# Veeam Backup & Replication

## Version 12

User Guide for VMware vSphere

July, 2025

© 2025 Veeam Software.

All rights reserved. All trademarks are the property of their respective owners.

No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language in any form by any means, without written permission from Veeam Software (Veeam). The information contained in this document represents the current view of Veeam on the issue discussed as of the date of publication and is subject to change without notice. Veeam shall not be liable for technical or editorial errors or omissions contained herein. Veeam makes no warranties, express or implied, in this document. Veeam may have patents, patent applications, trademark, copyright, or other intellectual property rights covering the subject matter of this document. All other trademarks mentioned herein are the property of their respective owners. Except as expressly provided in any written license agreement from Veeam, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

**NOTE**

Read the End User Software License Agreement before using the accompanying software programs. Using any part of the software indicates that you accept the terms of the End User Software License Agreement.

# Contents


CONTACTING VEEAM SOFTWARE ..... 18
ABOUT THIS DOCUMENT ..... 19
ABOUT VEEAM BACKUP & REPLICATION ..... 20
PLANNING AND PREPARATION ..... 23
Supported Platforms and Applications ..... 27
Deployment Scenarios ..... 39
Simple Deployment ..... 40
Advanced Deployment ..... 41
Distributed Deployment ..... 43
System Requirements ..... 44
Permissions ..... 74
Ports ..... 105
Naming Conventions ..... 184
SECURITY GUIDELINES ..... 185
General Security Considerations ..... 186
Securing Backup Infrastructure ..... 189
Security & Compliance Analyzer ..... 192
Malware Detection ..... 207
How Malware Detection Works ..... 208
Malware Detection Methods ..... 209
Configuring Malware Detection ..... 248
Viewing Malware Detection Events ..... 253
Managing Malware Status ..... 255
Communications Encryption ..... 258
Communications Encryption Standards ..... 259
Supported Cipher Suites and Protocols ..... 260
Trusted Certificates ..... 261
Data Encryption ..... 262
How Data Encryption Works ..... 263
How Data Decryption Works ..... 266
Multi-Factor Authentication ..... 274
Kerberos Authentication ..... 279
Using Group Managed Service Accounts ..... 283
LICENSING ..... 286
Types of Licenses ..... 290
Instance Consumption for Object Storage Backup, File Backup and File to Tape Jobs ..... 291
Obtaining and Renewing License ..... 295

Installing License ..... 297
Merging Licenses ..... 299
Viewing License Information ..... 301
Revoking License ..... 306
Removing License ..... 308
Exceeding License Limit ..... 311
License Expiration ..... 313
Updating License ..... 315
Updating License Manually ..... 316
Updating License Automatically ..... 318
Automatic License Usage Reporting ..... 321
DEPLOYMENT ..... 322
Installation ..... 323
Installing Veeam Backup & Replication ..... 324
Installing Veeam Backup & Replication Console ..... 346
Uninstalling Veeam Backup & Replication ..... 354
Upgrade and Update ..... 355
Upgrading to Veeam Backup & Replication 12.3 ..... 356
Updating Veeam Backup & Replication ..... 370
Veeam Backup & Replication Update Notifications ..... 376
Upgrading Veeam Backup & Replication Console ..... 378
Upgrading Infrastructure Components ..... 384
Silent Deployment ..... 386
Installing Veeam Backup & Replication in Silent Mode ..... 387
Installing Veeam Backup & Replication Console in Silent Mode ..... 396
Upgrading Veeam Backup & Replication in Silent Mode ..... 401
Upgrading Veeam Backup & Replication Console in Silent Mode ..... 406
Uninstalling Veeam Backup & Replication in Silent Mode ..... 410
Uninstalling Veeam Backup & Replication Console in Silent Mode ..... 413
GETTING STARTED WITH VEEAM BACKUP & REPLICATION ..... 416
Logging in to Veeam Backup & Replication ..... 417
Veeam Backup & Replication UI ..... 418
Main Menu ..... 419
Navigation Pane ..... 420
Ribbon and Tabs ..... 421
Views ..... 422
Working Area ..... 424
Job Filter ..... 425
Footer ..... 428
Changing Color Theme ..... 429

Infrastructure Icons ..... 430
Veeam Backup & Replication Services ..... 446
Veeam Installer Service ..... 449
Veeam Intelligence ..... 453
Resource Scheduling ..... 455
Limitation of Concurrent Tasks ..... 456
Limitation of Read and Write Data Rates for Backup Repositories ..... 459
Performance Bottlenecks ..... 460
Job Priorities ..... 462
CONFIGURING VEEAM BACKUP & REPLICATION ..... 464
Configuring General Settings ..... 465
Specifying I/O Settings ..... 466
Configuring Security Settings ..... 468
Specifying Email Notification Settings ..... 484
Specifying SNMP Settings ..... 494
Specifying Syslog Servers ..... 497
Specifying Other Notification Settings ..... 503
Specifying Session History Settings ..... 506
Managing Users and Roles ..... 507
Configuring Users ..... 508
Four-Eyes Authorization ..... 512
Managing Credentials ..... 516
Credentials Manager ..... 517
Cloud Credentials Manager ..... 525
Password Manager ..... 556
Key Management System Keys ..... 559
Managing Locations ..... 566
Creating and Assigning Locations to Infrastructure Objects ..... 569
Editing Locations ..... 571
Deleting Locations ..... 572
Exporting and Importing Locations List ..... 573
Managing Network Traffic ..... 574
Configuring Network Traffic Rules ..... 575
Managing Upload Streams ..... 582
Specifying Preferred Networks ..... 583
IPv6 Support ..... 585
Managing Logs ..... 586
Exporting Logs ..... 587
Configuring Global VM Exclusions ..... 593
Configuring Analytics View ..... 595

BACKUP INFRASTRUCTURE COMPONENTS ..... 598
Backup Server ..... 599
Backup & Replication Console ..... 600
Veeam Backup & Replication Configuration Database ..... 603
Managing Configuration Database ..... 604
Veeam Backup PowerShell Module ..... 650
Virtualization Servers and Hosts ..... 651
Adding VMware vSphere Servers ..... 653
Adding VMware Cloud Director Servers ..... 660
Adding Microsoft Windows Servers ..... 667
Adding Linux Servers ..... 675
Rescanning Servers ..... 684
Editing Server Settings ..... 685
Removing Servers ..... 686
General-Purpose Backup Proxies ..... 687
Adding General-Purpose Backup Proxies ..... 689
VMware Backup Proxies ..... 696
Requirements and Limitations for VMware Backup Proxies ..... 698
Transport Modes ..... 700
Adding VMware Backup Proxies ..... 715
Editing VMware Backup Proxy Settings ..... 724
Disabling and Removing VMware Backup Proxies ..... 725
VMware CDP Proxies ..... 727
Adding VMware CDP Proxies ..... 730
Cache Repositories ..... 738
Backup Repositories ..... 739
External Repositories ..... 740
Microsoft Windows Server ..... 772
Linux Server ..... 784
Hardened Repository ..... 796
SMB (CIFS) Share ..... 857
NFS Share ..... 868
Deduplicating Storage Appliances ..... 881
Backup Repositories with Rotated Drives ..... 922
Object Storage Repositories ..... 929
Managing Backup Repositories ..... 1091
Scale-Out Backup Repositories ..... 1109
Limitations for Scale-Out Backup Repositories ..... 1111
Immutability for Scale-Out Backup Repositories ..... 1115
Performance Tier ..... 1117

Capacity Tier ..... 1129
Archive Tier ..... 1157
Adding Scale-Out Backup Repositories ..... 1173
Managing Scale-Out Backup Repositories ..... 1188
Guest Interaction Proxies ..... 1215
Gateway Servers ..... 1218
Mount Servers ..... 1225
Veeam Data Mover Service ..... 1227
Veeam vPower NFS Service ..... 1229
WAN Accelerators ..... 1231
WAN Global Cache ..... 1233
WAN Accelerator Sizing ..... 1240
Adding WAN Accelerators ..... 1245
Removing WAN Accelerators ..... 1253
WAN Acceleration ..... 1254
Log Shipping Servers ..... 1258
Tape Servers ..... 1259
NDMP Servers ..... 1260
Veeam Backup Enterprise Manager ..... 1261
**BACKUP ..... 1262**
About Backup ..... 1263
How Backup Works ..... 1264
Backup Infrastructure for Backup ..... 1266
Backup Chain ..... 1269
Changed Block Tracking ..... 1302
Data Compression and Deduplication ..... 1304
Data Exclusion ..... 1308
VMware Tools Quiescence ..... 1317
Guest Processing ..... 1318
Microsoft SQL Server Log Backup ..... 1336
Oracle Log Backup ..... 1346
PostgreSQL WAL Files Backup ..... 1355
Backup Job Scheduling ..... 1363
Immutability for Backup Files ..... 1371
Health Check for Backup Files ..... 1372
Compact of Full Backup File ..... 1378
Backup Move ..... 1380
Resume on Disconnect ..... 1385
Snapshot Hunter ..... 1386
Creating Backup Jobs ..... 1388

Before You Begin ..... 1389
Step 1. Launch New Backup Job Wizard ..... 1390
Step 2. Specify Job Name and Description ..... 1391
Step 3. Select VMs to Back Up ..... 1392
Step 4. Exclude Objects from Backup Job ..... 1393
Step 5. Define VM Backup Order ..... 1396
Step 6. Specify Backup Storage Settings ..... 1397
Step 7. Configure Long-Term Retention ..... 1399
Step 8. Specify Advanced Backup Settings ..... 1401
Step 9. Specify Secondary Target ..... 1411
Step 10. Specify Guest Processing Settings ..... 1412
Step 11. Define Job Schedule ..... 1426
Step 12. Finish Working with Wizard ..... 1428
Encrypting Backup Jobs ..... 1429
Performing Active Full Backup ..... 1430
Quick Backup ..... 1432
Retention Policy for Quick Backups ..... 1433
Performing Quick Backup ..... 1434
Importing Backups Manually ..... 1435
Importing Encrypted Backups ..... 1437
Importing Transaction Logs ..... 1438
Importing Backup Files from Scale-Out Backup Repositories ..... 1439
Managing Backups ..... 1440
Viewing Backup Properties ..... 1441
Upgrading Backup Chain Formats ..... 1442
Moving Backups ..... 1444
Copying Backups ..... 1447
Exporting Backups ..... 1450
Detaching Backups from Jobs ..... 1458
Removing Backups from Configuration ..... 1459
Deleting Backups from Disk ..... 1460
Deleting Backups from Object Storage Repositories ..... 1462
Deleting Backups from Scale-Out Backup Repositories ..... 1463
Removing Missing Restore Points ..... 1464
Launching Background Retention ..... 1467
Disabling Background Retention ..... 1468
Managing Backup Jobs ..... 1469
Editing Job Settings ..... 1470
Cloning Jobs ..... 1472
Disabling and Deleting Jobs ..... 1473

Starting and Stopping Jobs ..... 1475
Starting and Stopping Transaction Log Backup Jobs ..... 1477
Retrying Jobs ..... 1479
Reconfiguring Jobs with Microsoft SQL Server VMs ..... 1481
Targeting Jobs to Another Repository ..... 1482
Reporting ..... 1483
Viewing Real-Time Statistics ..... 1484
Viewing History Statistics ..... 1487
Viewing Job Session Results ..... 1488
Viewing Job and Job Session Reports ..... 1490
**REPLICATION ..... 1492**
Considerations and Limitations ..... 1493
Backup Infrastructure for Replication ..... 1494
Replication Scenarios ..... 1496
How Replication Works ..... 1499
Replication Chain ..... 1500
Replica Seeding and Mapping ..... 1501
Replica from Backup ..... 1504
Creating Replication Jobs ..... 1507
Before You Begin ..... 1508
Step 1. Launch New Replication Job Wizard ..... 1509
Step 2. Specify Job Name and Description ..... 1510
Step 3. Select VMs to Replicate ..... 1512
Step 4. Specify Data Source ..... 1513
Step 5. Exclude Objects from Replication Job ..... 1514
Step 6. Specify VM Processing Order ..... 1517
Step 7. Specify Replica Destination ..... 1518
Step 8. Create Network Mapping Table ..... 1521
Step 9. Configure Re-IP Rules ..... 1522
Step 10. Specify Replication Job Settings ..... 1524
Step 11. Specify Advanced Replica Settings ..... 1525
Step 12. Specify Data Transfer Settings ..... 1531
Step 13. Define Seeding and Mapping Settings ..... 1533
Step 14. Specify Guest Processing Settings ..... 1535
Step 15. Define Job Schedule ..... 1547
Step 16. Finish Working with Wizard ..... 1549
Creating Replica Seeds ..... 1550
Managing Replicas ..... 1551
Viewing Replica Properties ..... 1552
Rescanning Replicas ..... 1553

Removing Replicas from Configuration ..... 1554
Deleting Replicas from Disk ..... 1555
Managing Replication Jobs ..... 1556
Retrying Replication Jobs ..... 1557
Editing Replication Jobs ..... 1558
Cloning Replication Jobs ..... 1559
Disabling and Deleting Replication Jobs ..... 1560
Failover and Failback for Replication ..... 1562
Failover Plans ..... 1564
Failover ..... 1574
Permanent Failover ..... 1581
Planned Failover ..... 1583
Failover Undo ..... 1589
Failback ..... 1591
Failback Commit ..... 1609
Failback Undo ..... 1611
**CONTINUOUS DATA PROTECTION (CDP) ..... 1613**
Backup Infrastructure for CDP ..... 1614
Considerations and Limitations ..... 1616
How CDP Works ..... 1618
CDP Replication Chain ..... 1622
Retention Policies ..... 1624
Guaranteed Delivery ..... 1627
Replica Seeding and Mapping ..... 1628
Installing I/O Filter ..... 1631
Step 1. Launch I/O Filter Management Wizard ..... 1632
Step 2. Select Clusters ..... 1633
Step 3. Apply Filter Settings ..... 1634
Step 4. Finish Working with Wizard ..... 1635
Updating and Uninstalling I/O Filter ..... 1636
Taking I/O Filter Ownership ..... 1638
Creating CDP Policies ..... 1639
Before You Begin ..... 1640
Step 1. Launch New CDP Policy Wizard ..... 1641
Step 2. Specify Policy Name and Advanced Settings ..... 1642
Step 3. Select VMs to Replicate ..... 1643
Step 4. Exclude Objects ..... 1644
Step 5. Specify VM Processing Order ..... 1646
Step 6. Select Replica Destination ..... 1647
Step 7. Configure Network Mapping ..... 1649

Step 8. Configure Re-IP Rules ..... 1650
Step 9. Configure Seeding and Mapping ..... 1652
Step 10. Specify Data Transfer and Replica Settings ..... 1654
Step 11. Specify Notification Settings ..... 1656
Step 12. Specify Replication Schedule ..... 1658
Step 13. Specify Guest Processing Settings ..... 1660
Step 14. Finish Working with Wizard ..... 1670
Creating Replica Seeds for CDP ..... 1671
Managing CDP Policies ..... 1672
Viewing Session Statistics and Results ..... 1673
Editing Policies ..... 1676
Cloning Policies ..... 1677
Disabling and Deleting Policies ..... 1679
Managing CDP Replicas ..... 1681
Viewing Replica Properties ..... 1682
Rescanning CDP Replicas ..... 1683
Removing Replicas from Configuration ..... 1684
Deleting Replicas from Disk ..... 1685
Failover and Failback for CDP ..... 1686
Failover ..... 1689
Failover Plans ..... 1695
Permanent Failover ..... 1696
Planned Failover ..... 1697
Failover Undo ..... 1703
Failback ..... 1704
Failback Commit ..... 1721
Failback Undo ..... 1723
**UNSTRUCTURED DATA BACKUP ..... 1724**
Backup Infrastructure for Unstructured Data Backup ..... 1725
Adding Unstructured Data Source ..... 1731
How Unstructured Data Backup Works ..... 1789
Data Structure in Backup, Archive and Secondary Repositories ..... 1790
Unstructured Data Backup Retention Scenarios ..... 1795
Unstructured Data Backups in Object Storage Repositories ..... 1802
Unstructured Data Backups in Immutable Repositories ..... 1805
Unstructured Data Backups in Scale-Out Repositories ..... 1809
Unstructured Data Backups in Deduplicating Storage Appliances ..... 1810
File Backup Integration with Storage Systems ..... 1812
Creating Backup Jobs for Protecting Unstructured Data ..... 1814
Creating File Backup Jobs ..... 1815

Creating Object Storage Backup Jobs ..... 1840
Managing Unstructured Data Backups ..... 1861
Viewing Unstructured Data Backup Properties ..... 1862
Copying Unstructured Data Backups ..... 1863
Starting New Backup Chain ..... 1865
Performing Health Check and Repair for Unstructured Data Backups ..... 1866
Detaching Unstructured Data Backups from Jobs ..... 1868
Deleting Unstructured Data Backups from Disk ..... 1869
Converting Backups from SMB or NFS Shares to NAS Filer Shares ..... 1870
Converting Backups from Non-Root to Root Shared Folders ..... 1872
Updating Source File Share Path for Backup Jobs with Secondary Target ..... 1873
Unstructured Data Recovery ..... 1875
File Share Data Recovery ..... 1876
Object Storage Data Recovery ..... 1911
**VEEAMZIP ..... 1937**
Creating VeeamZIP Backups ..... 1938
Encrypting VeeamZIP Backups ..... 1941
Managing and Restoring VeeamZIP Backups ..... 1942
**BACKUP COPY ..... 1943**
About Backup Copy ..... 1944
How Backup Copy Works ..... 1945
Backup Copy Architecture ..... 1947
Backup Copy Modes ..... 1950
Restore Point Selection ..... 1952
Transformation Processes ..... 1953
Backup Copy Window ..... 1954
Retention Policy for Backup Copy Jobs ..... 1955
Backup Copy Space Requirements ..... 1969
Health Check for Backup Files ..... 1970
Compact of Full Backup File ..... 1972
Backup Copy Job Mapping ..... 1973
Active Full Backup Copies ..... 1975
Automatic Job Retries ..... 1976
Creating Backup Copy Jobs for VMs and Physical Machines ..... 1978
Before You Begin ..... 1979
Step 1. Launch New Backup Copy Job Wizard ..... 1980
Step 2. Specify Job Name and Copy Mode ..... 1981
Step 3. Select Workloads to Process ..... 1982
Step 4. Exclude Objects from Backup Copy Job ..... 1984
Step 5. Define Processing Order ..... 1985

Step 6. Specify Target Repository and Retention Settings ..... 1986
Step 7. Map Backup File ..... 1988
Step 8. Specify Advanced Settings ..... 1990
Step 9. Specify Data Path Settings ..... 1997
Step 10. Define Backup Copy Window ..... 1998
Step 11. Finish Working with Wizard ..... 2000
Creating Backup Copy Jobs for HPE StoreOnce Repositories ..... 2001
Before You Begin ..... 2002
Step 1. Launch New Backup Copy Job Wizard ..... 2003
Step 2. Specify Job Name and Description ..... 2004
Step 3. Select Source and Target Repositories ..... 2005
Step 4. Specify Advanced Settings ..... 2007
Step 5. Define Backup Copy Window ..... 2011
Step 6. Finish Working with Wizard ..... 2012
Creating Backup Copy Jobs for Veeam Plug-ins ..... 2013
Encrypting Backup Copy Jobs ..... 2014
Linking Backup Jobs to Backup Copy Jobs ..... 2016
Managing Backups ..... 2017
Viewing Backup Properties ..... 2018
Upgrading Backup Chain Formats ..... 2020
Moving Backups ..... 2023
Copying Backups ..... 2026
Removing Backups ..... 2029
Removing Missing Restore Points ..... 2032
Managing Backup Copy Jobs ..... 2035
Editing Backup Copy Job Settings ..... 2036
Starting and Stopping Backup Copy Jobs ..... 2038
Creating Active Full Backups ..... 2041
Disabling and Deleting Backup Copy Jobs ..... 2042
Cloning Backup Copy Job ..... 2044
Reporting ..... 2046
**VM COPY ..... 2047**
Copying VMs ..... 2048
Before You Begin ..... 2049
Step 1. Launch VM Copy Job Wizard ..... 2050
Step 2. Specify Job Name and Description ..... 2051
Step 3. Select VMs to Copy ..... 2052
Step 4. Exclude Objects from VM Copy Job ..... 2053
Step 5. Specify Copy Destination ..... 2056
Step 6. Specify Guest Processing Settings ..... 2058

Step 7. Define Job Schedule ..... 2067
Step 8. Finish Working with Wizard ..... 2069
FILE COPY ..... 2070
Creating File Copy Jobs ..... 2071
Before You Begin ..... 2072
Step 1. Launch New File Copy Job Wizard ..... 2073
Step 2. Specify Job Name and Description ..... 2074
Step 3. Select Files and Folders to Be Copied ..... 2075
Step 4. Select Destination for Copying ..... 2076
Step 5. Define Job Schedule ..... 2077
Step 6. Finish Working with Wizard ..... 2079
Copying Files and Folders Manually ..... 2080
Managing Folders ..... 2081
Editing and Deleting Files ..... 2082
QUICK MIGRATION ..... 2083
Quick Migration Architecture ..... 2085
Migrating VMs ..... 2086
Before You Begin ..... 2087
Step 1. Launch Quick Migration Wizard ..... 2088
Step 2. Select VMs to Relocate ..... 2089
Step 3. Specify VM Destination ..... 2090
Step 4. Select Infrastructure Components for Data Transfer ..... 2092
Step 5. Finish Working with Wizard ..... 2094
Migrating First Class Disks (FCDs) ..... 2096
Before You Begin ..... 2097
Step 1. Launch FCD Quick Migration Wizard ..... 2098
Step 2. Select FCDs to Migrate ..... 2099
Step 3. Specify FCD Destination ..... 2100
Step 4. Finish Working with Wizard ..... 2101
RECOVERY VERIFICATION ..... 2102
SureBackup ..... 2103
How SureBackup Works ..... 2104
Backup Recovery Verification Tests ..... 2106
Application Group ..... 2111
Virtual Lab ..... 2121
SureBackup Job ..... 2144
XML Files with Machine Roles Description ..... 2169
Manual Recovery Verification ..... 2172
SureReplica ..... 2173
How SureReplica Works ..... 2174

Replica Recovery Verification Tests ... 2176
Application Group ... 2177
Virtual Lab Configuration ... 2178
SureBackup Job for VM Replicas ... 2184
ON-DEMAND SANDBOX ... 2187
On-Demand Sandbox for Storage Snapshots ... 2188
Mixed Scenarios ... 2190
Configuring On-Demand Sandbox ... 2191
DATA RECOVERY ... 2193
VM Recovery ... 2194
Instant Recovery to VMware vSphere ... 2195
Instant Recovery to Microsoft Hyper-V ... 2218
Instant Recovery to Nutanix AHV ... 2239
Entire VM Restore ... 2240
Staged Restore ... 2261
Restore to Amazon EC2 ... 2263
Restore to Microsoft Azure ... 2285
Restore to Google Compute Engine ... 2336
Restore to Nutanix AHV ... 2357
Restore to Proxmox VE ... 2358
Restore to OLVM and RHV ... 2359
Disk Recovery ... 2360
Instant Disk Recovery ... 2361
Instant First Class Disk (FCD) Recovery ... 2374
Virtual Disk Restore ... 2386
Disk Export ... 2396
Disk Publishing (Data Integration API) ... 2408
Disk Restore to oVirt KVM ... 2421
Item Recovery ... 2422
VM Files Restore ... 2423
Guest OS File Restore ... 2430
Application Item Restore ... 2478
Restoring Data from Encrypted Backups ... 2488
VMWARE CLOUD DIRECTOR SUPPORT ... 2489
Viewing VMware Cloud Director VMs ... 2490
Backup and Restore of vApps ... 2491
Backup for VMware Cloud Director ... 2492
Data to Back Up ... 2493
Cloud Director Backup Jobs ... 2495
Performing Backup of VMware Cloud Director VMs ... 2496

Managing Cloud Director Backups and Jobs ..... 2498
Creating VeeamZIP Files for VMware Cloud Director VMs ..... 2500
CDP for VMware Cloud Director ..... 2501
Backup Infrastructure for Cloud Director CDP ..... 2502
Requirements and Limitations ..... 2504
How Cloud Director CDP Works ..... 2505
Installing I/O Filter on VDCs ..... 2506
Updating and Uninstalling I/O Filter ..... 2511
Creating Cloud Director CDP Policies ..... 2513
Managing Cloud Director CDP Policies ..... 2540
Managing Cloud Director CDP Replicas ..... 2547
Failover and Failback for Cloud Director CDP ..... 2551
Replication for VMware Cloud Director ..... 2580
vApp Replica Seeding and Mapping ..... 2581
Network Mapping ..... 2583
Creating Cloud Director Replication Job ..... 2585
Creating vApp Replica Seeds ..... 2621
Managing Cloud Director Replication Jobs ..... 2622
Managing Cloud Director Replicas ..... 2624
Failover and Failback for Cloud Director ..... 2628
Data Recovery for VMware Cloud Director ..... 2657
VM Recovery ..... 2658
vApp Recovery ..... 2692
Item Recovery ..... 2706
**VMWARE CLOUD SUPPORT ..... 2707**
**TAPE DEVICES SUPPORT ..... 2708**
**STORAGE SYSTEM SNAPSHOT INTEGRATION ..... 2709**
**INTEGRATION WITH VEEAM BACKUP FOR AWS ..... 2710**
**INTEGRATION WITH VEEAM BACKUP FOR MICROSOFT AZURE ..... 2711**
**INTEGRATION WITH VEEAM BACKUP FOR MICROSOFT ENTRA ID ..... 2712**
**INTEGRATION WITH VEEAM BACKUP FOR GOOGLE CLOUD ..... 2713**
**INTEGRATION WITH VEEAM BACKUP FOR NUTANIX AHV ..... 2714**
**INTEGRATION WITH VEEAM BACKUP FOR ORACLE LINUX VIRTUALIZATION MANAGER AND RED HAT VIRTUALIZATION ..... 2715**
**INTEGRATION WITH VEEAM BACKUP FOR PROXMOX VE ..... 2716**
**INTEGRATION WITH KASTEN ..... 2717**
**VEEAM AGENT MANAGEMENT ..... 2718**
**VEEAM CLOUD CONNECT ..... 2719**
**ADVANCED VMWARE VSPHERE FEATURES ..... 2720**
VM Tags ..... 2721

Encrypted VMs ..... 2723
Storage Profiles ..... 2726
**VEEAM BACKUP & REPLICATION UTILITIES ..... 2727**
Extract Utility ..... 2728
Using Extract Utility in GUI ..... 2729
Using Extract Utility in Interactive Mode ..... 2730
Using Extract Utility from Command Line ..... 2731
Veeam VM Migrator Utility ..... 2735
Using Veeam VM Migrator Utility ..... 2736
Veeam Configuration Database Connection Utility ..... 2745
Using Veeam Configuration Database Connection Utility ..... 2746
Veeam Backup Validator ..... 2753
Using Veeam Backup Validator ..... 2754
Veeam Backup Configuration Tool ..... 2760
Using Veeam Backup Configuration Tool ..... 2761
**VEEAM BACKUP & REPLICATION EVENTS ..... 2764**

# Contacting Veeam Software

At Veeam Software we value feedback from our customers. It is important not only to help you quickly with your technical issues, but it is our mission to listen to your input and build products that incorporate your suggestions.

## Customer Support

Should you have a technical concern, suggestion or question, visit the Veeam Customer Support Portal to open a case, search our knowledge base, reference documentation, manage your license or obtain the latest product release.

## Company Contacts

For the most up-to-date information about company contacts and office locations, visit the Veeam Contacts Webpage.

## Online Support

If you have any questions about Veeam products, you can use the following resources:

- Full documentation set: veeam.com/documentation-guides-datasheets.html
- Veeam R&D Forums: forums.veeam.com

# About This Document

This user guide provides information about main features, installation and use of Veeam Backup & Replication in VMware vSphere environments. The document applies to version 12 and all subsequent versions until it is replaced with a new edition.

## Intended Audience

The user guide is intended for anyone who wants to use Veeam Backup & Replication. It is primarily aimed at VMware administrators, consultants, analysts and any other IT professionals using the product.

# About Veeam Backup & Replication

Veeam Backup & Replication is a comprehensive data protection and disaster recovery solution. With Veeam Backup & Replication, you can create image-level backups of virtual, physical and cloud machines and restore from them. Technology used in the product optimizes data transfer and resource consumption, which helps to minimize storage costs and the recovery time in case of a disaster.

Veeam Backup & Replication provides a centralized console for administering backup, restore and replication operations in all supported platforms (virtual, physical, cloud). Also, the console allows you to automate and schedule routine data protection operations and integrate with solutions for alerting and generating compliance reports.

This section contains an overview of Veeam Backup & Replication and solutions integrated with it.

## Main Features

Main functionality of Veeam Backup & Replication includes:

- Backup: creating image-level backups of virtual, physical, cloud machines and backups of file shares and object storage repositories.
- Restore: performing restore from backup files to the original or a new location.
  Veeam Backup & Replication offers a number of recovery options for various disaster recovery scenarios, including Instant Recovery, image-level restore, file-level restore, restore of application items and so on.
- Replication: creating an exact copy of a VM and maintaining the copy in sync with the original VM.
- Continuous Data Protection (CDP): replication technology that helps you protect mission-critical VMs and reach recovery point objective (RPO) up to seconds.
- Backup Copy: copying backup files to a secondary repository.
- Storage Systems Support: backing up and restoring VMs using capabilities of native snapshots created on storage systems. For more information, see the Storage System Snapshot Integration Guide.
- Tape Devices Support: storing copies of backups in tape devices. For more information, see the Tape Devices Support section in the Veeam Backup & Replication User Guide.
- Recovery Verification: testing VM backups and replicas before recovery.
- Scale-Out Backup Repositories: a repository system that allows you to distribute data between performance, capacity and archive tiers.

## Protected Objects

With Veeam Backup & Replication, you can back up and restore the following objects:

- **Virtual machines**:
  - VMware vSphere VMs
  - Microsoft Hyper-V VMs
  - Nutanix AHV VMs
  - oVirt VMs
  - Proxmox VE VMs

- **Cloud VMs**:
  - AWS EC2 instances
  - Microsoft Azure VMs
  - Google Cloud VMs
- Microsoft Entra ID tenant data
- Unstructured data:
  - File shares
  - Object storage repositories
- **Physical machines**. To back up machines running Windows, Linux or macOS operating systems, Veeam Backup & Replication uses backup agents installed on each computer. Veeam Backup & Replication operates as a centralized control center for deploying and managing Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for Mac, Veeam Agent for Oracle Solaris and Veeam Agent for IBM AIX. For details, see the Veeam Agent Management Guide.

# Protected Applications

Native functionality of Veeam Backup & Replication allows you to create application-consistent backups for:

- Microsoft SQL Server
- PostgreSQL
- Oracle Database
- Active Directory
- Microsoft Exchange
- Microsoft OneDrive
- Microsoft SharePoint

Also, you can install the following additional tools:

- Veeam Backup for Microsoft 365: for full protection of Microsoft applications.
- Veeam Plug-ins for Enterprise Applications: for integration of Veeam Backup & Replication with Oracle RMAN, SAP HANA Backint, and BR*Tools.

# Management and Reporting

Veeam Backup & Replication integrates with a set of solutions that provide reporting and management capabilities for enterprise environments:

- Veeam ONE: a solution that enables real-time monitoring, business documentation and management reporting for Veeam Backup & Replication, VMware vSphere and Microsoft Hyper-V.
- Veeam Backup Enterprise Manager: a management and reporting component that allows you to manage multiple Veeam Backup & Replication installations from a single web console.
- Management Pack for Veeam Backup & Replication: a component that integrates Veeam Backup & Replication infrastructure, services and jobs into Microsoft System Center Operations Manager.

- Veeam Recovery Orchestrator: a solution that orchestrates disaster recovery processes in VMware vSphere environments, supports one-click recovery for critical applications and sites, and provides features for documentation and testing.
- Veeam App for Splunk: a Splunk extension that allows you to monitor the health and security status of your Veeam backup infrastructure.
- Veeam App for Palo Alto Networks XSOAR: a Cortex XSOAR content pack that allows you to monitor various security activities in your Veeam backup infrastructure.

# Service Providers

If you are a service provider, you can use Veeam Service Provider Console to deliver Veeam-powered Backup-as-a-Service (BaaS) and Disaster Recovery-as-a-Service (DRaaS) services to your customers.

You can also use Veeam Backup & Replication to offer cloud repository as a service and disaster recovery as a service. For details, see Veeam Cloud Connect.

# Planning and Preparation

Infrastructure of Veeam Backup & Replication depends on the business needs and resources of your company. Before you install Veeam Backup & Replication, make sure that your backup infrastructure meet product hardware recommendations and system requirements. For more information, see these sections:

- Supported Platforms and Applications
- Deployment Scenarios
- System Requirements
- Permissions
- Ports
- Naming Conventions

Before you deploy Veeam Backup & Replication, consider the following tips and recommendations that may help you design your backup infrastructure:

1. Define protection scope.
2. Define RTO and RPO goals.
3. Select Veeam Backup & Replication features that you will need.
4. Plan how many copies of your data you need to store (3-2-1 rule).
5. Design Veeam Backup & Replication infrastructure.

## Step 1. Define Protection Scope

Define how many machines you need to protect and the amount of disk space the machines use.

After defining the protection scope, calculate how much of the total amount of data is actually changing on a daily basis. This information is required because of the mechanism of how Veeam Backup & Replication creates a backup chain. At the first run, Veeam Backup & Replication creates a full backup file; at the second and further runs, Veeam Backup & Replication creates an incremental backup file that contains only the blocks that has been changed since the last backup. As a result, the daily change rate has a significant impact on the backup window and the storage capacity needed to store the backups. As Veeam Backup & Replication creates image-level or block level backups, you need to know the daily change rate on the block level. For VMware vSphere or Microsoft Hyper-V, you can use Veeam ONE to measure and generate a report on the daily change rate of VMs.

As a result of this step, you can make a list of machines to be protected, including the data on which of the machines contain databases, which of the machines host business critical applications, and how much of the total amount of data is changing on these machines on a daily basis. This information will help you in further steps of deployment planning.

# Step 2. Define RPO and RTO

When you make a business continuity and disaster recovery plan, you must define two important parameters: Recovery Point Objective (RPO) and Recovery Time Objective (RTO).

- **RPO** defines a period during which you may accept to lose data. Basically, it is the age of the latest backup that will be used for recovery in case of a failure. It means that your company accepts that in case of a failure you may lose the data that has been accumulated since the latest restore point. RPO set by the policy of your company defines how often you need to create a recovery point. This will help you estimate how much storage you will need to store backups, how many copies of your data you need, and which Veeam Backup & Replication features are the most suitable for business needs of your company.
- **RTO** is related to downtime. RTO represents the amount of time from the beginning of an incident until all services are back online and available to users.

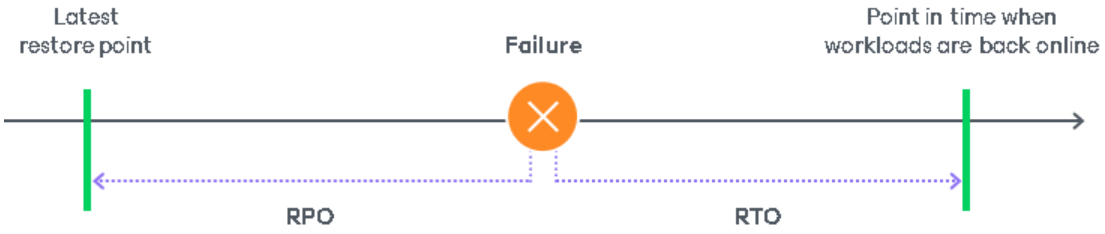


Define a list of your workloads grouped and organized by how fast they must be recoverable. Divide the list into categories. The higher the recovery priority, the lower the RTO will be required relative to the rest of your workloads.

# Step 3. Select Veeam Backup & Replication Features

Based on the analysis of your RTO and RPO, you can define your protection plan and select which features are the most suitable for your business needs. It is a common practice to divide servers and applications into categories and use different protection functionality for each category based on SLA (service level agreement). You can take the following table as a reference.

| | RPO: Seconds | RPO: Minutes | RPO: Hours (<24h) | RPO: Hours (24-48) |
|---|---|---|---|---|
| **RTO: Seconds** | Continuous Data Protection (for VMware vSphere) | Replication | | |
| **RTO: Minutes** | | Snapshot Orchestration (for VMware vSphere) | Backup | Backup Copy |
| **RTO: Hours** | | | | Tape Devices Support |

Apart from backup and replication options, the RTO also depends on the method of recovery and recovery verification. Veeam Backup & Replication offers a number of recovery options for various disaster recovery scenarios, including Instant Recovery, image-level restore, file-level restore, restore of application items and so on. For details, see the following sections:

- Restore: performing restore from backup files to the original or a new location.
- Recovery Verification: testing VM backups and replicas before recovery.

# Step 4. Plan How Many Copies of Data You Need (3-2-1 rule)

To build a successful data protection and disaster recovery plan, we recommend that you follow the 3-2-1 rule:

- 3: You must have at least three copies of your data: the original production data and two backups. The first copy of your data is the original production data. The second copy of your data is a backup created by a backup job. To create the third copy of data you can use Backup Copy or Backup to Tape jobs. Also, if you use cloud repositories (S3, Azure Blob, Google Cloud, IBM Cloud), you can copy backups to a capacity tier.
- 2: You must use at least two different types of media to store the copies of your data.
- 1: You must keep at least one backup off-site. For example, in the cloud or in a remote site. One of the repositories must be offline, air-gaped or immutable.

Veeam Backup & Replication provides integration with various types of repositories. Select where you want to store your backup files. For the full list of supported backup repositories, see Backup Repository, Tape Devices Support, Storage System Snapshot Integration Guide.

To plan the required space on repositories, you may also need to analyze for how long will you store the backups. Veeam Backup & Replication provides short-term and long-term (GFS) retention policies to effectively store the backup files.

# Step 5. Design Veeam Backup & Replication Infrastructure

Veeam Backup & Replication can be used in virtual environments of any size and complexity. The architecture of the solution supports on-site and off-site data protection, operations across remote sites and geographically dispersed locations. Veeam Backup & Replication provides flexible scalability and easily adapts to the needs of your virtual environment.

> **NOTE**
>
> Consider general security recommendations and recommendations for hardening specific backup infrastructure components when designing Veeam Backup & Replication infrastructure.

Before you install Veeam Backup & Replication, familiarize yourself with common deployment scenarios and plan your backup infrastructure layout. For details, see Deployment Scenarios.

The easiest way to start is to deploy a Veeam Backup & Replication server, one dedicated server for a backup proxy and one repository. While you keep adding backup jobs, add more proxies and repositories. Each backup infrastructure component has its own specifics and requirements that are described in the following sections of this guide:

- Veeam Backup & Replication Server
- VMware Backup Proxy
- Backup Repository

Also, note that in most cases, it is recommended to deploy Veeam Backup & Replication, Veeam Backup Enterprise Manager and Veeam ONE on separate servers.

# Supported Platforms and Applications

Veeam Backup & Replication supports the following virtualization platforms:

- VMware vSphere
- Azure Local (formerly Azure Stack HCI)
- Microsoft Hyper-V
- Proxmox VE
- Nutanix AHV
- Oracle Linux Virtualization Manager
- Red Hat Virtualization

> **NOTE**
>
> The information on this page is valid as of the date of the last page update.

# VMware vSphere Platform

## VMware vSphere Virtual Infrastructure

Veeam Backup & Replication supports the following VMware vSphere platforms.

| Specification | Requirement |
|---|---|
| **Platform** | • vSphere 8.0 (up to 8.0 U3)<br>• vSphere 7.x<br>• vSphere 6.x<br>• VMware Cloud Foundation (VCF)<br>This platform is supported as individual VMware software components. VMware components listed on this page can be part of VCF. For the information on the correspondence of VMware components to the VCF version, see this VMware KB article.<br>• Google Cloud VMware Engine (GCVE)<br>For more information on GCVE support, see this Veeam KB article.<br>• IBM Cloud for VMware Solutions<br>For more information on IBM Cloud for VMware Solutions support, see this Veeam KB article.<br>• Microsoft Azure VMware Solution (AVS)<br>For more information on AVS support, see this Veeam KB article.<br>• Oracle Cloud VMware Solution<br>For more information on Oracle Cloud VMware Solutions support, see this Veeam KB article.<br>• VMware Cloud on AWS<br>For more information on VMware Cloud on AWS support, see this Veeam KB article.<br>• VMware Cloud on Dell |
| **Hypervisor** | • ESXi 8.0 (up to 8.0 U3)<br>• ESXi 7.x<br>• ESXi 6.x<br>Free ESXi is not supported. Veeam Backup & Replication leverages vSphere and vStorage APIs that are disabled by VMware in free ESXi. |
| **Management Server (optional)** | • vCenter Server 8.0 (optional) (up to 8.0 U3)<br>• vCenter Server 7.x (optional)<br>• vCenter Server 6.x (optional) |

## Veeam Continuous Data Protection (CDP)

The following infrastructure requirements apply only when you protect VMs with Continuous Data Protection (CDP):

- VMware vSphere edition must be VMware vSphere Essentials Kits editions or higher.
- vCenter Server is required. Standalone ESXi hosts are not supported.
- Minimum ESXi version required is 6.5 U2.
- Minimum 16GB RAM is required for source and target ESXi hosts.
- Continuous Data Protection for VMware Cloud Director supports target ESXi hosts version 7.0 or later. Target hosts of version 6.5 and 6.7 are not supported.
- All hosts in a cluster must be of the same major version. For example, 7.x or 6.x (6.5, or 6.7, or a combination of 6.5 and 6.7 is supported). In turn, all clusters managed by the same vCenter Server must also be of the same major version.
- You can replicate data between vCenter Servers of different versions. However, if you replicate data within one vCenter Server, and this vCenter Server includes hosts of different versions (6.5 and 6.7), your CDP policy may fail. In this case, go to the vSphere Client and delete the *Veeam CDP Replication* VM storage policy applied to each VM included in the CDP policy. For more information on how to delete a VM storage policy, see VMware Docs. Then, in the Veeam Backup & Replication console, relaunch the CDP policy.
- The maximum number of disks per ESXi host is 500.
- Backup server, VMware CDP proxies, vCenter Server and ESXi hosts must be able to resolve each other DNS names.
- VMware Cloud on AWS is not supported.

For more information on Veeam CDP, its requirements and limitations, see Continuous Data Protection (CDP).

## VMware vSphere VMs

| Specification | Requirement |
|---|---|
| **Virtual Hardware** | • All types and versions of virtual hardware are supported, including 62 TB VMDK.<br>• Virtual machines with virtual NVDIMM devices, with virtual disks engaged in SCSI bus sharing or residing on PMem datastores are not supported for VM backup, because VMware does not support snapshotting such VMs. To protect such VMs, use Veeam Agent backup.<br>• RDM virtual disks in physical mode, independent disks, and disks connected through in-guest iSCSI initiator are not supported for VM backup, and are skipped from processing automatically. If backup of these disks is required, use Veeam Agent backup.<br><br>Network shares and mount points targeted to 3rd party storage devices are also skipped as these volumes/disks are not visible in the VM configuration file.<br><br>RDM virtual disks in virtual mode are supported to create backups based on VMware Changed Block Tracking technology, although there are some restrictions on the virtual disk restore operation. To learn more about them, see Restoring Virtual Disks. |
| **OS** | • All operating systems supported by VMware.<br>• Guest processing (which includes application-aware processing and indexing) is supported for Microsoft Windows 2008/Windows Vista or later except Nano Server, due to the absence of VSS framework.<br>• Guest processing (which includes application-aware processing and indexing) is supported for 32- bit and 64-bit versions of the following Linux operating systems:<br>  ○ CentOS 7.x<br>  ○ Debian 10.0 to 12.2<br>  ○ Oracle Linux 7 (UEK3) to 9 (UEK R7)<br>  ○ Oracle Linux 7 to 9 (RHCK)<br>  ○ RHEL 7.0 to 9.5<br>  ○ SLES 12 SP4 or later, 15 SP1 or later<br>  ○ Ubuntu: 18.04 LTS, 20.04 LTS, and 22.04 LTS<br><br>Note that persistent agent components can be installed only on 64-bit versions of these Linux operating systems. |

| Specification | Requirement |
| --- | --- |
| **Software** | • For Linux operating systems, GLIBC 2.12 or later. GLIBC is used for the following operations: application-aware processing, file-level restore to Linux guest OS, installing persistent agent components, and adding VM as a managed server.<br>• VMware Tools (optional, recommended). VMware Tools are required for the following operations: application-aware processing, file-level restore from Microsoft Windows guest OS, and SureBackup testing functions.<br>• Open VM Tools (OVT, optional). Open VM Tools are a set of services and modules used by VMware for interaction with VMs running Linux or other VMware supported Unix-like guest operating systems.<br>• All latest OS service packs and patches (required for application-aware processing). |

## Guest OS File Restore

File-level restore is supported for the following file systems, including Microsoft Windows Logical Disk Manager (LDM) dynamic disks and Linux Logical Volume Manager (LVM):

| OS | Supported File Systems |
| --- | --- |
| **Microsoft Windows** | • FAT, FAT32<br>• NTFS<br>• ReFS (ReFS is supported only if Veeam Backup & Replication is installed on Microsoft Windows Server 2012 or later)<br><br>Windows file-level restore to original location is supported for Microsoft Windows 2008/Windows Vista or later except Nano Server. |
| **Linux** | • ext2, ext3, ext4<br>• ReiserFS<br>• JFS<br>• XFS<br>• Btrfs<br>• NTFS<br><br>DRBD (Distributed Replicated Block Devices) are not supported. |
| **BSD** | UFS, UFS2 |
| **Mac** | HFS, HFS+ (volumes up to 2 TB) |

| OS | Supported File Systems |
| --- | --- |
| **OpenText OES (formerly Micro Focus OES)** | NSS<br><br>AD-enabled NSS volumes on Open Enterprise Server 2015 are supported. Besides restore of standard file and folder permissions, restore of NSS trustee rights on files and folders is supported.<br><br>File-level restore is supported for the following OSes:<br>• Open Enterprise Server (OES) versions 2015(SP1), 2018(SP1-SP3). Versions 2023 and 24.1 are supported with limitations. For more information, see this Veeam KB article.<br>• SUSE Linux Enterprise Server 11 SP1-SP4 with OES11(SP1-SP3)<br>• NetWare 6.5 (only restore to a new location is supported) |
| **Solaris** | • UFS<br>• ZFS (except any pool versions of Oracle Solaris)<br><br>The helper appliance uses module ZFSonLinux version 2.1.5. For this reason, Veeam Backup & Replication supports only those versions of pools and features that are available in ZFSonLinux version 2.1.5. |

For other requirements and limitations of guest OS file restore, see Requirements and Limitations.

### VMware Cloud Director

| Specification | Requirement |
| --- | --- |
| **VMware Cloud Director** | VMware Cloud Director 10.1 to 10.6<br><br>Note that if your backup server is based on Microsoft Windows Server 2012 R2 or an earlier version, VMware Cloud Director servers cannot be added or used. |

## Hyper-V Platform

Veeam Backup & Replication provides support for the Microsoft Hyper-V platform including Azure Local (formerly Azure Stack HCI). For more information on supported versions and VM requirements, see the Veeam Backup & Replication User Guide for Microsoft Hyper-V.

## Proxmox VE Platform

Veeam Backup for Proxmox VE is a solution developed for protection and disaster recovery tasks for Proxmox Virtual Environment (Proxmox VE). For more information on Veeam Backup for Proxmox VE features and system requirements, see the Veeam Backup for Proxmox VE User Guide.

# Nutanix AHV Platform

Veeam Backup for Nutanix AHV is a solution developed for protection and disaster recovery tasks for the Nutanix AHV environment. For more information on Veeam Backup for Nutanix AHV features and system requirements, see the Veeam Backup for Nutanix AHV User Guide.

# Oracle Linux Virtualization Manager and Red Hat Virtualization Platforms

Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization is a solution developed for protection and disaster recovery tasks for oVirt KVM environments. For more information on Veeam Backup for OLVM and RHV features and system requirements, see the Veeam Backup for OLVM and RHV User Guide.

# Supported Applications

You can create transactionally-consistent backups or replicas of VMs that run the following applications.

| Application | Requirement |
|---|---|
| **Microsoft Active Directory** | Veeam Backup & Replication supports backup of domain controllers for the following operating systems:<br>• Microsoft Windows Server 2025<br>• Microsoft Windows Server 2022<br>• Microsoft Windows Server 2019<br>• Microsoft Windows Server 2016<br>• Microsoft Windows Server 2012 R2<br>• Microsoft Windows Server 2012<br>• Microsoft Windows Server 2008 R2<br>• Microsoft Windows Server 2008<br>Minimum supported domain and forest functional level is Windows 2008. |
| **Microsoft Exchange** | Veeam Backup & Replication supports backup of the following Microsoft Exchange versions:<br>• Microsoft Exchange 2019<br>• Microsoft Exchange 2016<br>• Microsoft Exchange 2013 |

| Application | Requirement |
| --- | --- |
| **Microsoft SharePoint** | Veeam Backup & Replication supports backup of the following versions of Microsoft SharePoint Server (virtualized either on VMware or Hyper-V platform):<br>• Microsoft SharePoint Server Subscription Edition<br>• Microsoft SharePoint 2019<br>• Microsoft SharePoint 2016<br>• Microsoft SharePoint 2013<br><br>All editions are supported (Subscription, Foundation, Standard, Enterprise).<br><br>Restore of Microsoft SharePoint data may require an available staging Microsoft SQL Server. To learn how to configure this server, see the Staging SQL Server Settings section of the Veeam Explorers User Guide. |
| **Microsoft SQL Server** | Veeam Backup & Replication supports backup of the following Microsoft SQL Server versions:<br>• Microsoft SQL Server 2022 (only for Windows)<br>• Microsoft SQL Server 2019 (only for Windows)<br>• Microsoft SQL Server 2017 (only for Windows)<br>• Microsoft SQL Server 2016 SP2<br>• Microsoft SQL Server 2014 SP3<br>• Microsoft SQL Server 2012 SP4<br>• Microsoft SQL Server 2008 R2 SP3<br>• Microsoft SQL Server 2008 SP4<br><br>All editions of Microsoft SQL Server except LocalDB are supported.<br><br>The database whose logs you want to back up must use the *Full* or *Bulk-logged* recovery model. In this case, all changes of the Microsoft SQL Server state will be written to transaction logs, and you will be able to replay transaction logs to restore the Microsoft SQL Server. You can use the Microsoft SQL Server Management Studio to switch to one of these models. For more information, see Microsoft Docs.<br><br>Restore of Microsoft SQL Server data may require an available staging Microsoft SQL Server. To learn how to configure this server, see the Configuring Staging SQL Server section of the Veeam Explorers User Guide. |

| Application | Requirement |
| --- | --- |
| **Oracle on Windows OS** | Veeam Backup & Replication supports backup of Oracle Database running on the following Microsoft Windows operating systems:<br>• Microsoft Windows Server 2025<br>• Microsoft Windows Server 2022<br>• Microsoft Windows Server 2019<br>• Microsoft Windows Server 2016<br>• Microsoft Windows Server 2012 R2<br>• Microsoft Windows Server 2012<br>• Microsoft Windows Server 2008 R2<br>• Microsoft Windows Server 2008<br>Veeam Backup & Replication supports backup of the following Oracle Database versions on Microsoft Windows machines:<br>• Oracle Database 21c<br>• Oracle Database 19c<br>• Oracle Database 18c<br>• Oracle Database 12c Release 2<br>• Oracle Database 12c Release 1<br>• Oracle Database 11g Release 2<br>Starting from version 12.3 (build 12.3.0.310), Veeam Backup & Replication supports the backup of Windows-based VMs with several Oracle homes of different versions. |
| **Oracle on Linux OS** | Veeam Backup & Replication supports backup of Oracle Database running on the following Linux distributions:<br>• CentOS 7.x<br>• RHEL 7.0 to 9.5<br>• SLES 12 SP4 or later, 15 SP1 or later<br>• Oracle Linux 7 (UEK3) to 9 (UEK R7)<br>• Oracle Linux 7 to 9 (RHCK)<br>Veeam Backup & Replication supports backup of the following Oracle Database versions on Linux machines:<br>• Oracle Database 21c<br>• Oracle Database 19c<br>• Oracle Database 18c<br>• Oracle Database 12c Release 2<br>• Oracle Database 12c Release 1<br>• Oracle Database 11g Release 2<br>Starting from version 12.1 (build 12.1.0.2131), Veeam Backup & Replication supports the backup of Linux-based VMs with several Oracle homes of different versions. |

| Application | Requirement |
| --- | --- |
| **Oracle Database configuration** | Consider the following:<br>• Automatic Storage Management (ASM) is supported for Oracle 11g and later; requires *ASMlib* present.<br>• Oracle Real Application Clusters (RAC) are not supported within the image-level backup functionality. Use Veeam Plug-in for Oracle RMAN. For details, see the Veeam Plug-in for Oracle RMAN section of the Veeam Plug-ins for Enterprise Applications User Guide.<br>• Oracle Database Express Edition (XE) is supported for Windows-based machines only.<br>• [For Microsoft Windows-based Oracle machines] Configurations with different versions of Oracle Database deployed on the same server are not supported.<br>• To create Oracle database backups, all Oracle servers that use Data Guard must be added to the backup job.<br>• You can use Veeam Plug-in for Oracle RMAN to integrate RMAN with Veeam Backup & Replication repositories. For details, see the Veeam Plug-in for Oracle RMAN section of the Veeam Plug-ins for Enterprise Applications User Guide. |
| **PostgreSQL** | Veeam Backup & Replication supports backup of the following PostgreSQL versions on Linux machines:<br>• PostgreSQL 17<br>• PostgreSQL 16<br>• PostgreSQL 15<br>• PostgreSQL 14<br>• PostgreSQL 13<br>• PostgreSQL 12<br><br>Veeam Backup & Replication does not support backup of PostgreSQL clusters. |

**NOTE**

Consider that 32-bit Oracle running on 64-bit operating systems and Oracle Express Edition (XE) on Linux are not supported.

# File Servers

Veeam Backup & Replication supports backup of files and folders from the file servers managed by the following operating systems:

- 64-bit versions of the Microsoft Windows operating systems:
  - Microsoft Windows Server 2025
  - Microsoft Windows Server 2022

    - Microsoft Windows Server 2019
    - Microsoft Windows Server 2016
    - Microsoft Windows Server 2012 R2
    - Microsoft Windows Server 2012
    - Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)
    - Microsoft Windows 10 (versions 1909 to 22H2)
    - Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021)
- 64-bit versions of the Linux operating systems:
    - CentOS 7.x
    - Debian 10.0 to 12.2
    - Oracle Linux 7 (UEK3) to 9 (UEK R7)
    - Oracle Linux 7 to 9 (RHCK)
    - RHEL 7.0 to 9.3
    - SLES 12 SP4 or later, 15 SP1 or later
    - Ubuntu: 18.04 LTS, 20.04 LTS, and 22.04 LTS

# Unstructured Data Backup Support

Veeam Backup & Replication supports backup of data from the following sources:

- Object storage repositories:
    - S3 compatible object storage repositories.
    - Amazon S3 object storage.
    - Microsoft Azure Blob storage including Azure Data Lake Gen2 storage (HNS supported).
- Managed Microsoft Windows or Linux server
- Enterprise NAS systems: NetApp Data ONTAP (including Amazon FSx for NetApp ONTAP), Lenovo ThinkSystem DM/DG Series, Dell PowerScale (formerly Isilon), Nutanix Files Storage.
- NFS file shares.
- SMB file shares.

## Requirements and Limitations for File Shares

Consider the following requirements and limitations for file shares:

- Only 64-bit versions of operating systems are supported for managed Microsoft Windows or Linux server file share.
- Backup of file shares on Linux hosts added with single use credentials is not supported.
- NFS file share must run NFS protocol version 3 or 4.1. Parallel NFS (pNFS) is not supported.

- Network shares and files on them targeted to 3rd party storage devices may have difficulties being restored, or may not be restored at all. Such shares/files often rely upon specific software/OS filters to be recalled from the alternate storage location, which is not available when performing a file share data recovery. See your software vendor documentation to learn how to back up such files.
- Anonymous or AD/Kerberos authentication is not supported for access to file shares through NFS.

  In NFS settings of the source file share, you must explicitly specify what servers will have access to the file share.
- SMB file share must run on SMB version 1.x, 2.x or 3.x.
- To support the **VSS for SMB File Shares** feature, make sure that requirements listed in this Veeam KB article are met.
- To correctly back up SACL (Ownership) files and folders from the SMB file share and restore them:
  a. When you are specifying access settings for the SMB file share, select the **This share requires access credentials** check box.
  b. Make sure that the account you use to Unstructured Data Backup access the file share is either added to the **Backup Operators** group or has the **SeBackupPrivilege** and **SeRestorePrivilege** privileges in Windows Server on the file share.

### Requirements and Limitations for Object Storage

Consider the following limitations for object storage added as a source of unstructured data:

- Veeam Backup & Replication does not support backup from and restore to the following types of Azure Blob Storage: Azure Data Box Storage, Azure Stack Edge.
- Veeam Backup & Replication does not support backup from and restore to AWS Snowball Edge Storage.
- Veeam Backup & Replication does not support backup from and restore to Veeam Data Cloud Vault.

For more information on unstructured data backup, see the Unstructured Data Backup section.

## Network

Consider the following requirements and limitations:

- Names of port groups/segments/networks must be unique.
- VMware NSX-T 2.3 or later is supported with N-VDS for VMware vSphere and VMware Cloud on AWS/Dell.
- VMware NSX-T 3.0 or later is supported with VDS for VMware vSphere and VMware Cloud on AWS/Dell.
- VMware NSX-V is supported (see details on the support of vSphere and VMware Cloud version in VMware vSphere Virtual Infrastructure).

# Deployment Scenarios

Veeam Backup & Replication can be used in virtual environments of any size and complexity. The architecture of the solution supports on-site and off-site data protection, operations across remote sites and geographically dispersed locations. Veeam Backup & Replication provides flexible scalability and easily adapts to the needs of your virtual environment.

Before you install Veeam Backup & Replication, familiarize yourself with common deployment scenarios and carefully plan your backup infrastructure layout.

# Simple Deployment

The simple deployment scenario is good for small virtual environments or the purpose of the Veeam Backup & Replication evaluation. In this scenario, Veeam Backup & Replication and all services needed for data protection tasks are installed on a single Windows-based machine.

> **NOTE**
>
> If you decide to use the simple deployment scenario, it is recommended that you install Veeam Backup & Replication on a VM. This will enable you to use the Virtual appliance transport mode and, as a result, LAN-free data transfer. For details, see Transport Modes.

The machine where Veeam Backup & Replication is installed performs the following roles:

- Backup server that coordinates all jobs, controls their scheduling and performs other administrative activities.
- Default VMware backup proxy that handles job processing and transfers backup traffic.
- Default backup repository where backup files are stored. During installation, Veeam Backup & Replication checks volumes of the machine on which you install the product and identifies a volume with the greatest amount of free disk space. On this volume, Veeam Backup & Replication creates the `Backup` folder that is used as the default backup repository.
- Mount server that is needed for restoring of VM guest OS files.
- Guest interaction proxy that is needed for application-aware processing, guest file system indexing and transaction log processing.

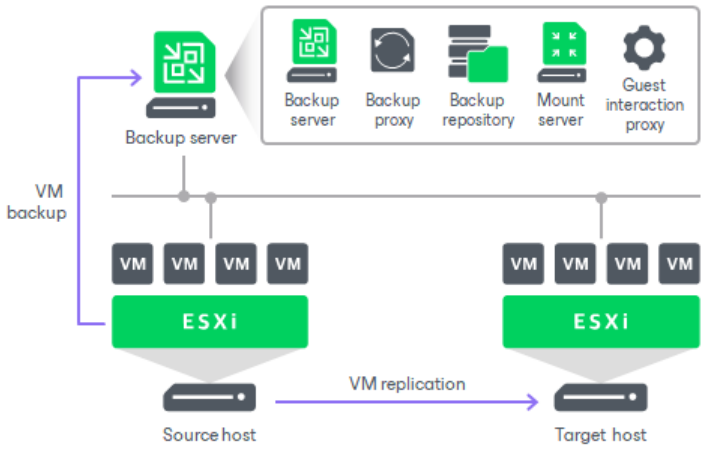


Veeam Backup & Replication is ready for use right after the installation. The only thing you must do is add VMware vSphere servers that you plan to use as source and target for backup, replication and other activities. For details, see Adding VMware vSphere Servers.

The drawback of a simple deployment scenario is that all data is handled and stored on the backup server. For medium-sized or large-scale environments, the capacity of a single backup server may not be enough. To take the load off the backup server and balance it throughout your backup infrastructure, we recommend that you use the advanced deployment scenario. For details, see Advanced Deployment.

# Advanced Deployment

In large-scale virtual environments with a large number of jobs, the load on the backup server is heavy. In this case, it is recommended that you use the advanced deployment scenario that moves the backup workload to dedicated backup infrastructure components. The backup server here functions as a "manager" for deploying and maintaining backup infrastructure components.

The advanced deployment includes the following components:

- Virtual infrastructure servers — VMware vSphere hosts used as source and target for backup, replication and VM copy.
- Backup server — a configuration and control center of the backup infrastructure.
- VMware backup proxy — a "data mover" component used to retrieve VM data from the source datastore, process it and deliver to the target.
- Backup repository — a location used to store backup files, VM copies and auxiliary replica files.
- Dedicated mount servers — component required for VM guest OS files and application items restore to the original location.
- Dedicated guest interaction proxies — components used to deploy the non-persistent runtime components or persistent agent components in Microsoft Windows VMs.

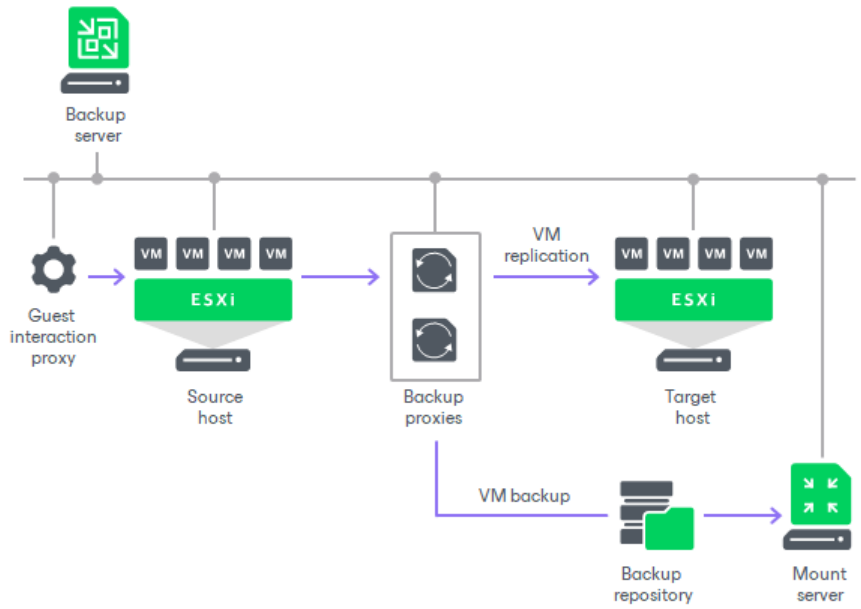


With the advanced deployment scenario, you can easily meet your current and future data protection requirements. You can expand your backup infrastructure horizontally in a matter of minutes to match the amount of data you want to process and available network throughput. Instead of growing the number of backup servers or constantly tuning job scheduling, you can install multiple backup infrastructure components and distribute the backup workload among them. The installation process is fully automated, which simplifies deployment and maintenance of the backup infrastructure in your virtual environment.

In virtual environments with several proxies, Veeam Backup & Replication dynamically distributes backup traffic among these proxies. A job can be explicitly mapped to a specific proxy. Alternatively, you can let Veeam Backup & Replication choose the most suitable proxy. In this case, Veeam Backup & Replication will check settings of available proxies and select the most appropriate one for the job. The proxy server to be used should have access to the source and target hosts as well as to the backup repository to which files will be written.

The advanced deployment scenario can be a good choice for backing up and replicating off-site. You can deploy a VMware backup proxy in the production site and another one in the disaster recovery (DR) site, closer to the backup repository. When a job is performed, backup proxies on both sides establish a stable connection, so this architecture also allows for efficient transport of data over a slow network connection or WAN.

To regulate backup load, you can specify the maximum number of concurrent tasks per proxy and set up throttling rules to limit proxy bandwidth. The maximum number of concurrent tasks can also be specified for a backup repository in addition to the value of the combined data rate for it.

Another advantage of the advanced deployment scenario is that it contributes to high availability — jobs can migrate between proxies if one of them becomes overloaded or unavailable.

# Distributed Deployment

The distributed deployment scenario is recommended for large geographically dispersed virtual environments with multiple Veeam Backup & Replication servers installed across different sites. These backup servers are federated under Veeam Backup Enterprise Manager — an optional component that provides centralized management and reporting for these servers through a web interface.

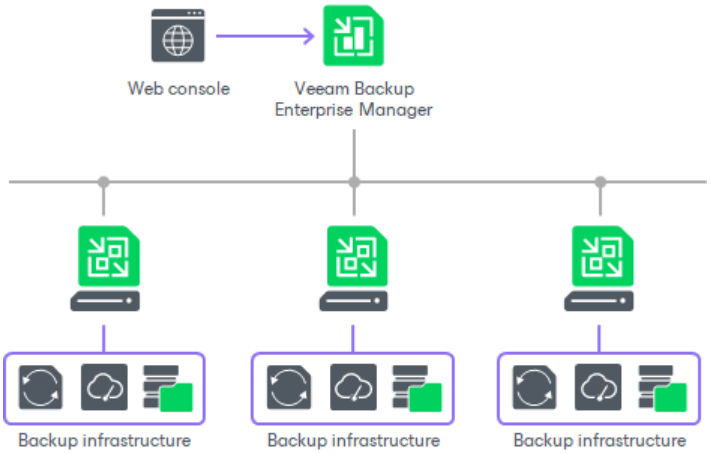


Veeam Backup Enterprise Manager collects data from backup servers and enables you to run backup and replication jobs across the entire backup infrastructure through a single web console, edit them and clone jobs using a single job as a template. It also provides reporting data for various areas (for example, all jobs performed within the last 24 hours or 7 days, all VMs engaged in these jobs and so on). Using indexing data consolidated on one server, Veeam Backup Enterprise Manager provides advanced capabilities to search for VM guest OS files in VM backups created on all backup servers (even if they are stored in backup repositories on different sites), and recover them in a single click. Search for VM guest OS files is enabled through Veeam Backup Enterprise Manager itself.

With flexible delegation options and security roles, IT administrators can delegate the necessary file restore or VM restore rights to authorized personnel in the organization – for example, allow database administrators to restore Oracle or Microsoft SQL Server VMs.

If you use Veeam Backup Enterprise Manager in your backup infrastructure, you do not need to install licenses on every backup server you deploy. Instead, you can install one license on the Veeam Backup Enterprise Manager server and it will be applied to all servers across your backup infrastructure. This approach simplifies tracking license usage and license updates across multiple backup servers.

In addition, VMware administrators will benefit from Veeam plug-in for vSphere Web Client that can be installed using Veeam Backup Enterprise Manager. They can analyze cumulative information on used and available storage space view and statistics on processed VMs, review success, warning, failure counts for all jobs, easily identify unprotected VMs and perform capacity planning for repositories, all directly from vSphere.

# System Requirements

Make sure that servers that you plan to use as backup infrastructure components meet the listed system requirements.

> **NOTE**
>
> The information on this page is valid as of the date of the last page update.

## Limitations and Recommendations

### Coexistence with Mission-Critical Production Servers

We do not recommend you to install Veeam Backup & Replication and its components on mission-critical machines in the production environment such as vCenter Server, Domain Controller, Microsoft Exchange Server, Small Business Server/ Windows Server Essentials and so on. If possible, install Veeam Backup & Replication and its components on dedicated machines. Backup infrastructure component roles can be co-installed.

### Microsoft Windows Server Core

You can assign roles of a backup proxy, backup repository, WAN accelerator, Veeam Cloud Connect infrastructure components and tape infrastructure components to machines running Microsoft Windows Server Core.
Installing Veeam Backup & Replication and Veeam Backup Enterprise Manager on a Windows OS machine without Desktop Experience (Core) is not supported.

### Windows Server IoT/Windows Storage Server Support

For information about support of Windows Server IoT/Windows Storage Server, see this Veeam KB article.

### Domain Member

The machine on which you plan to install Veeam Backup & Replication does not necessarily need to be a domain member. However, if you plan to restore Microsoft Exchange items from the Veeam Backup Enterprise Manager UI, you must install Veeam Backup Enterprise Manager on the domain member server from the Microsoft Active Directory forest in which Microsoft Exchange mailboxes are located.

### All-in-One Installations

For all-in-one installations, you can subtract 2 GB of memory resources from each but one role. These 2 GB are allotted to the OS itself, assuming each component is installed on the dedicated server.

### Unstructured Data Backup

Each of the following components for unstructured data backup may consume up to 4 GB RAM per task (in case of deduplicating storage appliances, up to 8 GB RAM): backup repository, general-purpose backup proxy, cache repository. Make sure you allocate enough memory resources for your installation. For all-in-one installations, where the server performs several roles, it must have enough memory resources for all components.

### Sharing Backup Infrastructure Components Across Veeam Installations

Using Linux-based shared backup infrastructure components across different Veeam installations is not supported.

As for non-Linux backup infrastructure components, we do not recommend using them shared across different Veeam installations due to several reasons:

- Veeam installations compete for resources.
- Backup components cannot simultaneously interact with Veeam Backup & Replication of different versions.
- Adding the same repository to different Veeam Backup & Replication installations may lead to corrupted backup and data in the database.

## Backup Server

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: x86-64 processor (minimum 4 cores recommended).<br><br>*Memory*: 8 GB RAM plus 500 MB RAM for each concurrent job. Memory consumption varies according to the number of VMs in the job, size of VM metadata, size of production infrastructure, and so on.<br><br>[For users with tape installations] For system requirements for large number of files in the file backup to tape job, see the Before You Begin section for file backup to tape.<br><br>*Disk Space*:<br>• 5 GB[1] for product installation and 4.5 GB for Microsoft .NET Framework 4.7.2 installation.<br>• 10 GB per 100 VM for guest file system catalog folder (persistent data).<br>• Additional free disk space for Instant Recovery cache folder (non-persistent data, at least 100 GB recommended).<br>• At least 10 GB for storing logs, although the disk space required for logging depends on the set of features used and may significantly increase. For details on logs locations, see Logging.<br>• When upgrading Veeam Backup & Replication, additional disk space is required. For details, see the Upgrade Checklist section.<br>• If you use a remote repository as a target for a configuration backup job, a temporary BCO file is created in the %TEMP% folder during the job (by default, this is C:\Windows\Temp). Therefore, additional free space is required for the temporary BCO file. For more information, see Creating Configuration Backups.<br><br>*Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported.<br><br>[1] Here and throughout this document GB is considered as 2^30 bytes, TB as 2^40 bytes. |

| Specification | Requirement |
| --- | --- |
| OS | 64-bit versions of the following Microsoft Windows operating systems are supported[1]:<br>• Microsoft Windows Server 2025<br>• Microsoft Windows Server 2022<br>• Microsoft Windows Server 2019<br>• Microsoft Windows Server 2016<br>• Microsoft Windows Server 2012 R2[2, 3]<br>• Microsoft Windows Server 2012[2, 3]<br>• Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>• Microsoft Windows 10 (versions 1909 to 22H2)<br>• Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021)<br>[1] Running Veeam backup server or any of Veeam backup infrastructure components on Insider versions of Microsoft Windows OS (both Client and Server) is not supported.<br>[2] VMware Cloud Director servers cannot be added or used in this version of Microsoft Windows OS.<br>[3] There may be issues when using Veeam Backup for Nutanix AHV with backup server running on this version of Microsoft Windows OS. For more information, see this Veeam KB article. |

| Specification | Requirement |
| --- | --- |
| **Configuration Database** | Local or remote installation of the following versions of PostgreSQL[1]:<br>• PostgreSQL 14.x<br>• PostgreSQL 15.x (PostgreSQL 15.12.1 is included in the Veeam Backup & Replication 12.3 setup)<br><br>The PostgreSQL instance must have UTF-8 as the default encoding for the database.<br><br>Local or remote installation of the following versions of Microsoft SQL Server[1]:<br>• Microsoft SQL Server 2022<br>• Microsoft SQL Server 2019<br>• Microsoft SQL Server 2017<br>• Microsoft SQL Server 2016<br>• Microsoft SQL Server 2014<br>• Microsoft SQL Server 2012<br><br>All editions of Microsoft SQL Server are supported. The usage of Microsoft SQL Server Express Edition is limited by the database size up to 10 GB. If you plan to have larger databases, use other editions of Microsoft SQL Server.<br><br>If you plan to use a database engine other than PostgreSQL 15.x, included in the Veeam Backup & Replication setup, you must install it yourself. If you want to use an already installed PostgreSQL instance for the configuration database, make sure the instance contains the default `postgres` database. If you allow the setup to install a new PostgreSQL instance, the `postgres` database will be created on the instance automatically. Since Veeam Backup & Replication connects to the `postgres` database to access the configuration database, do not rename the `postgres` database upon the installation.<br><br>Veeam Backup & Replication does not support Microsoft SQL Server database with case-sensitive collations.<br><br>Veeam Backup & Replication and Veeam Backup Enterprise Manager configuration databases can be deployed in Microsoft SQL Always On Availability Groups. For more information, see this Veeam KB article.<br><br>[1] Consider the following:<br>• Veeam Backup & Replication does not support PostgreSQL and Microsoft SQL Server installations on cloud database services (for example, Amazon Relational Database Service (RDS).<br>• We do not recommend sharing a local instance of PostgreSQL with any other services. It should be dedicated to host the backup server database only.<br>• If you are planning to use Veeam Backup & Replication with Veeam Cloud Connect, note that system requirements for the backup server are different. For more information, see the Veeam Cloud Connect Guide. |

| Specification | Requirement |
| --- | --- |
| **Software** | During setup, the system configuration check is performed to determine if all prerequisite software is available on the machine where you plan to install Veeam Backup & Replication. If some of the required software components are missing, the setup wizard will offer you to install missing software automatically. This refers to:<br>• Microsoft .NET Framework 4.7.2<br>• Microsoft ASP.NET Core Runtime 8.0<br>• Microsoft .NET Desktop Runtime 8.0<br>• Microsoft Edge WebView2 Runtime 130.0.2849.56 (not installed for Microsoft Windows Server 2012 and 2012 R2 due to the version incompatibility)<br>• Microsoft PowerShell 5.1<br>• Microsoft Report Viewer Redistributable 2015<br>• Microsoft SQL Server System CLR Types (both for SQL Server and PostgreSQL installations)<br>• Microsoft Universal C Runtime<br>• Microsoft Visual C++ 2015-2022 Redistributable 14.40.33810<br><br>The backup server installation also requires the automatic installation of the prerequisite software for Veeam Cloud Plug-Ins.<br><br>The following software must be installed manually:<br>• Windows Installer 4.5 |

Consider the following:

- If you plan to back up VMs running Microsoft Windows Server 2012 R2 or later, and Data Deduplication is enabled for some VM volumes, it is recommended that you deploy the Veeam Backup & Replication console and mount server on a machine running same or later version of Microsoft Windows Server with Data Deduplication feature enabled. Otherwise, some types of restore operations for these VMs (such as Microsoft Windows File-Level Recovery) may fail.
- Due to its limitations, Microsoft SQL Server Express Edition can only be used for evaluation purposes or in case of a small-scale production environment. For environments with a lot of VMs, it is necessary to install a fully functional commercial version of Microsoft SQL Server.

For more information, see the Backup Server section.

# Veeam Backup & Replication Console

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: x86-64 processor.<br><br>*Memory*: 8 GB RAM<br><br>*Disk Space*: 500 MB for product installation and 4.5 GB for Microsoft .NET Framework 4.7.2 installation.<br><br>*Network*: 1 Mbps connection to the backup server. High latency and low bandwidth impact user interface responsiveness. |
| **OS** | 64-bit versions of the following Microsoft Windows operating systems are supported:<br>• Microsoft Windows Server 2025<br>• Microsoft Windows Server 2022<br>• Microsoft Windows Server 2019<br>• Microsoft Windows Server 2016<br>• Microsoft Windows Server 2012 R2<br>• Microsoft Windows Server 2012<br>• Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>• Microsoft Windows 10 (versions 1909 to 22H2)<br>• Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021) |
| **Software** | During setup, the system configuration check is performed to determine if all prerequisite software is available on the machine where you plan to install the Veeam Backup & Replication Console. If some of the required software components are missing, the setup wizard will offer you to install missing software automatically. This refers to:<br>• Microsoft .NET Framework 4.7.2<br>• Microsoft .NET Desktop Runtime 8.0<br>• Microsoft Edge WebView2 Runtime 126.0.2592.81 (not installed for Microsoft Windows Server 2012 and 2012 R2 due to the version incompatibility)<br>• Microsoft PowerShell 5.1<br>• Microsoft Report Viewer Redistributable 2015<br>• Microsoft SQL Server System CLR Types (both for SQL Server and PostgreSQL installations)<br>• Microsoft Universal C Runtime<br><br>The following software must be installed manually:<br>• Windows Installer 4.5 |

For more information, see the Backup & Replication Console section.

# VMware Backup Proxy

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: x86-64 processor with 2 cores (vCPUs) minimum, plus 1 core (vCPU) for each 2 additional concurrent tasks. Using faster processors improves data processing performance. For more information, see Limitation of Concurrent Tasks.<br><br>*Memory*: 2 GB RAM plus 1 GB for each concurrent task. The actual size of memory required may be larger and depends on the amount of data to back up, machine configuration, and job settings. Using faster memory improves data processing performance.<br><br>*Disk Space*: 750 MB for Microsoft Windows-based proxies; 400 MB for Linux-based proxies.<br><br>*Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported. |
| **OS** | 64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition:<br>• Microsoft Windows Server 2025<br>• Microsoft Windows Server 2022<br>• Microsoft Windows Server 2019<br>• Microsoft Windows Server 2016<br>• Microsoft Windows Server 2012 R2<br>• Microsoft Windows Server 2012<br>• Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>• Microsoft Windows 10 (versions 1909 to 22H2)<br>• Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021)<br><br>64-bit versions of the following Linux distributions are supported. Note that bash shell and SSH are required.<br>• Almalinux 9.3<br>• CentOS 7.x<br>• Debian 10.0 to 12.2<br>• Oracle Linux 7 (UEK3) to 9 (UEK R7)<br>• Oracle Linux 7 to 9 (RHCK)<br>• RHEL 7.0 to 9.5<br>• Rocky Linux 9.0 to 9.5<br>• SLES 12 SP4 or later, 15 SP1 or later<br>• Ubuntu: 18.04 LTS, 20.04 LTS, 22.04 LTS, 24.04 LTS |

For more information, see the VMware Backup Proxies section.

# General-Purpose Backup Proxy

The following table shows the minimum system requirements for a general-purpose backup proxy used for unstructured data backup and Veeam Agent and storage system snapshot integration.

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: x86-64 processor with 2 cores (vCPUs) minimum. 4 cores (vCPUs) are recommended (2 cores are required) for each additional concurrent task. Using multi-core processors enhances data processing performance and enables the processing of more tasks simultaneously.<br><br>*Memory*: [For unstructured data backup] 4 GB RAM plus 4 GB RAM for each concurrent task.<br><br>[For Veeam Agent and storage system snapshot integration] 2 GB RAM plus 1 GB RAM for each concurrent task.<br><br>Using faster memory improves data processing performance. For all-in-one installations, where the server performs several roles, it must have enough memory resources for all components.<br><br>*Disk Space*: 300 MB.<br><br>*Network*: High latency and reasonably unstable WAN links are supported. |
| **OS** | 64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition:<br>• Microsoft Windows Server 2025<br>• Microsoft Windows Server 2022<br>• Microsoft Windows Server 2019<br>• Microsoft Windows Server 2016<br>• Microsoft Windows Server 2012 R2<br>• Microsoft Windows Server 2012<br>• Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>• Microsoft Windows 10 (versions 1909 to 22H2)<br>• Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021)<br><br>Backup proxies running Microsoft Windows Server 2012 R2 or later support the **VSS for SMB File Shares** feature. To use it, make sure that all requirements listed in this Veeam KB article are met.<br><br>For **object storage backup**, also 64-bit versions of the following Linux distributions are supported:<br>• CentOS 7.x<br>• Debian 10.0 to 12.2<br>• Oracle Linux 7 (UEK3) to 9 (UEK R7)<br>• Oracle Linux 7 to 9 (RHCK)<br>• RHEL 7.0 to 9.3<br>• SLES 12 SP4 or later, 15 SP1 or later<br>• Ubuntu: 18.04 LTS, 20.04 LTS, 22.04 LTS |

For more information, see the General-Purpose Backup Proxies section.

# Proxmox VE Worker

Workers process the backup workload and distribute the backup traffic when transferring data to backup repositories. If you deploy a worker using the default configuration, the following compute resources will be allocated:

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: 6 cores (vCPUs).<br>*Memory*: 6 GB RAM<br>*Disk Space*: 100 GB for product installation and logs. |

With the default configuration, the worker can handle up to 4 concurrent backup and restore tasks. While deploying a new worker or editing settings of an existing one, you can increase the maximum number of concurrent tasks. However, you must allocate 1 vCPU and 1 GB RAM for each additional task. When configuring the maximum number of concurrent tasks, you must also take into account the network traffic throughput in your virtual infrastructure.

For more information, see the Veeam Backup for Proxmox VE User Guide.

# VMware CDP Proxy

The following table shows the minimum system requirements for a VMware CDP proxy.

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: x86-64 processor with 4 cores (8 cores before Veeam Backup & Replication 12.1 (build 12.1.0.2131)) minimum. Using multi-core processors improves data processing performance and allows for more tasks to be processed concurrently. Enabling the encryption may cause the performance reduction: in this case, increase the vCPU count up to 2 times.<br>*Memory*: 8 GB RAM minimum (16 GB RAM before Veeam Backup & Replication 12.1 (build 12.1.0.2131)). Using more memory allows for longer peak write I/O periods before a CDP policy switches to the disk-based write I/O cache. Using faster memory improves data processing performance.<br>*Disk Space*: 300 MB plus disk-based write I/O cache (non-persistent data, at least 50 GB recommended). Larger cache allows for longer network downtime periods before a CDP policy switches to the CBT mode.<br>*Network*: 100 Mbps or faster. |

| Specification | Requirement |
| --- | --- |
| OS | 64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition:<br>• Microsoft Windows Server 2025<br>• Microsoft Windows Server 2022<br>• Microsoft Windows Server 2019<br>• Microsoft Windows Server 2016<br>• Microsoft Windows Server 2012 R2<br>• Microsoft Windows Server 2012<br>• Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>• Microsoft Windows 10 (versions 1909 to 22H2)<br>• Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021)<br>64-bit versions of the following Linux distributions are supported. Note that bash shell and SSH are required.<br>• CentOS 7.x<br>• Debian 10.0 to 12.2<br>• Oracle Linux 7 (UEK3) to 8.3 (UEK R6 U2)<br>• Oracle Linux 7 to 8.5 (RHCK)<br>• RHEL 7.0 to 9.5<br>• SLES 12 SP4 or later, 15 SP1 or later<br>• Ubuntu: 18.04 LTS, 20.04 LTS, 22.04 LTS |

For more information, see the VMware CDP Proxies section.

# Backup Repository

These requirements also apply to separate mount servers (they can be only Windows-based), gateway servers for unstructured data backup and deduplicating appliance-based repositories.

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: x86-64 processor. The number of cores depends on the concurrent task settings. For more information, see Limitation of Concurrent Tasks.<br><br>*Memory*: 4 GB RAM, plus not less than 1 GB RAM for each concurrently processed machine disk. For more information, see Limitation of Concurrent Tasks.<br><br>[For unstructured data backup] Not less than 4 GB RAM for each concurrently processed unstructured data source (file share or object storage); in case of deduplicating storage appliances, up to 8 GB RAM. Additionally, 1GB RAM is required for indexing each 200 million objects (files and folders). In case of all-in-one installations for unstructured data backup, where the server performs several roles, it must have enough memory resources for all components.<br><br>*Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported. |

| Specification | Requirement |
| --- | --- |
| OS | 64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition:<br>• Microsoft Windows Server 2025<br>• Microsoft Windows Server 2022<br>• Microsoft Windows Server 2019<br>• Microsoft Windows Server 2016<br>• Microsoft Windows Server 2012 R2<br>• Microsoft Windows Server 2012<br>• Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>• Microsoft Windows 10 (versions 1909 to 22H2)<br>• Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021)<br><br>64-bit versions of the following Linux distributions are supported:<br>• AlmaLinux 9.3<br>• CentOS 7.x<br>• Debian 10.0 to 12.2<br>• Oracle Linux 7 (UEK3) to 9 (UEK R7)<br>• Oracle Linux 7 to 9 (RHCK)<br>• RHEL 7.0 to 9.5<br>• Rocky Linux 9.0 to 9.5<br>• SLES 12 SP4 or later, 15 SP1 or later<br>• Ubuntu: 18.04 LTS, 20.04 LTS, 22.04 LTS, 24.04 LTS<br><br>Bash shell and SSH are required to deploy the management agent (SSH connection is not required for updating Veeam components and can be disabled afterwards). Perl is required only for non-persistent Veeam Data Movers. Check the full list of required Perl modules in this Veeam KB article.<br><br>For advanced XFS integration (fast clone), only the following 64-bit Linux distributions are supported:<br>• AlmaLinux 9.3<br>• Debian 10.0 to 12.2<br>• RHEL 8.2 to 9.5<br>• Rocky Linux 9.0 to 9.5<br>• SLES 15 SP2 to SP6<br>• Ubuntu 18.04 LTS, 20.04 LTS, 22.04 LTS, 24.04 LTS<br><br>For other distributions, XFS integration support is experimental, with kernel version 5.4 or later recommended. For more information, see this Veeam KB article. |

For more information, see the Backup Repositories section.

NOTE

Consider the following:

- If you plan to use a Microsoft Windows backup repository with Data Deduplication, make sure that you set up the Microsoft Windows server correctly. For more information, see this Veeam KB article.
- In rare cases, Veeam backups may become corrupted when stored on the exFAT file system. We strongly discourage using the exFAT filesystem for backup storage when attached to Windows 10 or Windows Server 2019 and later. For more information, see this Veeam KB article.

# Cache Repository

The following storage types can be used as a cache repository for unstructured data backup:

- Direct attached storage. You can add virtual and physical servers as cache repositories:
  - Microsoft Windows server (only 64-bit versions are supported).
  - Linux server (only 64-bit versions are supported).
- Network attached storage. You can add SMB (CIFS) Share or NFS Share as a cache repository.

| Specification | Requirement |
|---|---|
| **Hardware** | Cache repository hardware requirements depend on the target backup repository.<br><br>For direct attached storage, network attached storage (NAS) or deduplicating storage appliance as the backup target:<br>o CPU: x86-64 processor with 2 cores (vCPUs), plus 4 cores (vCPUs) for each concurrently processed unstructured data source.<br>o Memory: 4 GB RAM, plus 4 GB RAM for each concurrently processed unstructured data source.<br>o Network: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported.<br><br>For object storage as the backup target:<br>o CPU: x86-64 processor with 2 cores (vCPUs), plus 6 cores (vCPUs) for each concurrently processed unstructured data source.<br>o Memory: 4 GB RAM, plus 16 GB RAM for each concurrently processed unstructured data source.<br>o Network: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported. |

| Specification | Requirement |
| --- | --- |
| OS | 64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition:<br>○ Microsoft Windows Server 2025<br>○ Microsoft Windows Server 2022<br>○ Microsoft Windows Server 2019<br>○ Microsoft Windows Server 2016<br>○ Microsoft Windows Server 2012 R2<br>○ Microsoft Windows Server 2012<br>○ Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>○ Microsoft Windows 10 (versions 1909 to 22H2)<br>○ Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021)<br><br>64-bit versions of the following Linux distributions are supported:<br>○ AlmaLinux 9.3<br>○ CentOS 7.x<br>○ Debian 10.0 to 12.2<br>○ Oracle Linux 7 (UEK3) to 9 (UEK R7)<br>○ Oracle Linux 7 to 9 (RHCK)<br>○ RHEL 7.0 to 9.5<br>○ Rocky Linux 9.0 to 9.5<br>○ SLES 12 SP4 or later, 15 SP1 or later<br>○ Ubuntu: 18.04 LTS, 20.04 LTS, 22.04 LTS, 24.04 LTS<br><br>Bash shell and SSH are required to deploy the management agent (SSH connection is not required for updating Veeam components and can be disabled afterwards). Perl is required only for non-persistent Veeam Data Movers. Check the full list of required Perl modules in this Veeam KB article.<br><br>For advanced XFS integration (fast clone), only the following 64-bit Linux distributions are supported:<br>○ AlmaLinux 9.3<br>○ Debian 10.0 to 12.2<br>○ RHEL 8.2 to 9.5<br>○ Rocky Linux 9.0 to 9.5<br>○ SLES 15 SP2 to SP6<br>○ Ubuntu 18.04 LTS, 20.04 LTS, 22.04 LTS, 24.04 LTS<br><br>For other distributions, XFS integration support is experimental, with kernel version 5.4 or later recommended. For more information, see this Veeam KB article. |

For more information, see **Cache Repository** in the Backup Infrastructure for Unstructured Data Backup section.

### Cache Repository for Object Storage Repository

If the cache repository works with an unstructured data source being backed up to an object storage repository, it also processes active metadata that is heavily used during backup and restore operations. So you will require more disk space for the cache repository. The volume of this disk space depends on the number of the file versions on the data source and the number of unstructured data backup jobs that protect this data source. We recommend allocating not less than 1 GB of disk space for active metadata of each 1,000,000 file versions protected with 1 file backup job or object storage backup job. If you protect the same data source, for example, with 2 different backup jobs, the volume of metadata will double. For more information, see the Unstructured Data Backups in Object Storage Repositories section.

> **TIP**
>
> We strongly recommend utilizing a fast storage disk, for example, an SSD in the role of the cache repository used for working with an object storage repository.

## Tape Server

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: x86-64 processor with 2 cores (vCPUs) minimum. Using multi-core processors improves data processing performance and allows for more tasks to be processed concurrently.<br><br>*Memory*: 2 GB RAM plus 500MB for each concurrent task. Depending on the source of tape jobs, different entities are considered tasks: for machine backup to tape, a task covers a source job or a source chain if tape paralleling is enabled; for file backup to tape, a task covers an entire server or a file share. Restoring VMs directly from tape requires 400MB of RAM per 1TB of the restored virtual disk size. Tape cloning requires 1GB RAM for each concurrent task.<br><br>*Disk Space*: 300 MB, plus 10 GB for temporary data storage for backup and restore operations.<br><br>*Network*: 1 Gbps or faster.<br><br>For system requirements for backup to tape jobs, see the Before You Begin section for backup to tape.<br><br>For system requirements for large number of files in the file backup to tape job, see the Before You Begin section for file backup to tape. |

| Specification | Requirement |
| --- | --- |
| OS | 64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition:<br>o Microsoft Windows Server 2025<br>o Microsoft Windows Server 2022<br>o Microsoft Windows Server 2019<br>o Microsoft Windows Server 2016<br>o Microsoft Windows Server 2012 R2<br>o Microsoft Windows Server 2012<br>o Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>o Microsoft Windows 10 (versions 1909 to 22H2)<br>o Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021)<br><br>64-bit versions of the following Linux distributions are supported[1]:<br>o CentOS 7.x<br>o Debian 10.0 to 12.2<br>o Oracle Linux 7 (UEK3) to 9 (UEK R7)<br>o Oracle Linux 7 to 9 (RHCK)<br>o RHEL 7.0 to 9.5<br>o Rocky Linux 9.4, 9.5<br>o SLES 12 SP4 or later, 15 SP1 or later<br>o Ubuntu: 18.04 LTS, 20.04 LTS, 22.04 LTS, 24.04 LTS<br><br>[1] Note that the tape server requires root access rights. For this reason, hardened repository cannot be used as the tape server. |

For more information, see the Tape Servers section in the Tape Devices Support Guide.

## WAN Accelerator

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: x86-64 processor. Using multi-core processors improves data processing performance, and is highly recommended on WAN links faster than 10 Mbps.<br><br>*Memory*: 8 GB RAM. Using faster memory improves data processing performance.<br><br>*Disk Space*: Disk space requirements depend on the WAN Accelerator role. Source WAN Accelerator requires 20 GB per 1 TB of source data to store digests of data blocks of source VM disks. Disk space consumption is dynamic and changes as unique VMs are added to (or removed from) jobs with WAN Acceleration enabled. Target WAN Accelerator requires global cache size as defined by user (fixed amount). Disk space is reserved immediately upon selecting the WAN Accelerator as a target one in any job. For more information, see WAN Accelerator Sizing.<br><br>*Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported. |
| **OS** | 64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition:<br>○ Microsoft Windows Server 2025<br>○ Microsoft Windows Server 2022<br>○ Microsoft Windows Server 2019<br>○ Microsoft Windows Server 2016<br>○ Microsoft Windows Server 2012 R2<br>○ Microsoft Windows Server 2012<br>○ Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>○ Microsoft Windows 10 (versions 1909 to 22H2)<br>○ Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021) |

For more information, see the WAN Accelerators section.

> **NOTE**
>
> Global cache is not leveraged by source WAN accelerators, or by WAN accelerators operating in the **High bandwidth mode**, so it does not need to be allocated and populated in such cases.

## Backup Target

Backed up data can be stored in the following disk-based systems:

- Local (internal) storage of the backup repository
- Direct Attached Storage (DAS)

  The DAS must be connected to the backup repository, including external USB/eSATA drives, USB pass through and raw device mapping (RDM) volumes.

- Storage Area Network (SAN)

  The backup repository must be connected into the SAN fabric through hardware or virtual HBA, or software iSCSI initiator.

- Network Attached Storage (NAS)

  The NAS must be able to present its capacity as NFS share (protocol versions 3.0 and 4.1 only) or SMB (CIFS) share (any protocol versions). Using SMB (CIFS) protocol for non-continuously available (CA) file shares is not recommended for reliability reasons. Using consumer-grade NAS storage without an enterprise-grade RAID controller with battery-backed write cache (BBWC) is not recommended for reliability considerations.

- Veeam Data Cloud Vault
- Amazon S3
- Google Cloud Storage
- IBM Cloud Object Storage
- Microsoft Azure Blob Storage
- Wasabi Hot Cloud Storage
- Any S3-compatible object storage (on-premise appliance, or cloud storage provider)
- Dell Data Domain (DD OS version 7.3 to 8.3) with DDBoost license. Both Ethernet and Fibre Channel (FC) connectivity is supported.
- ExaGrid[1] (firmware version 5.0.0 or later)
- Fujitsu ETERNUS CS800[1] software version 3.4.0 or later
- HPE StoreOnce (firmware version 3.18.18 or later for Gen3 and 4.2.3 or later for Gen4) with Catalyst license

  Both Ethernet and Fibre Channel (FC) connectivity are supported. Note that HPE StoreOnce Federated Catalyst is not supported.

- Infinidat InfiniGuard[1] version 3.6 and later
- Quantum[1] (DXi software 3.4.0 or later)

  Supported Quantum DXi systems include DXi4700 (NAS configuration), DXi4700 (multi-protocol configuration), DXi 4800, DXi 6900, DXi 6900-S, DXi 9000, DXi 9100, and DXi V5000. FIPS-compliant operations mode requires DXi software 4.0 or later.

[1] These storage systems use the Veeam Transport Service. Make sure that they also meet system requirements for the backup repository.

Once backups are created, they can be copied (for redundancy) or offloaded (for long-term retention) to one of the following hot object storage types using the scale-out backup repository Capacity Tier:

- Veeam Data Cloud Vault
- Amazon S3 (including AWS Snowball Edge)
- Google Cloud Storage
- IBM Cloud Object Storage
- Microsoft Azure Blob storage (including Microsoft Azure Data Box)
- Wasabi Hot Cloud Storage

- Any S3-compatible object storage (on-premises appliance, or cloud storage provider)

Once backups are created on Amazon S3, Microsoft Azure Blob Storage, or S3-compatible object storage systems with the archiving extension of Smart Object Storage API, they can be further archived to one of the following respective cold object storage classes using the scale-out backup repository Archive Tier:

- Amazon S3 Glacier Instant Retrieval
- Amazon S3 Glacier Flexible Retrieval
- Amazon S3 Glacier Deep Archive
- Microsoft Azure Archive Tier
- Microsoft Azure Cold Tier
- Any S3-compatible object storage with data archiving enabled

For the full list of partner-tested solutions including primary backup storage solutions, S3-compatible object storage solutions and offline storage solutions, see this Veeam page.

For information on unstructured data backup target, see **Storage Repositories** in the Backup Infrastructure for Unstructured Data Backup section.

# Veeam CDP Source and Target Datastores

The following source and target datastores are supported for Veeam CDP:

- NFS on file storage
- VMFS on block storage
- VMFS on internal ESXi storage
- VSAN

  VSAN is supported for all hyper-converged infrastructure (HCI) appliances.
- vVol

  vVol is supported for the following vendors: NetApp, HPE Nimble, Pure Storage and HPE 3PAR. For the list of tested vendor product lines, see this Veeam KB article.

Support for HCI appliances is pending validation by Veeam. The documentation will be updated based on the testing results. For the updates, see CDP Requirements and Limitations.

> **NOTE**
>
> Consider the following:
>
> - Cisco HyperFlex 4.5 (2a) and later can be used as a source or target only with VMware vSphere 7.0 U2 and later. With the previous versions of VMware vSphere, Cisco HyperFlex is not supported.
> - vVol and VSAN of each vendor have limits for the number of storage objects. Once the limit is reached, CDP starts to fail because creation of new objects cannot be completed. To restore the CDP process, remove some objects from VSAN/vVol.
> - CDP performance depends on your datastore characteristics. For better performance, check that your target datastore is able to process the I/O workload produced on the source datastore.

# Tape

| Specification | Requirement |
|---|---|
| **Hardware** | The following types of tape libraries (including VTL) and standalone drives are supported:<br>○ LTO3-LTO9<br>○ IBM 3592 (TS1160 and TS1170)<br><br>Tape device must be directly attached to the backup server or to a tape proxy server through SAS, FC or iSCSI interface. Note that VMware does not support connecting tape libraries to ESXi for VM pass-through. |
| **Software** | ○ Tape devices without device-specific, vendor-supplied OEM drivers for Windows will appear in Windows Device Manager as Unknown or Generic and require enabling native SCSI commands mode.<br>○ If multiple driver installation modes are available for your tape device, use the one that allows for multiple open handles from a host to a drive to exist at the same time. Usually, such drivers are referred to as "non-exclusive".<br>○ No other backup server must be interacting with the tape device. |

For more information, see the Tapes section in the Tape Devices Support Guide.

# Storage Systems

Veeam Backup & Replication provides different types of snapshot integrations and these integrations support different storage systems:

- NAS integration with built-in storage systems
- VMware and Veeam Agent Integrations with built-in storage systems
- VMware and Veeam Agent Integrations with Universal Storage API integrated systems

> **NOTE**
>
> [For VMware integration] You may encounter scenarios when Veeam Backup & Replication and the storage system support different versions of the virtual platform. In such cases, we recommend using versions supported by the storage system. If other versions are used, Veeam support will assist only with issues unrelated to version compatibility.

## NAS Integration

NAS integration is possible with the following built-in storage systems.

### Dell PowerScale (formerly Isilon)

- Filer integration for NAS backup functionality
- NFS or SMB (CIFS) connectivity

- OneFS 8.1.2 to 9.9

### Fujitsu ETERNUS HX/AX

- NFS or SMB (CIFS) connectivity
- ONTAP 7-mode versions 8.2 up to 8.2.5
- ONTAP cluster-mode versions 9.1 to 9.16.1

### IBM N Series

- NFS, Fibre Channel or iSCSI connectivity
- ONTAP 7-mode versions 8.2 up to 8.2.5
- ONTAP cluster-mode versions 9.1 to 9.16.1

### Lenovo ThinkSystem DM/DG Series

- NFS or SMB (CIFS) connectivity
- ONTAP 7-mode versions 8.2 up to 8.2.5
- ONTAP cluster-mode versions 9.1 to 9.16.1

### NetApp ONTAP FAS/AFF/ASA, FlexArray (V-Series), Edge/Select/Cloud VSA

- NFS or SMB (CIFS) connectivity
- ONTAP 7-mode versions 8.2 up to 8.2.5
- ONTAP cluster-mode versions 9.1 to 9.16.1
- NetApp ASA r2 systems (ASA A1K, ASA A70, and ASA A90) are not supported

### Nutanix Files

- Filer integration for NAS backup functionality
- NFS or SMB (CIFS) connectivity
- Nutanix File Server 3.8.1.3 to 5.0

## VMware and Veeam Agent Integrations with Built-in Storage Systems

VMware and Veeam Agent integrations are possible with the following built-in storage systems.

### Cisco HyperFlex HX-Series

- VMware integration only
- NFS connectivity only
- HyperFlex 4.0(2x) or later (Backup from Storage Snapshots, Full Integration mode)
- Basic authentication is not supported for SSO users in HyperFlex

### Dell VNX, VNX2, VNXe and Unity XT, Unity

- NFS, Fibre Channel (FC) or iSCSI connectivity
- Dell VNX/VNX2 all OE versions are supported
- Dell VNXe OE versions 3.x
- Dell Unity XT/Unity OE versions 5.0 to 5.5

  OE version 5.5 requires TLS 1.2 usage. For more information, see this Veeam KB article.

### Fujitsu ETERNUS HX/AX

- NFS, Fibre Channel or iSCSI connectivity
- ONTAP 7-mode versions 8.2 up to 8.2.5
- ONTAP cluster-mode versions 9.1 to 9.16.1

### HPE 3PAR StoreServ

- Fibre Channel (FC) or iSCSI connectivity
- 3PAR OS versions 3.2.2 up to 3.3.2
- WSAPI 1.5 and later
- iSCSI VLAN tags are supported

### HPE Primera

- Fibre Channel (FC) or iSCSI (starting from OS versions 4.3 or later) connectivity
- OS versions 4.x

### HPE Alletra 9000

- Fibre Channel (FC) or iSCSI connectivity
- OS version 9.3 or later

### HPE Alletra MP

- Fibre Channel (FC) or iSCSI connectivity

  iSCSI is supported starting from OS version 10.3.
- OS version 10.2 or later

### HPE Nimble Storage AF-Series, HF-Series and CS-Series

- Fibre Channel (FC) or iSCSI connectivity
- Nimble OS 5.0 and later

### HPE Alletra 5000, 6000

- Fibre Channel (FC) or iSCSI connectivity

- OS version 6.1

### HPE StoreVirtual (formerly LeftHand/P4000 Series) and StoreVirtual VSA

- iSCSI connectivity only
- LeftHand OS versions 9.5 up to 12.8
- HPE SV3200 (LeftHand OS version 13) is not supported

### IBM N Series

- NFS, Fibre Channel or iSCSI connectivity
- ONTAP 7-mode versions 8.2 up to 8.2.5
- ONTAP cluster-mode versions 9.1 to 9.16.1

### Lenovo ThinkSystem DM/DG Series

- NFS, Fibre Channel or iSCSI connectivity
- ONTAP 7-mode versions 8.2 up to 8.2.5
- ONTAP cluster-mode versions 9.1 to 9.16.1

### NetApp ONTAP FAS/AFF/ASA, FlexArray (V-Series), Edge/Select/Cloud VSA

- NFS, Fibre Channel or iSCSI connectivity
- ONTAP 7-mode versions 8.2 up to 8.2.5
- ONTAP cluster-mode versions 9.1 to 9.16.1
- NetApp ASA r2 systems (ASA A1K, ASA A70, and ASA A90) are not supported.

## VMware and Veeam Agent Integrations with Universal Storage API Integrated Systems

To start working with Universal Storage API integrated systems, you must install storage system plug-ins. VMware and Veeam Agent integrations are possible with the following storage systems.

### DataCore SANsymphony

- Fibre Channel (FC) or iSCSI connectivity
- DataCore SANsymphony 10.0 PSP12 or later

### Dell PowerMax

- Fibre Channel (FC) or iSCSI connectivity
- Dell PowerMax/VMAX All Flash (PowerMax OS microcode family 5978 or later 5978 family updates and PowerMaxOS 10.1.0.0 or later)
- Unisphere for PowerMax 9.2.1.6 or later (versions 10.0.x are not supported)

### Dell PowerStore

- Fibre Channel (FC) or iSCSI connectivity
- Dell PowerStore (PowerStore OS 3.0 or later)

### Dell SC Series (formerly Compellent)

- Fibre Channel (FC) or iSCSI connectivity
- Storage Center OS 7.4.2 or later
- FluidFS volumes and Live Volumes are not supported

### Fujitsu ETERNUS AF and DX Series

- Fibre Channel (FC) or iSCSI connectivity
- ETERNUS AF series: AF250 S2, AF650 S2, AF150 S3, AF250 S3, AF650 S3
- ETERNUS DX series: DX60 S4, DX100 S4, DX200 S4, DX500 S4, DX600 S4, DX8900 S4, DX60 S5, DX100 S5, DX200 S5, DX500 S5, DX600 S5, DX900 S5, DX600 S6, DX900 S6, DX8900 S6
- Storage firmware version:
  - ETERNUS AF S2 and DX S4 series (except DX8900 S4): V10L88-1000 or later
  - ETERNUS AF S3 and DX S5 series, DX8900 S4: V11L30-5000 or later
  - ETERNUS DX S6 series: V12L10-0000 or later

### Hitachi VSP/VSP One Block

- Fibre Channel (FC) or iSCSI connectivity
- VSP E series (93-03-01-60/00 or later)
- VSP F series (88-07-01-x0/00 or later)
- VSP G series (88-07-01-x0/00 or later)
- VSP 5100 and VSP 5500 (90-05-01-00/00 or later)
- VSP 5200 and VSP 5600 (90-08-01-00/00 or later)
- VSP One Block series (A3-03-01-x0/01 or later)

### HPE XP

- Fibre Channel (FC) or iSCSI connectivity
- HPE XP8 (90-05-01-00/00 or later)

### IBM FlashSystem (formerly Spectrum Virtualize, includes IBM StorWize and IBM SVC)

- Fibre Channel (FC) or iSCSI connectivity
- IBM Spectrum Virtualize OS version 8.2 or later

- Policy-based replication relationships are not detected

  You will not be able to select this feature for backup and snapshot orchestration with secondary storage arrays.

- Policy-based high availability (PBHA) configurations are not supported

  Volumes covered by PBHA cannot be used for the storage integration feature.

### INFINIDAT Infinibox F-Series

- NFS, Fibre Channel (FC) or iSCSI connectivity
- InfiniBox 5.0 or later

> **NOTE**
>
> You must add to the backup infrastructure only one of the two InfiniBox storage arrays for which Active/Active Replication is configured, or exclude the replicating volumes on one of these arrays from rescan. For details on how to exclude volumes from rescan, see the Rescanning Storage Systems section in the Veeam Backup & Replication User Guide.

### NEC Storage M Series

- Fibre Channel (FC) or iSCSI connectivity
- M120, M320, M320F, M520, M720, M720F (Storage Control Software revision 1234 or later)

### NEC Storage V Series

- Fibre Channel (FC) or iSCSI connectivity
- V100, V110, V300 (93-04-21-XX or later), V310, V310F (A3-03-01-XX or later), V10e (88-08-09-XX or later)

### NetApp SolidFire/HCI

- iSCSI connectivity
- Element OS version 10.0 or later

### Pure Storage FlashArray

- NFS, Fibre Channel (FC) or iSCSI connectivity
- Purity 6.3 or later

### Tintri IntelliFlash (formerly Western Digital IntelliFlash, Tegile)

- NFS, Fibre Channel (FC) or iSCSI connectivity
- Tintri IntelliFlash 3.11 or later

# Gateway Server

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: x86-64 processor, 2 or more cores.<br><br>*Memory*: 4 GB RAM, plus up to 4 GB RAM for each concurrently processed machine, file share or object storage. For more information, see Limitation of Concurrent Tasks. For RAM allocation recommendations for unstructured data backup, see Limitations and recommendations for unstructured data backup.<br><br>*Disk space:* 750 MB for Microsoft Windows-based proxies; 400 MB for Linux-based proxies.<br><br>**Note**: If a unstructured data backup stored in an object storage does not have metadata in the cache repository, during the restore or health check operation this metadata will be downloaded to the gateway server. That can consume up to 80% of the gateway server disk space.<br><br>*Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported. |
| **OS** | 64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition:<br>○ Microsoft Windows Server 2025<br>○ Microsoft Windows Server 2022<br>○ Microsoft Windows Server 2019<br>○ Microsoft Windows Server 2016<br>○ Microsoft Windows Server 2012 R2<br>○ Microsoft Windows Server 2012<br>○ Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>○ Microsoft Windows 10 (versions 1909 to 22H2)<br>○ Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021)<br><br>For the communication with object storage repositories, external repositories and NFS backup repositories, you can use machines running 64-bit versions of the following Linux distributions:<br>○ CentOS 7.x<br>○ Debian 10.0 to 12.2<br>○ Oracle Linux 7 (UEK3) to 9 (UEK R7)<br>○ Oracle Linux 7 to 9 (RHCK)<br>○ RHEL 7.0 to 9.5<br>○ Rocky Linux 9.0 to 9.5<br>○ SLES 12 SP4 or later, 15 SP1 or later<br>○ Ubuntu: 18.04 LTS, 20.04 LTS, 22.04 LTS, 24.04 LTS |

For more information, see the Gateway Server section.

# Mount Server

| Specification | Requirement |
|---|---|
| **Hardware** | *CPU*: x86-64 processor with 2 cores (vCPUs) minimum, plus 1 core (vCPU) for each 2 additional concurrent tasks. Using faster processors improves data processing performance.<br><br>*Memory*: 4 GB RAM, plus not less than 1 GB RAM for each concurrently processed machine disk.<br><br>The following jobs consume not less than 400 MB RAM per guest VM on mount server:<br>○ Windows file level restore<br><br>The following jobs consume 1 GB RAM per guest VM disk on mount server + 100 MB RAM per VM:<br>○ SureBackup<br>○ Instant Recovery<br>○ Instant Disk Recovery<br><br>*Disk Space*: 1.4 GB for product installation and 4.5 GB for Microsoft .NET Framework 4.7.2 installation. If Microsoft .NET Framework 4.7.2 is not installed on the machine, Veeam Backup & Replication will install it automatically.<br><br>*Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported. |
| **OS** | 64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition:<br>○ Microsoft Windows Server 2025<br>○ Microsoft Windows Server 2022<br>○ Microsoft Windows Server 2019<br>○ Microsoft Windows Server 2016<br>○ Microsoft Windows Server 2012 R2<br>○ Microsoft Windows Server 2012<br>○ Microsoft Windows 11 (versions 21H2, 22H2, 23H2, 24H2)<br>○ Microsoft Windows 10 (versions 1909 to 22H2)<br>○ Microsoft Windows 10 LTS (versions LTSB 1607, LTSC 1809, LTSC 2021)<br><br>**Note**: If you plan to restore VM guest OS files from VMs running Microsoft Windows ReFS or from VMs with data deduplication enabled for some volumes, you must assign the mount server role to machines running specific OS versions. For more information, see ReFS and Data Deduplication subsections in Restoring VM Guest OS Files (FAT, NTFS or ReFS). |

For more information, see the Mount Server section.

## Helper Host

> **NOTE**
>
> Bash shell and SSH are required.

| OS | Version/Distribution |
|---|---|
| **Linux** | 64-bit versions of the following Linux distributions are supported:<br>○ CentOS 7.x<br>○ Debian 10.0 to 12.2<br>○ Oracle Linux 7 (UEK3) to 9 (UEK R7)<br>○ Oracle Linux 7 to 9 (RHCK)<br>○ RHEL 7.0 to 9.5<br>○ Rocky Linux 9.0 to 9.5<br>○ SLES 12 SP4 or later, 15 SP1 or later<br>○ Ubuntu: 18.04 LTS, 20.04 LTS, 22.04 LTS, 24.04 LTS |
| **IBM AIX** | The following versions of the operating system are supported:<br>IBM AIX 7.1 – 7.3 TL2 |
| **Oracle Solaris** | The following versions of the operating system are supported:<br>Oracle Solaris 10 1/13, 11.0 – 11.4 |

For more information, see Restore from Linux, Unix and Other File Systems.

## Veeam Backup Enterprise Manager Server

The machine where you plan to install Veeam Backup Enterprise Manager must meet the requirements listed in the System Requirements section of the Enterprise Manager User Guide.

# Veeam Plug-Ins

The machine where you plan to install Veeam plug-ins must meet the following requirements:

> **IMPORTANT**
>
> The Microsoft .NET Desktop Runtime and Microsoft ASP.NET Core Runtime must be of the same version (up to the minor version number). Otherwise, starting the plug-in will fail.

| Veeam Plug-In | Requirement |
|---|---|
| **AWS Plug-in for Veeam Backup & Replication version 12.8.0.xxx and later** | Microsoft ASP.NET Core Runtime 8<br>Microsoft .NET Desktop Runtime 8<br>For other system requirements of the plug-in, see the Veeam Backup for AWS User Guide. |
| **Microsoft Azure Plug-in for Veeam Backup & Replication version 12.7.0.xxx and later** | Microsoft ASP.NET Core Runtime 8<br>Microsoft .NET Desktop Runtime 8<br>For other system requirements of the plug-in, see the Veeam Backup for Microsoft Azure User Guide. |
| **Google Cloud Plug-in for Veeam Backup & Replication version 12.6.0.xxx and later** | Microsoft ASP.NET Core Runtime 8<br>Microsoft .NET Desktop Runtime 8<br>For other system requirements of the plug-in, see the Veeam Backup for Google Cloud User Guide. |
| **Nutanix AHV Plug-in for Veeam Backup & Replication version 12.7.0.xxx and later** | Microsoft ASP.NET Core Runtime 8<br>Microsoft .NET Desktop Runtime 8<br>For other system requirements of the plug-in, see the Veeam Backup for Nutanix AHV User Guide. |

IMPORTANT

| | |
|---|---|
| **oVirt KVM Plug-in for Veeam Backup & Replication version 12.6.0.xxx and later** | Microsoft ASP.NET Core Runtime 8¨<br>Microsoft .NET Desktop Runtime 8<br>For other system requirements of the plug-in, see the Veeam Backup for OLVM and RHV User Guide. |
| **Proxmox Virtual Environment Plug-in for Veeam Backup & Replication version 12.1.0.xxx and later** | Microsoft ASP.NET Core Runtime 8<br>Microsoft .NET Desktop Runtime 8<br>For other system requirements of the plug-in, see the Veeam Backup for Proxmox VE User Guide. |
| **Veeam Kasten Plug-in for Veeam Backup & Replication version 12.1.3.xxx and later** | Microsoft ASP.NET Core Runtime 8<br>Microsoft .NET Desktop Runtime 8<br>For other system requirements of the plug-in, see the Veeam Kasten for Kubernetes User Guide. |

# Veeam Explorers

- Veeam Explorer for Microsoft Active Directory
- Veeam Explorer for Microsoft Exchange
- Veeam Explorer for Microsoft SharePoint
- Veeam Explorer for Microsoft SQL Server
- Veeam Explorer for Microsoft Teams
- Veeam Explorer for Microsoft OneDrive for Business
- Veeam Explorer for MongoDB
- Veeam Explorer for Oracle
- Veeam Explorer for PostgreSQL
- Veeam Explorer for SAP HANA

# Permissions

Make sure the user accounts that you plan to use have permissions described in the following sections.

## Installing and Using Veeam Backup & Replication

The accounts used for installing and using Veeam Backup & Replication must have the following permissions.

| Account | Required Permission |
|---|---|
| **Setup Account** | The account used for product installation must have the local Administrator permissions on the target machine. |
| **Veeam Backup & Replication Console Permissions** | When you open the Veeam Backup & Replication console for the first time or after a cumulative patch is installed on the backup server, you must run the console under an account with the local Administrator permissions on the machine where the console is installed. In other cases, you can run it under an account that is a member of the Users group on the machine where the console is installed. However, you may require additional permissions to recover guest OS files of Microsoft Windows VMs.<br><br>[For recovery of Microsoft Windows VM guest OS files] If you plan to save files to a new location, the user who launched the Veeam Backup & Replication console does not have permissions to read and write data to the new location, and the mount point is located on the same machine as the Veeam Backup & Replication console, check that the user has the *SeBackupPrivilege* and *SeRestorePrivilege*. For more information on where mount points are created, see Mount Points and Restore Scenarios.<br><br>Accounts that are members of the Protected Users Active Directory group cannot be used to access the backup server remotely over the Veeam Backup & Replication console. For more information, see Microsoft Docs. |

| Account | Required Permission |
| --- | --- |
| Veeam Backup & Replication Services | The account used to run the Veeam Backup & Replication Services must be the LocalSystem account or must have local Administrator permissions on the backup server.<br><br>These privileges are required because the services need access to the Veeam registry configuration, which is available only to privileged users. Some services also perform disk management operations, open ports in firewalls and perform other operations that require a high permission level. |
| **Microsoft SQL Server (where the configuration database is stored)** | You require different sets of Microsoft SQL permissions in the following cases:<br>• Installation (remote or local): current account needs CREATE ANY DATABASE permission on the SQL server level. After database creation this account automatically gets a *db_owner* role and can perform all operations with the database. If the current account does not have this permission, a Database Administrator may create an empty database in advance and grant the *db_owner* role to the account that will be used for installing Veeam Backup & Replication.<br>• Upgrade: current account should have sufficient permissions for that database. To grant these permissions through role assignment, it is recommended that you use the account with *db_owner* role.<br>• Operation: the account used to run Veeam Backup Service requires *db_datareader* and *db_datawriter* roles as well as permissions to execute stored procedures for the configuration database on the Microsoft SQL Server. Alternatively, you can assign *db_owner* role for this database to the service account.<br><br>For more information, see Microsoft Docs. |
| **PostgreSQL** | The account used for installation, upgrade and operation requires *superuser* role. |

# Using Virtualization Servers and Hosts

The following permissions and roles are required to work with virtualization servers and hosts during data protection tasks.

| Component | Required Permission or Role |
|---|---|
| **Source/Target VMware vSphere Host** | Root permissions on the ESXi host. When adding the credentials, use the MACHINE\USER format for local accounts or DOMAIN\USER format for domain accounts.<br><br>If the vCenter Server is added to the backup infrastructure, an account that has administrative permissions is required.<br>You can either grant the Administrator role to the account or configure granular vCenter Server permissions for certain Veeam Backup & Replication operations in the VMware vSphere environment. For more information, see the Permissions Reference. |
| **VMware Cloud Director Server** | The account that you specify when adding a server must have system administrator privileges on VMware Cloud Director. You cannot use the organization administrator account to add the Cloud Director server. |
| **Source / Target Hyper-V host or cluster** | Administrator permissions. |
| **SCVMM** | Any SCVMM user. |
| **Windows Server added to the backup infrastructure** | The user account that you use to add a Microsoft Windows server must be in the local administrators group (on the server being added). When adding the credentials, use the MACHINE\USER format for local accounts or DOMAIN\USER format for domain accounts. |
| **Linux Server added to the backup infrastructure** | Permissions for the account that you specify when adding a Linux server differ depending on the role that you plan to assign to this server:<br>• Roles for which Veeam Data Movers must be persistent (backup proxy, hardened repository) require root or equivalent permissions. For the full list of roles, see Veeam Data Movers.<br>• Gateway server that communicates with NFS share requires root or equivalent permissions.<br>• Backup repository requires read and write permissions on the folder where backups will be stored. You will configure this folder at the Configure Backup Repository Settings step of the backup repository wizard.<br>• Other roles require read and write permissions on files and folders with which the server will work. |
| **SMB Backup Repository** | Read and write permission on the target folder and share. |

| Component | Required Permission or Role |
|---|---|
| **Dell Data Domain Deduplicating Storage Appliance** | DD Boost User. To specify the DD Boost User account settings, in Data Domain System Manager, open the **Data Management** > **DD Boost Settings** tab. |
| **HPE StoreOnce Deduplicating Storage Appliance** | Permissions on a Catalyst store where backup data will be kept. To check the client account permissions, in the HPE StoreOnce management console, select the Catalyst store and open the **Permissions** tab for it. |

# Performing Guest Processing

To use guest OS processing (application-aware processing, pre-freeze and post-thaw scripts, transaction log processing, guest file indexing and file exclusions), make sure to configure your accounts according to the requirements listed in this section. For more information on guest processing, see Guest Processing.

All user accounts used for guest processing of Windows VMs must have the following permissions:

- *Logon as a batch job* granted
- *Deny logon as a batch job* not set

Group Managed Service Accounts (gMSAs) must also have the following permissions:

- *Logon as a service* granted.
- *Deny logon as a service* not set.

Other permissions depend on applications that you back up. You can find permissions for backup operations in the following table. For restore operation permissions, see **Permissions** sections in the Veeam Explorers User Guide.

| Application | Required Permission |
|---|---|
| **Microsoft SQL Server** | To back up Microsoft SQL Server data, the user whose account you plan to use must be:<br>• Local Administrator on the target VM.<br>• System administrator (has the *Sysadmin* role) on the target Microsoft SQL Server.<br><br>If you need to provide minimal permissions, the account must be assigned the following roles and permissions:<br>• SQL Server instance-level role: *public* and *dbcreator*.<br>• Database-level roles and roles for the model system database: *db_backupoperator*, *db_denydatareader*, *public*;<br>for the master system database – *db_backupoperator*, *db_datareader*, *public*;<br>for the msdb system database – *db_backupoperator*, *db_datareader*, *public*, *db_datawriter*.<br>• Securables: *view any definition*, *view server state*, *connect SQL*. |

| Application | Required Permission |
| --- | --- |
| **Microsoft Active Directory** | To back up Microsoft Active Directory data, the account must be a member of the built-in *Administrators* group. |
| **Microsoft Exchange** | To back up Microsoft Exchange data, the account must have the local Administrator permissions on the machine where Microsoft Exchange is installed. |
| **Oracle** | The account specified at the **Guest Processing** step must be configured in the following way:<br>• For a Windows-based VM, the account must be a member of both the *Local Administrator* group and the *ORA_DBA* group (if OS authentication is used). In addition, if *ASM* is used, then such an account must be a member of the *ORA_ASMADMIN* group (for Oracle 12 and higher).<br>• For a Linux-based VM, the account must be a Linux user elevated to *root*. The account must have the `home` directory created.<br>To back up Oracle databases, you can specify the following options at the Oracle tab:<br>• Oracle account with SYSDBA privileges.<br>You can use, for example, the SYS Oracle account or any other Oracle account that has been granted SYSDBA privileges.<br>• Account specified for guest processing. That is, the **Use guest credentials** option selected.<br>In this case, the account that was specified at the **Guest Processing** step must be a member of the *ORA_DBA* group for a Windows-based VM and *OSASM*, *OSDBA* and *OINSTALL* groups for a Linux-based VM.<br>To perform guest processing for Oracle databases on Linux servers, make sure that the `/tmp` directory is mounted with the `exec` option. Otherwise, you will get an error with the permission denial. |
| **Microsoft SharePoint** | To back up a Microsoft SharePoint Server, the account must have the Farm Administrator role.<br>To back up Microsoft SQL databases of the Microsoft SharePoint Server, the account must have permissions required for Microsoft SQL Server backup operations. |

| Application | Required Permission |
| --- | --- |
| **PostgreSQL** | The account specified at the **Guest Processing** step must be a Linux user elevated to *root*. The account must have the home directory created.<br><br>For the directory specified as the temporary location for archive logs, the following permissions must be granted:<br>• The user running the PostgreSQL instance must have read, write, and execute (`rwx`) permissions.<br>• The user selected in the backup job settings must have read and execute (`rx`) permissions.<br><br>For more information, see PostgreSQL WAL Files Settings.<br><br>To back up PostgreSQL instances, the account must have the superuser privileges for the PostgreSQL instance. For more information, see PostgreSQL documentation.<br><br>**Note**: If you back up data using vSphere API, the account specified at the **Guest Processing** step must be a *root* Linux user. |

Consider the following general requirements when choosing a user account:

- [For guest OS file indexing] For Windows-based workloads, choose an account that has administrator privileges. For Linux-based workloads, choose an account of a root user or user elevated to root.
- To use networkless guest processing over VMware VIX/vSphere Web Services, you must specify one of the following accounts at the **Guest Processing** step of the backup wizard. Check that the account also has permissions listed in the table.
  - If Windows User Account Control (UAC) is enabled, specify Local Administrator (*MACHINE|Administrator*) or Domain Administrator (*DOMAIN|Administrator*) account.
  - If UAC is disabled, specify an account that is a member of the built-in Administrators group.
  - For Linux-based VMs, specify a root account.
- [For networkless guest processing over VMware VIX] To be able to perform more than 1000 guest processing operations, the user that you specify for guest processing must be logged into the VM at least once.
- [If you plan to use guest processing over network for workloads without listed applications] For Windows-based workloads, choose an account that has administrator privileges. For Linux-based workloads, choose an account of a root user or user elevated to root.
- When using Active Directory accounts, make sure to provide an account in the *DOMAIN|Username* format.
- When using local user accounts, make sure to provide an account in the *Username or HOST|Username* format.
- To process a Domain Controller server, make sure that you are using an account that is a member of the *DOMAIN|Administrators* group.
- To back up a Read-Only Domain controller, a delegated RODC administrator account is sufficient. For more information, see Microsoft Docs.

# Restoring to Amazon EC2

To restore workloads to Amazon EC2, it is recommended that the IAM user whose credentials you plan to use to connect to AWS has administrative permissions — access to all AWS actions and resources.

If you do not want to provide full access to AWS, you can grant to the IAM user a minimal set of permissions that will be sufficient for restore. For more information, see AWS IAM User Permissions.

# Adding Microsoft Azure Compute Accounts

Microsoft Azure Compute account is required to restore workloads to Microsoft Azure, add Azure archive storage and so on. For more information, see Microsoft Azure Compute Accounts.

The following permissions are required for adding a Microsoft Azure Compute account.

| Microsoft Entra ID (formerly Azure Active Directory) Application | Permissions |
|---|---|
| New<br>(select the **Create a new account** option at the Subscription step of the wizard) | The Microsoft Entra ID user account where the application will be created must have the following privileges:<br>• To register applications. For this, you can assign the *Global Administrator* privileges to the user or enable the **Users can register applications** option for the user in Azure portal. For details, see Microsoft Azure Docs.<br>• To assign a role on the subscription level for the registered application. For this, you can use the *Owner** role or if the *Owner* role cannot be used, you can create a custom role with minimal permissions. To learn how to create a custom role, see Creating Custom Role for Azure and Azure Stack Hub Accounts.<br>* - When you assign a privileged role, like *Owner,* using the Azure Portal, the default conditions can be added to this role assignment. With these default conditions, adding the Compute account will fail. To avoid this issue, select the **Allow user to assign all roles (highly privileged)** option. For more information on conditions, see Microsoft Docs. |

| Microsoft Entra ID (formerly Azure Active Directory) Application | Permissions |
| --- | --- |
| Existing<br>(select the **Use the existing account** option at the Subscription step of the wizard) | The application must have the *Contributor* role privilege for the selected subscription. If you restore workloads to Microsoft Azure Stack Hub and cannot use the *Contributor* role, you can create a custom role with minimal permissions. To learn how to create a custom role, see Creating Custom Role for Azure and Azure Stack Hub Accounts. |

# Adding Microsoft Azure Stack Hub Compute Accounts

A Microsoft Azure Stack Hub Compute account is required to restore workloads to Microsoft Azure Stack Hub. For more information, see Microsoft Azure Stack Hub Compute Accounts.

The following permissions are required to add a Microsoft Azure Stack Hub Compute account.

| Microsoft Entra ID (formerly Azure Active Directory) Application | Permissions |
| --- | --- |
| New<br>(select the **Create a new account** option at the Subscription step of the wizard) | The Microsoft Entra ID user account where the application will be created must have the following privileges:<br>• To register applications. For this, you can assign the *Global Administrator* privileges to the user or enable the **Users can register applications** option for the user in Azure portal. For details, see Microsoft Azure Docs.<br>• To assign a role on the subscription level for the registered application. For this, you can use the *Owner** role or if the *Owner* role cannot be used, you can create a custom role with minimal permissions. To learn how to create a custom role, see Creating Custom Role for Azure and Azure Stack Hub Accounts.<br>* - When you assign a privileged role, like *Owner,* using the Azure Portal, the default conditions can be added to this role assignment. With these default conditions, adding the Compute account will fail. To avoid this issue, select the **Allow user to assign all roles (highly privileged)** option. For more information on conditions, see Microsoft Docs. |

| Microsoft Entra ID (formerly Azure Active Directory) Application | Permissions |
| --- | --- |
| Existing<br>(select the **Use the existing account** option at the Subscription step of the wizard) | The application must have the *Contributor* role privilege for the selected subscription. If you restore workloads to Microsoft Azure Stack Hub and cannot use the *Contributor* role, you can create a custom role with minimal permissions. To learn how to create a custom role, see Creating Custom Role for Azure and Azure Stack Hub Accounts. |

# Adding Microsoft Azure Storage Accounts (Entra ID)

Microsoft Azure storage account with Microsoft Entra ID authorization is required to add Azure Blob storage, Azure archive storage, and so on. For more information, see Microsoft Azure Storage Accounts (Entra ID).

The following permissions are required to existing Microsoft Azure storage account with Microsoft Entra ID authorization:

- If you use a new Microsoft Entra ID (formerly Azure Active Directory) application (select the **Create a new account** option at the Subscription step of the wizard) when adding a Microsoft Azure Stack Hub Compute account, the Microsoft Entra ID user account where the Microsoft Entra ID application will be created must have the following privileges:

  a. To register applications. For this, you can assign the *Global Administrator* privileges to the user or enable the **Users can register applications** option for the user in Azure portal. For details, see Microsoft Azure Docs.

  b. To assign a role on the subscription level for the registered application. For this, you can use the *Owner* role or if the *Owner* role cannot be used, you can create a custom role with minimal permissions. To learn how to create a custom role, see Creating Custom Role for Azure and Azure Stack Hub Accounts.

- If you use an existing Microsoft Entra ID (formerly Azure Active Directory) application (select the **Use the existing account** option at the Account Type step of the wizard), the application must have the following role privileges for the selected storage account:

  - Storage Account Contributor
  - Storage Blob Data Contributor
  - Storage Blob Data Owner

  For more information, see Microsoft Docs.

# Using Object Storage Repositories

## General Amazon S3 and S3 Compatible Object Storage Permissions

Consider the following:

- Make sure the account you are using has access to Amazon buckets and folders.
- The `ListAllMyBuckets` permission is not required if you specify the bucket name explicitly at the **Bucket** step of the New Object Repository wizard.

Permissions for Amazon S3 or S3 compatible object storage depend on whether you use immutability and helper appliance settings.

> **NOTE**
>
> Consider the following:
>
> - S3 compatible object storage repositories use the same permissions as Amazon S3 object storage with the following exception: since you cannot setup helper appliance in the cloud, you do not need permissions for it. Therefore, S3 compatible object storage repositories require permissions which start with s3, for example, "`s3:ListBucket`". Permissions that start with ec2 can be skipped, for example, "`ec2:DescribeInstances`".
> - To deploy S3 compatible object storage in multiple bucket mode, you must add the `s3:CreateBucket` and `s3:DeleteBucket` permissions to the list of permissions that you are going to use.

### 1. Immutability Disabled and Helper Appliance not Used

The following permissions are required to use Amazon S3 or S3 compatible object storage with immutability disabled. A helper appliance is not used for health check operations.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
              "s3:ListBucket",
              "s3:GetBucketLocation",
              "s3:GetObject",
              "s3:PutObject",
              "s3:DeleteObject",
              "s3:ListAllMyBuckets",
              "s3:GetBucketVersioning",
             "s3:GetBucketObjectLockConfiguration"
            ],
            "Resource": "*"
        }
    ]
}
```

### 2. Immutability Disabled and New Helper Appliance Configured

The following permissions are required to use Amazon S3 or S3 compatible object storage with immutability disabled. For health check operations a new helper appliance is configured and the Amazon VPC, subnet and security group settings are set to *(Create new)* for the helper appliance settings.

```
{
   "Version": "2012-10-17",
   "Statement": [
   {
        "Effect": "Allow",
        "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation",
        "s3:GetObject",
        "s3:PutObject",
        "s3:DeleteObject",
        "s3:ListAllMyBuckets",
        "s3:GetBucketVersioning",
        "s3:GetBucketObjectLockConfiguration",
        "ec2:DescribeInstances",
        "ec2:CreateKeyPair",
        "ec2:DescribeKeyPairs",
        "ec2:RunInstances",
        "ec2:DeleteKeyPair",
        "ec2:DescribeVpcAttribute",
        "ec2:CreateTags",
        "ec2:DescribeSubnets",
        "ec2:TerminateInstances",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeImages",
        "ec2:DescribeVpcs",
        "ec2:CreateVpc",
        "ec2:CreateSubnet",
        "ec2:DescribeAvailabilityZones",
        "ec2:CreateRoute",
        "ec2:CreateInternetGateway",
        "ec2:AttachInternetGateway",
        "ec2:ModifyVpcAttribute",
        "ec2:CreateSecurityGroup",
        "ec2:DeleteSecurityGroup",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:DescribeRouteTables",
        "ec2:DescribeInstanceTypes"
            ],
        "Resource": "*"
        }
   ]
}
```

### 3. Immutability Disabled and Helper Appliance Configured Beforehand

The following permissions are required to use Amazon S3 or S3 compatible object storage with immutability disabled. Amazon VPC, subnet and security group settings for a helper appliance are configured beforehand.

```
{
   "Version": "2012-10-17",
   "Statement": [
   {
       "Effect": "Allow",
       "Action": [
       "s3:ListBucket",
       "s3:GetBucketLocation",
       "s3:GetObject",
       "s3:PutObject",
       "s3:DeleteObject",
       "s3:ListAllMyBuckets",
       "s3:GetBucketVersioning",
       "s3:GetBucketObjectLockConfiguration",
       "ec2:DescribeInstances",
       "ec2:CreateKeyPair",
       "ec2:DescribeKeyPairs",
       "ec2:RunInstances",
       "ec2:DeleteKeyPair",
       "ec2:DescribeVpcAttribute",
       "ec2:CreateTags",
       "ec2:DescribeSubnets",
       "ec2:TerminateInstances",
       "ec2:DescribeSecurityGroups",
       "ec2:DescribeImages",
       "ec2:DescribeVpcs"
           ],
       "Resource": "*"
       }
   ]
}
```

### 4. Immutability Enabled and Helper Appliance not Used

The following permissions are required to use Amazon S3 or S3 compatible object storage with immutability enabled. A helper appliance is not used for health check operations.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "s3:ListBucket",
                "s3:GetBucketLocation",
                "s3:GetObject",
                "s3:PutObject",
                "s3:DeleteObject",
                "s3:ListAllMyBuckets",
                "s3:GetBucketVersioning",
                "s3:GetBucketObjectLockConfiguration",
                "s3:ListBucketVersions",
                "s3:GetObjectVersion",
                "s3:GetObjectRetention",
                "s3:GetObjectLegalHold",
                "s3:PutObjectRetention",
                "s3:PutObjectLegalHold",
                "s3:DeleteObjectVersion"
                    ],
            "Resource": "*"
        }
    ]
}
```

### 5. Immutability Enabled and New Helper Appliance Configured

The following permissions are required to use Amazon S3 or S3 compatible object storage with immutability enabled. For health check operations a new helper appliance is configured and the Amazon VPC, subnet and security group settings are set to *(Create new)* for the helper appliance settings.

```
{
   "Version": "2012-10-17",
   "Statement": [
   {
        "Effect": "Allow",
        "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation",
        "s3:GetObject",
        "s3:PutObject",
        "s3:DeleteObject",
        "s3:ListAllMyBuckets",
        "s3:GetBucketVersioning",
        "s3:GetBucketObjectLockConfiguration",
        "s3:ListBucketVersions",
        "s3:GetObjectVersion",
        "s3:GetObjectRetention",
        "s3:GetObjectLegalHold",
        "s3:PutObjectRetention",
        "s3:PutObjectLegalHold",
        "s3:DeleteObjectVersion",
        "ec2:DescribeInstances",
        "ec2:CreateKeyPair",
        "ec2:DescribeKeyPairs",
        "ec2:RunInstances",
        "ec2:DeleteKeyPair",
        "ec2:DescribeVpcAttribute",
        "ec2:CreateTags",
        "ec2:DescribeSubnets",
        "ec2:TerminateInstances",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeImages",
        "ec2:DescribeVpcs",
        "ec2:CreateVpc",
        "ec2:CreateSubnet",
        "ec2:DescribeAvailabilityZones",
        "ec2:CreateRoute",
        "ec2:CreateInternetGateway",
        "ec2:AttachInternetGateway",
        "ec2:ModifyVpcAttribute",
        "ec2:CreateSecurityGroup",
        "ec2:DeleteSecurityGroup",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:DescribeRouteTables",
        "ec2:DescribeInstanceTypes"
            ],
        "Resource": "*"
        }
   ]
}
```

### 6. Immutability Enabled and Helper Appliance Configured Beforehand

The following permissions are required to use Amazon S3 or S3 compatible object storage with immutability enabled. Amazon VPC, subnet and security group settings for a helper appliance are configured beforehand.

```
{
    "Version": "2012-10-17",
    "Statement": [
    {
        "Effect": "Allow",
        "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation",
        "s3:GetObject",
        "s3:PutObject",
        "s3:DeleteObject",
        "s3:ListAllMyBuckets",
        "s3:GetBucketVersioning",
        "s3:GetBucketObjectLockConfiguration",
        "s3:ListBucketVersions",
        "s3:GetObjectVersion",
        "s3:GetObjectRetention",
        "s3:GetObjectLegalHold",
        "s3:PutObjectRetention",
        "s3:PutObjectLegalHold",
        "s3:DeleteObjectVersion",
        "ec2:DescribeInstances",
        "ec2:CreateKeyPair",
        "ec2:DescribeKeyPairs",
        "ec2:RunInstances",
        "ec2:DeleteKeyPair",
        "ec2:DescribeVpcAttribute",
        "ec2:CreateTags",
        "ec2:DescribeSubnets",
        "ec2:TerminateInstances",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeImages",
        "ec2:DescribeVpcs"
            ],
        "Resource": "*"
        }
    ]
}
```

For more information on permissions, see AWS Documentation.

## Amazon S3 Glacier and S3 Compatible with Data Archiving Object Storage Permissions

Permissions for Amazon S3 Glacier and S3 Compatible with data archiving depend on whether you use immutability and the archiver appliance settings:

### 1. Immutability Disabled and Archiver Appliance not Configured

The following permissions are required for Amazon S3 Glacier and S3 Compatible with data archiving object storage with immutability disabled. VPC, subnet and security group settings set to are set to *(Create new)* for the archiver appliance settings.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "s3:DeleteObject",
        "s3:PutObject",
        "s3:GetObject",
        "s3:RestoreObject",
        "s3:ListBucket",
        "s3:AbortMultipartUpload",
        "s3:GetBucketVersioning",
        "s3:ListAllMyBuckets",
        "s3:GetBucketLocation",
        "s3:GetBucketObjectLockConfiguration",
        "ec2:DescribeInstances",
        "ec2:CreateKeyPair",
        "ec2:DescribeKeyPairs",
        "ec2:RunInstances",
        "ec2:DeleteKeyPair",
        "ec2:DescribeVpcAttribute",
        "ec2:CreateTags",
        "ec2:DescribeSubnets",
        "ec2:TerminateInstances",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeImages",
        "ec2:DescribeVpcs",
        "ec2:CreateVpc",
        "ec2:CreateSubnet",
        "ec2:DescribeAvailabilityZones",
        "ec2:CreateRoute",
        "ec2:CreateInternetGateway",
        "ec2:AttachInternetGateway",
        "ec2:ModifyVpcAttribute",
        "ec2:CreateSecurityGroup",
        "ec2:DeleteSecurityGroup",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:DescribeRouteTables",
        "ec2:DescribeInstanceTypes"
      ],
      "Resource": "*"
    }
  ]
}
```

### 2. Immutability Disabled and Archiver Appliance Configured Beforehand

These permissions apply for Amazon S3 Glacier and S3 Compatible with data archiving object storage with immutability disabled. Amazon VPC, subnet and security group settings for an archiver appliance are configured beforehand.

```
{
 "Version": "2012-10-17",
 "Statement": [
   {
     "Sid": "VisualEditor0",
     "Effect": "Allow",
     "Action": [
       "s3:DeleteObject",
       "s3:PutObject",
       "s3:GetObject",
       "s3:RestoreObject",
       "s3:ListBucket",
       "s3:AbortMultipartUpload",
       "s3:GetBucketVersioning",
       "s3:ListAllMyBuckets",
       "s3:GetBucketLocation",
       "s3:GetBucketObjectLockConfiguration",
       "ec2:DescribeInstances",
       "ec2:CreateKeyPair",
       "ec2:DescribeKeyPairs",
       "ec2:RunInstances",
       "ec2:DeleteKeyPair",
       "ec2:DescribeVpcAttribute",
       "ec2:CreateTags",
       "ec2:DescribeSubnets",
       "ec2:TerminateInstances",
       "ec2:DescribeSecurityGroups",
       "ec2:DescribeImages",
       "ec2:DescribeVpcs"
     ],
     "Resource": "*"
   }
 ]
}
```

### 3. Immutability Enabled and Archiver Appliance not Configured

The following permissions are required for Amazon S3 Glacier and S3 Compatible with data archiving object storage with immutability enabled. VPC, subnet and security group settings set to are set to *(Create new)* for the archiver appliance settings.

```
{
 "Version": "2012-10-17",
 "Statement": [
   {
     "Sid": "VisualEditor0",
     "Effect": "Allow",
     "Action": [
       "s3:DeleteObject",
       "s3:PutObject",
       "s3:GetObject",
       "s3:RestoreObject",
       "s3:ListBucket",
       "s3:AbortMultipartUpload",
       "s3:GetBucketVersioning",
       "s3:ListAllMyBuckets",
       "s3:GetBucketLocation",
       "s3:GetBucketObjectLockConfiguration",
       "s3:PutObjectRetention",
       "s3:GetObjectVersion",
       "s3:PutObjectLegalHold",
       "s3:GetObjectRetention",
       "s3:DeleteObjectVersion",
       "s3:ListBucketVersions",
       "ec2:DescribeInstances",
       "ec2:CreateKeyPair",
       "ec2:DescribeKeyPairs",
       "ec2:RunInstances",
       "ec2:DeleteKeyPair",
       "ec2:DescribeVpcAttribute",
       "ec2:CreateTags",
       "ec2:DescribeSubnets",
       "ec2:TerminateInstances",
       "ec2:DescribeSecurityGroups",
       "ec2:DescribeImages",
       "ec2:DescribeVpcs",
       "ec2:CreateVpc",
       "ec2:CreateSubnet",
       "ec2:DescribeAvailabilityZones",
       "ec2:CreateRoute",
       "ec2:CreateInternetGateway",
       "ec2:AttachInternetGateway",
       "ec2:ModifyVpcAttribute",
       "ec2:CreateSecurityGroup",
       "ec2:DeleteSecurityGroup",
       "ec2:AuthorizeSecurityGroupIngress",
       "ec2:AuthorizeSecurityGroupEgress",
       "ec2:DescribeRouteTables",
       "ec2:DescribeInstanceTypes"
     ],
     "Resource": "*"
   }
 ]
}
```

### 4. Immutability Enabled and Archiver Appliance Configured Beforehand

These permissions apply for Amazon S3 Glacier and S3 Compatible with data archiving object storage with immutability enabled. Amazon VPC, subnet and security group settings for an archiver appliance are configured beforehand.

```
{
 "Version": "2012-10-17",
 "Statement": [
   {
     "Sid": "VisualEditor0",
     "Effect": "Allow",
     "Action": [
       "s3:DeleteObject",
       "s3:PutObject",
       "s3:GetObject",
       "s3:RestoreObject",
       "s3:ListBucket",
       "s3:AbortMultipartUpload",
       "s3:GetBucketVersioning",
       "s3:ListAllMyBuckets",
       "s3:GetBucketLocation",
       "s3:GetBucketObjectLockConfiguration",
       "s3:PutObjectRetention",
       "s3:GetObjectVersion",
       "s3:PutObjectLegalHold",
       "s3:GetObjectRetention",
       "s3:DeleteObjectVersion",
       "s3:ListBucketVersions",
       "ec2:DescribeInstances",
       "ec2:CreateKeyPair",
       "ec2:DescribeKeyPairs",
       "ec2:RunInstances",
       "ec2:DeleteKeyPair",
       "ec2:DescribeVpcAttribute",
       "ec2:CreateTags",
       "ec2:DescribeSubnets",
       "ec2:TerminateInstances",
       "ec2:DescribeSecurityGroups",
       "ec2:DescribeImages",
       "ec2:DescribeVpcs"
     ],
     "Resource": "*"
   }
 ]
}
```

## Read-Only Permissions for Amazon S3 and S3 Compatible Object Storage

The following permissions are required to connect to Amazon S3 or S3 compatible object storage from the second backup server in case you need to perform data recovery operations.

> **IMPORTANT**
>
> Read-only permissions are not supported for S3 compatible object storage with multiple buckets.

### 1. Immutability Enabled

These permissions apply for Amazon S3 or S3 compatible object storage with immutability enabled.

```
{
 "Version": "2012-10-17",
 "Statement": [
   {
     "Effect": "Allow",
     "Action": [
         "s3:ListBucket",
         "s3:GetBucketLocation",
         "s3:GetObject",
         "s3:ListAllMyBuckets",
         "s3:GetBucketVersioning",
         "s3:GetBucketObjectLockConfiguration",
         "s3:ListBucketVersions",
         "s3:GetObjectVersion",
         "s3:GetObjectRetention",
         "s3:GetObjectLegalHold"
     ],
     "Resource": "*"
   }
 ]
}
```

### 2. Immutability Disabled

These permissions apply for Amazon S3 or S3 compatible object storage with immutability disabled.

```
{
 "Version": "2012-10-17",
 "Statement": [
   {
     "Effect": "Allow",
     "Action": [
         "s3:ListBucket",
         "s3:GetBucketLocation",
         "s3:GetObject",
         "s3:ListAllMyBuckets",
         "s3:GetBucketVersioning",
         "s3:GetBucketObjectLockConfiguration"
     ],
     "Resource": "*"
   }
 ]
}
```

## S3 Batch Retrieval Permissions for Amazon S3 Glacier

The following permissions are required to perform standard accelerated retrieval that uses S3 Batch Operations:

```
"iam:GetRole",
"iam:CreateRole",
"iam:PutRolePolicy",
"iam:DeleteRolePolicy",
"iam:DeleteRole",
"iam:PassRole",
"s3:CreateJob",
"s3:DescribeJob”
```

## Permissions for Amazon EBS Encryption

If you have EBS encryption enabled in Amazon EC2, these permissions are necessary for a helper appliance to perform a health check of backup files in object storage and run the archiving job. You must grant the following permissions for the default encryption key to an IAM Role using a Key Policy.

```
"kms:Encrypt",
"kms:Decrypt",
"kms:ReEncrypt*",
"kms:GenerateDataKey*",
"kms:DescribeKey",
"kms:CreateGrant"
```

## Azure Archive Object Storage Permissions

The following permissions are required to use Azure Archive object storage.

```
{
 "properties": {
   "roleName": "CUSTOM_ROLE_MINIMAL_PERMISSIONS",
   "description": "CUSTOM_ROLE_MINIMAL_PERMISSIONS",
   "assignableScopes": [
     "/subscriptions/111111-1111-1111-0000-00000000000"
   ],
   "permissions": [
     {
       "actions": [
         "Microsoft.Authorization/*/read",
         "Microsoft.Compute/locations/*",
         "Microsoft.Compute/virtualMachines/*",
         "Microsoft.Compute/disks/read",
         "Microsoft.Compute/disks/write",
         "Microsoft.Compute/disks/delete",
         "Microsoft.Network/locations/*",
         "Microsoft.Network/networkInterfaces/*",
         "Microsoft.Network/networkSecurityGroups/join/action",
         "Microsoft.Network/networkSecurityGroups/read",
         "Microsoft.Network/networkSecurityGroups/write",
         "Microsoft.Network/networkSecurityGroups/delete",
         "Microsoft.Network/publicIPAddresses/join/action",
         "Microsoft.Network/publicIPAddresses/read",
         "Microsoft.Network/publicIPAddresses/write",
         "Microsoft.Network/publicIPAddresses/delete",
         "Microsoft.Network/virtualNetworks/read",
         "Microsoft.Network/virtualNetworks/write",
         "Microsoft.Network/virtualNetworks/subnets/join/action",
         "Microsoft.Storage/storageAccounts/listKeys/action",
         "Microsoft.Storage/storageAccounts/read",
         "Microsoft.Resources/deployments/*",
         "Microsoft.Resources/subscriptions/resourceGroups/read",
         "Microsoft.Resources/checkResourceName/action",
         "Microsoft.Resources/subscriptions/resourceGroups/write",
         "Microsoft.Resources/subscriptions/locations/read"
          ],
       "notActions": [],
       "dataActions": [],
       "notDataActions": []
     }
   ]
 }
}
```

## Google Cloud Object Storage Permissions

The following permissions are required to use Google Cloud Object Storage as an object storage repository.

### General Permissions to Add Google Cloud Object Storage

The following permissions are required to use Google Cloud object storage.

```
{
"storage.buckets.get",
"storage.buckets.list",
"storage.objects.create",
"storage.objects.delete",
"storage.objects.get",
"storage.objects.list"
}
```

Consider the following:

- The `storage.buckets.list` permission is not required if you specify the bucket name explicitly at the **Bucket** step of the New Object Repository wizard.
- The Owner IAM role does not necessarily grant the permissions required for working with Google Cloud Storage.

### Read-Only Permissions for Google Cloud Object Storage

The following permissions are required to connect to Google Cloud object storage from the second backup server in case you need to perform data recovery operations.

```
{
"storage.buckets.get",
"storage.buckets.list",
"storage.objects.get",
"storage.objects.list"
}
```

## Adding Object Storage as Unstructured Data Source

The following permissions are required for the account that you use to add Amazon S3 and S3 Compatible object storage as unstructured sources.

### 1. Permissions for Working with Bucket Objects

To be able to work with objects in buckets, the account for Amazon S3 and S3 Compatible object storage must have the following permissions:

```
HeadBucket
GetBucketLocation
ListBuckets
HeadObject
GetObject
GetObjectTagging
ListObjectsV1
ListObjectsV2
ListObjectVersions
```

### 2. Permissions for Restoring Objects and Perform Backup Objects Calls

To be able to perform restore, recovery, and backup object calls, the account for Amazon S3 and S3 Compatible object storage must have the following permissions:

```
PutObject
PutObjectTagging
DeleteObject
DeleteObjects
CreateMultipartUpload
CreateBucket
UploadPart
CompleteMultipartUpload
AbortMultipartUpload
```

### 3. Permissions for Getting Backup Bucket Properties

To be able to get backup bucket properties, the account for Amazon S3 and S3 Compatible object storage must have the following permissions:

```
GetBucketLifecycleConfiguration
GetBucketLogging
GetBucketMetricsConfiguration
ListBucketMetricsConfigurations
GetBucketNotificationConfiguration
GetBucketOwnershipControls
GetBucketPolicy
GetPublicAccessBlock
GetBucketReplication
GetBucketRequestPayment
GetBucketTagging
GetBucketVersioning
GetBucketWebsite
GetObjectLockConfiguration
```

# Storage Systems Snapshot Integration

## NetApp Data ONTAP/Lenovo Thinksystem DM/DG Permissions

The account used to connect to a NetApp ONTAP, IBM N, Fujitsu ETERNUS HX/AX, Lenovo ThinkSystem DM/DG storage system must have permissions described in this section. The commands are provided for the console, UI names may differ.

## 7-Mode

- login-http-admin
- api-system-*
- api-license-* (api-license-list-info)
- api-volume-*
- api-net-*
- api-options-*
- api-vfiler-*
- api-qtree-*
- api-nfs-*
- api-snapshot-*
- api-lun-*
- api-iscsi-*
- api-feature-*
- api-registry-*
- api-fcp-*
- api-file-*
- api-igroup-*
- api-clone-*
- api-snapvault-*
- api-snapmirror-*
- api-cf-*
- cli-options
- security-api-vfiler

## CDOT (VMware Integration)

| Command/Directory | Access/Query Level |
|---|---|
| DEFAULT | readonly |
| cluster | readonly |
| metrocluster | readonly |
| vserver fcp | readonly |
| volume file | readonly |
| lun igroup | all |
| vserver iscsi | all |
| network | readonly |
| system node | readonly |
| security | readonly |
| security login | readonly |
| set | readonly |
| snapmirror | all |
| system | readonly |
| version | readonly |
| volume qtree | readonly |
| lun | all |
| vserver nfs | all |

| Command/Directory | Access/Query Level |
| --- | --- |
| volume snapshot | all |
| volume | all |
| vserver | all |

### Only as SVM (VMware Integration)

| Command/Directory | Access/Query Level |
| --- | --- |
| DEFAULT | none |
| lun | all |
| lun igroup | all |
| network | readonly |
| security | readonly |
| security login | readonly |
| snapmirror | all |
| system | readonly |
| version | readonly |
| volume | all |
| volume file | readonly |
| volume qtree | all |
| volume snapshot | all |

| Command/Directory | Access/Query Level |
| --- | --- |
| vserver | all |
| vserver fcp | all |
| vserver iscsi | all |
| vserver nfs | all |

### CDOT (NAS Backup Integration)

| Command/Directory | Access/Query Level |
| --- | --- |
| DEFAULT | readonly |
| security | readonly |
| security login | readonly |
| volume snapshot | all |
| vserver | all |
| vserver nfs | all |

### Only as SVM (NAS Backup Integration)

| Command/Directory | Access/Query Level |
| --- | --- |
| DEFAULT | none |
| lun | readonly |
| network | readonly |
| security | readonly |

| Command/Directory | Access/Query Level |
| --- | --- |
| security login | readonly |
| snapmirror | readonly |
| version | readonly |
| volume | readonly |
| volume snapshot | all |
| vserver | all |

## CDOT (Veeam Agent Integration)

| Command/Directory | Access/Query Level |
| --- | --- |
| cluster | readonly |
| lun | all |
| metrocluster | readonly |
| network | readonly |
| system license | readonly |
| system node | readonly |
| version | readonly |
| volume | all |
| volume snapshot | all |
| vserver | all |

### Only as SVM (Veeam Agent Integration)

| Command/Directory | Access/Query Level |
|---|---|
| lun | all |
| network | readonly |
| version | readonly |
| volume | all |
| volume snapshot | all |
| vserver | all |

## Universal Storage API Integrated Systems Permissions

The account used to connect to a Universal Storage API integrated system must be assigned a necessary role in the storage system console and have a set of necessary permissions.

- For Dell PowerMax, the account must be assigned the Storage Administrator role.
- For Fujitsu ETERNUS AF and DX series, the account must be assigned the Software role.
- For NetApp SolidFire/HCI, the account must have the following permissions:
  - Volumes
  - Cluster Admins
- For Tintri IntelliFlash (formerly Western Digital IntelliFlash, Tegile), the account must be assigned the Veeam Admin Role.
- For DataCore, the account must have the following permissions:
  - General
  - Port
  - Host
  - Virtual disk
  - Snapshot
  - Physical disk
- For Hitachi VSP/VSP One Block, the account must be assigned the following roles:
  - Storage Administrator (View Only)
  - Storage Administrator (Provisioning)

    - Storage Administrator (Local Copy)
- For HPE XP, the account must be assigned the following roles:
    - Storage Administrator (View Only)
    - Storage Administrator (Provisioning)
    - Storage Administrator (Local Copy)
- For Dell PowerStore, the account must be assigned one of the following roles:
    - Administrator
    - Storage Administrator
    - Storage Operator
- For NEC Storage M Series, the account must be assigned the Administrator role.
- For NEC Storage V Series, the account must be assigned the following roles:
    - Storage Administrator (View Only)
    - Storage Administrator (Provisioning)
    - Storage Administrator (Local Copy)

For privileges required to integrate the unstructured data backup feature with Dell PowerScale (formerly Isilon), see Integration with Dell PowerScale in the Unstructured Data Backup section.

For storage systems not mentioned above, the account must have Administrator role.

# Ports

On backup infrastructure components, Veeam Backup & Replication automatically creates firewall rules for the required ports on Windows-based machines. If you are using a third-party firewall, these rules must be created manually. These rules allow components to communicate with each other.

> **IMPORTANT**
>
> Some Linux distributions also require firewall and security rules to be created manually. For details, see this Veeam KB article.

You can find the full list of the ports in this section.

## Backup Server

The following table describes network ports that must be opened to ensure proper communication of the backup server with backup infrastructure components.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| **Communication with Virtualization Servers** | | | | |
| Backup server | vCenter Server | TCP | 443 | Default port used for connections to vCenter Server.<br>**Note:** The backup server should have a direct connection to vCenter Server. HTTP/HTTPS proxy servers are not supported.<br>If you use VMware Cloud Director, make sure you open port 443 on underlying vCenter Servers. |
| | ESXi server | TCP | 443 | Default port used for connections to ESXi host.<br>This port is not required for VMware Cloud on AWS. |
| | | TCP | 902 | Port used for data transfer to ESXi host. It is also used during guest OS file recovery if you recover files from replicas.<br>This port is not required for VMware Cloud on AWS. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | VMware Cloud Director | TCP | 443 | Default port used for connections to VMware Cloud Director.<br>**Note:** The backup server should have a direct connection to VMware Cloud Director. HTTP/HTTPS proxy servers are not supported. |
| **Other Communications** | | | | |
| Backup server | PostgreSQL server hosting the Veeam Backup & Replication configuration database | TCP | 5432 | Port used for communication with PostgreSQL server on which the Veeam Backup & Replication configuration database is deployed. |
| | Microsoft SQL Server hosting the Veeam Backup & Replication configuration database | TCP | 1433 | Port used for communication with Microsoft SQL Server on which the Veeam Backup & Replication configuration database is deployed (if you use a Microsoft SQL Server default instance).<br>Additional ports may need to be open depending on your configuration. For more information, see Microsoft Docs. |
| | DNS server with forward/reverse name resolution of all backup servers | UDP | 53 | Port used for communication with the DNS Server. |
| | Veeam Update Notification Server | TCP | 443 | Default port used to download information about available updates from the Veeam Update Notification Server over HTTPS.<br>Veeam Update Notification Server endpoints:<br>• `dev.veeam.com`<br>• `vbrad.butler.veeam.com`<br>• `vbrce.butler.veeam.com` |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Veeam License Update Server | TCP | 443 | Default port used to automatically update license from the Veeam License Update Server over HTTPS. Veeam Threat Hunter and Veeam Data Cloud Vault also require this communication to work properly.<br>Veeam License Update Server endpoints:<br>• `vbr.butler.veeam.com`<br>• `autolk.veeam.com` |
| | | | 80 | Required for certificate validation when Veeam Backup & Replication connects to Veeam License Update Server to check if the new license is available and download it.<br>Certificate verification endpoints:<br>• `*.ss2.us`<br>• `*.amazontrust.com`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access** |
| | Certificate Revocation Lists | TCP | 80 or 443 | Veeam Backup & Replication requires access to the Certificate Revocation Lists (CRL) of the Certificate Authority (CA) that issued the certificate for each backup infrastructure component.<br>**Note**: The specific CRL endpoint that must be connected to depends on the CA that issued the certificate.<br>You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access** |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | KMS server | TCP | 5696 | Default port used for communication with Key Management System server. |
| | Veeam ONE Server | TCP | 2741 | Default port used for communication with Veeam ONE internal Web API.<br><br>Required for the Analytics view. For more information, see Configuring Analytics View. |
| | Veeam ONE Web Services | TCP | 1239 | Default port used by Veeam ONE Web Services.<br><br>Required for the Analytics view. For more information, see Configuring Analytics View. |
| | Backup server | TCP | 9501 (local) | Port used locally on the backup server for communication between Veeam Broker Service and Veeam services and components.<br><br>**Note**: Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |
| | Backup server | TCP | 6172 (local) | Port used to provide REST access to the Veeam Backup & Replication database.<br><br>**Note**: Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Backup server | TCP | 9393 (local) | Default port used by the Veeam Guest Catalog service for catalog replication. Can be customized during Veeam Backup & Replication installation.<br>**Note**: Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |
| Management client PC (remote access) | Backup server | TCP | 3389 | Default port used by Remote Desktop Services. If you use third-party solutions to connect to the backup server, other ports may need to be open. |
| REST client | Backup server | TCP | 9419 | Default port for communication with REST API service. |

# Backup & Replication Console

The following table describes network ports that must be opened to ensure proper communication with the Veeam Backup & Replication console.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Veeam Backup & Replication console | Backup server | TCP | 9392<br>9420 | Ports used by the Veeam Backup & Replication console to communicate with the backup server.<br>Note that both ports are required. |
| | | TCP | 9396 | Port used by the Veeam.Backup.UIService process for managing database connections. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 9401 | [Remote console only] Port used by the Veeam Backup & Replication console during Windows file-level recovery. Required to perform **Copy to** and **Mount to console** operations. |
| | | TCP | 10003 | [Remote console only] Port used by the Veeam Backup & Replication console to connect to the backup server only when managing the Veeam Cloud Connect infrastructure. |
| | Mount server | TCP | 2500 to 3300 | [Remote console only] Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br><br>This port is used if the mount server is not located on the console.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | Veeam AI Assistant (rest-ai.veeam.com) | TCP | 443 | Default port for communication with the Veeam AI Assistant service. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Nutanix AHV Plug-in for Veeam Backup & Replication | TCP | 8543 | Port used by the Veeam Backup & Replication console to communicate with Nutanix AHV Plug-in for Veeam Backup & Replication.<br><br>This port is opened by default after you install Veeam Backup & Replication. If you do not use the plug-in, it is still recommended that you keep this port open to speed up the Veeam Backup & Replication console loading time. |
| | oVirt KVM Plug-in for Veeam Backup & Replication | TCP | 8544 | Port used by the Veeam Backup & Replication console to communicate with oVirt KVM Plug-in for Veeam Backup & Replication.<br><br>This port is opened by default after you install Veeam Backup & Replication. If you do not use the plug-in, it is still recommended that you keep this port open to speed up the Veeam Backup & Replication console loading time. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Proxmox Virtual Environment Plug-in for Veeam Backup & Replication | TCP | 8545 | Port used by the Veeam Backup & Replication console to communicate with Proxmox Virtual Environment Plug-in for Veeam Backup & Replication.<br><br>This port is opened by default after you install Veeam Backup & Replication. If you do not use the plug-in, it is still recommended that you keep this port open to speed up the Veeam Backup & Replication console loading time. |
| | AWS Plug-in for Veeam Backup & Replication | TCP | 9402 | Port used by the Veeam Backup & Replication console to communicate with AWS Plug-in for Veeam Backup & Replication.<br><br>This port is opened by default after you install Veeam Backup & Replication. If you do not use the plug-in, it is still recommended that you keep this port open to speed up the Veeam Backup & Replication console loading time. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Google Cloud Plug-in for Veeam Backup & Replication | TCP | 9403 | Port used by the Veeam Backup & Replication console to communicate with Google Cloud Plug-in for Veeam Backup & Replication.<br><br>This port is opened by default after you install Veeam Backup & Replication. If you do not use the plug-in, it is still recommended that you keep this port open to speed up the Veeam Backup & Replication console loading time. |
| | Kasten Plug-in for Veeam Backup & Replication | TCP | 9404 | Port used by the Veeam Backup & Replication console to communicate with Kasten Plug-in for Veeam Backup & Replication.<br><br>This port is opened by default after you install Veeam Backup & Replication. If you do not use the plug-in, it is still recommended that you keep this port open to speed up the Veeam Backup & Replication console loading time. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Microsoft Azure Plug-in for Veeam Backup & Replication | TCP | 20443 | Port used by the Veeam Backup & Replication console to communicate with Microsoft Azure Plug-in for Veeam Backup & Replication.<br><br>This port is opened by default after you install Veeam Backup & Replication. If you do not use the plug-in, it is still recommended that you keep this port open to speed up the Veeam Backup & Replication console loading time. |

# Backup Proxy

The following table describes network ports that must be opened to ensure proper communication of backup proxies with other backup components. For more information about ports that must be opened between the backup proxy and specific backup repository, see Backup Repositories.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| **Communication with Backup Server** | | | | |
| Backup server | Backup proxy (Microsoft Windows) | TCP | 445<br>135 | Required for deploying Veeam Backup & Replication components. |
| | | TCP | 6160 | Default port used by Veeam Installer Service. |
| | | TCP | 6162 | Default port used by Veeam Data Mover Service. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br><br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |
| | Backup proxy (Linux) | TCP | 22 | Default SSH port used as a control channel. |
| | | TCP | 6160 | Default port used by Veeam Installer Service for Linux. |
| | | TCP | 6162 | Default port used by Veeam Data Mover Service.<br><br>You can specify a different port while adding the Linux server to the Veeam Backup & Replication infrastructure. Note that you can specify a different port only if there is no previously installed Veeam Data Mover on this Linux server. For more information, see Specify Credentials and SSH Settings. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Backup proxy | TCP | 2500 to 3300 | Default range of ports used as data transmission channels and for collecting log files. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6210 | Default port used by the Veeam Backup VSS Integration Service for taking a VSS snapshot during the SMB file share backup. |
| Backup proxy | Backup server | TCP | 2500 to 3300 | Default range of ports used for malware detection metadata transfer. |
| **Communication with Virtualization Servers** | | | | |
| Backup proxy | vCenter Server | TCP | 443 | Default VMware web service port that can be customized in vCenter settings. |
| | ESXi server | TCP | 902 | Default VMware port used for data transfer.<br>This port is not required for VMware Cloud on AWS. |
| | | TCP | 443 | Default VMware web service port that can be customized in ESXi host settings. Not required if vCenter connection is used.<br>This port is not required for VMware Cloud on AWS. |
| **Other Communications** | | | | |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup proxy | Gateway server | TCP | 2500 to 3300 | Default range of ports used as transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | Backup proxy | TCP | 2500 to 3300 | Default range of ports used as transmission channels for replication jobs. For every TCP connection that a job uses, one port from this range is assigned.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

# Gateway Server

The following table describes network ports that must be opened to ensure proper communication with gateway servers. For more information about ports that must be opened between the gateway server and specific backup repository, see Backup Repositories.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Gateway server (Microsoft Windows) | TCP | 445<br>135 | Required for deploying Veeam Backup & Replication components. |
| | | TCP | 6160 | Default port used by Veeam Installer Service. |
| | | TCP | 6162 | Default port used by Veeam Data Mover Service. |
| | Gateway server (Linux) | TCP | 22 | Default SSH port used as a control channel. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 6160 | Default port used by Veeam Installer Service for Linux. |
| | | TCP | 6162 | Default port used by Veeam Data Mover Service.<br><br>You can specify a different port while adding the Linux server to the Veeam Backup & Replication infrastructure. Note that you can specify a different port only if there is no previously installed Veeam Data Mover on this Linux server. For more information, see Specify Credentials and SSH Settings. |
| | Gateway server | TCP | 2500 to 3300 | Default range of ports used as transmission channels and for collecting log files. For every TCP connection that a job uses, one port from this range is assigned.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| Backup proxy | Gateway server | TCP | 2500 to 3300 | Default range of ports used as transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

# Backup Repositories

- Microsoft Windows/Linux-based Backup Repository
- NFS Backup Repository
- SMB Backup Repository

- Dell Data Domain System
- ExaGrid
- HPE StoreOnce
- Quantum DXi
- Fujitsu ETERNUS CS800
- Infinidat InfiniGuard
- Veeam Data Cloud Vault
- Object Storage Repository
- External Repository
- Archive Object Storage Repository

## Microsoft Windows/Linux-based Backup Repository

The following table describes network ports that must be opened to ensure proper communication with Microsoft Windows/Linux-based backup repositories.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Backup repository (Microsoft Windows) | TCP | 445<br>135 | Required for deploying Veeam Backup & Replication components. |
| | | TCP | 6160 | Default port used by Veeam Installer Service. |
| | | TCP | 6162 | Default port used by Veeam Data Mover Service. |
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br><br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Backup repository (Linux) | TCP | 22 | Default SSH port used as a control channel. |
| | | TCP | 6160 | Default port used by Veeam Installer Service for Linux. |
| | | TCP | 6162 | Default port used by Veeam Data Mover Service.<br>You can specify a different port while adding the Linux server to the Veeam Backup & Replication infrastructure. Note that you can specify a different port only if there is no previously installed Veeam Data Mover on this Linux server. For more information, see Specify Credentials and SSH Settings. |
| | | TCP | 2500 to 3300 | Default range of ports used as transmission channels and for collecting log files. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup repository (Linux) | Backup server | TCP | 2500 to 3300 | Default range of ports used as transmission channels for copy backup operations if the backup server is used as the target backup repository. These ports are also required for file copy operations between the Linux backup repository and the backup server.<br><br>For every TCP connection that a job uses, one port from this range is assigned.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | Veeam Update Repository | TCP | 443 | Used by backup repositories created with the Veeam Hardened Repository ISO to download security and operating system updates.<br><br>Veeam Update Repository endpoints:<br>• `repository.veeam.com` |
| Backup proxy | Backup repository | TCP | 2500 to 3300 | Default range of ports used as transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Source backup repository | Target backup repository | TCP | 2500 to 3300 | Default range of ports used as transmission channels for backup copy jobs and copy backup operations. For every TCP connection that a job uses, one port from this range is assigned.<br><br>If the backup copy job utilizes WAN accelerators, make sure that ports specific for WAN accelerators are opened.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## NFS Backup Repository

The following table describes network ports that must be opened to ensure proper communication with NFS shares added as backup repositories.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Gateway server or backup proxy | NFS backup repository | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service.<br><br>Also used as a transmission channel from the gateway server to the target NFS backup repository if a gateway server is specified explicitly in NFS backup repository settings. |
| Gateway server or backup proxy | NFS backup repository (NFS v3) | TCP, UDP | mountd_port | Dynamic port used for mountd service. Can be assigned statically. |
| | | TCP, UDP | statd_port | Dynamic port used for statd service. Can be assigned statically. |
| | | TCP, UDP | lockd_port | Dynamic port used for lockd service. Can be assigned statically. |

## SMB Backup Repository

The following table describes network ports that must be opened to ensure proper communication with SMB (CIFS) shares added as backup repositories.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Gateway server or backup proxy | SMB (CIFS) backup repository (Microsoft Windows) | TCP | 445 | Used as a transmission channel from the gateway server to the target SMB (CIFS) backup repository if a gateway server is specified explicitly in SMB (CIFS) backup repository settings. |

## Dell Data Domain System

For more information, see Dell Documents.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server or gateway server | Dell Data Domain | TCP | 111 | Port used to assign a random port for the mountd service used by NFS and DDBOOST. Mountd service port can be statically assigned. |
| | | TCP | 2049 | Main port used by NFS. Can be modified using the 'nfs set server-port' command. Command requires SE mode. |
| | | TCP | 2052 | Main port used by NFS MOUNTD. Can be modified using the 'nfs set mountd-port' command in SE mode. |

## ExaGrid

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | ExaGrid | TCP | 22 | Default command port used for communication with ExaGrid. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup proxy | ExaGrid | TCP | 2500 to 3300 | Default range of ports used for communication with the backup proxy.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## HPE StoreOnce

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server or gateway server | HPE StoreOnce | TCP | 9387 | Default command port used for communication with HPE StoreOnce. |
| | | | 9388 | Default data port used for communication with HPE StoreOnce. |

## Quantum DXi

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Quantum DXi | TCP | 22 | Default command port used for communication with Quantum DXi. |
| Backup proxy | Quantum DXi | TCP | 2500 to 3300 | Default range of ports used for communication with the backup proxy.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## Fujitsu ETERNUS CS800

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Fujitsu ETERNUS CS800 | TCP | 22 | Default command port used for communication with Fujitsu ETERNUS CS800. |
| Backup proxy | Fujitsu ETERNUS CS800 | TCP | 2500 to 3300 | Default range of ports used for communication with the backup proxy.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## Infinidat InfiniGuard

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Infinidat InfiniGuard | TCP | 22 | Default command port used for communication with Infinidat InfiniGuard. |
| Backup proxy | Infinidat InfiniGuard | TCP | 2500 to 3300 | Default range of ports used for communication with the backup proxy.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## Veeam Data Cloud Vault

The following table describes network ports and endpoints that must be opened to ensure proper communication with Veeam Data Cloud Vault. Note that a connection between the backup server and Veeam License Update Server is also required. For more information, see Backup Server.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Veeam Data Cloud Vault | TCP | 443 | Used to communicate with the Microsoft Azure object storage through the following endpoints:<br>• `<storage-account>.blob.core.windows.net`<br>• `<storage-account>.blob.storage.azure.net`<br>Consider that the `<storage-account>` part of the address must be replaced with the location name of your storage account. You can find the location name in the **Cloud Management** > **Vault Subscriptions** section of your Veeam Data Cloud account. |
| | | | 80 | Used to verify the certificate status through the following endpoints:<br>• `ocsp.digicert.com`<br>• `ocsp.msocsp.com`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access**<br>For more details, see also this Microsoft article. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup proxy (direct connection)/Gateway server | Veeam Data Cloud Vault | TCP | 443 | Used to communicate with the Microsoft Azure object storage through the following endpoints:<br>• `<storage-account>.blob.core.windows.net`<br>• `<storage-account>.blob.storage.azure.net`<br>Consider that the `<storage-account>` part of the address must be replaced with the location name of your storage account. You can find the location name in the **Cloud Management** > **Vault Subscriptions** section of your Veeam Data Cloud account. |
| | | | 80 | Used to verify the certificate status through the following endpoints:<br>• `ocsp.digicert.com`<br>• `ocsp.msocsp.com`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access**<br>For more details, see also this Microsoft article. |
| Gateway server | Backup proxy (direct connection)/Gateway server or backup server | TCP | 2500 to 3300 | Default range of ports used as transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## Object Storage Repository

The following table describes network ports and endpoints that must be opened to ensure proper communication with object storage repositories. For more information, see Object Storage Repository.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Gateway server | Backup proxy (direct connection)/Gateway server or backup server | TCP | 2500 to 3300 | Default range of ports used as transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| Backup proxy (direct connection)/Gateway server or backup server | Amazon S3 object storage | TCP | 443 | Used to communicate with the Amazon S3 object storage through the following endpoints:<br>• `*.amazonaws.com` (for both *Global* and *Government* regions)<br>• `*.amazonaws.com.cn` (for *China* region)<br>All AWS service endpoints are specified in the AWS documentation. |
| | | | 80 | Used to verify the certificate status through the following endpoints:<br>• `*.amazontrust.com`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access** |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Microsoft Azure object storage | TCP | 443 | Used to communicate with the Microsoft Azure object storage through the following endpoints:<br>• `<storage-account>.blob.core.windows.net` (for *Global* region)<br>• `<storage-account>.blob.storage.azure.net` (for *Global* region)<br>• `<storage-account>.blob.core.chinacloudapi.cn` (for *China* region)<br>• `<storage-account>.blob.core.usgovcloudapi.net` (for *Government* region)<br>Consider that the `<storage-account>` part of the address must be replaced with your actual storage account URL that can be found in the Azure management portal. |
| | | | 80 | Used to verify the certificate status through the following endpoints:<br>• `ocsp.digicert.com`<br>• `ocsp.msocsp.com`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access**<br>For more details, see also this Microsoft article. |
| | Google Cloud storage | TCP | 443 | Used to communicate with Google Cloud storage through the following endpoints:<br>• `storage.googleapis.com`<br>All cloud endpoints are specified in this Google article. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | | 80 | Used to verify the certificate status through the following endpoints:<br>• `ocsp.pki.goog`<br>• `pki.goog`<br>• `crl.pki.goog`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access** |
| | IBM Cloud object storage | TCP | Depends on device configuration | Used to communicate with IBM Cloud object storage. |
| | S3 compatible object storage | TCP | Depends on device configuration | Used to communicate with S3 compatible object storage. |

## External Repository

The following table describes network ports and endpoints that must be opened to ensure proper communication with external repositories. For more information, see External Repository.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Gateway server | Backup proxy (direct connection)/Gateway server or backup server | TCP | 2500 to 3300 | Default range of ports used as transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Gateway server or backup server | Amazon S3 object storage | TCP | 443 | Used to communicate with the Amazon S3 object storage through the following endpoints:<br>• `*.amazonaws.com` (for both *Global* and *Government* regions)<br>• `*.amazonaws.com.cn` (for *China* region)<br><br>All AWS service endpoints are specified in the AWS documentation. |
| | | | 80 | Used to verify the certificate status through the following endpoints:<br>• `*.amazontrust.com`<br><br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access** |
| | Microsoft Azure object storage | TCP | 443 | Used to communicate with the Microsoft Azure object storage through the following endpoints:<br>• `<storage-account>.blob.core.windows.net` (for *Global* region)<br>• `<storage-account>.blob.storage.azure.net` (for *Global* region)<br>• `<storage-account>.blob.core.chinacloudapi.cn` (for *China* region)<br>• `<storage-account>.blob.core.usgovcloudapi.net` (for *Government* region)<br><br>Consider that the `<storage-account>` part of the address must be replaced with your actual storage account URL that can be found in the Azure management portal. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | | 80 | Used to verify the certificate status through the following endpoints:<br>• `ocsp.digicert.com`<br>• `ocsp.msocsp.com`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access**<br>For more details, see also this Microsoft article. |
| | Google Cloud storage | TCP | 443 | Used to communicate with Google Cloud storage through the following endpoints:<br>• `storage.googleapis.com`<br>All cloud endpoints are specified in this Google article. |
| | | | 80 | Used to verify the certificate status through the following endpoints:<br>• `ocsp.pki.goog`<br>• `pki.goog`<br>• `crl.pki.goog`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access** |

## Archive Object Storage Repository

The following table describes network ports and endpoints that must be opened to ensure proper communication with object storage repositories used as a part of Archive Tier. For more information, see Archive Tier.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Gateway server or backup server | Amazon EC2 helper appliance | TCP | 443 | Used by default to communicate with the Amazon EC2 helper appliance through public/private IPv4 addresses of EC2 appliances.<br>If you use Amazon S3 Glacier object storage, the gateway server should have direct connection to AWS service endpoints. HTTP/HTTPS proxy servers are not supported.<br>If there is no gateway server selected, the backup server will be used as a gateway server. |
| | | TCP | 22 | Default SSH port used as a control channel. |
| | Microsoft Azure proxy appliance | TCP | 443 | Used by default to communicate with the Microsoft Azure helper appliance through public/private IPv4 addresses of Azure appliances.<br>If there is no gateway server selected, the backup server will be used as a gateway server. |
| | | TCP | 22 | Default SSH port used as a control channel. |
| Amazon EC2 helper appliance | Amazon S3 object storage | TCP | 443 | Used to communicate with the Amazon S3 object storage through the following endpoints:<br>• `*.amazonaws.com` (for both *Global* and *Government* regions)<br>• `*.amazonaws.com.cn` (for *China* region)<br>All AWS service endpoints are specified in the AWS documentation |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 80 | Used to verify the certificate status through the following endpoints:<br>• `*.amazontrust.com`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access** |
| Microsoft Azure proxy appliance | Microsoft Azure object storage | TCP | 443 | Used to communicate with the Microsoft Azure object storage through the following endpoints:<br>• `<storage-account>.blob.core.windows.net` (for *Global* region)<br>• `<storage-account>.blob.storage.azure.net` (for *Global* region)<br>• `<storage-account>.blob.core.chinacloudapi.cn` (for *China* region)<br>• `<storage-account>.blob.core.usgovcloudapi.net` (for *Government* region)<br>Consider that the `<storage-account>` part of the address must be replaced with your actual storage account URL that can be found in the Azure management portal. |
| | | TCP | 80 | Used to verify the certificate status through the following endpoints:<br>• `ocsp.digicert.com`<br>• `ocsp.msocsp.com`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access**<br>For more details, see also this Microsoft article. |

# Storage Systems

- Dell VNX(e) Storage
- Dell Unity XT, Unity Storage
- Dell PowerScale (Formerly Isilon) Storage
- HPE 3PAR StoreServ Storage
- HPE Alletra MP, Alletra 9000, Primera Storage
- HPE StoreVirtual (formerly LeftHand/P4000 Series) and StoreVirtual VSA Storage
- HPE Alletra 5000, Alletra 6000, Nimble Storage
- Lenovo ThinkSystem DM/DG Series Storage
- NetApp ONTAP Storage
- Nutanix Files Storage
- Universal Storage API Integrated System

## Dell VNX(e) Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | VNX File | TCP | 22 | Default command port used for communication with Dell VNX File over SSH. |
| | VNX Block<br>VNXe | TCP | 443 | Default port used for communication with Dell VNX Block/Dell VNXe over HTTPS and sending REST API calls. |
| Backup proxy | VNX Block<br>VNXe | TCP | 3260 | Default iSCSI target port. |
| | VNX File<br>VNXe | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |

## Dell Unity XT, Unity Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Dell Unity XT/Unity storage system | TCP | 443 | Default port used for communication with Dell Unity XT/Unity over HTTPS and sending REST API calls. |
| Backup proxy | Dell Unity XT/Unity storage system | TCP | 3260 | Default iSCSI target port. |
| | | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |

## Dell PowerScale (Formerly Isilon) Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Dell PowerScale storage system | TCP | 8080 | Default port used for communication with Dell PowerScale over HTTPS and sending REST API calls. |
| Backup proxy | Dell PowerScale storage system | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |
| | | TCP | 445 | Standard SMB port. |

## HPE 3PAR StoreServ Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | HPE 3PAR StoreServ storage system | TCP | 8008 | Default port used for communication with HPE 3PAR StoreServ over HTTP. |
| | | TCP | 8080 | Default port used for communication with HPE 3PAR StoreServ over HTTPS. |
| | | TCP | 22 | Default command port used for communication with HPE 3PAR StoreServ over SSH. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup proxy | HPE 3PAR StoreServ storage system | TCP | 3260 | Default iSCSI target port. |

## HPE Alletra MP, Alletra 9000, Primera Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | HPE Alletra MP/Alletra 9000/Primera storage system | TCP | 443 | Default port used for communication with HPE Alletra MP/Alletra 9000/Primera over HTTPS. |
| | | TCP | 22 | Default command port used for communication with HPE Alletra MP/Alletra 9000/Primera over SSH. |
| Backup proxy | HPE Alletra MP/Alletra 9000/Primera storage system | TCP | 3260 | Default iSCSI target port. |

## HPE StoreVirtual (formerly LeftHand/P4000 Series) and StoreVirtual VSA Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | HPE StoreVirtual/LeftHand/P4000 series storage system | TCP | 16022 | Default command port used for communication with HPE StoreVirtual/LeftHand/P4000 series. |
| Backup proxy | HPE StoreVirtual/LeftHand/P4000 series storage system | TCP | 3260 | Default iSCSI target port. |

## HPE Alletra 5000, Alletra 6000, Nimble Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | HPE Alletra 5000/Alletra 6000/Nimble storage system | TCP | 5392 | Default command port used for communication with HPE Alletra 5000/Alletra 6000/Nimble. |
| Backup proxy | HPE Alletra 5000/Alletra 6000/Nimble storage system | TCP | 3260 | Default iSCSI target port. |

## Lenovo ThinkSystem DM/DG Series Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Lenovo ThinkSystem DM/DG Series storage system | TCP | 80 | Default command port used for communication with Lenovo ThinkSystem DM/DG Series over HTTP. |
| | | TCP | 443 | Default command port used for communication with Lenovo ThinkSystem DM/DG Series over HTTPS. |
| Backup proxy | Lenovo ThinkSystem DM/DG Series storage system | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |
| | | TCP | 445 | Standard SMB port. |
| | | TCP | 3260 | Default iSCSI target port. |

## NetApp ONTAP Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | NetApp ONTAP | TCP | 80 | Default command port used for communication with NetApp ONTAP over HTTP. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | storage system | TCP | 443 | Default command port used for communication with NetApp ONTAP over HTTPS. |
| Backup proxy | NetApp ONTAP storage system | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |
| | | TCP | 445 | Standard SMB port. |
| | | TCP | 3260 | Default iSCSI target port. |

## Nutanix Files Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Nutanix Files storage system | TCP | 9440 | Default port used for communication with Nutanix Files and sending REST API calls. |
| Backup proxy | Nutanix Files storage system | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |
| | | TCP | 445 | Standard SMB port. |

## Universal Storage API Integrated System

The following tables describe network ports that must be opened to ensure proper communication with Universal Storage API integrated systems:

- DataCore SANsymphony
- Dell SC Series
- Dell PowerMax
- Dell PowerStore
- Fujitsu ETERNUS DX/AF
- Hitachi VSP/VSP One Block
- HPE XP
- IBM FlashSystem (formerly Spectrum Virtualize) Storage
- INFINIDAT InfiniBox
- NEC Storage M Series

- NEC Storage V Series
- NetApp SolidFire/HCI
- Pure Storage FlashArray
- Tintri IntelliFlash (formerly Western Digital IntelliFlash, Tegile)

### DataCore SANsymphony

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | DataCore SANsymphony storage system | TCP | 443 | Default command port used for communication with DataCore SANsymphony over HTTPS. |
| Backup proxy | DataCore SANsymphony storage system | TCP | 3260 | Default iSCSI target port. |

### Dell SC Series

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Dell SC Series storage system | TCP | 3033 | Default command port used for communication with Dell SC Series over HTTPS. |
| Backup proxy | Dell SC Series storage system | TCP | 3260 | Default iSCSI target port. |

### Dell PowerMax

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Dell PowerMax storage system | TCP | 8443 | Default command port used for communication with Dell PowerMax over HTTPS. |
| Backup proxy | Dell PowerMax storage system | TCP | 3260 | Default iSCSI target port. |

### Dell PowerStore

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Dell PowerStore storage system | TCP | 443 | Default command port used for communication with Dell PowerStore over HTTPS. |
| Backup proxy | Dell PowerStore storage system | TCP | 3260 | Default iSCSI target port. |

### Fujitsu ETERNUS DX/AF

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Fujitsu ETERNUS DX/AF storage system | TCP | 22 | Default command port used for communication with Fujitsu ETERNUS DX/AF over SSH. |
| Backup proxy | Fujitsu ETERNUS DX/AF storage system | TCP | 3260 | Default iSCSI target port. |

### Hitachi VSP/VSP One Block

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Hitachi VSP/VSP One Block storage system | TCP | 443 | Default command port used for communication with Hitachi VSP/VSP One Block over HTTPS. |
| Backup proxy | Hitachi VSP/VSP One Block storage system | TCP | 3260 | Default iSCSI target port. |

### HPE XP

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | HPE XP storage system | TCP | 443 | Default command port used for communication with HPE XP over HTTPS. |
| Backup proxy | HPE XP storage system | TCP | 3260 | Default iSCSI target port. |

### IBM FlashSystem (formerly Spectrum Virtualize) Storage

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | IBM FlashSystem storage system | TCP | 22 | Default command port used for communication with IBM FlashSystem over SSH. |
| Backup proxy | IBM FlashSystem storage system | TCP | 3260 | Default iSCSI target port. |

### INFINIDAT InfiniBox

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | INFINIDAT InfiniBox storage system | TCP | 443 | Default command port used for communication with INFINIDAT InfiniBox over HTTPS. |
| Backup proxy | INFINIDAT InfiniBox storage system | TCP | 3260 | Default iSCSI target port. |

## NEC Storage M Series

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | NEC Storage M Series storage system | TCP | 22 | Default command port used for communication with NEC Storage M Series over SSH. |
| Backup proxy | NEC Storage M Series storage system | TCP | 3260 | Default iSCSI target port. |

## NEC Storage V Series

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | NEC Storage V Series storage system | TCP | 443 | Default command port used for communication with NEC Storage V Series over HTTPS. |
| Backup proxy | NEC Storage V Series storage system | TCP | 3260 | Default iSCSI target port. |

## NetApp SolidFire/HCI

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | NetApp SolidFire/HCI storage system | TCP | 443 | Default command port used for communication with NetApp SolidFire/HCI over HTTPS. |
| Backup proxy | NetApp SolidFire/HCI storage system | TCP | 3260 | Default iSCSI target port. |

### Pure Storage FlashArray

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Pure Storage FlashArray system | TCP | 443 | Default command port used for communication with Pure Storage FlashArray over HTTPS. |
| Backup proxy | Pure Storage FlashArray system | TCP | 3260 | Default iSCSI target port. |

### Tintri IntelliFlash (formerly Western Digital IntelliFlash, Tegile)

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Tintri IntelliFlash system | TCP | 443 | Default command port used for communication with Tintri IntelliFlash over HTTPS. |
| Backup proxy | Tintri IntelliFlash system | TCP | 3260 | Default iSCSI target port. |
| | Tintri IntelliFlash system | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |

## Unstructured Data Backup Components

The following tables describe network ports that must be opened to ensure proper communication between unstructured data backup components.

- File Share Connections
- Cache Repository Connections
- Archive Repository Connections

## File Share Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup proxy | File server | TCP | 2500 to 3300 | Default range of ports used as transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | NAS filer (NetApp Data ONTAP or Lenovo ThinkSystem DM/DG Series storage system) | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |
| | | TCP | 445 | Standard SMB port. |
| | | TCP | 3260 | Default iSCSI target port. |
| | | TCP | 80, 443 | Used by NetApp SnapDiff when changed file tracking (CFT) is enabled. |
| | NAS filer (Dell PowerScale (formerly Isilon) or Nutanix Files storage system) | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |
| | | TCP | 445 | Standard SMB port. |
| Backup proxy or tape server | NFS share | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |
| | SMB share | TCP | 445 | Standard SMB port. |
| | Amazon S3 object storage | TCP | 443 | Used to communicate with the Amazon S3 object storage through the following endpoints:<br>• `*.amazonaws.com` (for both *Global* and *Government* regions)<br>• `*.amazonaws.com.cn` (for *China* region)<br>All AWS service endpoints are specified in the AWS documentation. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 80 | Used to verify the certificate status through the following endpoints:<br>• `*.amazontrust.com`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access** |
| | Microsoft Azure object storage | TCP | 443 | Used to communicate with the Microsoft Azure object storage through the following endpoints:<br>• `<storage-account>.blob.core.windows.net` (for *Global* region)<br>• `<storage-account>.blob.storage.azure.net` (for *Global* region)<br>• `<storage-account>.blob.core.chinacloudapi.cn` (for *China* region)<br>• `<storage-account>.blob.core.usgovcloudapi.net` (for *Government* region)<br>Consider that the `<storage-account>` part of the address must be replaced with your actual storage account URL that can be found in the Azure management portal. |
| | | TCP | 80 | Used to verify the certificate status through the following endpoints:<br>• `ocsp.digicert.com`<br>• `ocsp.msocsp.com`<br>Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. You can find the actual list of addresses in the certificate details in the following fields:<br>• **CRL Distribution Points**<br>• **Authority Information Access**<br>For more details, see also this Microsoft article. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | S3 compatible object storage | TCP | Depends on device configuration | Used to communicate with S3 compatible object storage. |

## Cache Repository Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup proxy | Cache repository | TCP | 2500 to 3300 | Default range of ports used as transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| Cache repository | Backup proxy | TCP | 2500 to 3300 | Default range of ports used as transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Primary or secondary backup repository | TCP | 2500 to 3300 | Default range of ports used as transmission channels for file share backup restore jobs. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

### Archive Repository Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Primary backup repository | Archive repository | TCP | 2500 to 3300 | Default range of ports used as transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## Tape Server

The following table describes network ports that must be opened to ensure proper communication with tape servers.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Tape server | TCP | 445, 135 | Required for deploying Veeam Backup & Replication components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 2500 to 3300 | Default range of ports used as data transmission channels and for collecting log files. For every TCP connection that a job uses, one port from this range is assigned.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6160 | Default port used by Veeam Installer Service. |
| | | TCP | 6162 | Default port used by Veeam Data Mover Service. |
| | | TCP | 6166 | Controlling port for RPC calls. |
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br><br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Tape server | Backup server | TCP | 2500 to 3300 | Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | Backup repository or gateway server | TCP | 2500 to 3300 | Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | NFS share | TCP, UDP | 111, 2049 | Standard NFS ports. Port 111 is used by the port mapper service. |
| | SMB share | TCP | 445 | Standard SMB port. |

# WAN Accelerator

The following table describes network ports that must be opened to ensure proper communication between WAN accelerators used in backup copy jobs and replication jobs.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | WAN accelerator (source and target) | TCP | 445, 135 | Required for deploying Veeam Backup & Replication components. |
| | | TCP | 6160 | Default port used by Veeam Installer Service. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 6162 | Default port used by Veeam Data Mover Service. |
| | | TCP | 6164 | Controlling port for RPC calls. |
| | | TCP | 6220 | Port used for traffic control (throttling) for tenants that use WAN accelerators.<br>This port is required only in the Veeam Cloud Connect infrastructure. |
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |
| WAN accelerator (target) | Backup repository (target) | TCP | 2500 to 3300 | Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is selected dynamically.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| WAN accelerator (source) | Backup repository (source) | TCP | 2500 to 3300 | Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is selected dynamically.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| WAN accelerator (source and target) | WAN accelerator (source and target) | TCP | 6164 | Controlling port for RPC calls. |
| | | TCP | 6165 | Default port used for data transfer between WAN accelerators. Ensure this port is open in firewall between sites where WAN accelerators are deployed. |

# Guest Processing Components

## Connections with Non-Persistent Runtime Components

The following tables describe network ports that must be opened to ensure proper communication of the backup server and backup infrastructure components with the non-persistent runtime components deployed inside the VM guest OS for application-aware processing and indexing.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | VM guest OS (Linux) | TCP | 22 | Default SSH port used as a control channel. |
| | Guest interaction proxy | TCP | 6190 | Used for communication with the guest interaction proxy. |
| | | TCP | 6290 | Used as a control channel for communication with the guest interaction proxy. |
| | | TCP | 445 | Port used as a transmission channel. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Guest interaction proxy | ESXi server | TCP | 443 | Default port used for connections to ESXi host.<br>[For VMware vSphere earlier than 6.5] Not required if vCenter connection is used. In VMware vSphere versions 6.5 and later, port 443 is required by vCenter Web Services. |

Network ports described in the following table are NOT required when working in networkless mode over VMware VIX/vSphere Web Services.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Guest interaction proxy | VM guest OS (Microsoft Windows) | TCP | 445<br>135 | Required to deploy the runtime coordination process on the VM guest OS. |
| | | TCP | 2500 to 3300 | Default range of ports used as transmission channels for log shipping.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br>Used by the runtime process deployed inside the VM for guest OS interaction (when working over the network, not over VIX API).<br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | VM guest OS (Linux) | TCP | 22 | Default SSH port used as a control channel. |
| | | TCP | 2500 to 3300 | Default range of ports used as transmission channels for log shipping.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| VM guest OS | Guest interaction proxy | TCP | 2500 to 3300 | Default range of ports used as transmission channels for log shipping.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## Connections with Persistent Agent Components

The following table describes network ports that must be opened to ensure proper communication of the backup server with the persistent agent components deployed inside the VM guest OS for application-aware processing and indexing.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | VM guest OS (Linux) | TCP | 6160 | Default port used by Veeam Installer Service for Linux. |
| | | TCP | 6162 | Default Management Agent port. Required if it is used as a control channel instead of SSH. |
| Guest interaction proxy | VM guest OS | TCP | 6160<br>11731 | Default port and failover port used by Veeam Installer Service. |
| | | TCP | 6173<br>2500 | Used by the Veeam Guest Helper for guest OS processing and file-level restore. |

# Log Shipping Components

The following tables describe network ports that must be opened to ensure proper communication between log shipping components.

- Log Shipping Server Connections
- MS SQL Guest OS Connections
- Oracle Guest OS Connections
- PostgreSQL Guest OS Connections

## Log Shipping Server Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Log shipping server | TCP | 445<br>135 | Required for deploying Veeam Backup & Replication components. |
| | | TCP | 6160 | Default port used by Veeam Installer Service. |
| | | TCP | 6162 | Default port used by Veeam Data Mover Service. |
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br><br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Log shipping server | Backup repository or gateway server | TCP | 2500 to 3300 | Default range of ports used for communication with a backup repository and transfer log backups.<br><br>By default, the log shipping server connects to the backup repository. However, if the target repository uses a gateway server, the connection will be established with that instead. For more information, see Gateway Servers.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## MS SQL Guest OS Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Guest interaction proxy | MS SQL VM guest OS | TCP | 445<br>135 | [Non-persistent runtime components only] Required for deploying Veeam Backup & Replication components including Veeam Log Shipper runtime component.<br><br>These ports are NOT required when working in networkless mode over VMware VIX/vSphere Web Services. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a guest OS.<br><br>These ports are NOT required when working in networkless mode over VMware VIX/vSphere Web Services.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 49152 to 65535 | [Non-persistent runtime components only] Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br><br>These ports are NOT required when working in networkless mode over VMware VIX/vSphere Web Services.<br><br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |
| | | TCP | 6160<br>11731 | [Persistent agent components only] Default port and failover port used by Veeam Installer Service. |
| | | TCP | 6167 | Used by the Veeam Log Shipping Service for preparing the database and taking logs. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| MS SQL VM guest OS | Guest interaction proxy | TCP | 2500 to 3300 | Default range of ports used for communication with a guest interaction proxy.<br>These ports are NOT required when working in networkless mode over VMware VIX/vSphere Web Services.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| MS SQL VM guest OS | Backup repository | TCP | 2500 to 3300 | Default range of ports used for communication with a backup repository and transfer log backups. Should be opened if log shipping servers are not used in the infrastructure and the MS SQL server has a direct connection to the backup repository.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| MS SQL VM guest OS | Log shipping server | TCP | 2500 to 3300 | Default range of ports used for communication with a log shipping server and transfer log backups.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## Oracle Guest OS Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Guest interaction proxy | Oracle VM guest OS (Microsoft Windows) | TCP | 445<br>135 | [Non-persistent runtime components only] Required for deploying Veeam Backup & Replication components including Veeam Log Shipper runtime component.<br><br>These ports are NOT required when working in networkless mode over VMware VIX/vSphere Web Services. |
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a guest OS.<br><br>These ports are NOT required when working in networkless mode over VMware VIX/vSphere Web Services.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 49152 to 65535 | [Non-persistent runtime components only] Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br><br>These ports are NOT required when working in networkless mode over VMware VIX/vSphere Web Services.<br><br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 6160<br>11731 | [Persistent agent components only] Default port and failover port used by Veeam Installer Service. |
| | | TCP | 6167 | Used by the Veeam Log Shipping Service for preparing the database and taking logs. |
| | Oracle VM guest OS (Linux) | TCP | 22 | [Non-persistent runtime components only] Default SSH port used as a control channel.<br>This port is NOT required when working in networkless mode over VMware VIX/vSphere Web Services. |
| | | TCP | 6162 | [Persistent agent components only] Default Management Agent port. Required if it is used as a control channel instead of SSH. |
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a guest OS.<br>These ports are NOT required when working in networkless mode over VMware VIX/vSphere Web Services.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| Oracle VM guest OS | Guest interaction proxy | TCP | 2500 to 3300 | Default range of ports used for communication with a guest interaction proxy.<br>These ports are NOT required when working in networkless mode over VMware VIX/vSphere Web Services.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Oracle VM guest OS | Backup repository | TCP | 2500 to 3300 | Default range of ports used for communication with a backup repository and transfer log backups. Should be opened if log shipping servers are not used in the infrastructure and the Oracle server has a direct connection to the backup repository.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| Oracle VM guest OS | Log shipping server | TCP | 2500 to 3300 | Default range of ports used for communication with a log shipping server and transfer log backups.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## PostgreSQL Guest OS Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Guest interaction proxy | PostgreSQL VM guest OS | TCP | 22 | [Non-persistent runtime components only] Default SSH port used as a control channel.<br>This port is NOT required when working in networkless mode over vSphere Web Services. |
| | | TCP | 6162 | [Persistent agent components only] Default Management Agent port. Required if it is used as a control channel instead of SSH. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a guest OS.<br><br>This port is NOT required when working in networkless mode over vSphere Web Services.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| PostgreSQL VM guest OS | Guest interaction proxy | TCP | 2500 to 3300 | Default range of ports used for communication with a guest interaction proxy.<br><br>This port is NOT required when working in networkless mode over vSphere Web Services.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| PostgreSQL VM guest OS | Backup repository | TCP | 2500 to 3300 | Default range of ports used for communication with a backup repository and transfer log backups. Should be opened if log shipping servers are not used in the infrastructure and the PostgreSQL server has a direct connection to the backup repository.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| PostgreSQL VM guest OS | Log shipping server | TCP | 2500 to 3300 | Default range of ports used for communication with a log shipping server and transfer log backups.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

# CDP Components

The following table describes network ports that must be opened to ensure proper communication of Veeam CDP components with other backup components.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| ESXi host (source) | CDP proxy (source) | TCP | 33032 | Default port used as a transmission channel to the source CDP proxy. |
| | ESXi host (source) | TCP | 33033 (local) | Port used locally on the source ESXi host for data transfer between I/O filter components.<br>**Note**: Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |
| | ESXi host (source) | TCP | 33035 (local) | Port used locally on the source ESXi host for data transfer between I/O filter components over shared-memory.<br>**Note**: Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | ESXi host (source) | TCP | 33036 | Port used on the source ESXi host for communication between CDP components over HTTPS without HTTP Reverse Proxy. |
| | ESXi host (source) | TCP | 33038 (local) | Port used locally on the source ESXi host for communication between CDP components over HTTPS.<br>**Note**: Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |
| | ESXi host (source) | TCP | 33039 (local) | Port used locally on the source ESXi host for control notifications between I/O filter components.<br>**Note**: Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |
| CDP proxy (source) | CDP proxy (target) | TCP | 33033 | Default port used as a transmission channel to the target CDP proxy. |
| | ESXi host (source and target) | TCP | 902 | Default VMware port used for data transfer. Used during initial synchronization and restore operations. |
| | vCenter Server (source and target) | TCP | 443 | Default VMware web service port that can be customized in vCenter settings. Used during initial synchronization and restore operations. |
| CDP proxy (target) | ESXi host (target) | TCP | 33032 | Default port used as a transmission channel to the target ESXi host. |
| | ESXi host (source and target) | TCP | 902 | Default VMware port used for data transfer. Used during initial synchronization and restore operations. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | vCenter Server (source and target) | TCP | 443 | Default VMware web service port that can be customized in vCenter settings. Used during initial synchronization and restore operations. |
| ESXi host (target) | ESXi host (target) | TCP | 33034 (local) | Port used locally on the target ESXi host for communication between the I/O filter components during failover.<br>**Note**: Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |
| | ESXi host (target) | TCP | 33036 | Port used on the target ESXi host for communication between CDP components over HTTPS without HTTP Reverse Proxy. |
| | ESXi host (target) | TCP | 33038 (local) | Port used locally on the target ESXi host for communication between CDP components over HTTPS.<br>**Note**: Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |
| Backup server | ESXi host (source and target) | TCP | 443 | Port used as a control channel. |
| | ESXi host (source and target) | TCP | 33035 | Port used to install I/O filter components on the ESXi hosts. |
| | vCenter Server (source and target) | TCP | 443 | Port used as a control channel. |
| | CDP proxy (source and target) | TCP | 6182 | Port used as a control channel. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Backup server | TCP | 9509 (local) | Port used locally on the backup server for communication between Veeam Backup Service and Veeam CDP Coordinator Service.<br>**Note:** Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |
| ESXi host (source and target) | Backup server | TCP | 33034 | Port used for communication with Veeam CDP Coordinator Service. |
| | | TCP | 33035 | Port used to install I/O filter components on the ESXi hosts. |
| vCenter Server (source and target) | Backup server | TCP | 33034 | Port used for communication with Veeam CDP Coordinator Service. |
| | | TCP | 33035 | Port used to install I/O filter components on the vCenter servers. |
| CDP proxy (source and target) | Backup server | TCP | 33034 | Port used for communication with Veeam CDP Coordinator Service. |

# Recovery Components

- Guest OS File Recovery
- Veeam vPower NFS Service
- SureBackup
- SureReplica Recovery Verification
- Veeam U-AIR
- Microsoft Active Directory Domain Controller Connections During Application Item Restore
- Microsoft Exchange Server Connections During Application Item Restore
- Microsoft SQL Server Connections During Application Item Restore
- Restore to Amazon EC2
- Restore to Google Cloud

- Restore to Microsoft Azure

## Guest OS File Recovery

The following table describes network ports that must be opened to ensure proper communication between components for guest OS file recovery.

- Mount Server Connections
- Helper Appliance Connections
- Helper Host Connections
- Guest OS Connections

## Mount Server Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Mount server | Backup server | TCP | 9401 | Used for communication with the Veeam Backup Service. |
| | Backup repository | TCP | 2500 to 3300 | Default range of ports used for communication with a backup repository.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | ESXi host | TCP | 443 | Default port used for connections to the ESXi host if Windows file-level restore is performed over VIX API. |
| | vCenter server | TCP | 443 | Default port used for connections to the vCenter server if Windows file-level restore is performed using vCenter Web Services. |
| | Veeam Signature Update Server | TCP | 443 | Default port used by Veeam Threat Hunter to download information about new malware signatures from the Veeam Signature Update Server over HTTPS.<br>Veeam Signature Update Server endpoints:<br>• `avupdate.veeam.com` |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Mount server | TCP | 6175 (local) | Used locally for communication with the Veeam Threat Hunter Service.<br>**Note**: Local ports do not require specific firewall rules. Make sure that this port is not used by another software. Otherwise, this can affect Veeam Backup & Replication functionality. |
| Backup server | Mount server | TCP | 445 | Required for deploying Veeam Backup & Replication components. |
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a mount server and for collecting log files.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6160 | Default port used by Veeam Installer Service including checking the compatibility between components before starting the recovery process. |
| | | TCP | 6162 | Default port used by Veeam Data Mover Service. |
| | | TCP | 6170 | Used for communication with a local or remote Mount Service. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br><br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the "RPC function call failed" error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |

## Helper Appliance Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Helper appliance | Backup repository | TCP | 2500 to 3300 | Default range of ports used for communication with a backup repository.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| Helper appliance | ESXi server | TCP | 443 | Default port used for connections to the ESXi host if restore is performed over VIX API/ Web Services.<br><br>[For VMware vSphere earlier than 6.5] Not required if vCenter connection is used. In VMware vSphere versions 6.5 and later, port 443 is required by vSphere Web Services. |
| Backup server | Helper appliance | TCP | 22 | Default SSH port used as a control channel. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a helper appliance.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| Mount server | Helper appliance | TCP | 22 | Default SSH port used as a control channel. |
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a helper appliance.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## Helper Host Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Helper host | Backup repository | TCP | 2500 to 3300 | Default range of ports used for communication with a backup repository.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Helper host | ESXi server | TCP | 443 | Default port used for connections to the ESXi host if restore is performed over VIX API/vSphere Web Services.<br>[For VMware vSphere earlier than 6.5] Not required if vCenter connection is used. In VMware vSphere versions 6.5 and later, port 443 must also be open to vSphere Web Services. |
| Backup server | Helper host | TCP | 22 | Default SSH port used as a control channel. |
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a helper host and for collecting log files.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6162 | Default port used by Veeam Data Mover Service. |
| | | TCP | 32768 to 60999 | Dynamic port range for Linux distributions. Used for communication with a helper host. For more information, see the Linux kernel documentation. |
| Mount server | Helper host | TCP | 22 | Default SSH port used as a control channel. |
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a helper host.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 32768 to 60999 | Dynamic port range for Linux distributions. Used for communication with a helper host. For more information, see the Linux kernel documentation. |

## Guest OS Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| VM guest OS (Linux/Unix) | Helper appliance | TCP | 21 | Default port used for protocol control messages if FTP server is enabled. |
| Helper appliance | VM guest OS (Linux/Unix) | TCP | 20 | Default port used for data transfer if FTP server is enabled. |
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a VM guest OS.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| Helper host | VM guest OS (Linux/Unix) | TCP | 2500 to 3300 | Default range of ports used for communication with a VM guest OS.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| Backup server | VM guest OS (Linux/Unix) | TCP | 22 | Default SSH port used as a control channel. |
| Mount server | | TCP | 445<br>135 | Required to deploy the runtime coordination process on the VM guest OS. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | VM guest OS (Microsoft Windows) | TCP | 6160<br>11731 | Default port and failover port used by Veeam Installer Service. |
| | | TCP | 6173<br>2500 | Used by the Veeam Guest Helper for guest OS processing and file-level restore if persistent agent components are deployed inside the VM guest OS. |
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the "RPC function call failed" error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |
| Backup server | VM guest OS | TCP | 2500 to 3300 | Default range of ports used for communication with a VM guest OS.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## Veeam vPower NFS Service

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Microsoft Windows | TCP | 6160 | Default port used by Veeam Installer Service. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | server with the mount server role running vPower NFS Service | TCP | 6161 | Default port used by the Veeam vPower NFS Service. |
| ESXi host | Microsoft Windows server with the mount server role running vPower NFS Service | TCP<br>UDP | 111 | Standard port used by the port mapper service. |
| | | TCP<br>UDP | 1058+ or 1063+ | Default mount port. The number of port depends on where the vPower NFS Service is located:<br>• 1058+: If the vPower NFS Service is located on the backup server.<br>• 1063+: If the vPower NFS Service is located on a separate Microsoft Windows machine.<br>If port 1058/1063 is occupied, the succeeding port numbers will be used. |
| | | TCP<br>UDP | 2049+ | Standard NFS port. If port 2049 is occupied, the succeeding port numbers will be used. |
| Backup repository or gateway server working with backup repository | Microsoft Windows server with the mount server role running vPower NFS Service | TCP | 2500 to 3300 | Default range of ports used as transmission channels during Instant Recovery, SureBackup or Linux file-level recovery.<br>For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Microsoft Windows server with the mount server role running vPower NFS Service | Backup repository or gateway server working with backup repository | TCP | 2500 to 3300 | Default range of ports used as transmission channels during Instant Recovery, SureBackup or Linux file-level recovery.<br>For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## SureBackup

The following table describes network ports that must be opened to ensure proper communication between SureBackup components.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Proxy appliance | TCP | 443 | Used for communication with the proxy appliance in the virtual lab. |
| | Applications on VMs in the virtual lab | — | — | Application-specific ports to perform port probing test. For example, to verify a DC, Veeam Backup & Replication probes port 389 for a response. |
| Internet-facing proxy server | VMs in the virtual lab | TCP | 8080 | Used to let VMs in the virtual lab access the Internet. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Microsoft Windows server with the mount server role running vPower NFS Service | Backup repository or gateway server working with backup repository | TCP | 2500 to 3300 | Default range of ports used as transmission channels during SureBackup.<br>For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | ESXi server | TCP | 443 | Default port used for connections to ESXi host. |
| Backup repository or gateway server working with backup repository | Microsoft Windows server with the mount server role running vPower NFS Service | TCP | 2500 to 3300 | Default range of ports used as transmission channels during SureBackup.<br>For every TCP connection that a job uses, one port from this range is assigned.<br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## SureReplica Recovery Verification

The following table describes network ports that must be opened to ensure proper communication between SureReplica components.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Proxy appliance | TCP | 443 | Used for communication with the proxy appliance in the virtual lab. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Applications on VMs in the virtual lab | — | — | Application-specific ports to perform port probing test. For example, to verify a DC, Veeam Backup & Replication probes port 389 for a response. |
| Internet-facing proxy server | VMs in the virtual lab | TCP | 8080 | Used to let VMs in the virtual lab access the Internet. |

## Veeam U-AIR

The following table describes network ports that must be opened to ensure proper communication of U-AIR wizards with other components.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| U-AIR wizards | Veeam Backup Enterprise Manager | TCP | 9394 | Used by default for communication with Veeam Backup Enterprise Manager. Can be customized during Veeam Backup Enterprise Manager installation. |

## Microsoft Active Directory Domain Controller Connections During Application Item Restore

The following table describes network ports that must be opened to ensure proper communication of the backup server with the Microsoft Active Directory VM during application-item restore.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Microsoft Active Directory VM guest OS | TCP | 135 | Used for communication between the domain controller and backup server. |
| | | TCP, UDP | 389 | LDAP connections. |
| | | TCP | 636, 3268, 3269 | LDAP connections. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later used by the runtime coordination process deployed inside the VM guest OS for application-aware processing (when working over the network, not over VIX API). For more information, see this Microsoft KB article.<br><br>**Note:** If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the "RPC function call failed" error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |

## Microsoft Exchange Server Connections During Application Item Restore

The following table describes network ports that must be opened to ensure proper communication of the Veeam backup server with the Microsoft Exchange Server system during application-item restore.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Microsoft Exchange 2003/2007 CAS Server | TCP | 80, 443 | WebDAV connections. |
| | Microsoft Exchange 2010/2013/2016/2019 CAS Server | TCP | 443 | Microsoft Exchange Web Services Connections. |

## Microsoft SQL Server Connections During Application Item Restore

The following table describes network ports that must be opened to ensure proper communication of the backup server with the VM guest OS system during application-item restore.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Microsoft SQL VM guest OS | TCP | 1433, 1434 and other | Used for communication with the Microsoft SQL Server installed inside the VM.<br>Port numbers depends on configuration of your Microsoft SQL server. For more information, see this Microsoft article. |
| | | UDP | 1434 | Used by the Microsoft SQL Server Browser service.<br>For more information, see this Microsoft article. |

## Restore to Amazon EC2

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server or backup repository | Helper appliance | TCP | 22 | Used as a communication channel to the helper appliance. |
| | | TCP | 443 | Default redirector port. You can change the port in helper appliance settings. For details, see the **Specify Helper Appliance** section in Restore to Amazon EC2. |

## Restore to Google Cloud

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server or backup repository | Helper appliance | TCP | 22 | Used as a communication channel to the helper appliance. |
| | | TCP | 443 | Default redirector port. You can change the port in helper appliance settings. For details, see the **Specify Helper Appliance** section in Restore to Google Cloud. |

## Restore to Microsoft Azure

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Helper appliance | TCP | 22 | Used by default as a communication channel to the helper appliance when restoring Linux workloads. Can be changed during helper appliance deployment. For details, see Configuring Helper Appliances. |
| | Microsoft Azure | TCP | 443 | Default management and data transport port required for communication with Microsoft Azure. |
| | Azure Windows VM agent distribution server | TCP | 443 | Used by Veeam Backup & Replication to install the Azure Windows VM agent on the restored VM through the following URLs:<br>• `go.microsoft.com`<br>• `aka.ms` (additional components required for the Azure Windows VM agent installation)<br>• `github.com` (additional components required for the Azure Windows VM agent installation)<br>• `objects.githubusercontent.com` (additional components required for the Azure Windows VM agent installation)<br>Consider that these URLs are subject to change. For more information, see this Microsoft article. |
| | Azure Stack Hub | TCP | 443, 30024 | Default management and data transport port required for communication with Azure Stack Hub. |
| Backup server or backup repository | Azure restore proxy appliance (former Azure proxy) | TCP | 443 | Default management and data transport port required for communication with the Azure restore proxy appliance. The port must be opened on the backup server and backup repository storing VM backups.<br>Can be changed in the settings of the Azure restore proxy appliance. For details, see Specify Credentials and Transport Port. |

# Veeam Backup Enterprise Manager

- Veeam Backup Enterprise Manager Connections

# Veeam Explorers

- Veeam Explorer for Microsoft Active Directory Connections
- Veeam Explorer for Microsoft Exchange Connections
- Veeam Explorer for Microsoft SharePoint and Veeam Explorer for Microsoft OneDrive for Business Connections
- Veeam Explorer for Microsoft SQL Server Connections
- Veeam Explorer for Microsoft Teams Connections
- Veeam Explorer for Oracle Connections
- Veeam Explorer for PostgreSQL Connections

# Veeam Cloud Connect

- Veeam Cloud Connect Connections

# Veeam Agents

## Veeam Agent for Microsoft Windows

- Connections for Veeam Agent for Microsoft Windows Operating in Managed Mode
- Connections for Veeam Agent for Microsoft Windows Operating in Standalone Mode

## Veeam Agent for Linux

- Connections for Veeam Agent for Linux Operating in Managed Mode
- Connections for Veeam Agent for Linux Operating in Standalone Mode

## Veeam Agent for Mac

- Connections for Veeam Agent for Mac Operating in Managed Mode
- Connections for Veeam Agent for Mac Operating in Standalone Mode

# Veeam Plug-ins for Enterprise Applications

- Veeam Plug-in for SAP HANA Connections
- Veeam Plug-in for Oracle RMAN Connections
- Veeam Plug-in for SAP on Oracle Connections
- Veeam Plug-in for Microsoft SQL Server Connections

# Veeam Plug-ins for Cloud Solutions

- AWS Plug-in for Veeam Backup & Replication
- Microsoft Azure Plug-in for Veeam Backup & Replication
- Google Cloud Plug-in for Veeam Backup & Replication

# Kasten

- Veeam Kasten Plug-in for Veeam Backup & Replication

# Virtualization Platforms

- Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization Connections
- Veeam Backup for Nutanix AHV Connections
- Veeam Backup for Proxmox VE Connections

# Nutanix Mine with Veeam

- Nutanix Mine with Veeam 4.0 Connections

# Other Connections

## NDMP Servers

The following table describes network ports that must be opened to ensure proper communication with NDMP servers.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Gateway server | NDMP server | NDMP | 10000 | Default port used to manage the NMDP server.<br>**Note:** The port range used for data transfer depends on your NDMP server configuration. For more information, contact your hardware vendor. |

## Mail Servers

The following table describes network ports that must be opened to ensure proper communication of the backup server with mail servers.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | SMTP server | TCP | 25 | Used by the SMTP server. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | | TCP | 587 | Used by the SMTP server if SSL is enabled. |
| | Gmail REST API (gmail.googleapis.com) | TCP | 443 | Used to communicate with Google Mail services. |
| | Microsoft Graph REST API (graph.microsoft.com, login.microsoftonline.com) | TCP | 443 | Used to communicate with Microsoft Exchange Online organizations. |

## Event Forwarding Components

The following table describes network ports that must be opened to ensure proper communication with event forwarding components.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | Syslog server | TCP<br>UDP | 514 | Default port used to communicate with the syslog server. |
| | | TLS | 6514 | Default port used to communicate with the syslog server over TLS. |

## Internet Connections

If you use an HTTP/HTTPS proxy server to access the Internet, make sure that WinHTTP settings are properly configured on Microsoft Windows machines with Veeam backup infrastructure components. For information on how to configure WinHTTP settings, see Microsoft Docs.

> **NOTE**
>
> Tenants cannot access Veeam Cloud Connect infrastructure components through HTTP/HTTPS proxy servers. For information on supported protocols for Veeam Cloud Connect, see the Ports section in the Veeam Cloud Connect Guide.

# Naming Conventions

Do not use Microsoft Windows reserved names for names of the backup server, managed servers, backup repositories, jobs, tenants and other objects created in Veeam Backup & Replication: CON, PRN, AUX, NUL, COM1, COM2, COM3, COM4, COM5, COM6, COM7, COM8, COM9, LPT1, LPT2, LPT3, LPT4, LPT5, LPT6, LPT7, LPT8 and LPT9.

If you plan to store backups on a repository operating in the per-machine mode, do not use Microsoft Windows reserved names for names of the virtual machines to back up.

If you use a reserved name, Veeam Backup & Replication may not work as expected. For more information on naming conventions in Microsoft Windows, see Microsoft Docs.

# Security Guidelines

When you set up the backup infrastructure, one thing that you must not overlook is security. The backup infrastructure can be potentially used as a backdoor to gain access to your systems and data. This section includes a number of security features and recommendations that will help you prevent potential security issues and reduce the risk of compromising sensitive data.

# General Security Considerations

General security considerations include best practices which help you to harden backup infrastructure, build a more secure environment, and mitigate risks of being compromised. Ensure that your backup infrastructure meet the common recommendations described in this section. For more information about hardening specific backup infrastructure components, see Securing Backup Infrastructure.

## Network

To secure the communication channel for backup traffic, consider the following recommendations:

- **Use network segmentation**. Create network segmentation policies to define network boundaries, control traffic between subnets and limit access to security-sensitive backup infrastructure components. Also, ensure that only ports used by backup infrastructure components are opened.
- **Isolate backup traffic**. Use an isolated network to transport data between backup infrastructure components — backup server, backup proxies, repositories and so on.
- **Disable outdated network protocols**. Check that the following protocols are disabled:
  - SSL 2.0 and 3.0 as they have well-known security vulnerabilities and are not NIST-approved. For more information, see NIST guidelines.
  - TLS 1.0 and 1.1 if they are not needed. For more information, see NIST guidelines.
  - LLMNR and NetBIOS broadcast protocols to prevent spoofing and man-in-the-middle (MITM) attacks.
  - SMB 1.0 protocol as it has a number of serious security vulnerabilities including remote code execution. For more information, see this Microsoft article.

## User Roles and Permissions

Administrator privileges on a backup server or a backup proxy allow the user to access other backup infrastructure components. If an attacker gains such permissions, they can destroy most of the production data, backups, and replicas, as well as compromise other systems in your environment. To mitigate risks, use the principle of the least privilege. Provide the minimal required permissions needed for the accounts to run. For more information, see Permissions.

## File System

Do not add paths writable by untrusted users to the PATH environment variable. This recommendation applies to Microsoft Windows-based and Linux-based backup infrastructure components as well as protected virtual and physical machines. A potential attacker may exploit this vulnerability to execute malware or access sensitive data. For more information, see this CWE article.

## Security Audit

Perform regular security audits to estimate your backup infrastructure by security criteria and understand if it is compliant with best practices, industry standards, or federal regulations.

The most possible causes of a credential theft are missing operating system updates and use of outdated authentication protocols. To mitigate risks, ensure that all software and hardware running backup infrastructure components are updated regularly. If the latest security updates and patches are installed on backup infrastructure servers, this will reduce the risk of exploiting vulnerabilities by attackers. Note that you should work out an update management strategy without a negative impact on production environment.

> **NOTE**
>
> You can subscribe to Veeam security advisories published in the Veeam Knowledge Base to stay up to date with the latest security updates.

# Microsoft Windows Server

To secure Microsoft Windows-based backup infrastructure components, consider the following recommendations:

- **Use operating system versions with Long Term Servicing Channel (LTSC)**. For these versions Microsoft provides extended support including regular security updates. For more information, see this Microsoft article.
- **Turn on Microsoft Defender Firewall with Advanced Security**. Set up rules for inbound and outbound connections according to your infrastructure and Microsoft best practices. For more information, see this Microsoft article.
- **Disable remote services if they are not needed**:
  - Remote Desktop Service
  - Remote Registry service
  - Remote PowerShell
  - Windows Remote Management service

> **NOTE**
>
> A backup server also requires additional configuration described in section Securing Backup Infrastructure.

# Linux Server

To secure Linux-based backup infrastructure components, consider the following recommendations:

- **Use operating system versions with long-term support (LTS).** LTS versions of popular community-based and commercial Linux distributions have extended support including regular security updates.
- **Choose strong encryption algorithms for SSH**. To communicate with Linux servers deployed as a part of the backup infrastructure, Veeam Backup & Replication uses SSH. Make sure that for the SSH tunnel you use a strong and proven encryption algorithm, with sufficient key length. For more information, see this section. Also, ensure that private keys are kept in a highly secure place and cannot be uncovered by a third-party.

  > **NOTE**
  >
  > For the Linux hardened repository, instead of SSH Veeam Backup & Replication uses SHA256RSA self-signed certificates with 2048-bit RSA key.

- **Avoid using password authentication to connect to remote servers over SSH**. Using key-based SSH authentication is generally considered more secure than using password authentication and helps averting man-in-the-middle (MITM) attacks. The private key is not passed to the server and cannot be captured even if a user connects to a fake server and accepts a bad fingerprint.

> **NOTE**
>
> A Linux hardened repository requires a specific security configuration. For more information, see Hardened Repository.

# Securing Backup Infrastructure

This section includes recommendations for hardening specific backup infrastructure components in addition to general security considerations.

## Infrastructure Planning

For large environments, adding the backup server and other backup infrastructure components to a management domain in a separate Active Directory forest is the best practice for building the most secure infrastructure.

For medium-sized and small environments, backup infrastructure components can be placed to a separate workgroup. If you want to use specific Veeam Backup Enterprise Manager features, for example, SAML authentication or restore of Microsoft Exchange items, you can add this component to the domain.

In both cases, backup infrastructure components should be placed to a separate network where applicable. Also, it is recommended to use the hardened backup repository.

## Backup Server

To secure the backup server, consider the following recommendations:

- **Restrict outbound connections**. To enable product update check, automatic license update, and license usage reporting, the backup server must be connected to the internet and be able to send requests to servers on the internet. Allow only HTTPS connections to the Veeam Update Notification Server (`dev.veeam.com`), Veeam License Update Servers (`vbr.butler.veeam.com`, `autolk.veeam.com`), and Microsoft WSUS servers or Microsoft Update sites.
- **Restrict inbound connections**. Inbound connectivity to backup servers from the internet must not be allowed. If you want to manage backup servers remotely over the Internet, you can deploy the Veeam Backup & Replication console on a jump server. Service providers who want to manage backup servers remotely can use the Veeam Backup Remote Access functionality. For more information, see the Using Remote Access Console section in the Veeam Cloud Connect Guide.

  > **NOTE**
  >
  > The account used for RDP access must not have local Administrator privileges on the jump server, and you must never use the saved credentials functionality for RDP access or any other remote console connections. To restrict users from saving RDP credentials, you can use Group Policies. For more information, see this article.

- **Encrypt backup traffic**. By default, Veeam Backup & Replication encrypts network traffic transferred between public networks. To ensure secure communication of sensitive data within the boundaries of the same network, encrypt backup traffic also in private networks. For more information, see Enabling Traffic Encryption.
- **Use multi-factor authentication**. Enable multi-factor authentication (MFA) in the Veeam Backup & Replication console to protect user accounts with additional user verification. For more information, see Multi-Factor Authentication.
- **Use self-signed TLS certificates generated by Veeam Backup & Replication**. This type of certificates is recommended for establishing a secure connection from backup infrastructure components to the backup server. For more information, see Generating Self-Signed Certificate.

- **Reduce the number of user sessions opened for a long time**. Set the idle timeout to automatically log off users. To do this, go to **Users and Roles**, select the **Enable auto log off after <number> min of inactivity** check box, and set the number of minutes.
- **Restrict untrusted Linux VMs and Linux servers to connect to the backup server**. Enable a manual SSH fingerprint verification for machines that do not meet specific conditions. For more information, see Linux Hosts Authentication.
- **Use the recommended Access Control List (ACL) for the custom installation folder**. If you specify a custom installation folder for Veeam Backup & Replication, use the recommended ACL configuration to prevent privilege escalation and arbitrary code execution (ACE) attacks. Remove all inherited permissions from this folder. Then, add the following permissions:
    - Administrators: Full control, applies to this folder, subfolders and files
    - SYSTEM: Full control, applies to this folder, subfolders and files
    - CREATOR OWNER: Full control, applies to subfolders and files only
    - Users: Read & Execute, applies to this folder, subfolders and files

# Veeam Backup & Replication Database

The Veeam Backup & Replication configuration database stores credentials of user accounts required to connect to virtual servers and other systems in the backup infrastructure. All passwords stored in the database are encrypted. However, a user with administrator privileges on the backup server can decrypt passwords which is a potential threat.

To secure the Veeam Backup & Replication configuration database, consider the following recommendations:

- **Restrict user access to the database**. Check that only authorized users can access the backup server and the server that hosts the Veeam Backup & Replication configuration database (if the database runs on a remote server).
- **Encrypt data in configuration backups**. Enable data encryption for configuration backup to secure sensitive data stored in the configuration database. For details, see Creating Encrypted Configuration Backups. Also, ensure that the repository for configuration backups is not located in the same network with the backup server.

# Backup Repositories

To secure data stored in backups and replicas, consider the following recommendations:

- **Follow the 3-2-1 rule**. To build a successful data protection, use the 3-2-1 rule when designing your backup infrastructure. For more information, see Plan How Many Copies of Data You Need (3-2-1 rule).
- **Ensure physical security of all data storage components**. All devices including backup repositories, proxies, and gateway servers must be physically located in an access-controlled area.
- **Restrict user access to backups and replicas**. Check that only authorized users have permissions to access backups and replicas on target servers.
- **Encrypt data in backups**. Use Veeam Backup & Replication built-in encryption to protect data in backups. For more information, see Data Encryption.
- **Encrypt SMB traffic**. If you use SMB shares in your backup infrastructure, enable SMB signing to prevent NTLMv2 relay attacks. Also, enable SMB encryption.

- **Enable immutability for backups**. To protect backup files from being modified or deleted, you can make them immutable. The feature is supported for any tier of scale-out backup repository.
- **Use offline media to keep backup files in addition to virtual storage**. For more information, see Backup Repositories with Rotated Drives and the Tape Devices Support Guide.
- **Ensure security of mount servers**. Machines performing roles of mount servers have access to the backup repositories and ESXi hosts which make them a potential source of vulnerability. Check that all required security recommendations are applied to these backup infrastructure components.

## Veeam Backup Enterprise Manager

To secure Veeam Backup Enterprise Manager server, consider the following recommendations:

- **Install Veeam Backup & Replication server and Veeam Backup Enterprise Manager on different machines**. Deploy Veeam Backup Enterprise Manager on a server different from the Veeam Backup & Replication server to reduce the risk of a key change attack.
- **Enable encryption password loss protection**. This provides an alternative way to decrypt the data if a password for encrypted backup or tape is lost. Even if passwords are lost due to unauthorized access, you can restore lost data using **Veeam Backup Enterprise Manager**. For more information, see Password Loss Protection and Managing Encryption Keys.
- **Use the recommended Access Control List (ACL) for the custom installation folder**. If you specify a custom installation folder for Veeam Backup Enterprise Manager, use the recommended ACL configuration to prevent privilege escalation and arbitrary code execution (ACE) attacks. Remove all inherited permissions from this folder. Then, add the following permissions:
  - Administrators: Full control, applies to this folder, subfolders and files
  - SYSTEM: Full control, applies to this folder, subfolders and files
  - CREATOR OWNER: Full control, applies to subfolders and files only
  - Users: Read & Execute, applies to this folder, subfolders and files

## Veeam Cloud Connect

Veeam Cloud Connect secures communication between the provider side and tenant side with TLS. If an attacker obtains a provider's private key, backup traffic can be eavesdropped and decrypted. The attacker can also use the certificate to impersonate the provider (man-in-the-middle attack). To mitigate risks, Veeam Cloud Connect providers must ensure that the TLS certificate is kept in a highly secure place and cannot be uncovered by a third-party.

# Security & Compliance Analyzer

Veeam Backup & Replication provides a built-in tool to ensure that your backup server configuration follows security best practices for Veeam backup infrastructure components based on Microsoft Windows Server and Linux operating systems.

To perform a security check, open the **Home** tab and click **Security & Compliance** on the ribbon. After that, the **Security & Compliance Analyzer** window opens and the security check starts automatically.

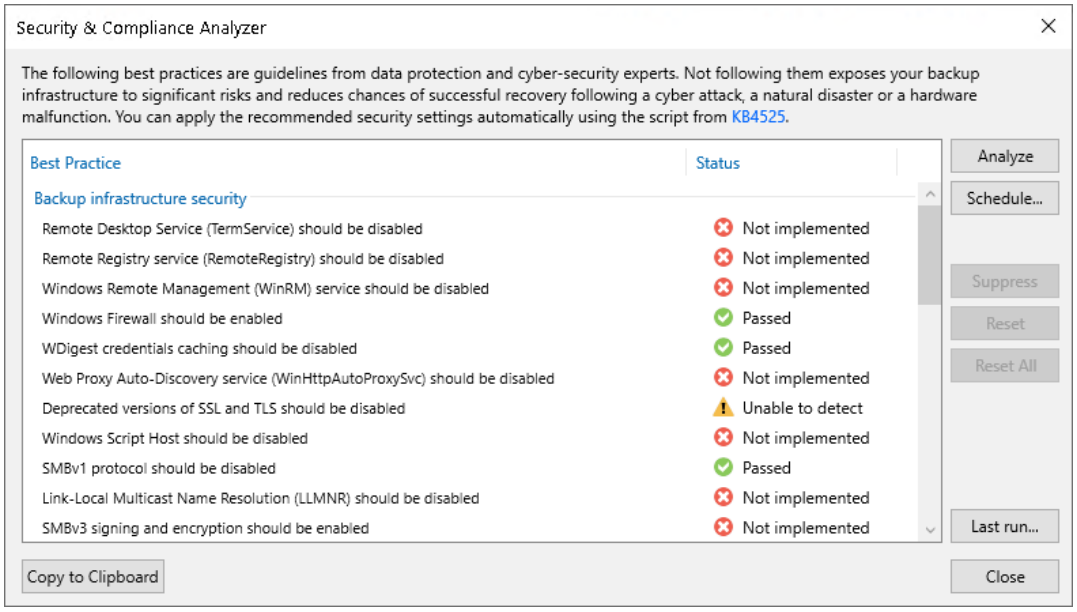


Configuration parameters that have recommended settings will have the *Passed* status. Parameters that have the *Not implemented* status should be revised in terms of your backup infrastructure. You can set them up as recommended or exclude specific parameters from the checklist.

To see the last scan results, click **Last run**. For more information about the results, you can view the log file. The path by default: `C:\ProgramData\Veeam\Backup\Job.BestPracticesAnalyzer.log`.

> **TIP**
>
> Run Security & Compliance Analyzer regularly, especially after you made significant changes in the backup infrastructure. To configure scan scheduling, see this section.

## Configuration Parameters

Security & Compliance Analyzer checks configuration parameters both for the operating system and Veeam products. You can implement these recommendations manually or use the automatic configuration script provided by Veeam. For more information, see this KB article.

| Parameter | Check Condition | Notes |
|---|---|---|
| **Backup Infrastructure Security** | | |

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| Remote Desktop Services (TermService) should be disabled | The Remote Desktop Services service is not running. The **Startup type** parameter is set to *Disabled*. | Remote services should be disabled if they are not needed. Note that for the Veeam Cloud Connect infrastructure, this parameter must be enabled if the SP uses Remote Desktop Protocol (RDP) to connect to the tenant backup server. For more information, see Remote Desktop Connection to Tenant. |
| Remote Registry service (RemoteRegistry) should be disabled | The Remote Registry service is not running. The **Startup type** parameter is set to *Disabled*. | Remote services should be disabled if they are not needed. |
| Windows Remote Management (WinRM) service should be disabled | The Windows Remote Management (WS-Management) service is not running. The **Startup type** parameter is set to *Disabled*. | Remote services should be disabled if they are not needed. |
| Windows Firewall should be enabled | The following PowerShell command returns *True* for `Domain`, `Public`, and `Private` firewall profiles:<br>`Get-NetFirewallProfile \| Format-Table Name, Enabled` | Microsoft Defender Firewall with Advanced Security should be turned on. Also, rules for inbound and outbound connections should be configured according to your infrastructure and Microsoft best practices. For more information, see this Microsoft article. |
| WDigest credentials caching should be disabled | The value of the `HKLM\System\CurrentControlSet\Control\SecurityProviders\WDigest\UseLogonCredential` registry key is set to *0*. | WDigest credentials caching stores cleartext credentials in Windows RAM. To reduce the risk of credential dumping attacks, the setting should be disabled with a registry value. For more information, see this Microsoft article. |

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| Web Proxy Auto-Discovery service (WinHttpAutoProxySvc) should be disabled | The WinHTTP Web Proxy Auto-Discovery service is not running. The **Startup type** parameter is set to *Disabled*.<br><br>The value of the `HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Internet Settings\WinHttp\DisableWpad` registry key is set to *1*. | The Web Proxy Auto-Discovery (WPAD) protocol provides automatic discovery of web proxy configuration. If this feature is not used in the backup infrastructure, the WinHTTP Web Proxy Auto-Discovery Service should be disabled to prevent man-in-the-middle (MITM) attacks. |

| | | |
|---|---|---|
| Deprecated versions of SSL and TLS should be disabled | Values of the following registry keys are set to *1*:<br>• `HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\SSL 2.0\Client\DisabledByDefault`<br>• `HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\SSL 2.0\Server\DisabledByDefault`<br>• `HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\SSL 3.0\Client\DisabledByDefault`<br>• `HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\SSL 3.0\Server\DisabledByDefault`<br>• `HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\TLS 1.0\Client\DisabledByDefault`<br>• `HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\TLS 1.0\Server\DisabledByDefault`<br>• `HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\TLS 1.1\Client\DisabledByDefault`<br>• `HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\TLS 1.1\Server\DisabledByDefault`<br>Values of the following registry keys are set to *0*:<br>• `HKLM\SYSTEM\CurrentControlSet\Cont` | Outdated network protocols SSL 2.0 and 3.0 should be disabled as they have well-known security vulnerabilities and are not NIST-approved. Also, TLS 1.0 and 1.1 should be disabled if they are not needed. For more information, see NIST guidelines.<br>Note that this parameter will have the *Passed* or *Not implemented* status only if specific registry keys with specific values exist. For more information, see this Microsoft article. If the registry key does not exist, the parameter will have the *Unable to detect* status.<br>If the registry key existence cannot be checked for some reason, the parameter will also have the *Not implemented* status. |

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| | rol\SecurityProviders\SCHANNEL\Protocols\SSL 2.0\Client\Enabled<br>• HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\SSL 2.0\Server\Enabled<br>• HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\SSL 3.0\Client\Enabled<br>• HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\SSL 3.0\Server\Enabled<br>• HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\TLS 1.0\Client\Enabled<br>• HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\TLS 1.0\Server\Enabled<br>• HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\TLS 1.1\Client\Enabled<br>• HKLM\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\TLS 1.1\Server\Enabled | |

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| Windows Script Host should be disabled | The value of the `HKLM\SOFTWARE\Microsoft\Windows Script Host\Settings\Enabled` registry key is set to *0*. | Windows Script Host should be disabled to prevent script-based malware attacks.<br>Consider the following:<br>• Make sure that this service is not used by backup infrastructure components you plan to install on the backup server. If there are any (for example, PostgreSQL database), install these components first, then disable the service. To update these components, you need to enable the service temporarily.<br>• To run specific verification scripts, you need to enable the service. For more information, see Custom Verification Scripts. |
| SMBv1 protocol should be disabled | The following PowerShell command returns *False*:<br>`Get-SmbServerConfiguration \| Select EnableSMB1Protocol` | Outdated network protocol SMB 1.0 should be disabled as it has a number of serious security vulnerabilities including remote code execution. For more information, see this Microsoft article. |
| Link-Local Multicast Name Resolution (LLMNR) should be disabled | The `HKLM\SOFTWARE\Policies\Microsoft\Windows NT\DNSClient\EnableMultiCast` registry key exists. The value of the key is set to *0*. | Outdated broadcast protocol Link-Local Multicast Name Resolution (LLMNR) should be disabled to prevent spoofing and man-in-the-middle (MITM) attacks.<br>Note that this parameter will have the *Passed* or *Not implemented* status only if specific registry keys with specific values exist. If the registry key does not exist, the parameter will have the *Unable to detect* status.<br>If the registry key existence cannot be checked for some reason, the parameter will also have the *Not implemented* status. |

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| SMBv3 signing and encryption should be enabled | The following PowerShell command returns *True* for all specified parameters:<br>`Get-SmbServerConfiguration \| select RequireSecuritySignature, EncryptData, EnableSecuritySignature` | If SMB shares are used in the backup infrastructure, SMB signing and encryption should be enabled to prevent NTLMv2 relay attacks. For more information, see these Microsoft articles: Configure SMB Signing with Confidence, SMB security enhancements. |
| Local Security Authority Server Service (LSASS) should be set to run as a protected process | The value of the `HKLM\SYSTEM\CurrentControlSet\Control\Lsa\RunAsPPL` registry key is set to *1* or *2*. | The protection for the Local Security Authority (LSA) process should be configured properly to prevent code injection and credential theft attacks. For more information, see this Microsoft article.<br>If the registry key existence cannot be checked for some reason, the parameter will also have the *Not implemented* status. |
| NetBIOS protocol should be disabled on all network interfaces | The value of `HKLM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\Interfaces\Tcpip_{GUID}\NetbiosOptions` registry keys is set to *2*. | NetBIOS should be disabled to reduce the risk of data theft attacks through shared folders.<br>If the registry key existence cannot be checked for some reason, the parameter will also have the *Not implemented* status. |
| **Product Configuration** | | |
| MFA for the backup console should be enabled | In the **Users and Roles** > **Security** settings, the **Enable multi-factor authentication (MFA)** check box is selected. | Multi-factor authentication (MFA) should be enabled for the Veeam Backup & Replication console to protect user accounts with additional user verification. For more information, see Multi-Factor Authentication. |

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| Immutable or offline (air gapped) media should be used | At least one of the following components is added to the Veeam Backup & Replication console and actively used:<br>• Backup repository with enabled immutability<br>• Backup repository with rotated drives<br>• Tape device | Immutable repositories should be used to protect backup files from being modified or deleted. For more information, see Immutability.<br>Offline media should be used to keep backup files in addition to virtual storage devices. For more information, see Backup Repositories with Rotated Drives and Tape Devices Support. |
| Password loss protection should be enabled | In Veeam Backup Enterprise Manager settings, the **Enable encryption password loss protection** check box is selected. | Password loss protection should be enabled on Veeam Backup Enterprise Manager to provide an alternative way to decrypt the data if a password for encrypted backup or tape is lost. For more information, see Managing Encryption Keys. |
| Backup server should not be a part of the production domain | The backup server is in a workgroup. | For large environments, it is recommended to add the backup server and other backup infrastructure components to a management domain in a separate Active Directory forest. For medium-sized and small environments, backup infrastructure components can be placed to a separate workgroup.<br>Note that this parameter will have the *Passed* status only if the backup server is not joined to any domain. In other cases, it will have the *Unable to detect* status because there is no way to identify the production domain automatically. |
| Email notifications should be enabled | In the global email notification settings, the **Enable e-mail notifications** check box is selected. | Email notifications should be enabled to monitor job statuses. For more information, see Specifying Email Notification Settings. |

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| All backups should have at least one copy (the 3-2-1 backup rule) | At least one of the following jobs or components exists in the Veeam Backup & Replication console:<br>• Backup copy job<br>• Scale-out backup repository with the copy mode<br>• Archive Tier | To be compliant with the 3-2-1 rule, at least one backup copy job should be created, or a scale-out backup repository with the copy mode or archive tier should be added. For more information, see Plan How Many Copies of Data You Need (3-2-1 rule). |
| Reverse incremental backup mode is deprecated and should be avoided | In the backup job settings, the incremental backup method is selected. | The reverse incremental backup method should not be used as it produces the heaviest I/O impact on the backup storage compared to other backup methods. For more information, see Backup Methods. |
| Unknown Linux servers should not be trusted automatically | In the **Options** > **Security** settings, the **Add unknown hosts to the list manually** option is selected in the **Linux hosts authentication** section. | Untrusted Linux VMs and Linux servers must be allowed to connect to the backup server only using manual SSH fingerprint verification. For more information, see Linux Host Authentication. |
| The configuration backup must not be stored on the backup server | In the configuration backup settings, the default backup repository or any other folder on the backup server are not selected as target backup repository. | The configuration backup must not be stored on the backup server or on the default backup repository to be able to recover its configuration in case of failure. For more information, see Configuration Backup. |
| Host to proxy traffic encryption should be enabled for the Network transport mode | For VMware backup proxy that is used the Network transport mode, the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box is selected. | If a VMware backup proxy uses the Network transport mode, it is recommended to transfer VM data over an encrypted TLS connection. For more information about this configuration and its limitations, see Choose Server. |

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| Hardened repositories should not be hosted in virtual machines | The hardened repository added to the Veeam Backup & Replication console is not hosted on a virtual machine. | To reduce the attack surface, the hardened repository should be hosted on a physical machine with local storage. For more information about hardened repository requirements, see Requirements and Limitations. |
| Network traffic encryption should be enabled in the backup network | All global network traffic rules have the **Encrypt network traffic** check box selected. | Network traffic encryption should be enabled in the backup network to ensure secure communication of sensitive data not only between public networks but also between private ones. For more information, see Enabling Traffic Encryption. |
| Linux servers should have password-based authentication disabled | Linux servers added to the Veeam Backup & Replication console do not use standard accounts. | Key-based SSH authentication is generally considered more secure than password-based authentication. The private key is not passed to the server and cannot be captured even if a user connects to a fake server and accepts a bad fingerprint. This helps averting man-in-the-middle (MITM) attacks. |
| Backup services should be running under the LocalSystem account | The Veeam Backup Service runs under a LocalSystem account. | The account used to run Veeam services should be a LocalSystem account. |

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| Configuration backup should be enabled and use encryption | In the configuration backup settings, the following check boxes are selected:<br>• The **Enable configuration backup to the following repository** check box.<br>• The **Enable configuration backup file encryption** check box. | Configuration backup should be enabled to reduce the risk of data loss and manage the Veeam Backup & Replication configuration database easier. For more information, see Configuration Backup and Restore.<br><br>Data encryption for configuration backup should be enabled to secure sensitive data stored in the configuration database. For more information, see Creating Encrypted Configuration Backups. |
| Credentials and encryption passwords should be rotated at least annually | Passwords of the user accounts added to the Credentials Manager, Cloud Credentials Manager, and Password Manager were changed less than 365 days ago. | For all user accounts added to the Credentials Manager, Cloud Credentials Manager and Password Manager, passwords should be changed at least once a year. |
| Hardened repositories should have the SSH Server disabled | Hardened repositories added to the Veeam Backup & Replication console are not available through SSH connection. | SSH connection is necessary only for the deployment of Veeam Data Mover. For security purposes, after adding the hardened repository to the backup infrastructure, the SSH connection should be disabled for the user account used to connect to the Linux server or for the server itself. |
| S3 Object Lock in the Governance mode does not provide true immutability | Immutable object storage repositories added to the Veeam Backup & Replication console use the Compliance retention mode. | The Compliance retention mode should be used for object storage repositories with immutability enabled. This is a more secure option compared to the Governance retention mode. For more information about immutability for object storage repositories, see this section. For more information about retentions modes, see this Amazon article. |

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| Backup jobs to cloud repositories should use encryption | In the backup job settings, if the cloud repository is selected as a backup repository, the **Enable backup file encryption** check box is also selected. | To reduce the cloud attack surface, job-level encryption should be enabled. For more information, see Job Encryption. |
| Latest product updates should be installed | In the **Options** > **Notifications** settings, the **Check for product and hypervisor updates periodically** check box is selected. | Veeam Backup & Replication should be updated regularly. Major releases and cumulative patches usually contain bug fixes, performance enhancements, and new features. |
| PostgreSQL server should be configured with recommended settings | In the `postgresql.conf` file, the following parameters have specific values:<br>• `max_connections = 3000`<br>• `max_wal_senders = 0` | PostgreSQL should have optimal run-time settings to operate correctly.<br>For more information about configuration file, see PostgreSQL documentation. |
| Hardened repositories should not be used as backup proxy servers | The hardened repository added to the Veeam Backup & Replication console is not used as a VMware backup proxy. | A VMware backup proxy requires VMware VDDK components to be installed. To reduce the risk of attacks through VMware VDDK vulnerabilities, a hardened repository should have only one role assigned.<br>For more information about hardened repositories, see Hardened Repository. |

| | | |
|---|---|---|
| Backup encryption password length and complexity recommendations should be followed | Encryption passwords meet Veeam requirements for password complexity. | To minimize the possibility of unauthorized access, encryption passwords should meet Veeam requirements for password complexity:<br>• The password is at least 12 characters long.<br>• The password contains at least:<br>○ One uppercase character<br>○ One lowercase character<br>○ One special character<br>○ One numeric character<br><br>If you want to perform a security check with custom values according to your company password policy, see this KB article.<br><br>If you implement recommended settings for configuration parameters, click **Analyze** to perform a security check again. Make sure that the status changed to *Passed*. |

## Excluding Parameters from Checklist

You can skip security check for specific parameters. For example, if you use Remote Desktop Service to connect to Veeam Backup & Replication and do not need to disable it, exclude this parameter from the checklist. To do this, perform the following steps:

1. Select a parameter and click **Suppress**.
2. [Optional] Leave a comment in the **Note** field.
3. Click **OK**.

Excluded parameters are displayed in the **Suppressed** section. To restore

default settings for the selected parameter and return it to the checklist, click **Reset**. If you want to return all excluded parameters to the checklist, click **Reset All**.

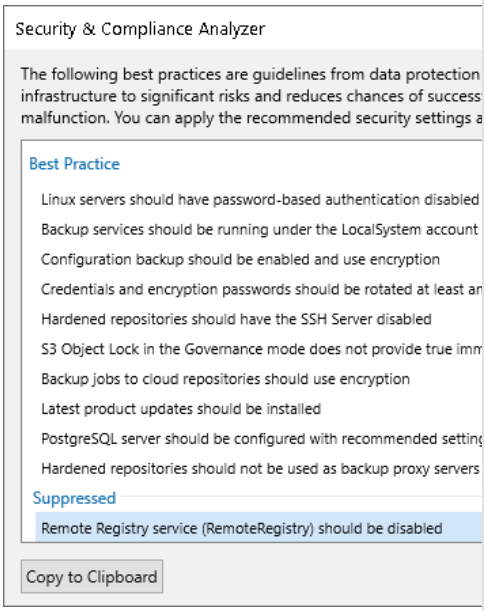


## Scan Scheduling

To configure daily scan scheduling, do the following:

1. Click **Schedule**.
2. Select the **Scan the backup infrastructure daily at** check box and specify the time.
3. If you want to receive scan results by email, select the **Send scan results to the following recipients** check box and specify one or several email addresses separated with a semicolon. You can use global notification settings or specify custom notification settings as required.
4. Click **OK**.

| Parameter | Check Condition | Notes |
| --- | --- | --- |
| | | 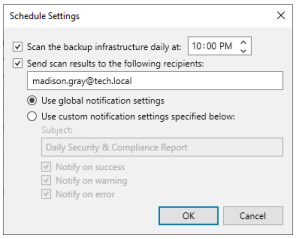 |

# Malware Detection

You can use built-in or third-party malware detection methods to scan backup data and get information about suspicious activity or infected objects. The functionality includes:

- Detecting malware activity in guest indexing data and data stream
- Performing scan backup with different scan engines — Veeam Threat Hunter, third-party antivirus software, or YARA
- Performing secure restore with different scan engines — Veeam Threat Hunter, third-party antivirus software, or YARA
- Integration with third-party malware protection solutions through Veeam Incident API
- Viewing malware detection events
- Receiving daily and immediate reports about malware detection events
- Managing the malware status of the machines marked as *Suspicious* or *Infected*
- Marking specific restore points as *Infected* or *Clean*

For more details about licensing support for specific malware detection features, see Veeam Data Platform Feature Comparison.

## Requirements and Limitations

Malware detection has the following requirements and limitations:

- Malware detection is only supported by specific platforms and applications. For more details, see the supported scenarios of the specific malware detection method:
  - Guest Indexing Data Scan
  - Inline Scan
  - Scan Backup
  - Secure Restore
- Only users with the Veeam Backup Administrator role have full access to the functionality. Users with other roles can view malware detection events and machines marked as *Suspicious* or *Infected*.
- When using Cloud Connect, malware detection is managed by a tenant. A service provider cannot scan tenant data for malware or view the malware detection status.

# How Malware Detection Works

Malware detection is managed by the Veeam Data Analyzer Service. The service restarts once a day at 12:00 AM and starts a new malware detection session. During the session, the Veeam Data Analyzer Service performs the following operations:

- Checks for updates for the list of known suspicious files, extensions and indicators of compromise. For more information, see Configuring Guest Indexing Data Scan .
- Sends an email notification about all malware detection events that were created within the last 24 hours. For more information, see Notifications.
- Initiates a scan session using a specific malware detection method if there is new backup data that needs to be scanned. Otherwise, the service waits for new data to appear.

To transfer malware detection metadata, allow incoming connections from the backup proxy to the backup server on ports 2500 to 3300. For more information, see Ports.

If malware activity is detected, the Veeam Data Analyzer Service does the following:

1. Creates a malware detection event.
2. Marks the machine and the restore point where malware activity was detected for the first time as *Suspicious* or *Infected*.

> **NOTE**
>
> All next restore points created by the original backup job and any additional jobs (backup copy job, backup to tape job, and so on) that include the scanned machine will also be marked as *Suspicious* or *Infected* until the machine is marked as clean. For more information, see Managing Malware Status.
>
> The malware status of machines and restore points is displayed only in the Veeam Backup & Replication console. If you perform restore operations using standalone applications, for example, Veeam Agent for Microsoft Windows, information about the malware status will not be available.

# Malware Detection Methods

Veeam Backup & Replication supports the following malware detection methods:

| Malware detection method | Scan objects | Notes |
|---|---|---|
| File system activity analysis | Guest indexing data | During the backup job, detects the following malware activity:<br>• Known suspicious files and extensions<br>• Indicators of compromise<br>• Deleted files<br>• Extension changes<br>Marks objects as *Suspicious*.<br>For more information, see Guest Indexing Data Scan. |
| Inline entropy analysis | Blocks in a data stream | During the backup job, detects the following malware activity:<br>• Encrypted files<br>• Onion links<br>• Ransom notes<br>Marks objects as *Suspicious*.<br>For more information, see Inline Scan. |
| Signature-based detection (Veeam Threat Hunter) | Restore points | During the Scan Backup session, finds the last clean restore point. For more information, see Veeam Threat Hunter for Scan Backup<br>During the restore session with the Secure Restore option, detects malware activity. For more information, see Veeam Threat Hunter for Secure Restore.<br>During the SureBackup job, detects malware activity. For more information, see SureBackup Job.<br>Marks objects as *Infected*. |

| Malware detection method | Scan objects | Notes |
|---|---|---|
| Third-party antivirus software | Restore points | During the Scan Backup session, finds the last clean restore point. For more information, see Antivirus Scan for Scan Backup.<br><br>During the restore session with the Secure Restore option, detects malware activity as specified in the antivirus configuration file. For more information, see Antivirus Scan for Secure Restore.<br><br>During the SureBackup job, detects malware activity. For more information, see SureBackup Job.<br><br>Marks objects as *Infected*. |
| Rule-based detection (YARA) | Restore points | During the Scan Backup session, does one of the following:<br>• Finds the last clean restore point<br>• Analyzes the content for specific information<br><br>For more information, see YARA Scan for Scan Backup.<br><br>During the restore session with the Secure Restore option, detects malware activity as specified in the YARA rule. For more information, see YARA Scan for Secure Restore.<br><br>During the SureBackup job, detects malware activity. For more information, see SureBackup Job.<br><br>Marks objects as *Infected*. |
| Third-party malware protection solution | Depends on the configuration of the malware protection solution | Uses Veeam Incident API to send a request about detected malware activity to Veeam Backup & Replication.<br><br>Marks objects as *Infected*.<br><br>For more information, see Veeam Backup & Replication REST API Reference. |

# Guest Indexing Data Scan

To scan guest indexing data, Veeam Backup & Replication uses file system activity analysis. The following malware activity can be detected:

- Known Suspicious Files and Extensions
- Indicators of Compromise
- Multiple Deleted Files
- Multiple Extension Changes

# Supported Scenarios

Consider the following:

- You can only scan guest indexing data when backing up the following machines:
  - VMware VMs including VMware Cloud Director VMs
  - Hyper-V VMs
  - Nutanix AHV VMs
  - Machines with Veeam Agent for Microsoft Windows
- Detection of "sleeping" malware is not supported by this method.

## How Guest Indexing Data Scan Works

A guest indexing data scan is started in the following situations:

- When the backup job with guest file system indexing enabled is complete and indexing data is saved in the `VBRCatalog` folder on the backup server.
- If the Veeam Data Analyzer Service gets new indexing data after the service starts.
- If you import backups with the **Import guest file system index data to the catalog** check box selected.

If malware activity is detected, the Veeam Data Analyzer Service will create a malware detection event and mark objects as *Suspicious*.

> **NOTE**
>
> If you upgrade to Veeam Backup & Replication 12.3 from versions older than 12.1, old indexing data will not be scanned.

If you disable guest indexing data scan for a period of time and enable it again, indexing data created during this time will be scanned in the following cases:

- When the next backup job with guest file system indexing enabled is complete.
- When the Veeam Data Analyzer Service restarts.

Note that this process may increase the load on the backup server, depending on the size of the indexing data.

### Known Suspicious Files and Extensions

To detect suspicious files and extensions commonly used by malware, the Veeam Data Analyzer Service compares the guest indexing data of a restore point with the `SuspiciousFiles.xml` file. If files have extensions or names that match any in the `SuspiciousFiles.xml` file, a malware detection event is created.

**NOTE**

Do not edit the `SuspiciousFiles.xml` file directly. You can customize the list of suspicious files and extensions in the following way:

- Add custom suspicious files and extensions.
- Exclude suspicious files and extensions.
- Exclude files and folders.
- Export and import custom lists.

# Adding Custom Suspicious Files and Extensions

To add custom files and extensions that should be marked as suspicious, perform the following steps:

1. From the main menu, select **Malware Detection** > **File Detection** and click **Suspicious files to monitor**.
2. Click **Add** next to the **Suspicious files** field.
3. Specify a file extension or a file name with or without extension. You can also use `*` and `?` wildcard characters. For example:

```
bot.txt
bot
b?t
*.avi
```

**NOTE**

Malware detection scan is case-insensitive. You do not need to add extensions or file names with different cases, for example, bot and Bot.

4. Click **OK**.

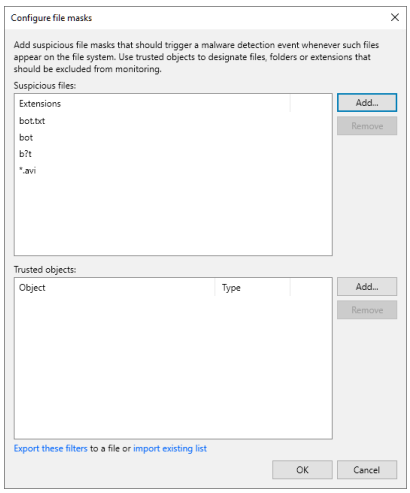

# Excluding Suspicious Files and Extensions

To exclude a file name or file extension listed in the `SuspiciousFiles.xml` file and ignore it during the scan, do the following:

1. From the main menu, select **Malware Detection** > **File Detection** and click **Suspicious files to monitor**.
2. Click **Add** > **Extension** next to the **Trusted objects** field.
3. Specify a file name or file extension as it is listed in the `SuspiciousFiles.xml` file (the `FileMask` tag). For example:

```
*.cryptoboss
```

4. Click **OK**.

> **NOTE**
>
> You can also add files and extensions to the trusted list from the **Event Details** window. For more information, see Viewing Malware Detection Events.

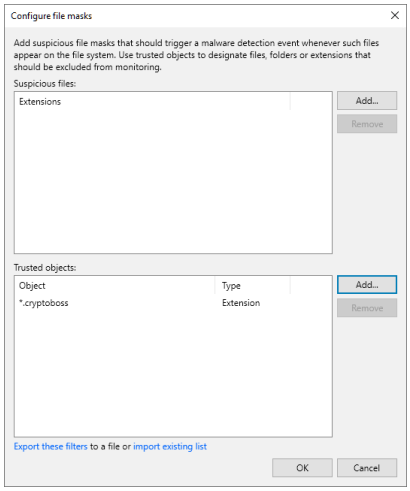


# Excluding Files and Folders

To ignore a specific file or a folder during the scan, do the following:

1. From the main menu, select **Malware Detection** > **File Detection** and click **Suspicious files to monitor**.
2. Click **Add** > **Path** next to the **Trusted objects** field.
3. Specify a path to the file or folder. For example:

```
C:\Users\user1\New Folder\
C:\Users\user1\file.txt
/home/user1/Folder/
/home/user1/file.md
```

4. Click **OK**.

NOTE

Consider the following:

- A path to a folder must include the last slash (" /" or "\" ) symbol.
- A Linux path must start with a slash ("/") symbol.
- A Windows path must start with a drive letter. For example, "C:\".
- UNC paths are not supported.
- Wildcard characters are not supported.
- Excluding files and folders is applied only to the following malware activity types:
  - Known suspicious files and extensions
  - Renamed files
  - Deleted files

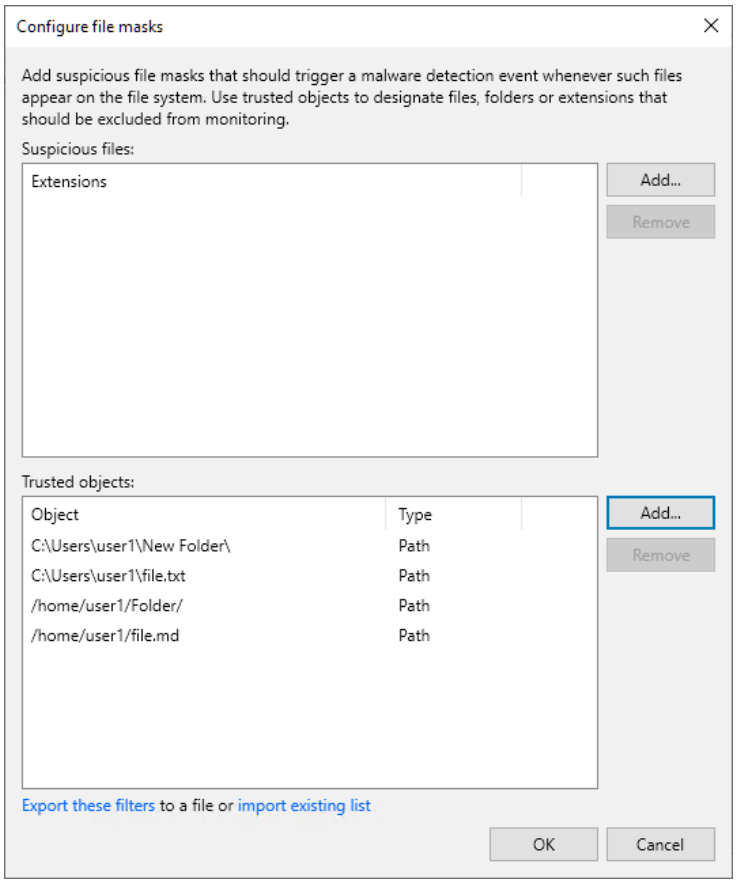


# Exporting and Importing Custom Lists

You can export and import the list of custom files and extensions to/from a file in the XML format.

To export the list, do the following:

1. From the main menu, select **Malware Detection** > **File Detection** and click **Suspicious files to monitor**.
2. Click **Export these filters**.
3. Click **Browse** and select the folder to save the list.
4. Specify the name of the file and click **Save**.

5. Click **OK**.

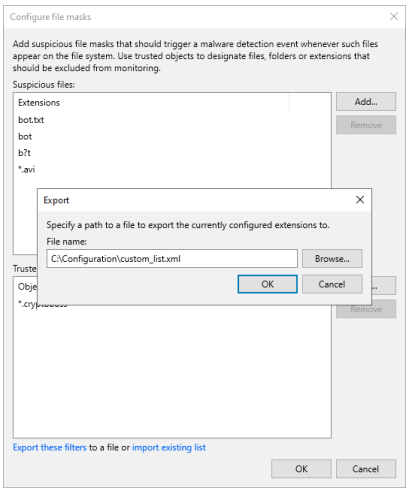


To import the list, do the following:

1. From the main menu, select **Malware Detection** > **File Detection** and click **Suspicious files to monitor**.
2. Click **Import existing list**.

   > **IMPORTANT**
   >
   > The import operation will override files and extensions you specified earlier. It is recommended to export your current list before you import a new one.

3. Click **Browse** and select the folder where the file is located.
4. Select the file and click **Open**.
5. Click **OK**.

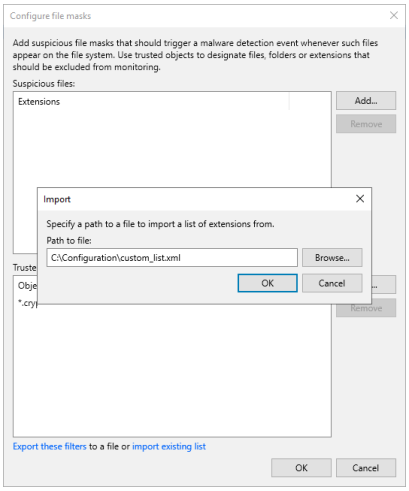


A valid XML Schema for the list must contain the following elements:

| Element | Description |
|---|---|
| `RansomwareExclusions` | Root element. |

| Element | Description |
| --- | --- |
| Includes | Child element of the RansomwareExclusions element. Contains custom files and extensions that must be marked as suspicious. Can be empty. |
| Excludes | Child element of the RansomwareExclusions element. Contains files and extensions listed in the SuspiciousFiles.xml file that must be ignored during the guest indexing data scan. Can be empty. |
| Item (Includes/Excludes) | Child element of the Includes or Excludes element. Each Item contains a file extension or a file name with or without extension that must be marked as suspicious or ignored.<br>Consider the following:<br>• File names and file extensions can include * and ? wildcard characters.<br>• File names and file extensions must not contain the following characters: <, >, :, ", \, /, \|.<br>• Use & to escape an ampersand (&).<br>Note that files and extensions you want to exclude must be listed in the SuspiciousFiles.xml file. |
| IgnoredPaths | Child element of the RansomwareExclusions element. Contains specific files and folders that must be ignored during the guest indexing data scan. Can be empty. |
| Item (IgnoredPaths) | Child element of the IgnoredPaths element. Each Item contains a path to the specific file or folder that must be ignored.<br>The element also has a specific Type attribute. Possible values:<br>• Absolute — Absolute paths<br>Consider the following:<br>• A path to the folder must include the last slash (" /" or "\" ) symbol.<br>• File and folder names can contain spaces.<br>• Paths must not contain the following characters: <, >, ", \|, *, ?.<br>• Use & to escape an ampersand (&). |

Example:

```
<RansomwareExclusions>
   <Includes>
       <Item>bot.txt</Item>
       <Item>*.avi</Item>
       <Item>bot1&amp;2.txt</Item>
   </Includes>
   <Excludes>
       <Item>*.cryptoboss</Item>
   </Excludes>
   <IgnoredPaths>
       <Item Type="Absolute">C:\Users\user1\New Folder\</Item>
       <Item Type="Absolute">C:\Users\user1\file.txt</Item>
       <Item Type="Absolute">/home/user1/Folder/</Item>
       <Item Type="Absolute">/home/user1/file.md</Item>
   </IgnoredPaths>
</RansomwareExclusions>
```

## Indicators of Compromise

Indicators of compromise are non-malware programs. However, their unexpected presence on a system can indicate a security risk. Indicators of compromise are specified in the `SuspiciousFiles.xml` file. They are selected from and categorized using the MITRE ATT&CK Matrix.

To detect indicators of compromise, the Veeam Data Analyzer Service compares guest indexing data from two restore points with the `SuspiciousFiles.xml` file. The later restore point will be compared with the earliest restore point from the same 25-hour period. If there is no restore point, the scan will use the most recent restore point from the same 30-day period. If an indicator of compromise is present only in the later restore point, a malware detection event is created.

> **NOTE**
>
> The indicators of compromise list cannot be customized. To receive new indicators of compromise, select the **Update malware definitions automatically** check box. For more information, see Configuring Guest Indexing Data Scan.

You can exclude specific indicators of compromise from monitoring. To disable an indicator of compromise, do the following:

1. From the main menu, select **Malware Detection** > **File Detection** and click **Tactics to monitor**.
2. Select an indicator of compromise you want to disable.

3. Click **Disable**.

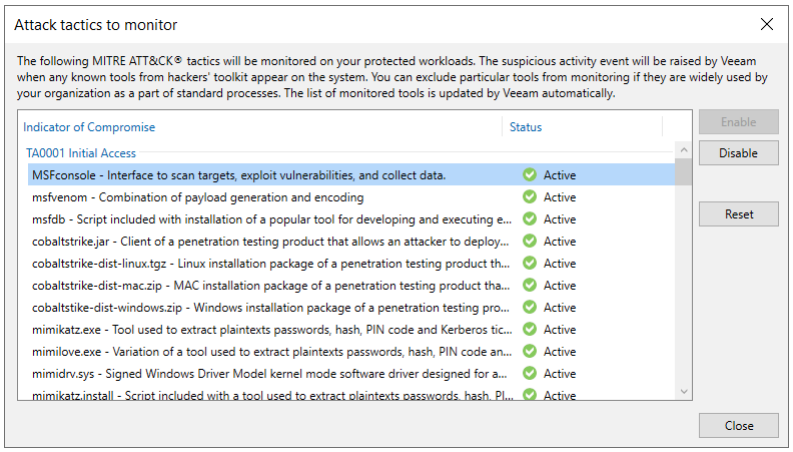


To include an indicator of compromise in monitoring again, do the following:

1. From the main menu, select **Malware Detection** > **File Detection** and click **Tactics to monitor**.
2. Select an indicator of compromise you want to enable.
3. Click **Enable**.

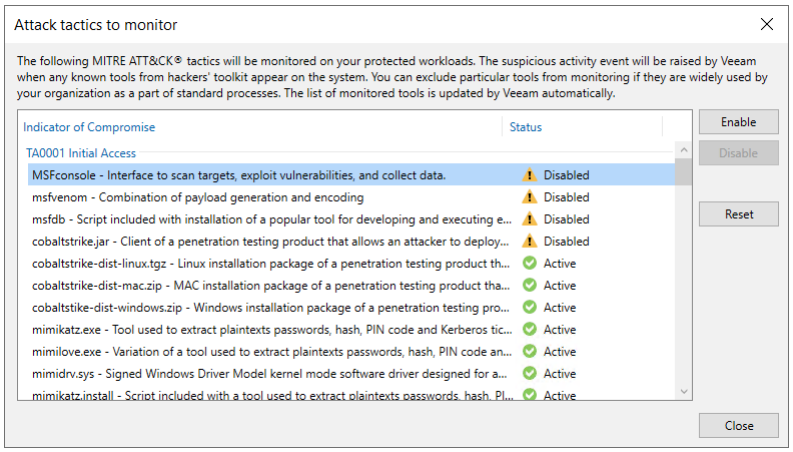


To restore default settings for all indicators of compromise, click **Reset**.

## Multiple Deleted Files

To detect the deletion of multiple files, the Veeam Data Analyzer Service compares guest indexing data from two restore points. The later restore point will be compared with the earliest restore point from the same 25-hour period. If there is no restore point, the scan will use the most recent restore point from the same 30-day period. If the deletion of multiple files is present only in the later restore point, a malware detection event is created.

A malware detection event will be created if the following conditions are met:

- At least 100 files with a specific extension have been deleted.
- The percentage of deleted files compared to the total amount of files with this extension is more than 50%.

Tracked extensions can be found in the `TrackedFiles.xml` file. The default path is `C:\Program Files\Veeam\Backup and Replication\Backup\TrackedFiles.xml`.

You can configure tracked file extensions in the following way:

- Change the file thresholds.
- Add and remove file extensions.
- Add ignored folders.

# Changing File Thresholds

To change scan sensitivity, open the `TrackedFiles.xml` file and specify new values for the following thresholds:

- `Thresholdpercent` — the percentage of files with the specified extension that must be deleted before a malware detection event is created. The default value is 50.
- `Thresholdfiles` — the minimum number of files with the specified extension that must be deleted before a malware detection event is created. The default value is 100.

Make sure that you configure the values correctly. During a scan, both thresholds must be reached to trigger a malware detection event.

> **NOTE**
>
> When a malware detection event is marked as a false positive, the file deletion detection threshold will be automatically increased by 5 percent to a maximum of 90 percent.

# Adding and Removing File Extensions

To start tracking a custom file extension, open the `TrackedFiles.xml` file and add the `FileMaskData` element with the following attributes:

```
<FileMaskData>
     <FileMask>.%fileextension%</FileMask>
     <ThresholdPercent>%1-100%</ThresholdPercent>
     <ThresholdFiles>%integer%</ThresholdFiles>
</FileMaskData>
```

To stop tracking specific file extensions, delete the `FileMaskData` element for the file extension in the `TrackedFiles.xml` file. You can also remove default extensions.

# Adding Ignored Folders

By default, ignored folders are inherited from the `SuspiciousFiles.xml` file and located at the beginning of the `TrackedFiles.xml` file. To add a specific folder to the ignored paths:

1. Open the `TrackedFiles.xml` file.
2. Go to the `IgnoredPaths` element.
3. Add a new `Path` attribute with the folder name as shown in the example below.

```
<IgnoredPaths>
   <Path>$Recycle.Bin</Path>
   <Path>System Volume Information</Path>
   <Path>~Example Folder Name~</Path>
</IgnoredPaths>
```

### Multiple Extension Changes

To detect multiple extension changes, the Veeam Data Analyzer Service compares guest indexing data from two restore points. The later restore point will be compared with the earliest restore point from the same 25-hour period. If there is no restore point, the scan will use the most recent restore point from the same 30-day period. If multiple extension changes are present only in the later restore point, a malware detection event is created.

A malware detection event will be created in the following cases:

- If the following conditions are met:
  - At least 200 files on the system have a new extension.
  - These extensions are not specified in the `SuspiciousFiles.xml` file.
- If one of the following conditions is met:
  - The total number of new unique extensions exceeds 5000.
  - The most common extension of changed files in a folder represents less than 1% of all extensions in this folder.
  - The most common extension of changed files in a folder is greater than or equal to 80% of all extensions in this folder.

Each folder that appears in both restore points and contains at least one file with a new extension is compared.

## Configuring Guest Indexing Data Scan

To enable the guest indexing data scan, do the following:

1. From the main menu, select **Malware Detection**.
2. In the **File Detection** tab, select the **File system activity analysis** check box. This option enables the scan for known suspicious files and extensions, multiple deleted files, and multiple extension changes. For more information, see the following sections:
   - Known Suspicious Files and Extensions
   - Multiple Deleted Files
   - Multiple Extension Changes

   > **NOTE**
   >
   > This functionality is enabled by default when you install or upgrade to Veeam Backup & Replication 12.3.
3. Make sure that you enable guest file system indexing for the required backup job. For more information, see Specify Guest Processing Settings.

To scan guest indexing data for indicators of compromise, select the **File system activity analysis** and **Indicators of compromise detection** check boxes. For more information, see Indicators of Compromise.

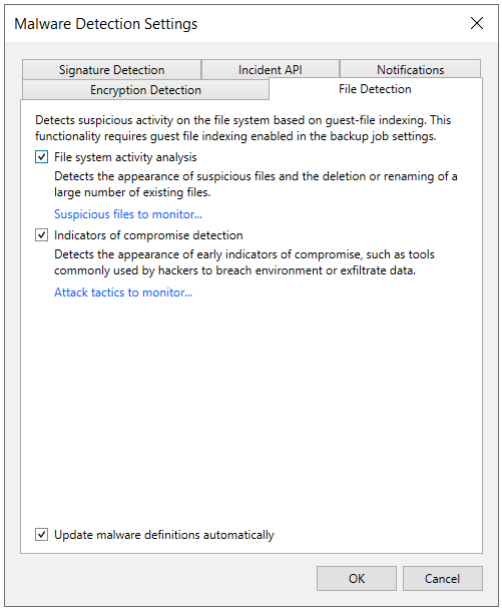


To keep the list of suspicious files and extensions as well as indicators of compromise up to date, do the following:

1. From the main menu, select **Malware Detection**.
2. In the **File Detection** tab, select the **Update malware definitions automatically** check box.

Veeam Backup & Replication will communicate with the Veeam Update Server (`vbr.butler.veeam.com`) daily and download the latest version of the `SuspiciousFiles.xml` file. By default, this occurs once a day at 12:00 AM.

> **NOTE**
>
> If your backup server has limited internet access, you can manually update the `SuspiciousFiles.xml` file. For more information, see this Veeam KB article.

## Logs and Indexing Data Retention

You can view the log of malware activity in the **Event Details** window. For more details, see Viewing Malware Detection Events.

Alternatively, the information about malware activity is stored in malware detection logs. The path by default: `C:\ProgramData\Veeam\Backup\Malware_Detection_Logs`. Logs older than 7 days are automatically archived on Mondays.

By default, guest indexing data is stored for 14 days. If Veeam Backup & Replication is connected to Veeam Backup Enterprise Manager, consider the following, the length of the retention period will change depending on how Veeam Backup Enterprise Manager is deployed:

- If Veeam Backup & Replication is installed on the same server as Veeam Backup Enterprise Manager, the guest index retention period will depend on your Veeam Backup Enterprise Manager settings. For more information, see Configuring Retention Settings for Index and History.

- If Veeam Backup & Replication and Veeam Backup Enterprise Manager are installed on separate servers, all guest indexing data will be stored:
  - In the Veeam Backup Catalog folder for 14 days.
  - In the Veeam Backup Catalog folder on the Veeam Backup Enterprise Manager server for the retention period specified in your Veeam Backup Enterprise Manager settings.

# Inline Scan

To scan blocks in a data stream, Veeam Backup & Replication uses inline entropy analysis. During the backup job, the following malware activity can be detected:

- Files encrypted by malware. A malware detection event will be created if the amount of encrypted data exceeds scan sensitivity limits.
- Text artifacts created by malware:
  - V3 onion addresses that consist of 56 symbols in the `[a-z2-7]{56}.onion` format. For example, *vykenniek4sagugiayj3z32rpyrinoadduprjtdy4wharue6cz7zudid.onion*.
  - Ransomware notes created by Medusa and Clop.

  A malware detection event will be created if a new restore point contains more onion addresses or ransomware notes than the previous restore point selected for comparison. If both restore points contain the same number of onion addresses or ransomware notes, a malware detection event will not be created. For more details, see How Inline Scan Works.

> **NOTE**
>
> Inline scan is disabled by default when you install or upgrade to a new version of Veeam Backup & Replication. If you want to use this functionality, consider the following:
>
> - Scanning may increase CPU usage (10-15% on average) on the backup proxy, depending on the workload type and amount of data.
> - Scanning may increase RAM usage on the backup server, depending on the amount of data. To analyze 1 million files on a machine, inline scan requires approximately 50 MB RAM.

## Supported Scenarios

When performing backup jobs, you can scan blocks in a data stream from the following machines:

- VMware VMs including VMware Cloud Director VMs
- Hyper-V VMs
- Nutanix AHV VMs
- Machines with Veeam Agent for Microsoft Windows operating in the managed by backup server mode (volume-level backup only)

## Requirements and Limitations

The inline scan has the following requirements and limitations:

- Scanning is supported only for simple volumes and for the following file systems: NTFS, ext4, ext3, ext2.

- Scanning dynamic disks and disks encrypted by BitLocker is not supported.
- To store ransomware data, you need enough disk space on the backup server. The disk space calculation is based on the following data:
    - The number of machines.
    - Used disk space per machine.
    - The number of restore points per machine.

  Storing ransomware data per machine requires approximately 270 KB of disk space on the backup server per each 10 GB of used disk space multiplied by the number of restore points.

  For example, a machine has 200 GB of used space and 10 restore points. Storing ransomware data for this machine requires 54 MB (270 KB * 20 * 10 restore points).
- Text artifacts will be detected only if the following conditions are met:
    - The block size of the file system is 4 KB.
    - Text file has the UTF-8 encoding.
    - Text file is not stored in the Master File Table (MFT).
- Detection of "sleeping" malware is not supported.
- Some file types may be unintentionally marked as suspicious during inline scan, for example, Linux packages with LZMA compression, files encrypted with Windows EFS, specific ISO files, and so on. If you have such files, you can mark related malware detection events as false-positive. For more information, see Managing Malware Status.

## How Inline Scan Works

For inline scan, malware detection works in the following way:

1. During the backup job, Veeam Backup & Replication analyzes data blocks metadata and saves ransomware data in the temporary folder on the backup proxy. A file in the RIDX format is created for each disk and contains the following information:
    - Disk metadata (disk name, creation time, disk size, used size, sector size, partition table)
    - Ransomware data for each data block (encrypted data, file types, onion addresses, ransomware notes)

    > **NOTE**
    >
    > If LZMA headers are found, they will be excluded from encrypted data calculation to decrease the number of false positive events.
2. When the backup job is complete, ransomware data is saved in the `VBRCatalog` folder on the backup server. By default, the path is `%volume%:\VBRCatalog\Index\Machines\%machine_name%\%date%%guid%\ransomwareidx`. The Veeam Guest Catalog Service notifies the Veeam Data Analyzer Service about new data that needs to be scanned.
3. The Veeam Data Analyzer Service checks last scan results in the `RansomwareIndexAnalyzeState.xml` file located in the `VBRCatalog` folder and initiates a new inline scan. The scan is also initiated if the Veeam Data Analyzer Service gets new indexing data after the service starts.

4. During the scan, the Veeam Data Analyzer Service compares a new restore point with the earliest one created for the last 25 hours. For example, two restore points were created 10 and 5 hours ago. The new restore point will be compared with the restore point created 10 hours ago.

   If the previous restore point was not created for the last 25 hours, the service tries to find the nearest restore point created for the last 30 days. For example, two restore points were created 2 days and 10 days ago. The new restore point will be compared with the restore point created 2 days ago.

5. The Veeam Data Analyzer Service compares the last and previous RIDX files and updates the `RansomwareIndexAnalyzeState.xml` file. If malware activity is detected, the service will create a malware detection event and mark objects as *Suspicious*.

   If the previous RIDX file is not found, the Veeam Data Analyzer Service will perform a full disk read operation to create a RIDX file. In this case, the job session will last longer than usual but the size of the incremental backup file will not be affected. During this operation the Changed Block Tracking (CBT) option will not be used. For more information about the option, see Changed Block Tracking.

   A full disk read operation will also be performed if you add a new disk to the VM.

   > **NOTE**
   >
   > The first RIDX file is used as a source for the first scan session and will not be analyzed after creation.

## Enabling Inline Scan

To enable inline scan, do the following:

1. From the main menu, select **Malware Detection**.
2. In the **Encryption Detection** tab, select the **Enable inline entropy analysis** check box.
3. Specify the scan sensitivity depending on your backup data and backup infrastructure capabilities. The default value is *Normal*.

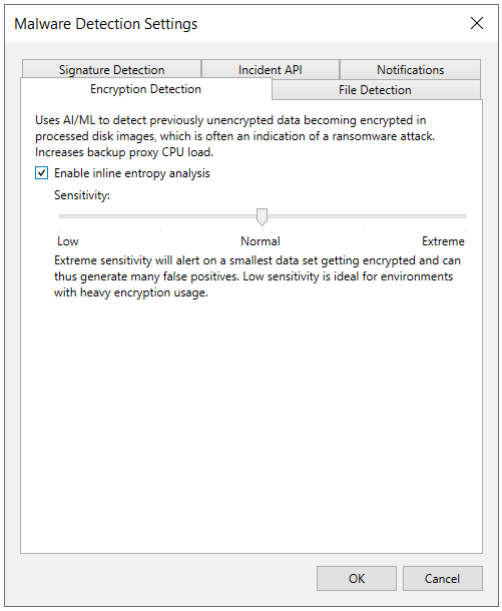

# Scan Backup

Scan Backup allows you to perform the following operations:

- Find the last clean restore point after a recent malware attack.
- Find the last clean restore point if the date of the malware attack is unknown.
- Find some specific information, for example, sensitive data.

To scan restore points, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Scan Backup.
- Third-party antivirus software. For more information, see Antivirus Scan for Scan Backup.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Scan Backup.

# Supported Scenarios

You can run the Scan Backup session for the following backups:

- Image-level virtual machine backups and backup copies of Microsoft Windows VMs (VMware, Hyper-V, Cloud Director, Nutanix AHV, OLVM and RHV, Proxmox VE).
- Physical machine backups and backup copies (Microsoft Windows only).

The following entities are not supported:

- Image-level virtual machine backups and backup copies of Linux VMs.
- Backups stored in the Veeam Cloud Connect repository.
- Backups stored in the archive tier of the scale-out backup repository.
- Storage snapshots.

## How Scan Backup Works

For Scan Backup session, malware detection works in the following way:

1. Veeam Backup & Replication mounts disks of the machine that you plan to scan to the mount server.
2. On the mount server, Veeam Backup & Replication runs the Veeam Mount Service to perform the following steps:
   a. Mount machine disks from backups to the mount server under the `C:\VeeamFLR\<machinename>` folder.
   b. Initiate a new scan session.
3. If you search for the last clean restore point using Veeam Threat Hunter, third-party antivirus software, or YARA, consider the following:
   a. If a clean restore point is found, the Scan Backup session will be finished with the *Success* status. The malware detection event will not be created.

b. If a clean restore point is not found, the Scan Backup session will be finished with the *Failed* status. The malware detection event will be created for each restore point. Objects will be marked as *Infected*.

4. If you check the restore point for sensitive data using YARA, consider the following:

   a. If sensitive data is found, the Scan Backup session will be finished with the *Failed* status.

   b. If sensitive data is not found, the Scan Backup session will be finished with the *Success* status.

   In both cases, the malware detection event will not be created.

By default, the mount server role is assigned to the backup server or a backup repository. However, you can assign the mount server role to any 64-bit Microsoft Windows machine in your backup infrastructure. For example, you may want to run the malware detection scan on a different server for security reasons. For more information about mount server deployment and requirements, see Mount Servers.

## Configuring Scan Backup Session

To run the Scan Backup session, do the following:

1. Open the **Scan Backup** window by doing one of the following:
   - Open the **Inventory** view and select the **Malware Detection** node. Select the required machine and click **Scan Backup** on the ribbon. Alternatively, right-click the machine and select **Scan backup**.
   - Open the **Home** view and select the **Backups** node. Select the job and required machine, and click **Scan Backup** on the ribbon. Alternatively, right-click the machine and select **Scan backup**.
2. Specify the scan mode you want to use:
   - Find the last clean restore point.
   - Find the last clean restore point in range.
   - Scan content of all restore points in range.
3. Specify the scan engine you want to use:
   - To use Veeam Threat Hunter as a scan engine, select the Scan restore points with Veeam Threat Hunter check box. For more information, see Veeam Threat Hunter for Scan Backup.

     If you use third-party antivirus software as an alternative to Veeam Threat Hunter, select the **Scan restore points with the existing antivirus software** check box. For more information, see Antivirus Scan for Scan Backup.
   - To use YARA as a scan engine, select the **Scan restore points with the following YARA rule** check box and specify the YARA file located in the Veeam Backup & Replication product folder. For more information, see YARA Scan for Scan Backup.
   - To use YARA and Veeam Threat Hunter or third-party antivirus software simultaneously, select both check boxes.
4. Configure the scan range. You can specify the following options:
   - Scan all restore points, from the most recent restore point to the oldest one.
   - Scan restore points created during a specific time period.

   If you want to continue the Scan Backup session after the first malware or the first piece of specific information is found, select the **Continue scanning all remaining files after the first occurrence** check box.

5. Click **OK**.

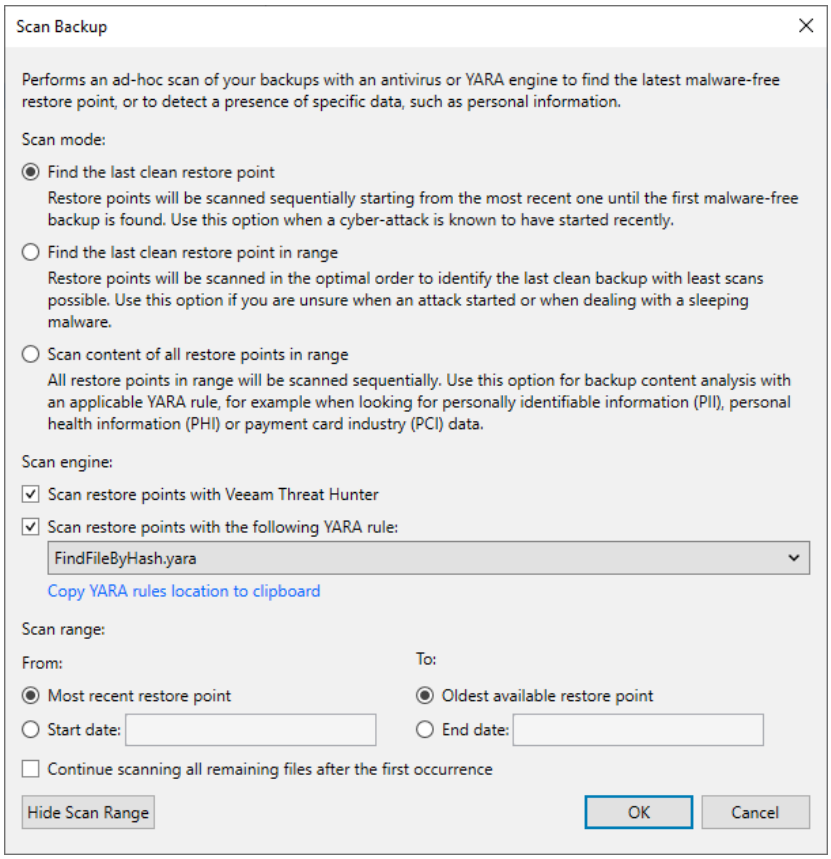


# Veeam Threat Hunter for Scan Backup

For the Scan Backup session, you can run a Veeam Threat Hunter scan to find the last clean restore point. To do this, select the **Scan restored points with Veeam Threat Hunter** check box in the **Scan Backup** window.

## How Veeam Threat Hunter Works

Veeam Threat Hunter is a signature-based scan engine provided by Veeam. It is used as an alternative to third-party antivirus software to scan the restore points. The Veeam Threat Hunter Service is automatically installed on a mount server and runs in the background.

During the Scan Backup session, the Veeam Threat Hunter scan works in the following way:

1. On the mount server, Veeam Backup & Replication runs the Veeam Mount Service to perform the following steps:
   a. Mount machine disks from backups to the mount server under the `C:\VeeamFLR\<machinename>` folder.
   b. Initiate a Veeam Threat Hunter scan.

   > **NOTE**
   >
   > Veeam Threat Hunter checks all disks simultaneously.

2. If Veeam Threat Hunter does not find a clean restore point, the Scan Backup session will be finished with the *Failed* status. The malware detection event will be created for each restore point. Objects will be marked as *Infected*.

3. If Veeam Threat Hunter finds a clean restore point, the Scan Backup session will be finished with the *Success* status. The malware detection event will not be created.

> **NOTE**
>
> Consider the following:
>
> - Veeam Threat Hunter checks updates for malware signatures before running the scan, but not more often than every 1 hour. Note that the initial malware signature update may take longer than the subsequent updates.
> - By default, Veeam Threat Hunter checks all files on disks. If you want to add exclusions, see this KB article.
> - If you deploy a new installation of Veeam Backup & Replication, Veeam Threat Hunter will be selected as a default scan engine in the malware detection settings. The Veeam Threat Hunter Service will be automatically installed on a mount server when you add it to the backup infrastructure.
> - If you upgrade to Veeam Backup & Replication 12.3 or later, the Veeam Threat Hunter Service will be automatically installed on a mount server after the upgrade. For backward compatibility, third-party antivirus software will be selected as a default scan engine in the malware detection settings.

## Requirements and Limitations

Veeam Threat Hunter has the following requirements and limitations:

- To get malware signature updates, a backup server and a mount server must have an internet connection. You need to allow the following incoming connections:
  - From the backup server to the Veeam License Update Server on port 443. For more information, see backup server connections.
  - From the mount server to the Veeam Signature Update Server on port 443. For more information, see mount server connections.

> **NOTE**
>
> Some files may not be scanned because of the file access problems caused by the conflict between Veeam Threat Hunter and third-party antivirus software. If you use Veeam Threat Hunter as a default scan engine, make sure that there is no third-party antivirus software installed on the mount server. Alternatively, you can add exclusions for Veeam Backup & Replication to the third-party antivirus software configuration. For more details, see this KB article.

## Viewing Veeam Threat Hunter Scan Results

Results of the Veeam Threat Hunter scan are available in the Scan Backup session statistics.

To view Veeam Threat Hunter scan results, do one of the following:

- Open the **Home** view. In the inventory pane, select **Last 24 hours**. In the working area, double-click the necessary Scan Backup session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

- Open the **History** view. In the inventory pane, select **Jobs**. In the working area, double-click the necessary Scan Backup session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

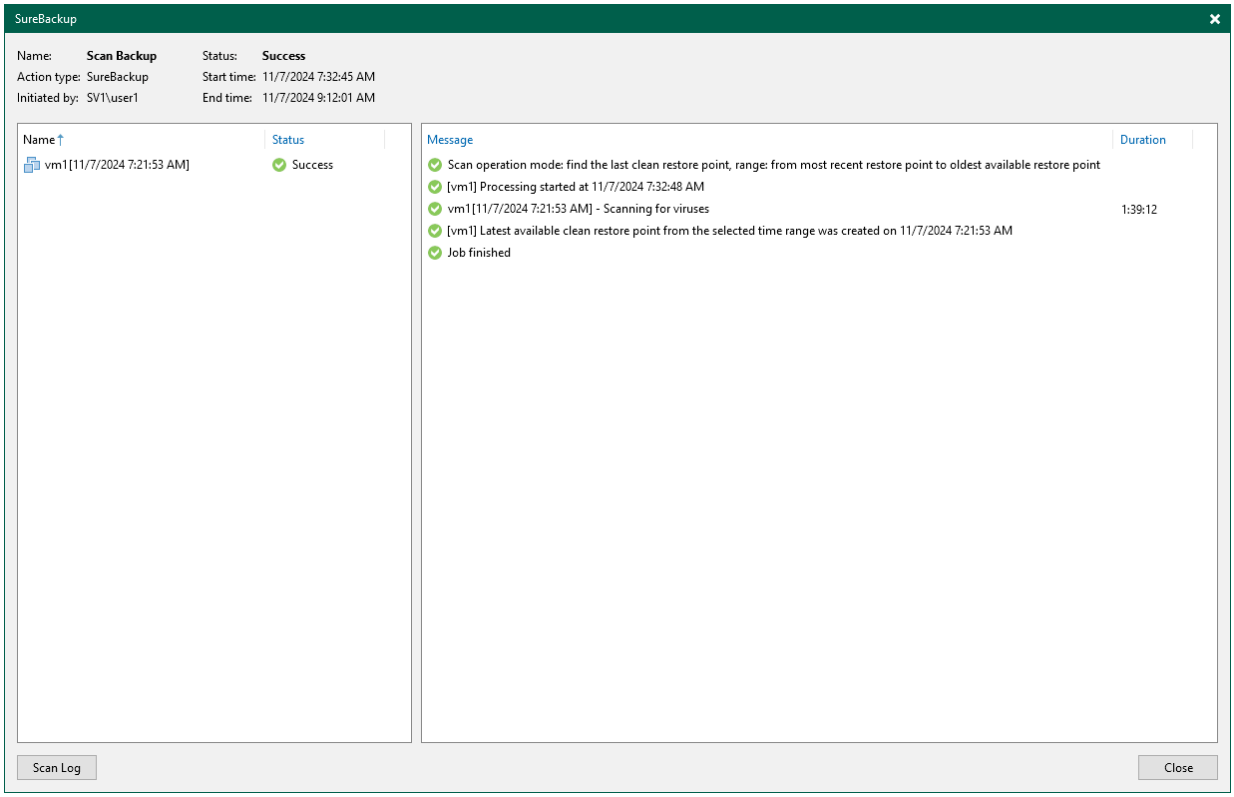


# Antivirus Scan for Scan Backup

For the Scan Backup session, you can run an antivirus scan to find the last clean restore point. To do this, select the **Scan restore points with the existing antivirus software** check box in the **Scan Backup** window.

## How Antivirus Scan Works

During the Scan Backup session, the antivirus scan works in the following way:

1. On the mount server, Veeam Backup & Replication runs the Veeam Mount Service to perform the following steps:
   a. Mount machine disks from backups to the mount server under the `C:\VeeamFLR\<machinename>` folder.
   b. Initiate an antivirus scan.
2. If the antivirus does not find a clean restore point, the Scan Backup session will be finished with the *Failed* status. The malware detection event will be created for each restore point. Objects will be marked as *Infected*.

3. If the antivirus finds a clean restore point, the Scan Backup session will be finished with the *Success* status. The malware detection event will not be created.

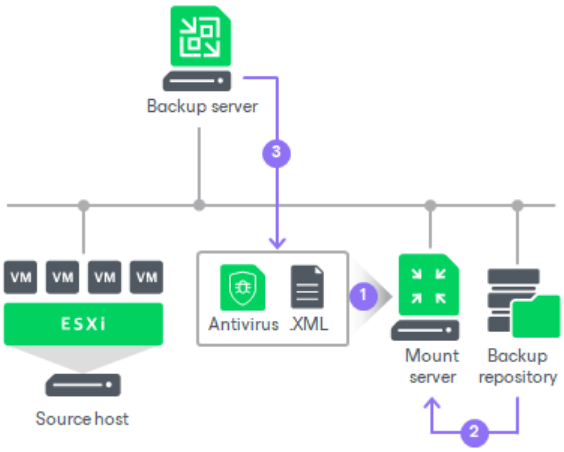


## Requirements and Limitations

The antivirus scan has the following requirements and limitations:

- The antivirus software must be installed on the mount server and support the command line interface (CLI).
- The antivirus configuration file must be located on the mount server and properly configured. For details, see Antivirus Configuration File.

  > **NOTE**
  >
  > If the antivirus is not installed or the configuration file is improperly configured, the Scan Backup session will fail.

- For ESET NOD32 Antivirus version 9 and earlier, the **Continue scanning all remaining files after the first occurrence** option is supported with limitations. If the option is not selected, the antivirus scan will continue for the volume where malware activity was detected. Other volumes will not be scanned.
- If you install several antivirus software on the mount server, the Scan Backup session may fail because of the antivirus conflict and file access problems. To avoid such issues, use only one antivirus software.

## Viewing Antivirus Scan Results

Results of the antivirus scan are available in the Scan Backup session statistics.

To view antivirus scan results, do one of the following:

- Open the **Home** view. In the inventory pane, select **Last 24 hours**. In the working area, double-click the necessary Scan Backup session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

- Open the **History** view. In the inventory pane, select **Jobs**. In the working area, double-click the necessary Scan Backup session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

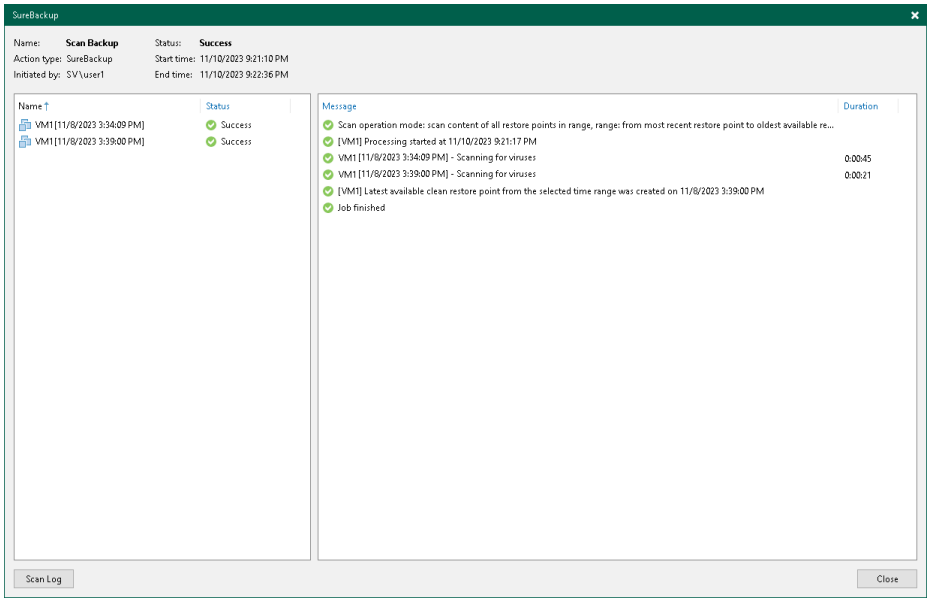


To view the detailed log of the antivirus scan, click the **Scan Log** button at the bottom of the window with the Scan Backup session statistics. Veeam Backup & Replication will display the most recent logs in a file of 1 MB in size.

Full logs of the scan are stored on the mount server in the following folder: `C:\ProgramData\Veeam\Backup\FLRSessions\Windows\FLR__<machinename>_\Antivirus`.

# YARA Scan for Scan Backup

For Scan Backup session, you can run a YARA scan to perform the following operations:

- Find the last clean restore point.
- Analyze the content for specific information, for example, sensitive data.

To perform the YARA scan during the Scan Backup session, do the following:

1. In the **Scan Backup** window, enable the **Scan restore points with the following YARA rule** option.
2. Specify the YARA file located in the Veeam Backup & Replication product folder. The path by default: `C:\Program Files\Veeam\Backup and Replication\Backup\YaraRules`. The YARA file must have the `.yara` or `.yar` extension.

For more information on how to create a YARA rule, see YARA documentation.

## How YARA Scan Works

During the Scan Backup session, the YARA scan works in the following way:

1. On the mount server, Veeam Backup & Replication runs the Veeam Mount Service to perform the following steps:
   a. Mount machine disks from backups to the mount server under the `C:\VeeamFLR\<machinename>` folder.
   b. Initiate a new YARA scan.

2. If you search for the last clean restore point, consider the following:

   a. If a clean restore point is found, the Scan Backup session will be finished with the *Success* status. The malware detection event will not be created.

   b. If a clean restore point is not found, the Scan Backup session will be finished with the *Failed* status. The malware detection event will be created for each restore point. Objects will be marked as *Infected*.

      If you do not want to create a malware detection event for a YARA rule, you can add a `SuppressMalwareDetectionNotification` tag to the name of the rule. For example:

      ```
      rule SearchFileHash : SuppressMalwareDetectionNotification
      ```

      In this case, the malware detection event will not be created but the Scan Backup session will be finished with the *Warning* status.

3. If you check the restore point for sensitive data, consider the following:

   a. If sensitive data is found, the Scan Backup session will be finished with the *Failed* status.

   b. If sensitive data is not found, the Scan Backup session will be finished with the *Success* status.

   In both cases, the malware detection event will not be created.

## Viewing YARA Scan Results

Results of the YARA scan are available in the Scan Backup session statistics.

To view YARA scan results, do one of the following:

- Open the **Home** view. In the inventory pane, select **Last 24 hours**. In the working area, double-click the necessary Scan Backup job. Alternatively, you can select the job and click **Statistics** on the ribbon or right-click the job and select **Statistics**.
- Open the **History** view. In the inventory pane, select **Jobs**. In the working area, double-click the necessary Scan Backup job. Alternatively, you can select the job and click **Statistics** on the ribbon or right-click the job and select **Statistics**.

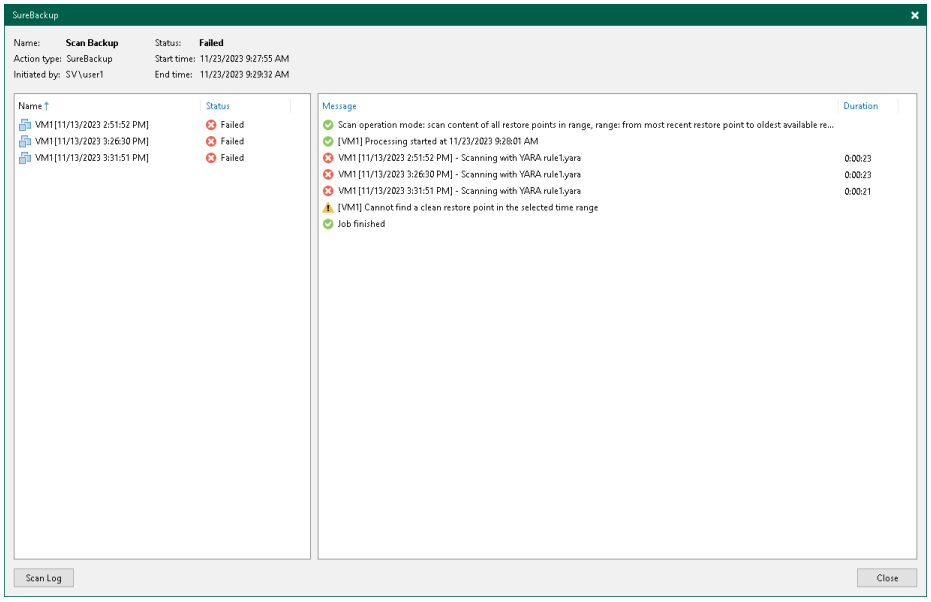

To view the detailed log of the YARA scan, click the **Scan Log** button at the bottom of the window with the Scan Backup job statistics. Veeam Backup & Replication will display the most recent logs in a file of 1 MB in size.

Full logs of the scan are stored on the mount server in the following folder: `C:\ProgramData\Veeam\Backup\FLRSessions\Windows\FLR__<machinename>_\Antivirus`.

# Secure Restore

Secure restore allows you to check restore points for malware activity before restoring the machine to the production environment. To scan restore points, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

# Supported Scenarios

Secure restore is available for the following operations:

- Instant Recovery
- Instant Disk Recovery
- Virtual Disks Restore
- Entire VM Restore
- Restore to Microsoft Azure
- Restore to Amazon EC2
- Restore to Google Compute Engine
- Disk Export

Consider the following:

- You can perform secure restore only for Microsoft Windows machines.
- Veeam Backup & Replication does not perform malware scan for disks or volumes that cannot be mounted to the mount server. For example, Storage Spaces disks or ReFS volumes (if ReFS is not supported by the mount server OS) are skipped from the scan and restored in a regular way.

## How Secure Restore Works

For secure restore, malware detection works in the following way:

1. On the mount server, Veeam Backup & Replication runs the Veeam Mount Service to perform the following steps:
   a. Mount machine disks from backups to the mount server under the `C:\VeeamFLR\<machinename>` folder.
   b. Initiate a new scan session.

2. If malware activity is not detected, Veeam Backup & Replication will restore the machine or its disks to the target location. The malware detection event will not be created.
3. If malware activity is detected, Veeam Backup & Replication will perform the following steps:
    a. Abort the restore process or restore the machine or its disks with restrictions depending on secure restore settings.
    b. Create the malware detection event and mark objects as *Infected*.

By default, the mount server role is assigned to the backup server or a backup repository. However, you can assign the mount server role to any 64-bit Microsoft Windows machine in your backup infrastructure. For example, you may want to run the malware detection scan on a different server for security reasons. For more information about mount server deployment and requirements, see Mount Servers.

> **TIP**
>
> You can also scan machines for malware regularly within a SureBackup job. For information on how to enable the malware scan for a SureBackup job, see the Settings step of the SureBackup job wizard.

# Veeam Threat Hunter for Secure Restore

For restore session, you can run a Veeam Threat Hunter scan to detect malware activity.

## How Veeam Threat Hunter Works

Veeam Threat Hunter is a signature-based scan engine provided by Veeam. It is used as an alternative to third-party antivirus software to scan the restore points. The Veeam Threat Hunter Service is automatically installed on a mount server and runs in the background.

During the restore session, Veeam Threat Hunter scan works in the following way:

1. On the mount server, Veeam Backup & Replication runs the Veeam Mount Service to perform the following steps:
    a. Mount machine disks from backups to the mount server under the `C:\VeeamFLR\<machinename>` folder.
    b. Initiate a new scan session.

    > **NOTE**
    >
    > Veeam Threat Hunter checks all disks simultaneously.

2. If Veeam Threat Hunter does not detect malware activity, Veeam Backup & Replication will restore the machine or its disks to the target location. The malware detection event will not be created.
3. If Veeam Threat Hunter detects malware activity, Veeam Backup & Replication will perform the following steps:
    a. Abort the restore process or restore the machine or its disks with restrictions depending on secure restore settings.
    b. Create the malware detection event and mark objects as *Infected*.

NOTE

Consider the following:

- Veeam Threat Hunter checks updates for malware signatures before running the scan, but not more often than every 1 hour. Note that the initial malware signature update may take longer than the subsequent updates.
- By default, Veeam Threat Hunter checks all files on disks. If you want to add exclusions, see this KB article.
- If you deploy a new installation of Veeam Backup & Replication, Veeam Threat Hunter will be selected as a default scan engine in the malware detection settings. The Veeam Threat Hunter Service will be automatically installed on a mount server when you add it to the backup infrastructure.
- If you upgrade to Veeam Backup & Replication 12.3 or later, the Veeam Threat Hunter Service will be automatically installed on a mount server after the upgrade. For backward compatibility, third-party antivirus software will be selected as a default scan engine in the malware detection settings.

## Requirements and Limitations

Veeam Threat Hunter has the following requirements and limitations:

- To get malware signature updates, a backup server and a mount server must have an internet connection. You need to allow the following incoming connections:
  - From the backup server to the Veeam License Update Server on port 443. For more information, see backup server connections.
  - From the mount server to the Veeam Signature Update Server on port 443. For more information, see mount server connections.

NOTE

Some files may not be scanned because of the file access problems caused by the conflict between Veeam Threat Hunter and third-party antivirus software. If you use Veeam Threat Hunter as a default scan engine, make sure that there is no third-party antivirus software installed on the mount server. Alternatively, you can add exclusions for Veeam Backup & Replication to the third-party antivirus software configuration. For more details, see this KB article.

## Performing Veeam Threat Hunter Scan

To perform Veeam Threat Hunter scan during the restore session, do the following at the **Secure Restore** step of the restore wizard:

1. Enable the **Scan restore points with Veeam Threat Hunter** option.
2. Specify the behavior scenario if malware activity is found. For more information about available options, see the following sections:
   - Secure Restore settings for Instant Recovery
   - Secure Restore settings for Instant Disk Recovery
   - Secure Restore settings for Entire VM Restore
   - Secure Restore settings for Virtual Disk Restore

- Secure Restore settings for Disk Export
- Secure Restore settings for Restore to Microsoft Azure
- Secure Restore settings for Restore to Amazon EC2
- Secure Restore settings for Restore to Google Compute Engine

3. If you want to continue Veeam Threat Hunter scan after the first malware is found, select the **Continue scanning all remaining files after the first occurrence** check box.

## Viewing Veeam Threat Hunter Scan Results

Results of the Veeam Threat Hunter scan are available in restore session statistics.

To view restore session statistics, do one of the following:

- Open the **Home** view. In the inventory pane, select **Last 24 hours**. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.
- Open the **History** view. In the inventory pane, select **Restore**. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

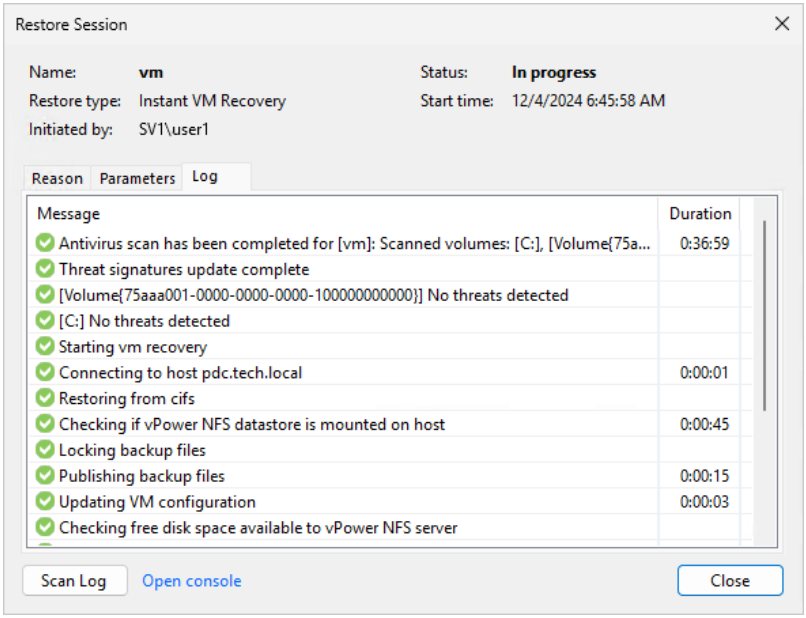


To view the detailed log of the malware scan, click the **Scan Log** button at the bottom of the window with restore session statistics. Veeam Backup & Replication will display the most recent logs in a file of 1 MB in size.

Full logs of the scan are stored on the mount server in the following folder: `C:\ProgramData\Veeam\Backup\FLRSessions\Windows\FLR__<machinename>_\Antivirus`.

# Antivirus Scan for Secure Restore

For the restore session, you can run an antivirus scan to detect malware activity.

## How Antivirus Scan Works

During the restore session, antivirus scan works in the following way:

1. On the mount server, Veeam Backup & Replication runs the Veeam Mount Service to perform the following steps:
   a. Mount machine disks from backups to the mount server under the `C:\VeeamFLR\<machinename>` folder.
   b. Initiate a new scan session.
2. If the antivirus does not detect malware activity, Veeam Backup & Replication will restore the machine or its disks to the target location. The malware detection event will not be created.
3. If the antivirus detects malware activity, Veeam Backup & Replication will perform the following steps:
   a. Abort the restore process or restore the machine or its disks with restrictions depending on secure restore settings.
   b. Create the malware detection event and mark objects as *Infected*.

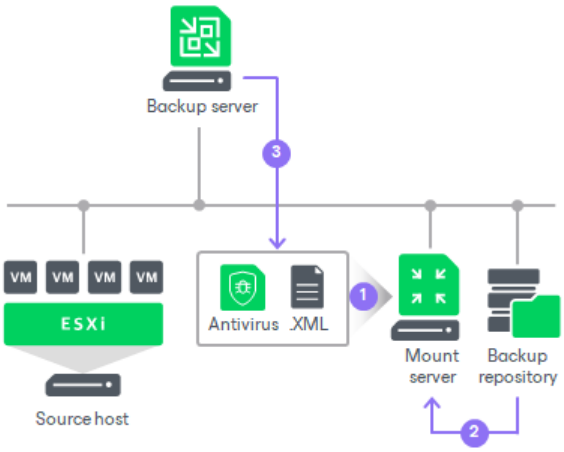


## Requirements and Limitations

The antivirus scan has the following requirements and limitations:

- The antivirus software must be installed on the mount server and support the command line interface (CLI).
- The antivirus configuration file must be located on the mount server and properly configured. For details, see Antivirus Configuration File.
- If you install several antivirus software on the mount server, the restore session may fail because of the antivirus conflict and file access problems. To avoid such issues, use only one antivirus software.

## Antivirus Configuration File

The antivirus software that you plan to use for scanning backups is described in the `AntivirusInfos.xml` file. By default, the file contains predefined settings for the following antivirus software:

- Symantec Protection Engine
- ESET
- Windows Defender

- Bitdefender Endpoint Security Tools
- Trellix (formerly McAfee) — Experimental support

Veeam Backup & Replication creates the `AntivirusInfos.xml` file in the `%ProgramFiles%\Common Files\Veeam\Backup and Replication\Mount Service` folder on every machine with the mount server role. During restore session, Veeam Backup & Replication reads settings from the file and triggers the antivirus to scan backup files. If you use several antivirus software on the mount server, Veeam Backup & Replication will trigger the antivirus whose configuration is defined first in the file.

- See the default configuration file

```
<Antiviruses>
  <!-- Symantec -->
  <AntivirusInfo Name='Symantec' IsPortableSoftware='false' ExecutableFilePath=
'Veeam.Backup.Antivirus.Scan.exe' CommandLineParameters='/p:%Path%' RegPath='HK
EY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\symcscan' ServiceName='symcs
can' ThreatExistsRegEx='Threat\s+found' IsParallelScanAvailable='false'>
     <ExitCodes>
        <ExitCode Type='Success' Description='No threats detected'>0</ExitCode>
        <ExitCode Type='Error' Description='Invalid command line argument'>1</E
xitCode>
        <ExitCode Type='Error' Description='Antivirus scan was completed with e
rrors'>2</ExitCode>
        <ExitCode Type='Error' Description='Antivirus scan was canceled'>4</Exi
tCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>3</Ex
itCode>
     </ExitCodes>
  </AntivirusInfo>
  <!-- Eset -->
  <AntivirusInfo Name='Eset File Security' IsPortableSoftware='true' Executable
FilePath='%ProgramFiles%\ESET\ESET File Security\ecls.exe' CommandLineParameter
s='%Path% /clean-mode=None /no-symlink' RegPath='' ServiceName='' ThreatExistsR
egEx='threat\s*=\s*["&apos;](?!is OK["&apos;])[^"&apos;]+["&apos;]' IsParallelS
canAvailable='false'>
     <ExitCodes>
        <ExitCode Type='Success' Description='No threats detected'>0</ExitCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>1</Ex
itCode>
        <ExitCode Type='Warning' Description='Some files were not scanned'>10</
ExitCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>50</E
xitCode>
        <ExitCode Type='Error' Description='Antivirus scan was completed with e
rrors'>100</ExitCode>
     </ExitCodes>
  </AntivirusInfo>
  <AntivirusInfo Name='ESET Antivirus' IsPortableSoftware='true' ExecutableFile
Path='%ProgramFiles%\ESET\ESET Security\ecls.exe' CommandLineParameters='%Path%
/clean-mode=None /no-symlink' RegPath='' ServiceName='' ThreatExistsRegEx='thre
at\s*=\s*["&apos;](?!is OK["&apos;])[^"&apos;]+["&apos;]' IsParallelScanAvailab
le='false'>
     <ExitCodes>
        <ExitCode Type='Success' Description='No threats detected'>0</ExitCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>1</Ex
itCode>
        <ExitCode Type='Warning' Description='Some files were not scanned'>10</
ExitCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>50</E
xitCode>
        <ExitCode Type='Error' Description='Antivirus scan was completed with e
rrors'>100</ExitCode>
     </ExitCodes>
  </AntivirusInfo>
  <!-- Windows Defender -->
  <AntivirusInfo Name='Windows Defender' IsPortableSoftware='false' ExecutableF
ilePath='%ProgramFiles%\Windows Defender\mpcmdrun.exe' CommandLineParameters='-
Scan -ScanType 3 -File %Path% -DisableRemediation -BootSectorScan' RegPath='HKE
Y_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WinDefend' ServiceName='WinDe
fend' ThreatExistsRegEx='Threat\s+information' IsParallelScanAvailable='false'>
     <ExitCodes>
```

```
        <ExitCode Type='Success' Description='No threats detected'>0</ExitCode>
        <ExitCode Type='Error' Description='Antivirus scan was completed with e
rrors'>2</ExitCode>
        <ExitCode Type='Infected' Description='Virus threat was detected'>2</Ex
itCode>
     </ExitCodes>
  </AntivirusInfo>
  <!-- Bitdefender Endpoint Security Tools -->
  <AntivirusInfo Name='Bitdefender Endpoint Security Tools' IsPortableSoftware=
'true' ExecutableFilePath='%ProgramFiles%\Bitdefender\Endpoint Security\product
.console.exe' CommandLineParameters= ' /c FileScan.OnDemand.RunScanTask custom
path=%Path%' RegPath='' ServiceName='' ThreatExistsRegEx='Remaining issues:\s[1
-9]\d*|Resolved issues:\s[1-9]\d*' IsParallelScanAvailable='false'>
     <ExitCodes>
        <ExitCode Type='Success' Description='Command executed successfully'>0<
/ExitCode>
        <ExitCode Type='Error' Description='Invalid Parameter'>87</ExitCode>
        <ExitCode Type='Error' Description='Bad Arguments'>160</ExitCode>
        <ExitCode Type='Error' Description='Function Failed - an error occurred
while executing the command'>1627</ExitCode>
        <ExitCode Type='Infected' Description='A threat was detected on the sys
tem'>-526</ExitCode>
     </ExitCodes>
  </AntivirusInfo>
  <!-- Trellix (formerly McAfee, experimental support) -->
  <AntivirusInfo Name='Trellix Command Line Scanner' IsPortableSoftware='true'
ExecutableFilePath='<FULL_PATH>\scan.exe' CommandLineParameters='%Path% /NOMEM
/ALL /SUB /PROGRAM /RECURSIVE /UNZIP /ANALYZE /MANALYZE /PANALYZE' RegPath='' S
erviceName='' ThreatExistsRegEx='Found' IsParallelScanAvailable='false'>
     <ExitCodes>
        <ExitCode Type='Success' Description='No viruses found'>0</ExitCode>
        <ExitCode Type='Infected' Description='Virus found'>13</ExitCode>
        <ExitCode Type='Infected' Description='Virus found in memory'>10</ExitC
ode>
        <ExitCode Type='Error' Description='Self-integrity check failed'>15</Ex
itCode>
        <ExitCode Type='Error' Description='DAT file not found'>8</ExitCode>
        <ExitCode Type='Error' Description='There has been a problem with scan'
>6</ExitCode>
        <ExitCode Type='Error' Description='DAT file integrity check failed'>2<
/ExitCode>
     </ExitCodes>
  </AntivirusInfo>
</Antiviruses>
```

# Antivirus Configuration File Structure

The `AntivirusInfos.xml` file contains the following elements:

- `Antiviruses`. Encapsulates the file with antivirus settings.
- `AntivirusInfo`. Describes the antivirus software.

› See AntivirusInfo attributes

| Attribute | Description |
|---|---|
| **Name** | Specifies the antivirus name. Veeam Backup & Replication will display this name in restore session logs. |
| **IsPortableSoftware** | Indicates if antivirus software is portable:<br>• If you set this attribute to `True`, Veeam Backup & Replication will treat the antivirus software as portable. Before performing secure restore, Veeam Backup & Replication will verify if the antivirus executable file exists. The path to the file is specified by the ExecutableFilePath attribute.<br>• If you set this attribute to `False`, Veeam Backup & Replication will treat the antivirus software as non-portable. Before performing secure restore, Veeam Backup & Replication will verify if the antivirus registry value exists and if the antivirus service is running. The key is specified by the RegPath attribute. The service name is specified by the ServiceName attribute. |
| **ExecutableFilePath** | Specifies the full path to the antivirus executable file.<br>**Note:** Some antivirus software uses separate installation folders for different versions. Make sure that you add the full path to the antivirus executable file you use. |
| **CommandLineParameters** | Specifies antivirus commands that you want to execute during the scan. Make sure that the antivirus supports the specified commands. For example, the list of commands for ESET is available in this ESET KB article.<br>**Note:** The `%Path%` variable is required for this attribute. During secure restore, Veeam Backup & Replication substitutes this variable for the path to the folder with mounted disks (`C:\VeeamFLR\<machinename>`). |
| **ServiceName** | Specifies the name of the antivirus service. The service must be responsible for data scanning. The attribute value can be an empty string if **IsPortableSoftware** = True and **ExecutableFilePath** is specified. |
| **RegPath** | Specifies the registry value of the anitivirus service. The attribute value can be an empty string if **IsPortableSoftware** = True and **ExecutableFilePath** is specified. |

| Attribute | Description |
|---|---|
| **ThreatExistsRegEx** | Specifies regular expressions. A regular expression is a sequence of characters that form a search pattern. Veeam Backup & Replication will search the antivirus output messages for the specified regular expression. If any of the output messages match the expression, Veeam Backup & Replication will notify you on detected threat.<br><br>**Note:** You must have a good understanding of the regular expression language to specify this attribute properly. For more information on the regular expression language, see Microsoft Docs. |
| **IsParallelScanAvailable** | Indicates if the antivirus will run multiple jobs to scan files on mounted disks simultaneously.<br><br>If you set this attribute to `True`, Veeam Backup & Replication will lock the antivirus to perform the scan for the current restore session. The antivirus will not be available for other sessions with enabled secure restore until the scan completes.<br><br>The default value for antivirus lock time-out is 24 hours. If the scan does not complete after this period, Veeam Backup & Replication will finish other restore sessions as specified in the restore wizard: abort restore sessions or restore machines (or its disks) with restrictions.<br><br>**Note:** You can change the lock time-out using registry values. For more information, contact Veeam Support.<br><br>If the antivirus CLI does not support multiple scan jobs, set this attribute to `False`. |

- `ExitCodes`. Encapsulates messages that Veeam Backup & Replication displays on scan results.
- `ExitCode`. Describes the subject and the body of the message that Veeam Backup & Replication displays on scan results.

### See ExitCode attributes

| Attribute | Description |
|---|---|
| **Type** | Specifies the subject of the message that Veeam Backup & Replication will display on scan results:<br>• `Success`<br>• `Infected`<br>• `Warning`<br>• `Error` |
| **Description** | Specifies the body of the message that Veeam Backup & Replication will display on scan results. |

# Customizing Antivirus Configuration File

If you want to scan machine data with other antivirus software, make sure that it supports the command line interface (CLI). Then, add configuration for this software to the `AntivirusInfos.xml` file. The configuration must contain the `AntivirusInfo` element with all nested elements and attributes. For more information, see Antivirus Configuration File Structure.

**NOTE**

Consider the following:

- If you made changes to the antivirus configuration file, you do not need to restart Veeam services on the backup server — Veeam Backup & Replication will perform the next malware scan with new settings.
- During the upgrade, the customized `AntivirusInfos.xml` file is replaced by the default one. Do not forget to make necessary changes to it.

**TIP**

You can distribute the XML configuration file among other mount servers in your backup infrastructure using Veeam PowerShell. For more information, see the Copy-VBRAntivirusConfigurationFile section in the Veeam PowerShell Reference.

## Performing Antivirus Scan

To perform the antivirus scan during the restore session, do the following at the **Secure Restore** step of the restore wizard:

1. Enable the **Scan restore points with the existing antivirus software** option.
2. Specify the behavior scenario if malware activity is found. For more information about available options, see the following sections:
   - Secure Restore settings for Instant Recovery
   - Secure Restore settings for Instant Disk Recovery
   - Secure Restore settings for Entire VM Restore
   - Secure Restore settings for Virtual Disk Restore
   - Secure Restore settings for Disk Export
   - Secure Restore settings for Restore to Microsoft Azure
   - Secure Restore settings for Restore to Amazon EC2
   - Secure Restore settings for Restore to Google Compute Engine
3. If you want to continue the antivirus scan after the first malware is found, select the **Continue scanning all remaining files after the first occurrence** check box.

   **NOTE**

   For ESET NOD32 Antivirus version 9 and earlier, the **Continue scanning all remaining files after the first occurrence** option is supported with limitations. If the option is not selected, the antivirus scan will continue for the volume where malware activity was detected. Other volumes will not be scanned.

Note that if the antivirus is not installed or the configuration file is improperly configured, Veeam Backup & Replication will display a warning. In that case, to pass the step with secure restore settings, you can do one of the following:

- Set up the antivirus software properly.
- Clear the **Scan restore points with the existing antivirus software** option.
- Use Veeam Threat Hunter. For more information, see Veeam Threat Hunter for Secure Restore.
- Use YARA scan. For more information, see YARA Scan for Secure Restore.

### Viewing Antivirus Scan Results

Results of the antivirus scan are available in restore session statistics.

To view restore session statistics, do one of the following:

- Open the **Home** view. In the inventory pane, select **Last 24 hours**. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.
- Open the **History** view. In the inventory pane, select **Restore**. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

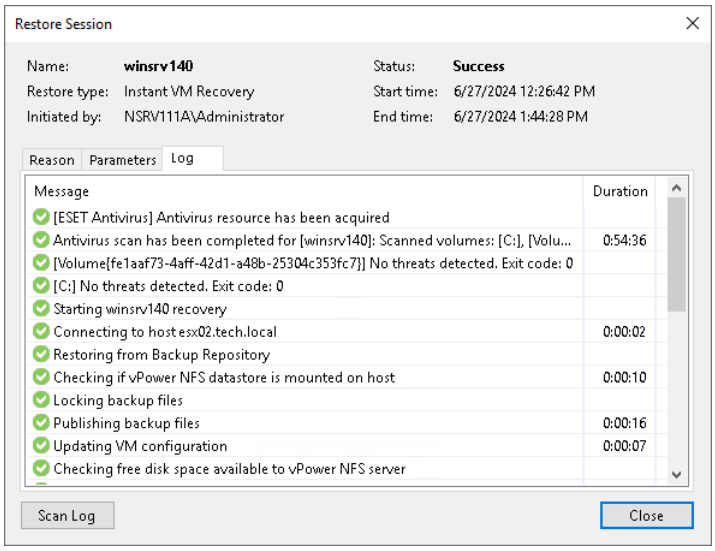


To view the detailed log of the malware scan, click the **Scan Log** button at the bottom of the window with restore session statistics. Veeam Backup & Replication will display the most recent logs in a file of 1 MB in size.

Full logs of the scan are stored on the mount server in the following folder: `C:\ProgramData\Veeam\Backup\FLRSessions\Windows\FLR__<machinename>_\Antivirus`.

## YARA Scan for Secure Restore

For the restore session, you can run a YARA scan to detect malware activity.

## How YARA Scan Works

During the secure restore, YARA scan works in the following way:

1. On the mount server, Veeam Backup & Replication runs the Veeam Mount Service to perform the following steps:
   a. Mount machine disks from backups to the mount server under the `C:\VeeamFLR\<machinename>` folder.
   b. Initiate a new scan session.
2. If malware activity is not detected, Veeam Backup & Replication will restore the machine or its disks to the target location. The malware detection event will not be created.
3. If malware activity is detected, Veeam Backup & Replication will perform the following steps:
   a. Abort the restore process or restore the machine or its disks with restrictions depending on secure restore settings.
   b. Create the malware detection event and mark objects as *Infected*.

   If you do not want to create a malware detection event for a YARA rule, you can add a `SuppressMalwareDetectionNotification` tag to the name of the rule. For example:

   ```
   rule SearchFileHash : SuppressMalwareDetectionNotification
   ```

   In this case, the malware detection event will not be created but the restore session will be finished with the *Warning* status.

## Performing YARA Scan

To perform the YARA scan during the restore session, do the following at the **Secure Restore** step of the restore wizard:

1. Enable the **Scan the restore point with the following YARA rule** option.
2. Specify the YARA file located in the Veeam Backup & Replication product folder. The path by default: `C:\Program Files\Veeam\Backup and Replication\Backup\YaraRules`. The YARA file must have the `.yara` or `.yar` extension. For more information on how to create a YARA rule, see YARA documentation.
3. Specify the behavior scenario if malware activity is found. For more information about available options, see the following sections:
   - Secure Restore settings for Instant Recovery
   - Secure Restore settings for Instant Disk Recovery
   - Secure Restore settings for Entire VM Restore
   - Secure Restore settings for Virtual Disk Restore
   - Secure Restore settings for Disk Export
   - Secure Restore settings for Restore to Microsoft Azure
   - Secure Restore settings for Restore to Amazon EC2
   - Secure Restore settings for Restore to Google Compute Engine

4. If you want to continue the YARA scan after the first malware is found, select the **Continue scanning all remaining files after the first occurrence** check box.

Note that if the YARA rule is not found, Veeam Backup & Replication will display a warning. In that case, to pass the step with secure restore settings, you can do one of the following:

- Check if the YARA file is located in the Veeam Backup & Replication product folder, has the proper syntax and the `.yara` or `.yar` extension.
- Clear the **Scan the restore point with the following YARA rule** option.
- Use Veeam Threat Hunter or third-party antivirus sofware. For more information, see Veeam Threat Hunter for Secure Restore and Antivirus Scan for Secure Restore.

## Viewing YARA Scan Results

Results of the YARA scan are available in restore session statistics.

To view restore session statistics, do one of the following:

- Open the **Home** view. In the inventory pane, select **Last 24 hours**. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.
- Open the **History** view. In the inventory pane, select **Restore**. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

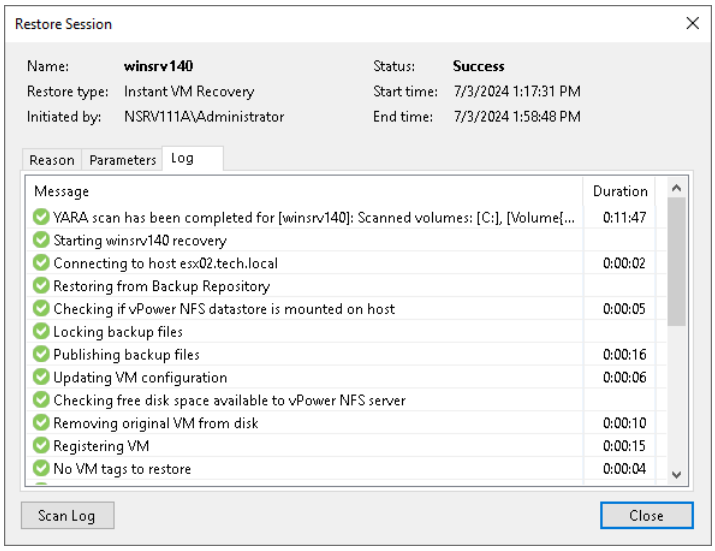


To view the detailed log of the YARA scan, click the **Scan Log** button at the bottom of the window with restore session statistics. Veeam Backup & Replication will display the most recent logs in a file of 1 MB in size.

Full logs of the scan are stored on the mount server in the following folder: `C:\ProgramData\Veeam\Backup\FLRSessions\Windows\FLR__<machinename>_\Antivirus`.

# Configuring Malware Detection

This section describes malware detection settings.

## File Detection

In the **File Detection** tab, you can configure guest indexing data scan settings. For more information, see the Configuring Guest Indexing Data Scan section.

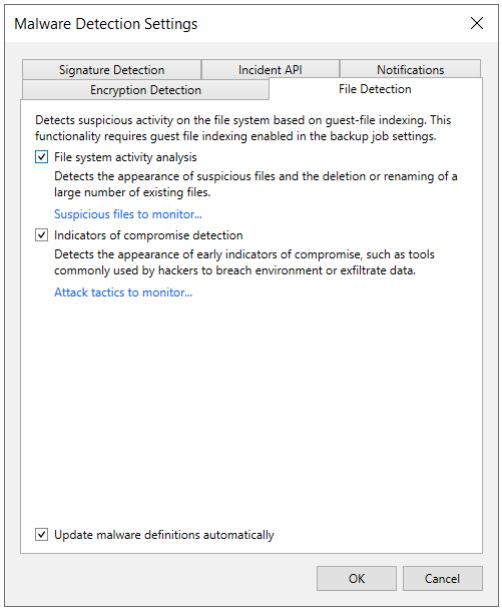


## Encryption Detection

In the **Encryption Detection** tab, you can configure inline scan settings. For more information, see the Enabling Inline Scan section.

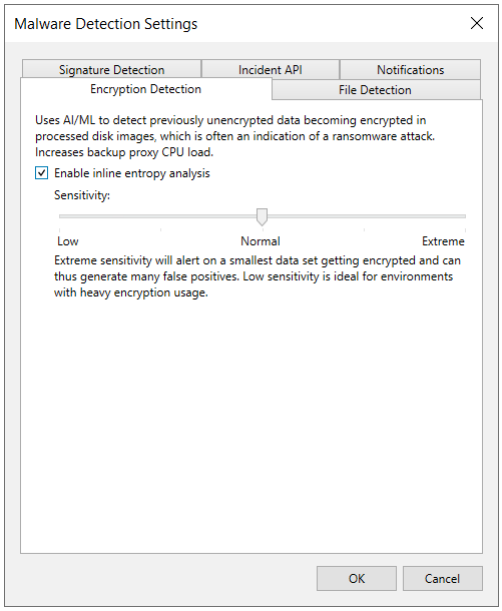

# Signature Detection

You can configure default engines to scan restore points. To do this, perform the following steps:

1. From the main menu, select **Malware Detection** > **Signature Detection**.
2. If you want to use Veeam Threat Hunter, select **Veeam Threat Hunter**. To include additional archive and compound files, such as RAR, ZIP, CHM, installers, and so on in the scan list, select the **Scan content of archive files** check box. Note that this option may increase scan time. For more information about Veeam Threat Hunter, see Veeam Threat Hunter for Scan Backup and Veeam Threat Hunter for Secure Restore.
3. If you want to use third-party antivirus software, select **Bring your own antivirus**. For more information about antivirus scan, see Antivirus Scan for Scan Backup and Antivirus Scan for Secure Restore.

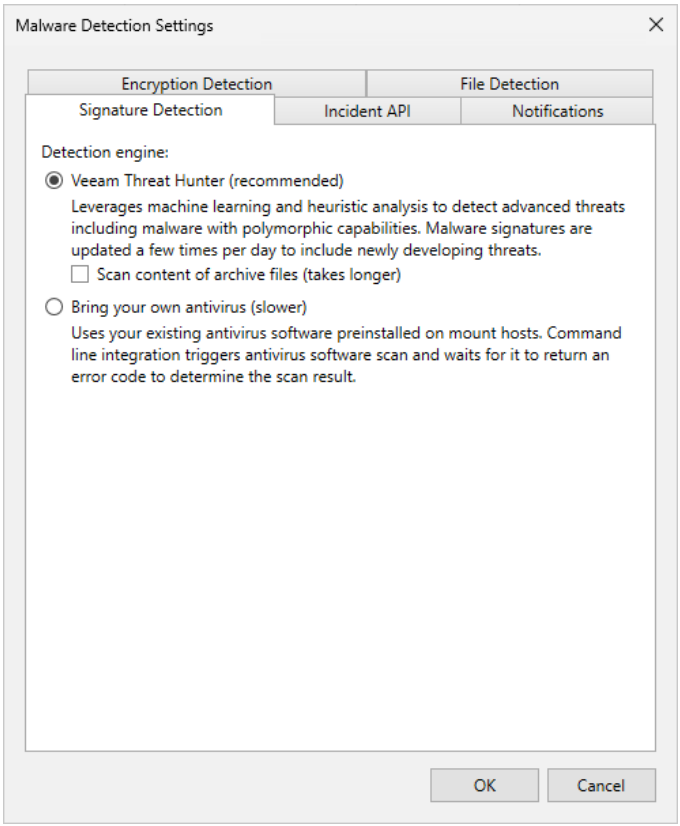


# Veeam Incident API

You can run quick backup when Veeam Backup & Replication gets a Veeam Incident API response that triggers a malware detection event. To do this, perform the following steps:

1. From the main menu, select **Malware Detection** > **Incident API**.
2. Select the **Perform a Quick Backup upon receiving an external event** check box.

> **NOTE**
>
> If there are several machines with malware detection events triggered by Veeam Incident API, quick backup will process these machines one by one.

For more information about quick backup, see Quick Backup. For more information about Incident API methods, see Veeam Backup & Replication REST API Reference.

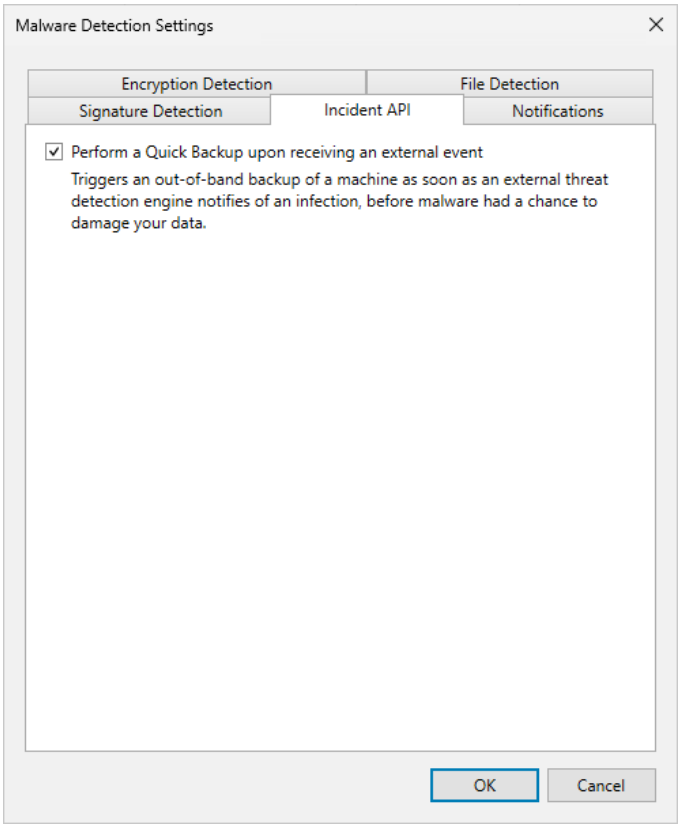


# Notifications

You can receive SNMP traps or email notifications about created malware detection events. To do this, perform the following steps:

1. From the main menu, select **Malware Detection** > **Notifications**.
2. If you want to receive SNMP traps, select the **Send SNMP notifications for detected malware** check box. For more information, see Configuring Global SNMP Settings.
3. If you want to receive email notifications, select the **Send e-mail notifications to the following recipients** check box and specify one or several email addresses separated with a semicolon. You can use global notification settings or specify custom notification settings as required. For more information, see Configuring Global Email Notification Settings.

   Veeam Backup & Replication sends the following email notifications:

   - Daily report is sent when the Veeam Data Analyzer Service restarts the malware detection session (once a day at 12:00 AM) or when you restart the service manually. The report contains consolidated data about all malware detection events that were created within the last 24 hours.

- Immediate report is sent each time a malware detection event is created. If several events were created in 10 seconds, one report with consolidated data will be sent.

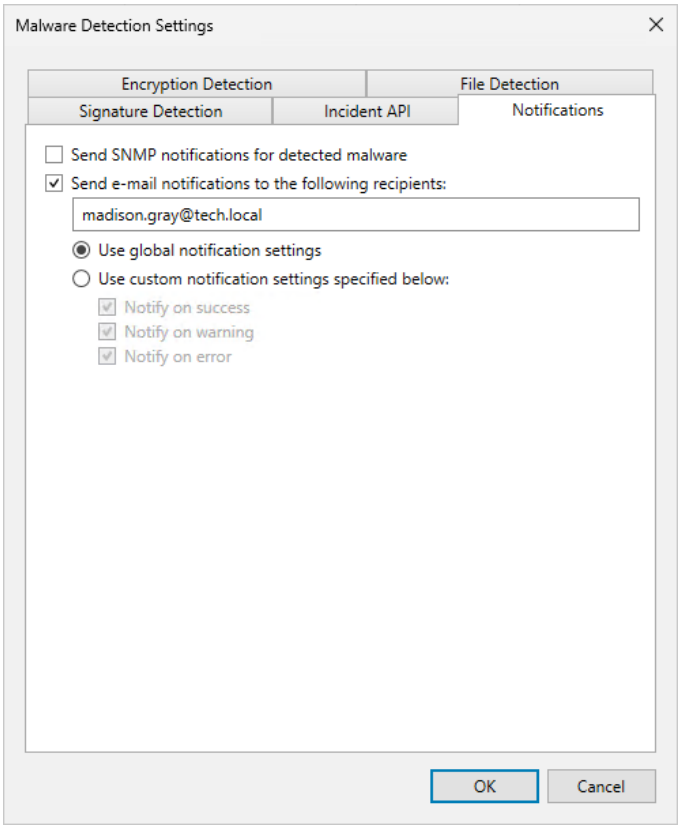


# Malware Exclusions

To exclude specific machines from the malware detection scan, perform the following steps:

1. Open the **Malware Detection Exclusions** window by doing one of the following:
   a. From the main menu, select **Global Exclusion** > **Malware Exclusions**.
   b. Open the **Inventory** view, select the **Malware Detection** node, and click **Exclusions** on the ribbon.
2. Click **Add** and select **VMware vSphere VMs** or **Physical and cloud machines**.
3. Select the object from the list and click **Add**.
4. Click **Note** to provide a description for future reference.
5. Click **OK**.

You can also add the machine to the exclusions list when you mark it as clean. For more information, see Managing Malware Status.

> **NOTE**
>
> Malware exclusions are applied only to guest indexing data scan and inline scan and do not affect scan using Veeam Threat Hunter, third-party antivirus software, or YARA.

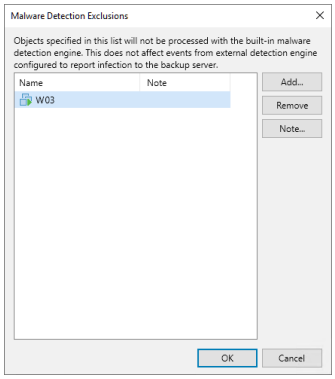
Malware Detection Exclusions
Objects specified in this list will not be processed with the built-in malware detection engine. This does not affect events from external detection engine configured to report infection to the backup server.
Name
Note
W03
Add...
Remove
Note...
OK
Cancel

# Viewing Malware Detection Events

Information about malware detection sessions is displayed in the **History** view under the **Malware Detection** node. The **Malware Detection** > **Malware Events** subnode displays events created as a result of the following operations:

- Performing a backup job with the guest indexing data scan
- Performing a backup job with the inline scan
- Performing a Scan Backup session
- Performing a restore session with the Secure Restore option
- Sending Veeam Incident API requests

> **NOTE**
>
> The **Activity date** column displays the creation date and time of the restore point where malware activity was detected for the first time.

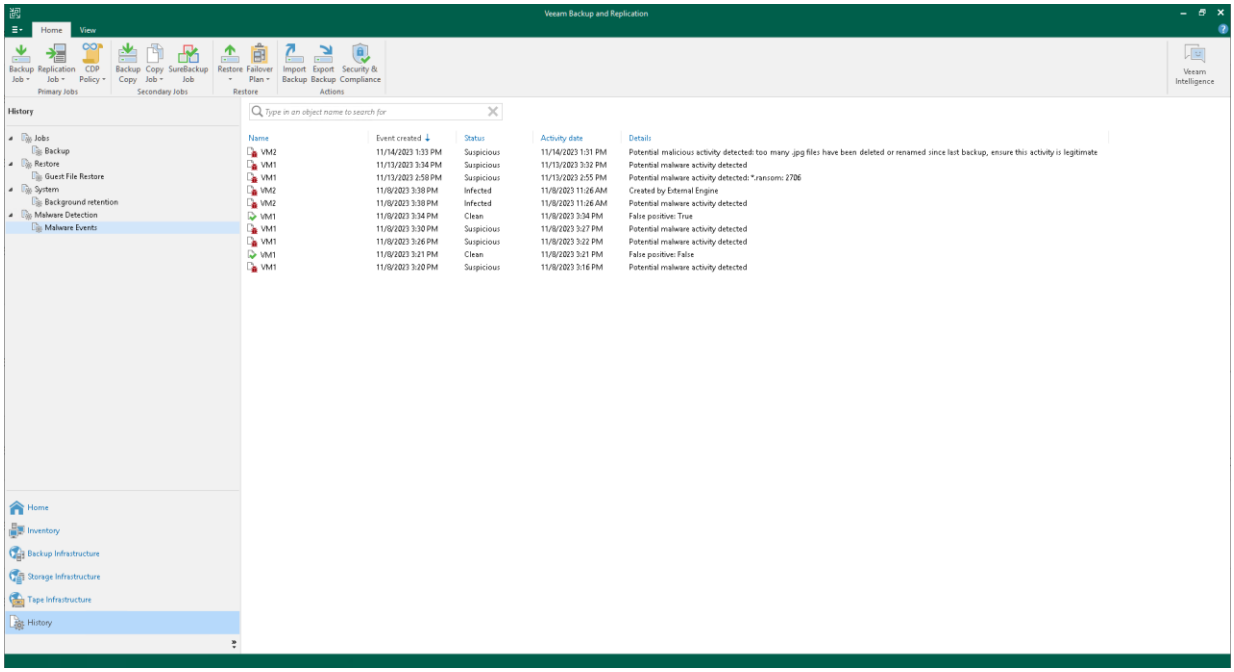


To view detailed information related to specific malware detection event, do one of the following:

- In the working area, double-click the necessary malware detection event.
- Select the malware detection event and click **Details** on the ribbon or right-click the session and select **Details**.

Events created during the guest indexing data scan include additional information:

- Paths to the log files that contain the list of detected suspicious files and extensions.
- Paths to the log files that contain the list of multiple files deleted by malware.

In the **Event Details** window, you can add detected suspicious files and extensions to the trusted list and ignore them during the next guest indexing data scans. To do this, click **Exclude these extensions from monitoring**. For more information on how to manage the list of suspicious files and extensions, see Known Suspicious Files and Extensions.

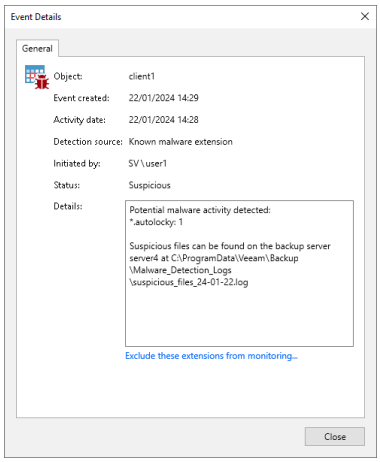

# Managing Malware Status

Veeam Backup & Replication allows you to manage the malware status of machines and specific restore points.

## Marking Machines as Clean

All machines marked by malware detection as suspicious or infected can be found in the **Inventory** view under the **Malware Detection** node.

If you cleaned the machine from malware or the malware detection event was false positive, you can mark the machine as clean. To do this, perform the following steps:

1. Right-click one or more machines and select **Mark as clean**. Alternatively, click **Mark as Clean** on the ribbon.
2. If machines was cleaned from malware, specify the reason and click **Yes**. The malware status of the machine will be automatically updated. Previous restore points will be left with the *Suspicious* or *Infected* status. All next restore points will not be marked as suspicious or infected unless a new malware detection event is created.
3. If the malware detection event was false positive, specify the reason, select the **Mark restore points affected by corresponding detection events as clean** check box, and click **Yes**. The malware status of the machine will be automatically updated. Previous restore points will be marked as clean. Next restore points will not be marked as suspicious or infected unless a new malware detection event is created.

If you want to disable next malware detection scans for machines that you mark as clean, select the **Exclude the workload from malware detection** check box. The machines will be added to the exclusions list. For more information, see Malware Exclusions.

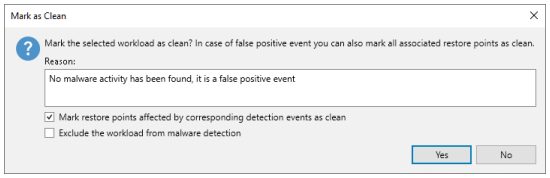


## Managing Malware Status of Specific Restore Points

If you know that a specific machine is infected but the malware detection scan did not detect any suspicious activity, you can manually change the malware status of the specific restore point. To do this, perform the following steps:

1. Open the **Home** view and navigate to the **Backups** node.
2. Right-click the job name and select **Properties**.
3. [Optional] If all backup data is backed up to a single storage and you want to change the malware status of the specific restore point for a specific machine, select the machine in the **Objects** window.

4. In the **Files** window, right-click the restore point and select **Malware** > **Mark as infected**.

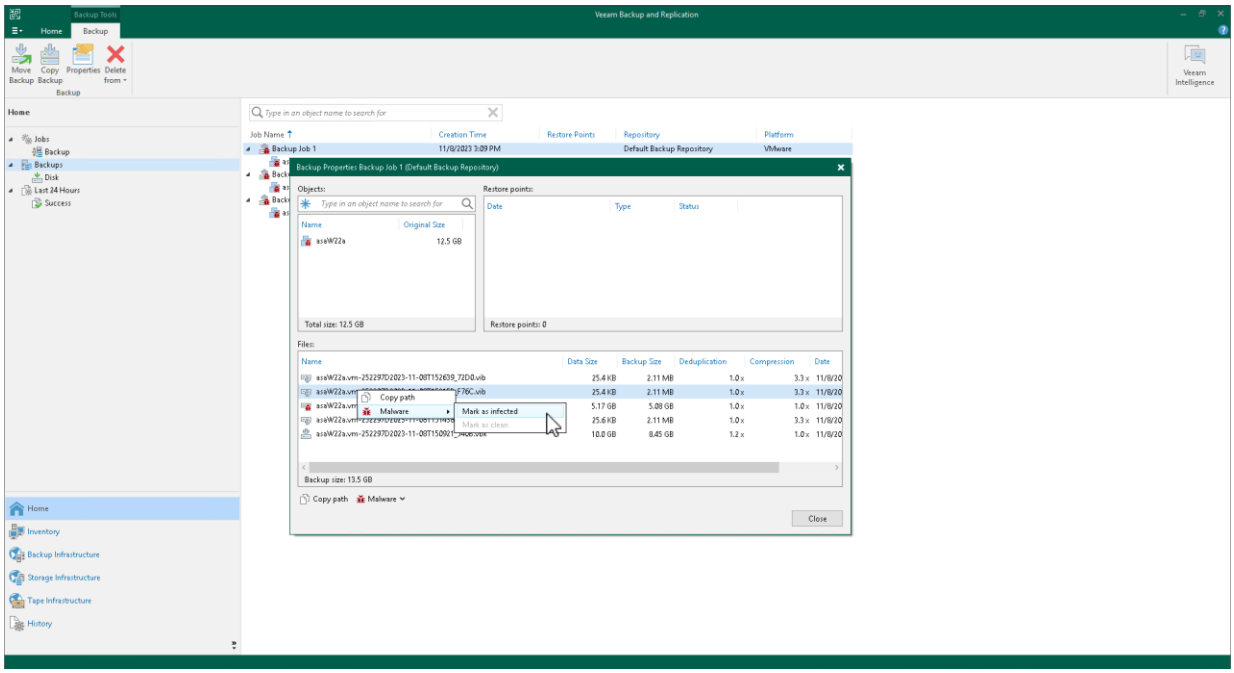


If you know that a specific restore point is not infected, you can manually change the malware status of the specific restore point. To do this, perform the following steps:

1. Open the **Home** view and navigate to the **Backups** node.
2. Right-click the job name and select **Properties**.
3. [Optional] If all backup data is backed up to a single storage and you want to change the malware status of the specific restore point for a specific machine, select the machine in the **Objects** window.
4. In the **Files** window, right-click the restore point and select **Malware** > **Mark as clean**.

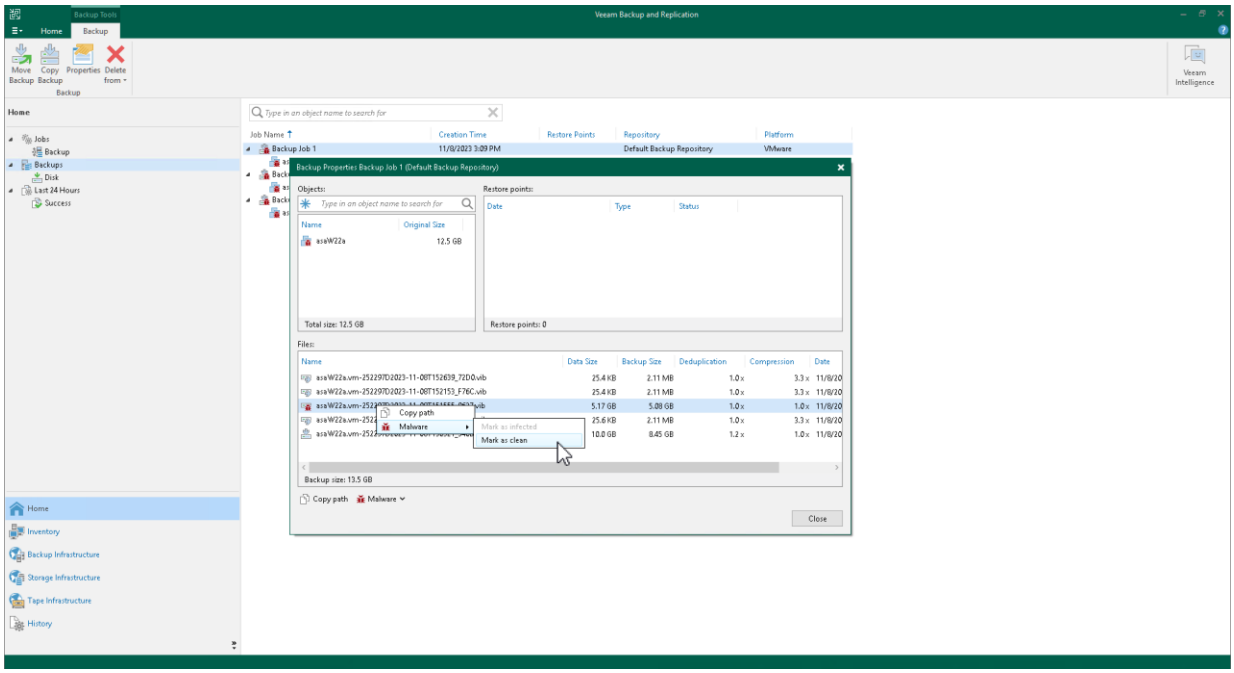

NOTE

If you set the malware status of the restore point manually, it will have a higher priority. The result of the malware detection scan will not affect the malware status of this restore point even if the malware detection event is created.

# Communications Encryption

To secure communication between backup infrastructure components, Veeam Backup & Replication uses the following mechanisms:

- TLS certificates to establish a trusted management connection between components. For more information, see Backup Server Certificate.
- TLS encrypted traffic to transfer backup data between components. For more information, see Enabling Traffic Encryption.

# Communications Encryption Standards

To encrypt communication, Veeam Backup & Replication supports the following libraries, modules, and algorithms:

- SHA-256 — for digital signature generation and SSH fingerprint verification.
- SHA-1 — for HMAC generation, backward compatibility, and certificate thumbprint generation.
- OpenSSL, cryptographic libraries provided by the operating system — for random number generation.
- Veeam Cryptographic Module — for Linux-based components and services. This module is also used for Veeam Data Mover Service installed on Microsoft Windows-based machines.
- Microsoft Crypto API — for other Microsoft Windows-based components and services.
- Microsoft Base Cryptographic Provider. For more information, see Microsoft Docs.
- Microsoft Enhanced RSA and AES Cryptographic Provider. For more information, see Microsoft Docs.
- Microsoft Enhanced Cryptographic Provider. For more information, see Microsoft Docs.

> **NOTE**
>
> If you need Veeam Cryptographic Module and Microsoft Crypto API to be compliant with the Federal Information Processing Standards (FIPS 140), enable FIPS compliance as described in section FIPS Compliance.

Veeam Backup & Replication encrypts certificates stored in the configuration database using Data Protection API (DPAPI) mechanisms. For more information, see Microsoft Docs.

# Supported Cipher Suites and Protocols

Veeam backup infrastructure components support the following TLS versions:

- TLS 1.3 including mutual TLS (mTLS) is partially supported by backup infrastructure components installed on Microsoft Windows Server 2022. PowerShell components and components using OpenSSL do not support TLS 1.3.
- TLS 1.2 including mutual TLS (mTLS).

> **NOTE**
>
> For security reasons, disable outdated protocols TLS 1.0 and 1.1 if they are not needed. For more information, see this Microsoft article.

To avoid negotiation problems between Veeam Backup & Replication and a Microsoft Windows server, ensure that both sides of communication support the same cipher suites.

To avoid negotiation problems between Veeam Backup & Replication and a Linux server, the latter should use ciphers, Key Exchange (KEX) algorithms, and MAC algorithms compatible with SSH libraries supported by Veeam Backup & Replication:

| Algorithms | Supported values |
|---|---|
| Ciphers | Recommended: *aes128-ctr*, *aes192-ctr*, *aes256-ctr*, *aes128-gcm@openssh.com*, *aes256-gcm@openssh.com*, *chacha20-poly1305@openssh.com*<br><br>Supported for backward compatibility: *3des-cbc*, *3des-ctr*, *aes128-cbc*, *aes192-cbc*, *aes256-cbc*, *arcfour*, *arcfour128*, *arcfour256*, *blowfish-cbc*, *blowfish-ctr*, *cast128-cbc*, *twofish-cbc*, *twofish128-cbc*, *twofish128-ctr*, *twofish192-cbc*, *twofish192-ctr*, *twofish256-cbc*, *twofish256-ctr* |
| KEX algorithms | Recommended: *diffie-hellman-group-exchange-sha256*, *diffie-hellman-group14-sha256*, *diffie-hellman-group15-sha512*, *diffie-hellman-group16-sha512*, *ecdh-sha2-nistp256*, *ecdh-sha2-nistp384*, *ecdh-sha2-nistp521*, *curve25519-sha256*, *curve25519-sha256@libssh.org*<br><br>Supported for backward compatibility: *diffie-hellman-group1-sha1*, *diffie-hellman-group14-sha1*, *diffie-hellman-group-exchange-sha1* |
| MAC algorithms | Recommended: *hmac-sha2-256*, *hmac-sha2-512*<br><br>Supported for backward compatibility: *hmac-md5*, *hmac-md5-96*, *hmac-sha-256-96*, *hmac-sha1-96*, *hmac-sha2-512-96*, *hmac-ripemd160*, *hmac-ripemd160@openssh.com*, *hmac-sha1* |

Ensure that your SSH configuration on the Linux server allows you to use at least one cipher, KEX algorithm, and MAC algorithm from the table above. You can run the following command to verify the list of allowed algorithms:

```
sudo sshd -T | grep "\(ciphers\|macs\|kexalgorithms\)"
```

# Trusted Certificates

There are several root and intermediate certificates necessary for the Veeam Backup & Replication to operate correctly. Removal of these certificates from the backup server may limit the functionality of Veeam Backup & Replication or may cause it to fail.

In most cases, these certificates are already installed on Microsoft Windows machines. Some Microsoft Windows installations do not contain needed certificate authorities as trusted certificates, or have non-current certificates. This may happen on servers with locked down security settings, or servers with no internet access or if the latest updates are not installed.

Make sure the following certificates are installed on the backup server:

- Root certificates:
  - https://www.digicert.com/CACerts/DigiCertAssuredIDRootCA.crt (DigiCert Assured ID Root CA)
  - https://www.digicert.com/CACerts/DigiCertHighAssuranceEVRootCA.crt (DigiCert High Assurance EV Root CA)
  - https://support.globalsign.com/customer/portal/articles/1426602-globalsign-root-certificates (install R1 and R3 certificates)
  - https://files.entrust.com/root-certificates/entrust_g2_ca.cer (Entrust Root Certification Authority — G2)
  - https://files.entrust.com/root-certificates/CSBR1.cer (Entrust Code Signing Root Certification Authority — CSBR1)
- Intermediate certificates:
  - https://www.digicert.com/CACerts/DigiCertEVCodeSigningCA-SHA2.crt (DigiCert EV Code Signing CA — SHA2)
  - https://files.entrust.com/subca-certificates/EVCS2-CSBR1-crosscert.cer (Entrust Extended Validation Code Signing CA — EVCS2)

If your backup server does not have internet access, download certificate files from another computer.

# Data Encryption

Data security is an important part of the backup strategy. You must protect your information from unauthorized access, especially if you back up sensitive VM data to off-site locations or archive it to tape. To keep your data safe, use data encryption.

Data encryption transforms data to an unreadable, scrambled format. If encrypted data is intercepted, it cannot be unlocked and read by the eavesdropper. Only intended recipients who know the secret key can reverse encrypted information to a readable format.

This section contains information about encrypting backup files. For more information about encrypting communication between backup infrastructure components, see Communications Encryption.

# How Data Encryption Works

In Veeam Backup & Replication, data encryption works at the following levels:

- Job encryption
- Storage encryption

When you enable data encryption, Veeam Backup & Replication encrypts backup files with data encryption keys, transfers encrypted data to the target side and stores it in a backup repository or a tape device.

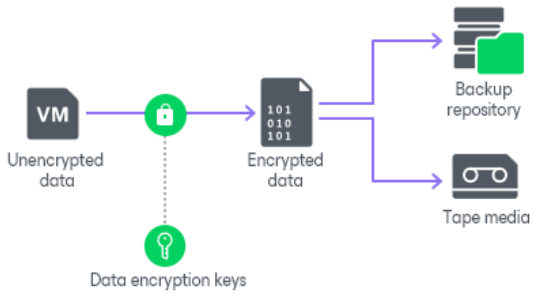


## Data Encryption Keys

To encrypt each type of backup files, Veeam Backup & Replication uses specific data encryption keys. These keys are encrypted with the secret key based on the password or KMS key and stored in the configuration database. Data encryption keys are not passed to the target side except for the following situations:

- You run a backup copy job over WAN accelerators.
- You perform health check for the encrypted backup files.
- You use tapes with enabled encryption. Depending on the encryption level, data encryption keys are passed to the tape library (hardware encryption) or to the tape server (software encryption).

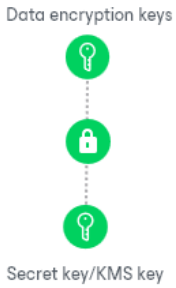


When you enable encryption for the job or storage, you can select the encryption method you want to use — a password that will generate secret keys or an external KMS server that will generate KMS keys.

For secret keys, consider the following:

- Veeam Backup & Replication generates a symmetric secret key based on the password you specify in the encryption options. To create a strong password, follow the recommendations:
  - The password must be at least 12 characters long.
  - The password must contain uppercase and lowercase characters.
  - The password must be a mixture of alphabetic, numeric and punctuation characters.

    - The password must significantly differ from the password you used previously.
    - The password must not contain any real information related to you, for example, date of birth, your pet's name, your logon name, and so on.
    - The password hint must not contain the password.
    - The password must be updated regularly.
- To generate a secret key, Veeam Backup & Replication uses 600,000 HMAC-SHA256 iterations and a 512-bit salt according to Recommendation for Password-Based Key Derivation.
- To encrypt data encryption keys with secret keys, Veeam Backup & Replication uses the 256-bit AES with a 256-bit key length in the CBC mode. For more information, see Advanced Encryption Standard (AES).

For KMS keys, consider the following:

- Asymmetric KMS keys are managed and rotated by an external KMS server. Veeam Backup & Replication only sends requests to the KMS server. For more information about KMS keys, see Key Management System Keys.
- To encrypt data encryption keys with KMS keys, Veeam Backup & Replication uses the RSA algorithm with a 4096-bit key length. For more information, see RSA Cryptography Specifications.

**NOTE**

Consider the following:

- To protect data encryption keys, passwords, secret keys, KMS keys and other sensitive information stored in the configuration database, Veeam Backup & Replication uses Data Protection API (DPAPI) mechanisms. For more information, see Microsoft Docs.
- If you have lost or forgotten a password or you cannot use KMS keys due to a KMS server failure, you will not be able to recover data from backups or tapes unless you use Enterprise Manager keys in the encryption process. For more information, see Password Loss Protection.

# Data Encryption and Deduplication

Data encryption has a negative effect on the deduplication ratio if you use a deduplicating storage appliance as a target. Veeam Backup & Replication uses different data encryption keys for every job session. For this reason, encrypted data blocks sent to the deduplicating storage appliances appear as different, even though they may contain duplicate data. If you want to achieve a higher deduplication ratio, you can disable data encryption. If you still want to use encryption, you can enable the encryption feature on the deduplicating storage appliance itself.

# Data Encryption and Compression

If data compression and data encryption are enabled for a job, Veeam Backup & Replication compresses VM data first and after that encrypts the compressed data blocks. Both operations are performed at the source side.

If the **Decompress backup data blocks before storing** check box is selected in the backup repository settings, Veeam Backup & Replication does not compress VM data before encryption. Therefore, in the job statistics, you may observe a higher amount of transferred data (the **Transferred** counter) as compared to a job for which encryption is disabled. For details on job statistics, see Viewing Real-Time Statistics.

# Job Encryption

Job-level encryption is supported for the following jobs:

- Backup jobs. For more information, see Encrypting Backup Jobs.
- Backup copy jobs. For more information, see Encrypting Backup Copy Jobs.
- File backup jobs. For more information, see the Storage Settings section in the job configuration.
- Object storage backup jobs. For more information, see the Storage Settings section in the job configuration.
- VeeamZIP sessions. For more information, see Encrypting VeeamZIP Backups.
- Configuration backup jobs. For more information, see Creating Encrypted Configuration Backups.

> **NOTE**
>
> If you plan to store backups in Veeam Data Cloud Vault, you must enable job encryption.

# Storage Encryption

Storage-level encryption is supported for the following backups:

- Backup files stored to capacity tier. For more information, see Encryption for Capacity Tier.
- Backup files stored to archive tier. For more information, see Encryption for Archive Tier.
- Backup files created by standalone applications. For more information, see Encrypting Standalone Application Backups in Backup Repositories.
- Backup files stored to external repositories. For more information, see Encryption.
- Backup files stored to tape media pools. For more information, see Tape Encryption.

# How Data Decryption Works

To perform operations with backup files, you need to decrypt them first. Backup files can be decrypted automatically or manually:

- If you encrypt data encryption keys with KMS keys, backup files are decrypted automatically when you import them. In some cases, you need to decrypt backup files manually. For more information, see How KMS Works.
- If you encrypt data encryption keys with secret keys, backup files are decrypted automatically only when the following conditions are met:
  - You encrypt and decrypt the backup file on the same backup server that uses the same configuration database.
  - The backup is not removed from the Veeam Backup & Replication console.

  In other cases, you need to enter a password you specified in the encryption options.

Veeam Backup & Replication transfers encrypted data to the source side and performs data decryption. Decrypted files appear under the **Backups** > **Disk (Imported)** node in the inventory pane.

For more information on how to decrypt backups, see Decrypting Backups. If you want to restore data from encrypted backups, see the following sections:

- Restoring Data from Encrypted Backups
- Restoring Encrypted Data from Tapes in the Tape Device Support Guide

## Decrypting Backups

If you encrypt data encryption keys with KMS keys, for more information about decrypting data, see Using KMS Keys.

If you encrypt data encryption keys with secret keys, do the following:

1. In the inventory pane, select the **Backups** > **Disk (Encrypted)** node.
2. In the working area, select an encrypted backup file and click **Specify Password** on the ribbon. Alternatively, you can right-click the backup and select **Decrypt backup**.
3. In the **Password** field, enter the password for the backup file. The password must be the same as the password that was used to encrypt the backup file.

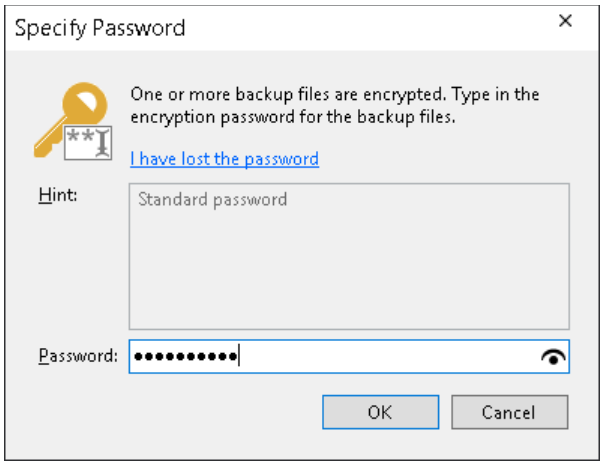

If you enter the correct password, Veeam Backup & Replication will decrypt the backup file. The decrypted backup will be moved under the **Backups** > **Disk (Imported)** node in the inventory pane. You can perform operations with the backup file in a regular manner.

Consider the following:

- If data encryption keys are stored in the configuration database, you do not need to specify the password. In this case, data decryption is performed in the background.
- If the password has changed several times or you have changed the encryption method to KMS keys, you need to specify them in the following manner:
  - If you select a metadata file (VBM) for import, you must specify the latest KMS key or password that was used to encrypt data encryption keys.
  - If you select a full backup file (VBK) for import, you must specify the whole set of KMS keys and passwords that were used to encrypt data encryption keys.

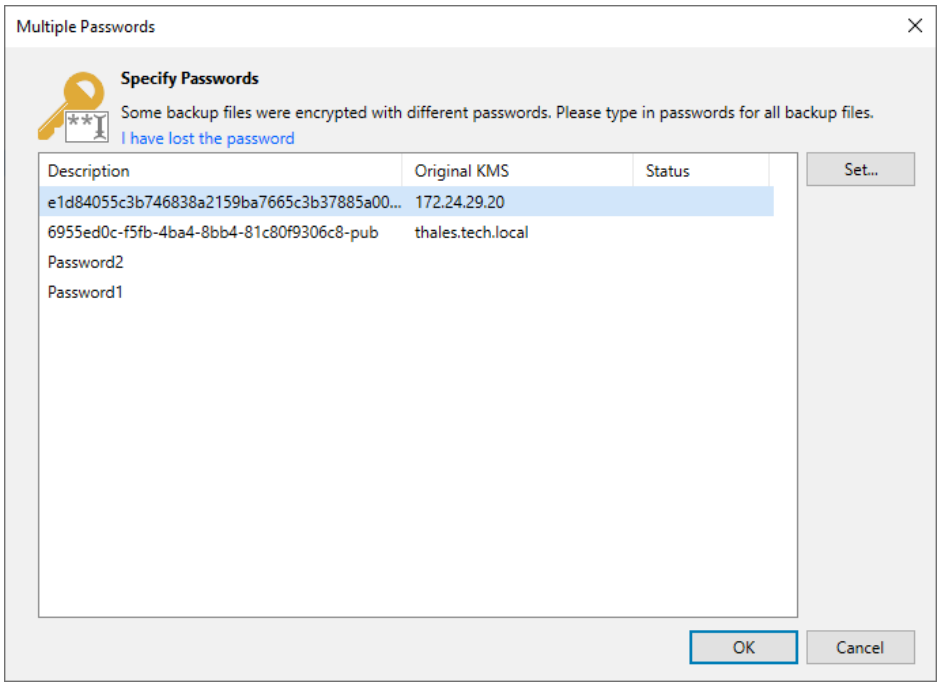


- If you have lost or forgotten a password, you can issue a request to Veeam Backup Enterprise Manager and restore data from an encrypted file using Enterprise Manager keys. This option works only if the password loss protection is enabled. For more information, see Password Loss Protection.

# Password Loss Protection

If you have lost or forgotten a password or you cannot use KMS keys due to a KMS server failure, you can restore encrypted data using password loss protection provided by Veeam Backup Enterprise Manager. This solution allows you to include Enterprise Manager keyset in the encryption process and use it to decrypt backup files. For more information, see Enterprise Manager Keys.

Password loss protection has the following requirements and limitations:

- To use password loss protection, you must have a paid license. For more details about licensing support, see Veeam Data Platform Feature Comparison.
- Backup servers that you use for data decryption must be connected to the same instance of Veeam Backup Enterprise Manager. If you connect the backup server to several instances of Veeam Backup Enterprise Manager, this may cause unexpected behavior, and the decryption process may fail.

- Password loss protection must be enabled on Veeam Backup Enterprise Manager. For more details, see the Managing Encryption Keys section in the Veeam Backup Enterprise Manager Guide.

For more information on how to decrypt backups with Enterprise Manager keys, see Decrypting Backups With Enterprise Manager Keys.

## Enterprise Manager Keys

Enterprise Manager keys make up a pair of matching keys — a keyset. Keysets are created and managed on the Veeam Backup Enterprise Manager server. During installation of Veeam Backup Enterprise Manager, the setup automatically generates a new keyset containing a public Enterprise Manager key and a private Enterprise Manager key. You can use Veeam Backup Enterprise Manager to create new Enterprise Manager keysets, activate them, import and export keysets and specify retention for their lifetime. For more information, see the Managing Encryption Keys section in the Veeam Backup Enterprise Manager User Guide.

If you enable password loss protection, the public Enterprise Manager key is automatically sent to the backup server and stored in the configuration database. When Veeam Backup & Replication encrypts backup files, it uses a secret key or a KMS key and a public Enterprise Manager key simultaneously.

> **NOTE**
>
> Consider the following:
>
> - To generate Enterprise Manager keys, Veeam Backup Enterprise Manager uses the RSA algorithm with a 4096-bit key length. For more information, see RSA Cryptography Specifications.
> - To encrypt data encryption keys with Enterprise Manager keys, Veeam Backup & Replication uses the RSA algorithm with a 4096-bit key length.
> - To protect public Enterprise Manager keys stored in the configuration database, Veeam Backup & Replication uses Data Protection API (DPAPI) mechanisms. For more information, see Microsoft Docs.

The private Enterprise Manager key is used to decrypt data encrypted with the public Enterprise Manager key. Private Enterprise Manager keys are not distributed and stored only on Veeam Backup Enterprise Manager.

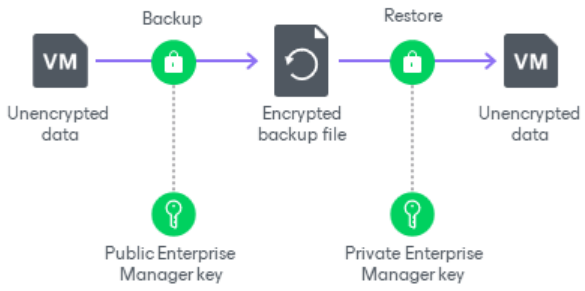


For Enterprise Manager keysets, consider the following recommendations:

- Create and activate new Enterprise Manager keysets regularly.
- Export your Enterprise Manager keysets and keep them in a safe place. If your installation of Veeam Backup Enterprise Manager goes down for some reason, you will lose private Enterprise Manager keys. As a result, you will not be able to use the password loss protection.
- If a matching private Enterprise Manager key is not found in the Veeam Backup Enterprise Manager configuration database, the decryption process will fail. In such situations, you can import the private key manually. For more information, see the Exporting and Importing Enterprise Manager Keyset section in the Veeam Backup Enterprise Manager Guide.

# Decrypting Backups With Enterprise Manager Keys

To decrypt backup files with Enterprise Manager keys, you need to issue a request to Veeam Backup Enterprise Manager using the following wizards:

- The **Encryption Key Restore** wizard on the backup server.
- The **Password Recovery** wizard on the Veeam Backup Enterprise Manager server.

> **NOTE**
>
> To generate a request for data restore from a backup server, Veeam Backup & Replication uses the RSA algorithm with a 2048-bit key length. For more information, see RSA Cryptography Specifications.

To complete the operation, perform the following steps:

1. Create a request for key restore.
2. Process the request in Veeam Backup Enterprise Manager.
3. Complete the key restore process.

## Step 1. Create Request for Key Restore

This procedure is performed by the Veeam Backup Administrator on the backup server.

To create a request for key restore, do the following:

1. Import encrypted backup to the Veeam Backup & Replication console.
2. Select the imported backup under the **Backups** > **Disk (Encrypted)** node and click **Specify Password** on the ribbon or right-click the backup and select **Decrypt backup**.
3. In the **Specify Password** window, click the **I have lost the password** link. Veeam Backup & Replication will launch the **Encryption Key Restore** wizard.

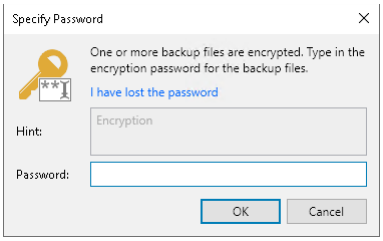


4. At the **Request** step of the **Encryption Key Restore** wizard, review the generated request for data recovery. Copy the request to the clipboard or save it to a text file.
5. Send the copied request by email or pass it in any other way to the Veeam Backup Enterprise Manager Administrator.

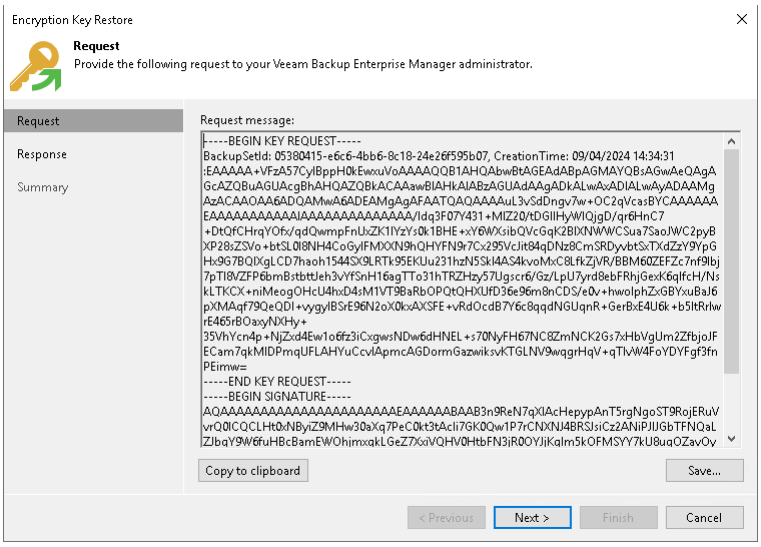


### TIP

You can close the **Encryption Key Restore** wizard on the backup server and start it again when you receive a response from the Veeam Backup Enterprise Manager Administrator.

## Step 2. Process Request in Veeam Backup Enterprise Manager

This procedure is performed by the Veeam Backup Enterprise Manager Administrator on the Veeam Backup Enterprise Manager server.

To process the key restore request, do the following:

1. Copy the obtained request to the clipboard.
2. In the **Configuration** > **Settings** > **Key Management** section, click **Password Recovery** to open the **Password Recovery** wizard.
3. Paste the request that you have received from the Veeam Backup Administrator.

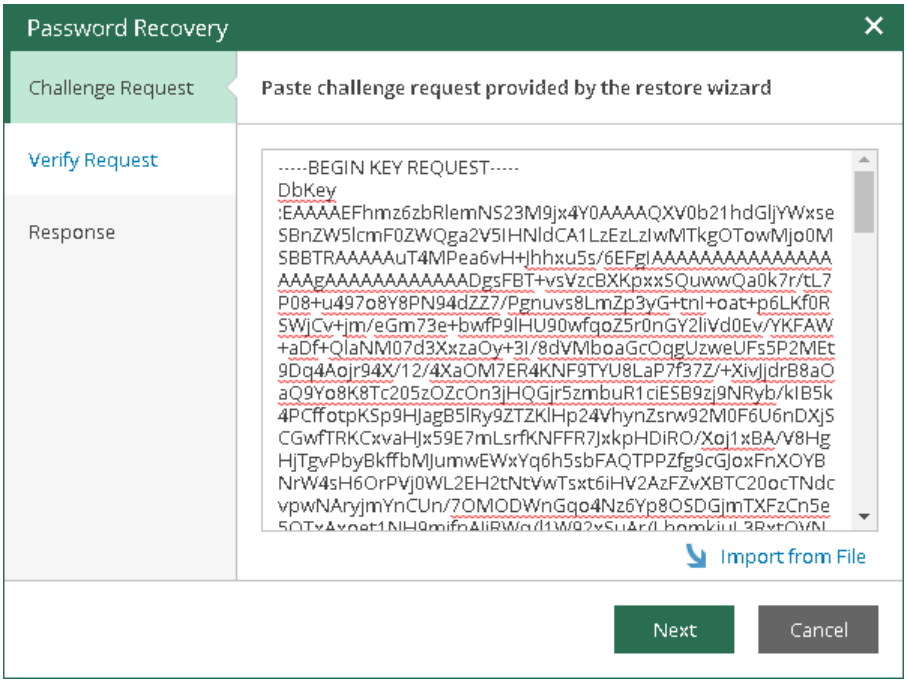


4. Follow the next steps of the wizard. At the **Response** step of the wizard, copy the text displayed in the wizard to the clipboard.

5. Send the copied response by email or pass it in any other way to the Veeam Backup Administrator working on the backup server.

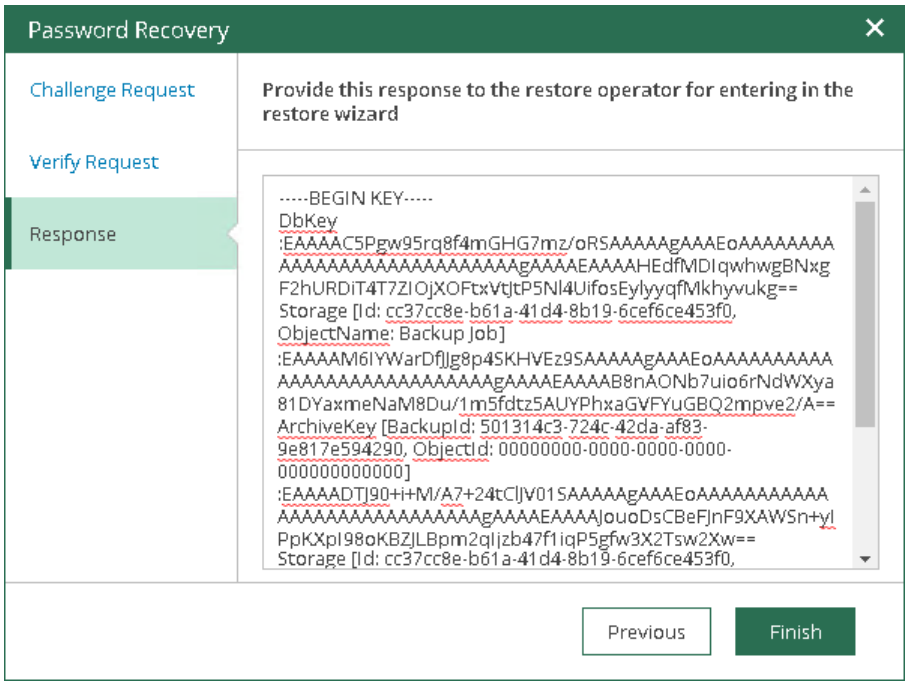

## Step 3. Complete Key Restore Process

This procedure is performed by the Veeam Backup Administrator on the backup server.

To complete the key restore process, do the following:

1. In Veeam Backup & Replication, return to the **Encryption Key Restore** wizard.
2. Enter the copied response to the text window at the **Response** step of the **Encryption Key Restore** wizard.
3. Follow the next steps of the wizard. At the last step, click **Finish**. Veeam Backup & Replication will retrieve the information required for decryption from the response, apply it to the encrypted file and unlock the file content.

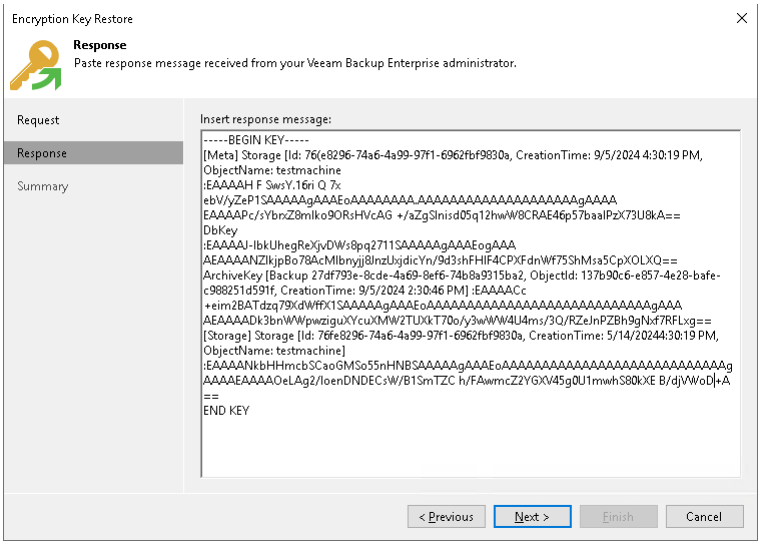

# Multi-Factor Authentication

Veeam Backup & Replication supports multi-factor authentication (MFA) for additional user verification. A one-time password (OTP) generated in a mobile authenticator application is used as a second verification method. Combined with login and password credentials, it creates a more secure environment and protects user accounts from being compromised.

The feature includes:

- Enabling/disabling MFA for all users
- Disabling MFA for service accounts
- Resetting MFA for specific users

> **TIP**
>
> Multi-factor authentication is compatible with mobile authentication applications that support RFC4226 and RFC6238.

## Requirements and Limitations

MFA has the following requirements and limitations:

- Only users with the *Veeam Backup Administrator* role can manage MFA.
- MFA is not supported in the Veeam Backup & Replication Community Edition.
- MFA is not natively supported for Veeam Backup Enterprise Manager. It can be used with a third party identity provider specified in the SAML authentication settings.
- User groups are not supported. You can enable MFA only for user accounts.
- MFA is not supported for non-interactive connections used by the following applications and backup infrastructure components:
  - Veeam Backup & Replication REST API
  - Veeam Backup Enterprise Manager (for communication with the Veeam Backup & Replication server)
  - Veeam ONE agent (for communication with the Veeam Backup & Replication server)
  - Veeam Backup Validator

  To avoid connection issues, you must disable MFA for the accounts used to run these applications and backup infrastructure components. For more information, see Disabling MFA for Service Accounts.
- MFA is not supported for PowerShell (either interactive logon or non-interactive connections). To use PowerShell cmdlets with Veeam Backup PowerShell Module or Microsoft Windows PowerShell, run the Veeam Backup & Replication console or Microsoft Windows PowerShell under the service account with disabled MFA.
- MFA also must be disabled for the account in the following cases:
  - To restore the configuration database properly, run the Veeam Backup & Replication console or Veeam Backup Configuration Restore application under the service account with disabled MFA.
  - To upgrade the remote Veeam Backup & Replication console properly, run it under the service account with disabled MFA.

- If a service provider (SP) uses Veeam Service Provider Console and wants to use multi-factor authentication on the SP backup server, they must set up a service account in Veeam Backup & Replication. For more information, see this Veeam KB article.
- Mobile push notifications are not supported. You can get an OTP code only in the mobile authenticator application.

# How MFA Works

Veeam Backup & Replication supports the following scenario for MFA:

- A user logs in to the Veeam Backup & Replication console.
- Veeam Backup & Replication checks if MFA is enabled and configured for the user:
  a. **MFA is enabled but not configured**. The user gets the instruction how to set up MFA. Veeam Backup & Replication generates a secret key which is used once for the initial setup in the mobile authenticator application. The hash of the secret key is also saved in the configuration database.

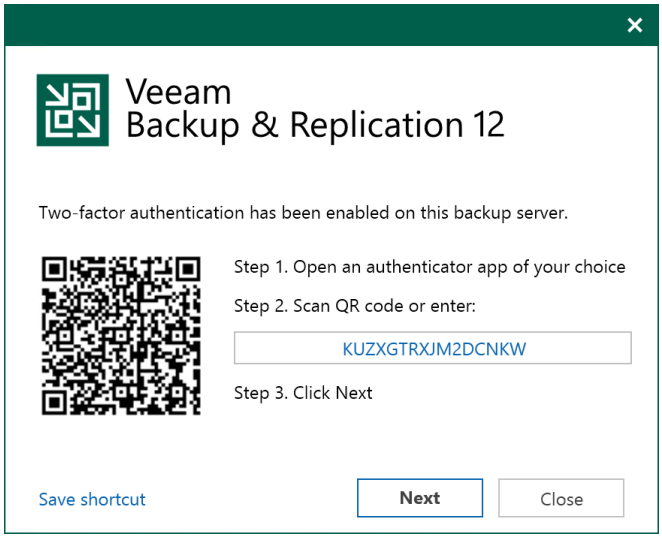


  b. **MFA is enabled and configured**. Each time the user logs in they should enter a 6-digit confirmation code generated in the mobile authenticator application. Veeam Backup & Replication checks if the code is valid and, in case of success, starts a user session.

  If there are more than 5 unsuccessful attempts, the user can reopen the console and try to log in again after waiting for at least one minute. If the problem persists, the backup administrator can reset MFA by request.

> **IMPORTANT**
>
> The code confirmation works when there is no time shifting between the mobile authenticator application and the Veeam Backup & Replication server. Ensure that they are synchronized with the UTC time. Otherwise, the authentication will fail.

If Veeam Explorers and other applications (except for Veeam Backup PowerShell Module) are started from the console, they do not require additional authentication.

# Enabling MFA

To enable the feature for all users:

1. Log in to the Veeam Backup & Replication console as an administrator.
2. Go to **Users and Roles** > **Security**.
3. Remove user groups from the list if there are any. Leave only specific users.
4. Select the **Enable multi-factor authentication (MFA)** check box.
5. Click **OK**.

> **IMPORTANT**
>
> If you generate a new backup server certificate while multi-factor authentication is enabled, any Veeam Backup & Replication consoles connected to the backup server must be restarted to avoid connection issues.

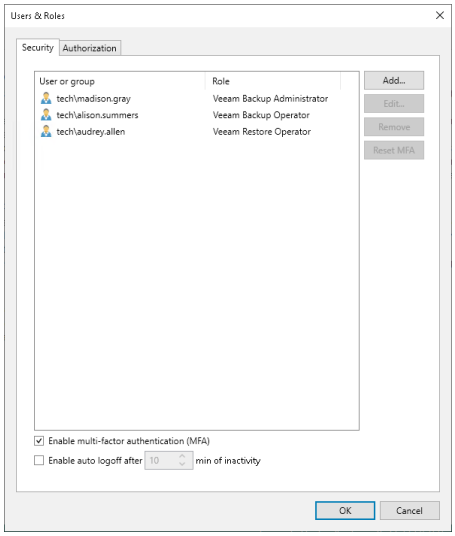


# Resetting MFA for Specific User

The backup administrator can reset MFA by user request if they have authentication issues, lose or change a mobile device with the mobile authentication application, and so on.

To reset MFA for a specific user:

1. Log in to the Veeam Backup & Replication console as an administrator.
2. Go to **Users and Roles** > **Security**.

3. Select the user and click **Reset MFA**. The next time the user logs in they will get the instruction how to set up MFA.

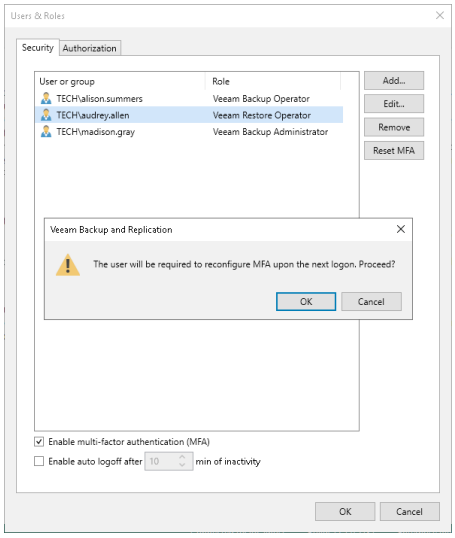


# Disabling MFA

To disable the feature for all users:

1. Log in to the Veeam Backup & Replication console as an administrator.
2. Go to **Users and Roles** > **Security**.
3. Clear the **Enable multi-factor authentication (MFA)** check box.
4. Click **OK**.

# Disabling MFA for Service Accounts

If you cannot use MFA due to limitations described in section Requirements and Limitations, you can disable this feature for specific service accounts used to run applications and backup infrastructure components.

To disable the feature for service accounts:

1. Log in to the Veeam Backup & Replication console as an administrator.
2. Go to **Users and Roles** > **Security**.
3. Select the service account and click **Edit**.
4. Select the **This is a service account (disables two-factor authentication)** check box.

5. Click **OK**.

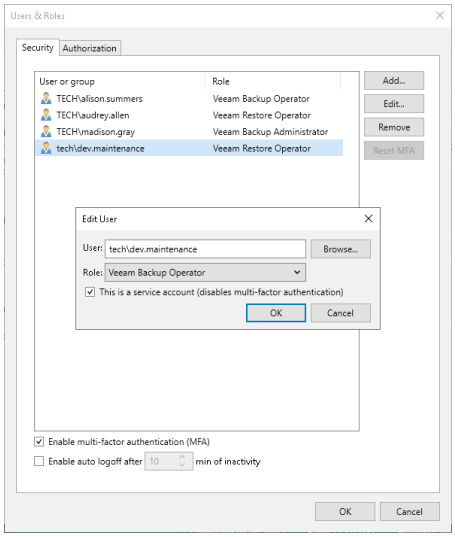
Users & Roles
Security
Authorization
User or group
Role
TECH\alison.summers
Veeam Backup Operator
TECH\audrey.allen
Veeam Restore Operator
TECH\madison.gray
Veeam Backup Administrator
tech\dev.maintenance
Veeam Restore Operator
Add...
Edit...
Remove
Reset MFA
Edit User
User:
tech\dev.maintenance
Browse...
Role:
Veeam Backup Operator
This is a service account (disables multi-factor authentication)
OK
Cancel
Enable multi-factor authentication (MFA)
Enable auto logoff after
10
min of inactivity
OK
Cancel

# Kerberos Authentication

Veeam Backup & Replication supports Kerberos authentication for all components of the backup infrastructure including Veeam Explorers, Veeam Agents, Veeam Plug-ins, and Veeam Backup Enterprise Manager. You can build the backup infrastructure in the following environments:

- Kerberos authentication is the primary domain authentication protocol, NTLM is supported for compatibility. This configuration is used by default starting from Microsoft Windows 2000 Server.
- Kerberos is the only domain authentication protocol, NTLM is disabled (more secure).

> **NOTE**
>
> Kerberos authentication is supported only in Microsoft Active Directory-based environments.

## How Kerberos Works

Unlike NTLM, Kerberos uses Ticket Granting Tickets (TGT) issued by the Key Distribution Center (KDC). TGT files contain a session key, the key's expiration date, and a user's IP address. They are encrypted and have limited validity period (10 hours by default). This authentication mechanism protect users from man-in-the-middle (MITM) attacks.

Veeam Backup & Replication supports the standard Kerberos authentication scenario:

1. A client sends a request to the KDC, which is located on a domain controller and uses Microsoft Active Directory as the account database. The request contains user details and information about the Veeam service the client wants to access.
2. The KDC verifies the client's request and issues a TGT.
3. The client uses the TGT to send a request to the KDC Ticket Granting Service (TGS) and get a TGS ticket. The request also contains the Service Principal Name (SPN) of the service.
4. The client uses the TGS ticket to send a request to the server.
5. The server verifies the client's request and provides access to the service for a limited period specified in the Kerberos configuration.

For more information about Kerberos authentication, see this Microsoft article.

## Requirements and Limitations

Kerberos authentication has the following requirements and limitations:

- The client and the server must belong to the same domain, or a trust relationship must exist between domains.
- All backup infrastructure components must be added to the Veeam Backup & Replication console using FQDN.
- Veeam backup infrastructure servers must resolve FQDNs.
- FQDN must be used when connecting to the remote Veeam Backup & Replication console.
- The hostname length for all Windows-based backup infrastructure components and VM guest OSes must not exceed 15 characters.

- The maximum time difference between the client and the domain controller must be 5 minutes to protect the backup infrastructure from relay attacks.
- NFS is supported by Kerberos starting from version 4.1.
- For guest OS processing, consider the following:
    - Local accounts do not support Kerberos authentication. To authenticate with Microsoft Windows guest OS using Kerberos, specify an Active Directory account.
    - If you use networkless application-aware guest processing through VIX API/vSphere Web Services, the guest OS must still have access to the domain controller. Otherwise, Kerberos authentication will not work.

For SPNs, consider the following aspects:

- Each Veeam service must have two SPNs registered with the Active Directory in the following formats:
    - `{ServiceName}/{FQDN}`, for example, `VeeamBackupSvc/vbrserver01.tech.local`
    - `{ServiceName}/{NetBIOSName}`, for example, `VeeamBackupSvc/VBRSERVER01`
- For the following services, SPNs are registered automatically each time they start:
    - Veeam Backup Service (VeeamBackupSvc)
    - Veeam Backup Enterprise Manager Service (VeeamEnterpriseManagerSvc)
    - Veeam Cloud Connect Service (VeeamCloudConnectSvc)
    - Veeam Cloud Gateway Service (VeeamGateSvc)
    - Veeam CDP Coordinator Service (VeeamCdpSvc)
    - Veeam CDP Proxy Service (VeeamCdpProxySvc)
    - Veeam Guest Catalog Service (VeeamCatalogSvc)
    - Veeam Distribution Service (VeeamDistributionSvc)
    - Veeam Mount Service (VeeamMountSvc)
    - Veeam Broker Service (VeeamBrokerSvc)
    - Veeam Hyper-V Integration Service (VeeamHvIntegrationSvc)
    - Veeam Data Mover Service (VeeamTransportSvc)
    - Veeam WAN Accelerator Service (VeeamWANSvc)
    - Veeam vPower NFS Service (VeeamNFSSvc)
    - Veeam Backup VSS Integration Service (VeeamFilesysVssSvc)
    - Veeam Installer Service (VeeamDeploySvc)
    - Veeam Tape Service (VeeamTapeSvc)
    - Veeam Log Shipping Service (VeeamLogShipperSvc)
    - Veeam Agent for Microsoft Windows Service (VeeamAgentWindows)

- Veeam Guest Helper Service (VeeamGuestHelperSvc)

  For services running under the LocalSystem account, SPNs are mapped to the Active Directory computer objects. For services, running under a dedicated Active Directory service account, SPNs are mapped to the Active Directory user objects.

  Note that If for any reason the SPN registration fails, the service will continue working, but there may be authentication issues in Kerberos-only environments.

- If you want to register SPN manually, you can use the `setspn` tool. For more information about manual SPN registration, see this Microsoft article. To disable automatic SPN registration, contact Veeam Customer Support.

  Note that Veeam Platform Services (AWS, Azure, Google Cloud) and Veeam Explorer Recovery Service do not need to be registered as they do not have SPNs.

- If you want to use another account for running a Veeam service, you must manually remove the existing SPN from the Active Directory beforehand. Otherwise, the SPN registration with the new account will fail.

> **IMPORTANT**
>
> If you use persistent agents for guest OS processing and want to upgrade Veeam Backup & Replication to version 12, there may be issues with application-aware processing in a Kerberos-only environment. To mitigate risks, register SPNs for Guest Helper Service using the following format: `{VeeamGuestHelperSvc}/{FQDN}` and `{VeeamGuestHelperSvc}/{NetBIOSName}`. For more information, see this Veeam KB article.

# Configuring Kerberos Environments

If you use default Microsoft Windows configuration for Kerberos authentication with NTLM fallback, configure the **Network security: LAN Manager Authentication Level** security policy setting to send NTLMv2 responses only. For more information, see this Microsoft article.

> **NOTE**
>
> Using NTLM increases the attack surface of your backup infrastructure. To build a more secure environment, disable NTLM and leave Kerberos the only domain authentication protocol.

To configure a Kerberos-only environment, perform the following steps:

1. **Enable an NTLM audit**. Start to collect and analyze NTLM audit events to determine which applications and services send NTLM requests. Reconfigure or update them to use Kerberos. If an application or service does not support Kerberos, replace it with an alternative one. For more information about auditing NTLM usage, see this Microsoft article.
2. **Enable Kerberos extended logging mode**. By default, Veeam Backup & Replication saves logs only for failed SPN registrations. It is recommended to collect information about all SPN registrations during initial Veeam Backup & Replication deployment and NTLM audit troubleshooting. To enable extended logging for SPNs used for Kerberos authentication between backup infrastructure components, see this Veeam KB article.
3. **Restrict NTLM traffic**. When all applications and services are configured to use Kerberos and there are no NTLM audit events, apply **Network Security: Restrict NTLM** Active Directory group policies to restrict NTLM traffic and authentication. You can also configure exclusions for specific hosts if they still require NTLM authentication to work properly. For more information about restricting NTLM usage, see this Microsoft article.

NOTE

To prevent Kerberos environments from Kerberoasting attacks, do the following:

- Make sure that you use strong encryption algorithms allowed for Kerberos. For more information, see this Microsoft article.
- Prevent Kerberos change password that uses RC4 secret keys. For more information, see this Microsoft article.

# Using Group Managed Service Accounts

A Group Managed Service Account (gMSA) is the type of domain account configured on the server. It does not need the administrator to manage the password as this role is performed by the Microsoft Windows operating system. Randomly generated complex passwords are automatically changed every 30 days which reduce the risk of brute force and dictionary attacks. For more information about gMSAs, see this Microsoft article.

You can use gMSAs to only run guest processing tasks.

> **NOTE**
>
> For application-aware processing, using a gMSA is supported for backups or replicas of VMs that run Microsoft Active Directory (domain controllers), Microsoft Exchange, Microsoft SQL Server, and Oracle 12c Release 2 and later. You cannot back up or replicate VMs that run Microsoft SharePoint with the gMSA.

## Requirements and Limitations

gMSAs have the following requirements and limitations:

- gMSAs are applicable to Microsoft Windows Server 2012 and later. For more information about operating system and Microsoft Active Directory requirements, see this Microsoft article.

  > **NOTE**
  >
  > Consider that this section describes the gMSA configuration compatible with most environments. To reduce the number of possible issues, it is recommended to perform all required steps to configure the domain controller and implement the gMSA.

- 
- gMSAs are not supported for backups of the Linux target machines joined to the Active Directory domain.
- When using gMSAs for guest processing, consider the following:
  - For backup and replication jobs:
    - A guest interaction proxy must be joined to the Active Directory domain where the gMSA was created.
    - The gMSA must also have *Logon as a service* permission granted.
  - For CDP policies, the backup server must be joined to the Active Directory domain where the gMSA was created.
- Veeam Explorers do not support data recovery using gMSAs.
- Since gMSAs require a connection to the domain controller, these accounts work only over network.
- If you use gMSAs to manage the restore process of VM guest OS files, consider Requirements and Limitations.
- If you back up a machine using a gMSA, both the guest interaction proxy and the target machine must have network access to the domain controllers and be in the same domain to obtain the gMSA password. On the target machine the gMSA must be added to the Administrators group (local or domain). Domain Administrator permissions are only required for Microsoft Active Directory backups, for other supported applications local Administrator permissions are sufficient.

> **IMPORTANT**
>
> Consider that granting Domain Administrator permissions to gMSAs makes them a potential source of vulnerability.

## Before You Begin

Before you start using gMSAs, configure the domain controller:

- The domain controller requires a root key to generate gMSA passwords. Ensure that the Key Distribution Service is enabled on the domain controller and use the following command in Microsoft Windows PowerShell to generate a root key:

  ```
  Add-KdsRootKey -EffectiveImmediately
  ```

  Wait until the Active Directory replication is finished or force it manually. For more information about creating KDS root keys, see this Microsoft article.

- To enable gMSA support in Microsoft Windows PowerShell, use the following commands:

  ```
  Install-WindowsFeature RSAT-AD-PowerShell,NET-Framework-Features | Out-Nul
  l;
  Import-Module ServerManager;
  Import-Module ActiveDirectory;
  ```

## Creating gMSA

To create a gMSA, use the `New-ADServiceAccount` cmdlet on the domain controller. For example:

```
$gMSAName = 'DOMAIN\gmsa01'
$gMSAGroupName = 'gMSAComputerAccountsGroup'
$gMSADNSHostName = 'gmsa01.srv.local'
New-ADServiceAccount -Name $gMSAName -DNSHostName $gMSADNSHostName -PrincipalsA
llowedToRetrieveManagedPassword $gMSAGroupName -Enabled $True
```

Consider the following:

- In the `DNSHostName` parameter, specify the FQDN of the gMSA.
- For the `PrincipalsAllowedToRetrieveManagedPassword` parameter specify the AD group containing computer accounts which will use the gMSA. As an alternative, specify computer accounts separated by comma.
- By default, created gMSAs are shown in the **Managed Service Accounts** container.

For more information about all cmdlet parameters, see this Microsoft article.

> **NOTE**
>
> To provide a more secure environment, use separate gMSAs for critical backup infrastructure components.

To refresh AD group membership after you create a gMSA, run the following command on the machines on which you plan to install the gMSA:

```
C:\WINDOWS\system32\klist.exe -lh 0 -li 0x3e7 purge
```

Alternatively, you can reboot these machines.

## Installing gMSA

To install a gMSA on the server or the target machine, perform the following steps:

- Run the `Install-ADServiceAccount` cmdlet.

  ```
  Install-ADServiceAccount "DOMAIN\gmsa01$"
  ```

- Ensure that the gMSA was successfully installed.

  ```
  Test-ADServiceAccount "DOMAIN\gmsa01$"
  ```

3. Add the gMSA to the local *Administrators* group.

   ```
   Add-LocalGroupMember -Group "Administrators" -Member "DOMAIN\gmsa01$"
   ```

   **NOTE**

   For domain controller VMs, add the gMSA to the domain *Administrators* group.

For guest processing, install the gMSA on the guest OS machine and the guest interaction proxy.

## Using gMSA

To run guest processing tasks with the gMSA, do the following:

1. Make sure that the following services run under the LocalSystem account:
   - The Veeam Backup Service on the backup server.
   - The Veeam Data Mover Service on the guest interaction proxy.
2. Add the gMSA account type to the Credentials Manager. For more information, see Group Managed Service Accounts.
3. Select the gMSA when specifying guest OS credentials for jobs or policies.

# Licensing

To work with Veeam Backup & Replication, you must obtain a license key and install it on the backup server. If you do not install the license key, the product will operate in the Veeam Backup & Replication Community (free) Edition. For more information, see Veeam Backup & Replication Community Edition.

Veeam licenses Veeam Backup & Replication in 3 ways: per socket, per instance, per capacity.

You can use instance and socket licenses together. For more information, see Merging Licenses.

> **NOTE**
>
> Veeam Backup & Replication consumes licenses only to back up data, restore does not require licenses. You can restore VMs and data no matter how many available licenses you have or how many VMs you are restoring.

For specific details on Veeam Agents licensing, see the following user guides:

- If you work with Veeam Agents operating in the managed mode, see the Licensing Requirements section in the Veeam Agent Management Guide.
- If you work with Veeam Agents operating in the standalone mode, see the user guide for the Veeam Agent depending on the operating system of the protected computer. For example, if you work with Veeam Agent for Microsoft Windows, see the Managing License section in the Veeam Agent for Microsoft Windows User Guide.

## Socket Licensing

With the socket licensing model, Veeam Backup & Replication is licensed by the number of CPU sockets on protected hosts. For more information, see Veeam Licensing Policy.

A license is required for every occupied motherboard socket as reported by the hypervisor API.

License is required only for source hosts — hosts on which VMs that you back up or replicate reside. Target hosts (for replication and migration jobs) do not need to be licensed.

The socket license assignment happens automatically as soon as you start a backup or replication job for VMs on a specific source host. You can revoke licenses from licensed hosts and re-apply them to other objects if needed. For more information, see Revoking License.

> **NOTE**
>
> If you use a socket license that was obtained for an earlier version of Veeam Backup & Replication, Veeam Software adds up to 6 gift (built-in) instances free of charge to your license scope. You can use these instances to protect any type of supported workloads except VMware and Hyper-V VMs — they are covered by the licensed CPU sockets on virtualization hosts.
>
> If the number of licensed sockets is less than 6, you can use the number of instances that equals the number of licensed sockets. For example, if the number of licensed sockets is 5, you can use 5 instances. If the number of licensed sockets is 100, you can use 6 instances.
>
> Note that starting with Veeam Backup & Replication 12, the gift (built-in) instances are disabled if the perpetual license support expiration date is reached.

# Instance Licensing

Veeam Backup & Replication can be licensed by the number of instances. Instances are units (or tokens) that you can use to protect your virtual, physical or cloud-based workloads. For more information, see Veeam Licensing Policy.

You must obtain a license with the total number of instances for workloads that you plan to protect in Veeam Backup & Replication.

Workloads that have been processed in the past 31 days are considered protected. Every protected workload consumes instances from the license scope. The number of instances that a workload requires depends on the workload type and product edition. If there were no backups created for a workload (for example, for a VM) during the last 31 days, the instances previously consumed to protect it are automatically revoked into the pool of available license instances.

This licensing model allows you to obtain a license with a certain number of instances without knowing in advance what types of workloads you plan to protect. When a need arises, you can revoke instances from a protected workload, and reuse them to protect other workloads regardless of the workload type.

Veeam Backup & Replication keeps track of instances consumed by protected workloads. If the number of consumed instances exceeds the license limit, Veeam Backup & Replication displays a warning when you open the Veeam Backup & Replication console. For more information, see Exceeding License Limit.

Consider the following:

- VM templates are regarded as protected VMs and consume license instances.
- VMs and unstructured data sources processed with backup copy and tape jobs are not regarded as protected VMs and data sources and do not consume license instances. These types of jobs provide an additional protection level for VMs and unstructured data sources that are already protected with backup jobs.
- VMs processed by snapshot-only jobs are regarded as protected VMs and consume license instances. Veeam Backup & Replication will revoke instances from these VMs if you re-add a storage array to the backup infrastructure.
- For more information on how Veeam Backup & Replication calculates license instances to consume when protecting unstructured data sources, see Instance Consumption for Object Storage Backup, File Backup and File to Tape Jobs.

# Capacity Licensing

Veeam Backup & Replication can be licensed by the capacity of protected data. TBs of front-end storage capacity pack are units that you can use to protect your non-deduplicated and uncompressed front-end source data in unstructured data sources (file servers, file shares, NAS filers, object storage repositories). For more information, see Veeam Licensing Policy.

You must obtain a license with the total capacity for source data that you plan to protect in Veeam Backup & Replication.

Data sources that have been processed in the past 31 days are considered protected. Every protected data source consumes capacity from the license scope.

This licensing model allows you to obtain a license with a certain license capacity without knowing in advance what types of unstructured data sources you plan to protect. When a need arises, you can revoke license capacity from a protected data source, and reuse it to protect other data sources regardless of the workload type.

Veeam Backup & Replication keeps track of capacity consumed by protected data sources. If the number of consumed capacity licenses exceeds the license limit, Veeam Backup & Replication displays a warning when you open the Veeam Backup & Replication console. For more information, see Exceeding License Limit.

Consider the following:

- Veeam Backup & Replication rounds the protected amount of data for each unstructured data source down to 1 TB.
- Capacity license is consumed in 1 TB chunks.
- When your unstructured data sources are protected with file backup to tape or object storage backup to tape jobs, Veeam backup files are excluded from the capacity consumption calculation. These files have the following extensions: VAB, VBM, VBK, VIB, VRB, VSB, VLB, VSM, VLM, VOM, VACM, VASM, VSOURCE, VSOURCETEMP, VSTORE, VSTORETEMP, VSLICE, VBASKET, VLIST, VCACHE, VBLOB, BCO, ADB.
- If different data sources are protected with different file backup or object storage backup jobs or with different file backup to tape or object backup to tape jobs, Veeam Backup & Replication rounds the protected amount to 500 GB separately for each data source and calculates the number of instances required to protect each of them. After that, it sums up the total number of license instances to consume for unstructured data backup or for unstructured data backup to tape.
- If the same data source is protected with more than one unstructured data backup job, to calculate the size of the protected amount of data Veeam Backup & Replication first sums the size of the source data protected with all the jobs. After that, it rounds the overall protected amount of data down to 500 GB and calculates the license capacity to consume for the unstructured data backup.
- Unstructured data sources processed with backup copy and backup to tape jobs are not regarded as protected data sources and do not consume the license capacity. These types of jobs provide an additional protection level for unstructured data sources that are already protected with backup jobs.
- Veeam Backup & Replication calculates the protected amount of data for each data source during every run of the unstructured data backup job or unstructured data backup to tape job that protects data on this data source and keeps the result for 30 days. To calculate the license capacity to consume for the data source protection, Veeam Backup & Replication takes the largest protected amount of data on the data source within the last 30 days.

  If the size of the protected data reduces and does not increase or the data source is removed from the unstructured data backup job or unstructured data backup to tape job, after 30 days Veeam Backup & Replication recalculates the protected amount of data and automatically revokes the excessively consumed license capacity. You can manually revoke the licenses without waiting for 30 days, as described in the Revoking License section. During the next unstructured data backup job or unstructured data backup to tape job run, Veeam Backup & Replication will recalculate the license capacity consumption as of the current date.
- If an unstructured data backup job protects several file shares residing on the same NAS device (the same share root or NAS filer), the approach to calculating license consumption depends on how the file shares are added to the infrastructure:
  - If you have added the whole share root (`\\root\`) to the infrastructure and the file backup job protects shares `\\root\share1` and `\\root\share2`, Veeam Backup & Replication first sums the protected amount of both file shares, rounds it down to 500 GB, and then calculates the license consumption to consume for file backup support.
  - If you have separately added shares `\\root\share1` and `\\root\share2` to the infrastructure and the file backup job protects both of them, Veeam Backup & Replication first rounds the protected amount for each file share down to 500 GB, separately calculates the license capacity to consume for each file share, and then sums the license capacity to calculate the total license capacity to consume for file backup support.

- Unstructured data backup to an object storage requires a license. Thus, this feature is not supported in the Veeam Backup & Replication Community (free) Edition. For details, see Veeam Editions Comparison.
- Unstructured data restore from an object storage does not require a license. Thus, this feature is supported in the Veeam Backup & Replication Community (free) Edition. For details, see Veeam Editions Comparison.

# Types of Licenses

Veeam Software offers the following types of licenses for Veeam Backup & Replication:

## Paid Licenses

- **Subscription license** — license that expires at the end of the subscription term. The Subscription license term is normally 1-3 years from the date of license issue.
- **Perpetual license** — permanent license. The support and maintenance period included with the license is specified in months or years. Typically, one year of basic support and maintenance is included with the Perpetual license.
- **Rental license** — license with the license expiration date set according to the chosen rental program (normally 1-12 months from the date of license issue). The Rental license can be automatically updated upon expiration.

  Rental licenses are provided to Veeam Cloud & Service Providers (VCSPs) only. For more information, see the Rental License section in the Veeam Cloud Connect Guide.

## Free Licenses

- **Evaluation license** — license used for product evaluation. The Evaluation license is valid for 30 days from the moment of product download.
- **NFR license** — license used for product demonstration, training and education. The person to whom the license is provided agrees that the license is not for resell or commercial use.
- **Promo license** — license that grants additional instances. You can install it only on top of an existing Perpetual or Subscription license (primary license). The primary license can have any units (only sockets, only instances, or both). Number of additional instances and duration of promo period are decided by a sales representative.

# Instance Consumption for Object Storage Backup, File Backup and File to Tape Jobs

For unstructured data backup and unstructured data backup to tape, there are the following peculiarities in calculating the number of license instances to consume:

- Veeam Backup & Replication rounds the protected amount of data for each file share or object storage down to the nearest 500 GB.
- One license instance covers 500 GB of the protected amount of data.
- When your unstructured data sources are protected with file backup to tape or object storage backup to tape jobs, Veeam backup files are excluded from the instance consumption calculation. These files have the following extensions: VAB, VBM, VBK, VIB, VRB, VSB, VLB, VSM, VLM, VOM, VACM, VASM, VSOURCE, VSOURCETEMP, VSTORE, VSTORETEMP, VSLICE, VBASKET, VLIST, VCACHE, VBLOB, BCO, ADB.
- If different data sources are protected with different file backup or object storage backup jobs or with different file backup to tape or object backup to tape jobs, Veeam Backup & Replication rounds the protected amount to 500 GB separately for each data source and calculates the number of instances required to protect each of them. After that, it sums up the total number of license instances to consume for unstructured data backup or for unstructured data backup to tape.

  For example, if the amount of data is 1100 GB and it is protected by 2 backup jobs simultaneously, Veeam Backup & Replication will round it down to 1000 GB and multiply the amount of data by 2. As a result, a total of 4 license instances will be consumed.
- If the same data source is protected with more than one unstructured data backup job, to calculate the size of the protected amount of data Veeam Backup & Replication first sums the size of the source data protected with all the jobs. After that, it rounds the overall protected amount of data down to 500 GB and calculates the number of license instances to consume for the unstructured data backup support.
- Unstructured data sources processed with backup copy and backup to tape jobs are not regarded as protected data sources and do not consume the license capacity. These types of jobs provide an additional protection level for unstructured data sources that are already protected with backup jobs.
- Veeam Backup & Replication calculates the protected amount of data for each data source during every run of the unstructured data backup job or unstructured data backup to tape job.

  If the size of the protected data reduces and does not increase or the data source is removed from the unstructured data backup job or unstructured data backup to tape job, after 30 days Veeam Backup & Replication recalculates the protected amount of data and automatically revokes the excessively consumed license instances. You can manually revoke the licenses without waiting for 30 days, as described in the Revoking License section. During the next unstructured data backup job or unstructured data backup to tape job run, Veeam Backup & Replication will recalculate the license instance consumption as of the current date.

  If backup job contains two data sources and you remove one of them, Veeam Backup & Replication will recalculate the protected amount of data and revoke the excessively consumed license instances after 30 days. But if you add the previously removed data source to another backup job, Veeam Backup & Replication will consume twice as many license instances because the data source still exists in the backup of the previous backup job.

- If an unstructured data backup job protects several file shares residing on the same NAS device (the same share root or NAS filer), the approach to calculating license consumption depends on how the file shares are added to the infrastructure:
  - If you have added the whole share root (`\\root\`) to the infrastructure and the file backup job protects shares `\\root\share1` and `\\root\share2`, Veeam Backup & Replication first sums the protected amount of both file shares, rounds it down to 500 GB, and then calculates the number of license instances to consume for file backup support.

    **NOTE**

    If a file share is accessible through two different protocols (for example, SMB and NFS), Veeam Backup & Replication recognizes it as 2 individual shares. You can exclude one of them from processing in the job's settings. For more information, see Step 3. Select Files and Folders to Back Up.
  - If you have separately added shares `\\root\share1` and `\\root\share2` to the infrastructure and the file backup job protects both of them, Veeam Backup & Replication first rounds the protected amount for each file share down to 500 GB, separately calculates the number of license instances to consume for each file share, and then sums the license instances to calculate the total number of license instances to consume for file backup support.
- Unstructured data backup to an object storage requires a license. Thus, this feature is not supported in the Veeam Backup & Replication Community (free) Edition. For more information, see Veeam Editions Comparison.
- Unstructured data restore from an object storage does not require a license. Thus, this feature is supported in the Veeam Backup & Replication Community (free) Edition. For more information, see Veeam Editions Comparison.

# Examples

**Case 1**

If the protected amount of data for the file share is 499 GB or less, Veeam Backup & Replication rounds it down to 0 GB. In this case, protection of this file share will not consume license instances.

If the protected amount of data for the file share is 590 GB or 990 GB, Veeam Backup & Replication rounds it down to 500 GB. In this case, protection of this file share will consume 1 license instance.

**Case 2**

If the protected amount of data for the file share is 1000 GB after rounding, Veeam Backup & Replication divides this amount by 500 GB to calculate the number of instances to consume:

1000 GB / 500 GB = 2 license instances

**Case 3**

You have 2 file shares File Share 1 (990 GB) and File Share 2 (890 GB) each protected with a separate file backup job. In this case, Veeam Backup & Replication rounds the protected amount of each file share down to 500 GB, calculates the number of license instances required to protect each of the file shares. After that, it sums the calculated number of license instances required to protect the shares:

990 GB ~ 500 GB = 1 license instance

890 GB ~ 500 GB = 1 license instance

1 + 1 = 2 — protection of 2 file shares with separate file backup jobs consumes 2 license instances.

### Case 4

You have a file share with the protected amount of data 1490 GB. Veeam Backup & Replication runs the file backup job and after rounding down the amount, it calculates that it will consume 2 license instances for protecting this file share:

1490 GB ~ 1000 GB = 2 license instances

Two days later, the size of the file share increases to 1510 GB. Veeam Backup & Replication runs the file backup job and recalculates the number of license instances to consume based on the increased size of the NAS share:

1510 GB ~ 1500 GB = 3 license instances

Two days later, the size of the file share decreases back to 1490 GB and does not increase any more. Although the protected amount of data decreases, for the next 30 days Veeam Backup & Replication uses value 1510 GB as a basis to calculate the consumption of license instances.

30 days later, Veeam Backup & Replication runs the file backup job and recalculates the number of instances to consume taking into account that the largest protected amount of data within the last 30 days is 1490 GB. After that, protection of the file share starts consuming 2 instances again.

### Case 5

You have a file share (100 GB) protected with 2 file backup jobs (1 and 2). During the 1st run of job 1, Veeam Backup & Replication rounds the protected amount down to 0 GB and calculates that the protection of this file share with job 1 does not consume license instances. After that the file share size increases to 270 GB. During the 1st run of job 2 which is scheduled after the increase of the file share, Veeam Backup & Replication calculates amount of data protected with all file backup jobs protecting this file share:

100 GB (run 1 of job 1) + 270 GB (run 1 of job 2) = 370 GB

Veeam Backup & Replication rounds the protected amount down to 0 GB and calculates that the protection of this file share with backup job 1 and 2 does not consume license instances. After that the file share size remains 270 GB. During the 2nd run of job 1, Veeam Backup & Replication calculates amount of data protected with all file backup jobs protecting this file share:

270 GB (run 2 of job 1) + 270 GB (run 1 of job 2) = 540 GB

Veeam Backup & Replication rounds the protected amount down to 500 GB and calculates that the protection of this file share with backup job 1 and 2 now consumes 1 license instance.

### Case 6

You have a single NAS device with 2 NFS file shares residing on it: `\\root\share1` (490 GB) and `\\root\share2` (600 GB). You have added the root server folder (`\\root\`) of this NAS device as an NFS file share to the infrastructure. The file shares `\\root\share1` and `\\root\share2` are added to one file backup job. In this case, Veeam Backup & Replication sums the protected amount of both file shares, rounds it down to 500 GB, and then calculates the number of license instances:

490 GB + 600 GB = 1090 GB ~ 1000 GB = 2 — protection of 2 file shares in this case consumes 2 license instances.

### Case 7

You have a single NAS device, but you have added 2 of its shared folders `\\root\share1` (490 GB) and `\\root\share2` (600 GB) as separate file shares to the inventory. The file shares `\\root\share1` and `\\root\share2` are added to one file backup job. In this case, Veeam Backup & Replication first rounds the protected amount for each file share down to 500 GB, separately calculates the number of license instances to consume for each file share, and then sums the license instances:

490 GB ~ 0 GB = 0 license instances

600 GB ~ 500 GB = 1 license instance

0 + 1 = 1 — protection of 2 file shares in this case consumes 1 license instance.

# Obtaining and Renewing License

You can obtain an Evaluation or paid license for the product when you download the product from the Veeam website.

## Obtaining Paid License

To obtain a paid license, refer to the Veeam Backup & Replication Pricing page.

To obtain a Perpetual Instance license, find a reseller.

## Obtaining Evaluation License

To obtain an Evaluation license:

1. Sign in to veeam.com.
2. On the Download Veeam products page, click the product link.
3. In the **Get trial key** section, click the **Request Trial Key** link to download the Evaluation license.

## Renewing License

To renew your maintenance plan, contact Veeam Renewals Team at renewals@veeam.com.

If you have a Perpetual or Subscription license, you can also renew your license contract online.

To renew the license:

1. From the main menu, select **License**.
2. In the **License Information** window, click **Renew**.

Veeam Backup & Replication will forward you to the Renewals page of Veeam website, where you can select your new maintenance plan. When your contract is renewed, you have to update your license.

Note that the **Renew** option is subject to restrictions. If online renewal is not possible for your account, you will be redirected to the Renewal Request page. There you will be able to submit a request for Veeam Renewals Team.

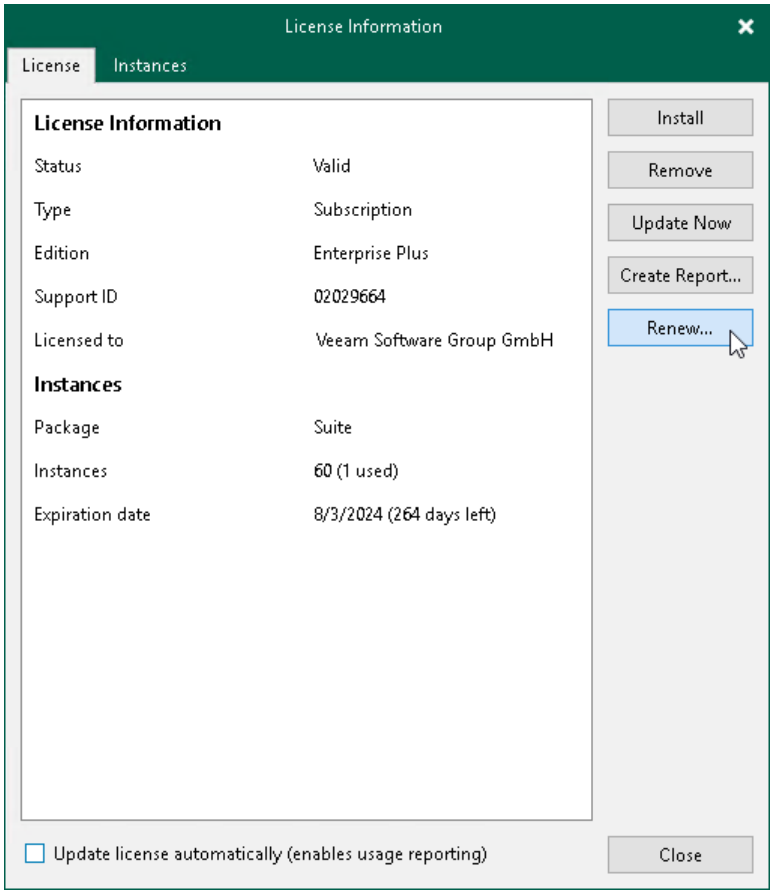

# Installing License

When you install Veeam Backup & Replication, you are asked to specify a path to the license file. If you do not specify a path to the license file, Veeam Backup & Replication will run in the Veeam Backup & Replication Community (free) Edition. For more information, see Veeam Backup & Replication Community Edition.

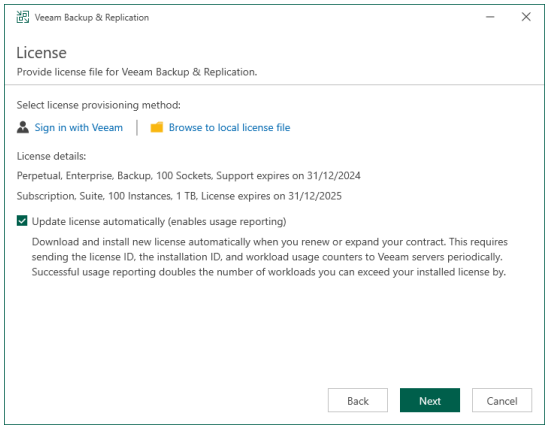


You can install or change the license after product installation:

1. From the main menu, select **License**.

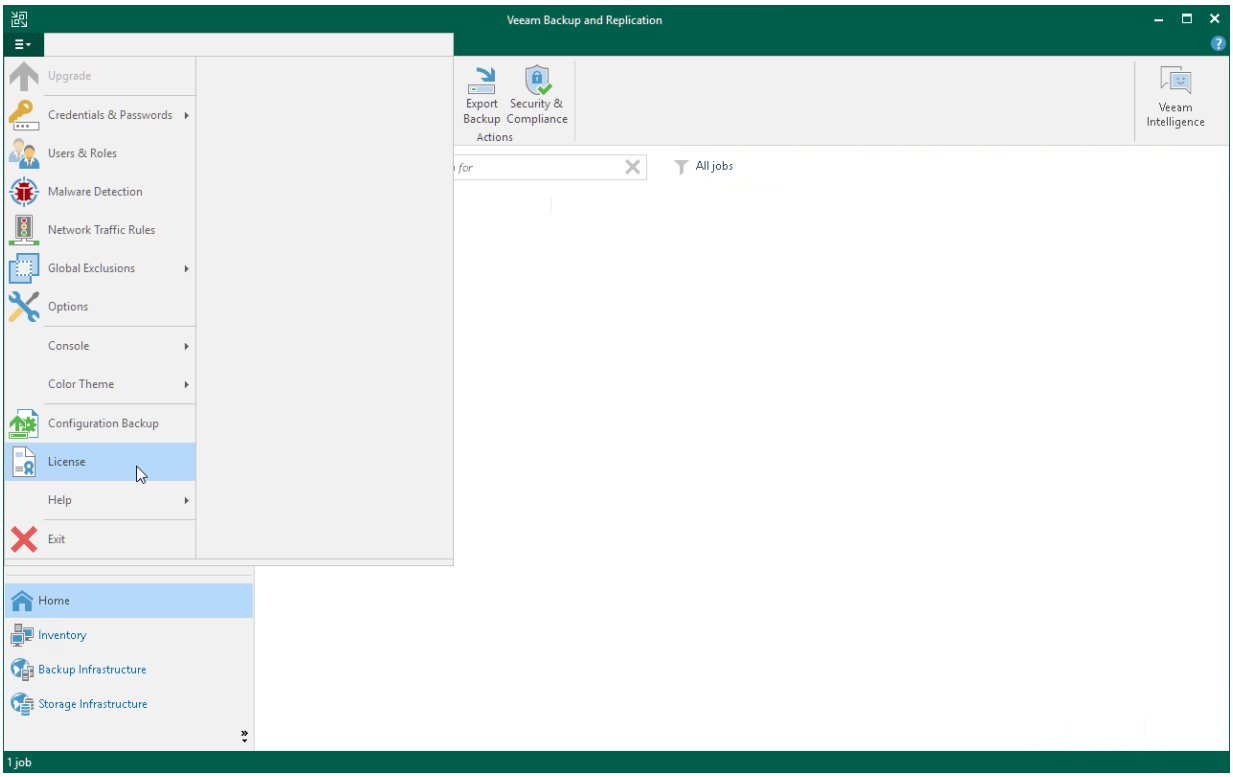

2. In the **License Information** window, click **Install**.

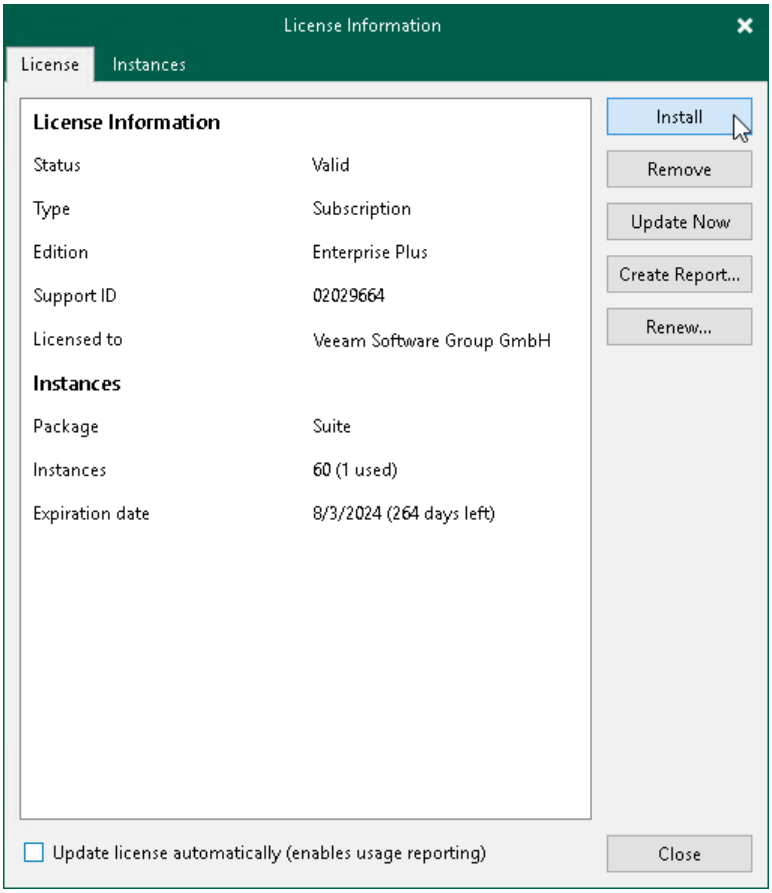


3. Browse to the LIC file and click **Open**.

# Licenses in Veeam Backup Enterprise Manager

If backup servers are connected to Veeam Backup Enterprise Manager, Veeam Backup Enterprise Manager collects information about all licenses installed on backup servers. When Veeam Backup Enterprise Manager replicates databases from backup servers, it also synchronizes license data: checks if the license installed on the backup server coincides with the license installed on the Veeam Backup Enterprise Manager server. If the licenses do not coincide, the license on the backup server is automatically replaced with the license installed on the Veeam Backup Enterprise Manager server.

For information on Veeam Backup Enterprise Manager license, see the Licensing section of the Veeam Backup Enterprise Manager Guide.

# Merging Licenses

Merging licenses is an option for customers who have a Perpetual Socket license. If you have a Perpetual Socket license, and want to also protect, for example, your cloud or physical workloads, or work with Veeam plug-ins, you can obtain an instance license and merge it with the socket license.

> **TIP**
>
> Merging licenses is a built-in mechanism that allows you to merge two different license types (Socket and Instance). If you want to merge several licenses of the same type (Socket and Socket, Instance and Instance) to obtain a single license key file, you can use the merge tool in the Customer Portal. For more information on merging licenses in the Customer Portal, see this Veeam KB article. For more information on merge rules and exceptions for the Customer Portal merging process, see the License Key Merge section of the Veeam Licensing Policy.

Under the merged license, the following workloads will consume a Socket license:

- VMware vSphere and Microsoft Hyper-V VMs
- Application servers protected by Veeam Plug-ins (only for application servers running in VMs of a virtual infrastructure registered with Veeam Backup & Replication)
- VMs protected by Veeam Agents (only for VMs of a virtual infrastructure registered with Veeam Backup & Replication)

Other workloads are processed per instance.

## License Types Available for Merging

You can merge licenses of the following types:

- Perpetual Socket license and Subscription Instance license
- Perpetual Socket license and Perpetual Instance license

## License Packages Available for Merging

You can merge licenses of the following packages:

| Type of license | Essentials Socket | Backup Socket | Suite Socket | ONE Socket |
|---|---|---|---|---|
| **Essentials Instance** | ○ | ○ | ○ | ○ |
| **Backup Instance** | ● | ● | ● | ○ |
| **Suite Instance** | ● | ● | ● | ○ |

| Type of license | Essentials Socket | Backup Socket | Suite Socket | ONE Socket |
|---|---|---|---|---|
| ONE Instance | ○ | ○ | ○ | ○ |

# Support Terms in Merged Licenses

If the licenses that you want to merge have different support expiration date, the merged license will take the support that expires first.

# Merging Licenses

Before you merge licenses, check the following prerequisites:

- The license type and package allow merging.
- The company names are identical in both licenses. Company name check is case sensitive.

To merge licenses, install a new license over the already installed license. For more information, see Installing License.

> **IMPORTANT**
>
> If the license types do not allow merging, the newly installed license will replace the previous license.

# Splitting Licenses

To split licenses, contact Veeam Customer Support.

# Viewing License Information

You can view details of the installed license in the **License Information** window.

To open the **License Information** window, from the main menu select **License**.

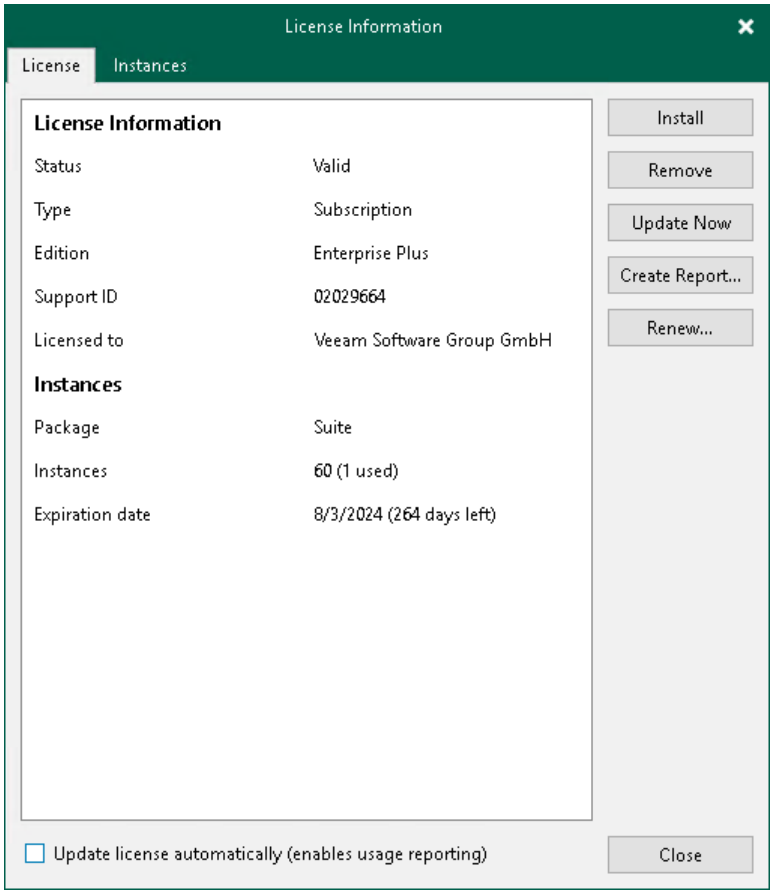


The following details are available for the current license:

- **Status** – license status (*Valid*, *Invalid*, *Expired*, *Not Installed*, *Warning*, *Error*).
- **Type** – license type (*Perpetual*, *Subscription*, *Rental*, *Evaluation*, *NFR*, *Free*).
- **Edition** – license edition (*Community*, *Standard*, *Enterprise*, *Enterprise Plus*).
- **Support ID** – support ID required for contacting Veeam Support.
- **Licensed to** – name of a person or organization to which the license was issued.
- **Cloud Connect Provider** – shows if you can use Veeam Backup & Replication to offer cloud repository as a service and disaster recovery as a service to your customers (*Enterprise*, *Yes*, *No*). For more information on Veeam Cloud Connect, see Veeam Cloud Connect Guide.
- **Package**[1] – Veeam license pack: *Essentials*, *Backup*, *Suite*, *ONE*.
- **Sockets** – number of sockets that you can use to protect workloads.
- **Instances** – number of instances that you can use to protect workloads.
- **Promo instances** – number of additional instances granted by the Promo license.
- **Expiration date** – date when the license expires.

- **Support expiration date** – date when expires the support and maintenance specified by Veeam Licensing Policy. Valid for Perpetual Socket and Perpetual Instance licenses.
- **Promo expiration date** – date when the Promo license expires.
- **Total instances including promo** – number of all available instances, regular and promo added up.
- **Capacity** – protected front end capacity (in TB) for unstructured data backup.

To enable automatic license update, select the **Update license automatically** check box. For more information, see Updating License Automatically.

# Viewing Information on Sockets

With socket licenses, Veeam Backup & Replication applies a license to the virtualization host on which the processed VMs reside.

To view to which objects the license is currently applied, open the **Sockets** tab.

For peculiarities of socket licensing, see the Licensing section.

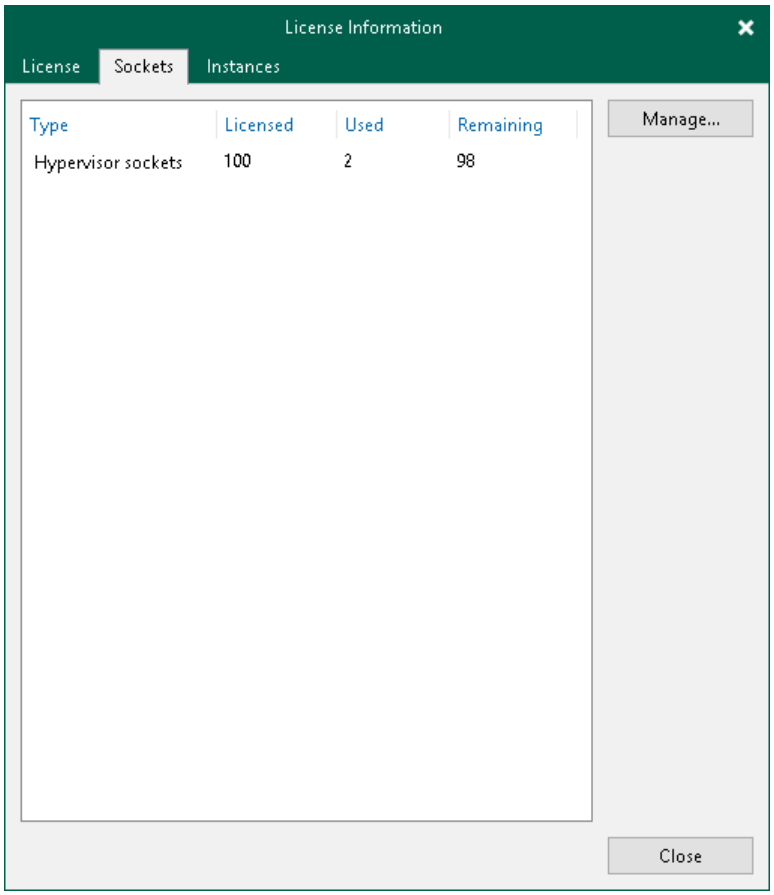


# Viewing Information on Instances

With instance licenses, Veeam Backup & Replication applies a license to a protected workload. The number of license instances that a protected workload consumes depends on the workload type and product edition. For details, see Veeam Licensing Policy.

To view to which objects the license is currently applied, open the **Instances** tab.

For peculiarities of instance licensing, see the Licensing section.

By default, Veeam Backup & Replication allows Veeam Agents to connect to the Veeam backup server and consume instances in the license. If you do not want Veeam Agents to consume instances, clear the **Allow unlicensed agents to consume instances** check box. For more information on Veeam Agents licensing, see the following user guides:

- If you work with Veeam Agents operating in the managed mode, see the Licensing Requirements section in the Veeam Agent Management Guide.
- If you work with Veeam Agents operating in the standalone mode, see the user guide for the Veeam Agent depending on the operating system of the protected computer. For example, if you work with Veeam Agent for Microsoft Windows, see the Managing License section in the Veeam Agent for Microsoft Windows User Guide.

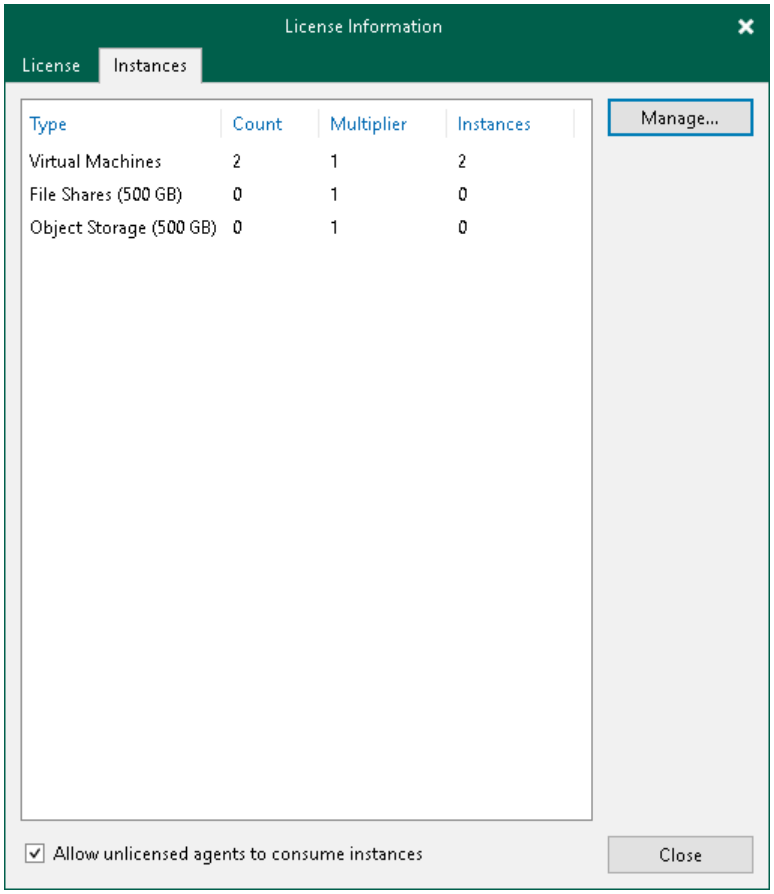


# Viewing Information on Capacity

With a capacity license, Veeam Backup & Replication applies the license to protected unstructured data sources. For details, see Veeam Licensing Policy.

To view which data sources currently consume the license capacity, open the **Capacity** tab.

For peculiarities of capacity licensing, see the Licensing section.

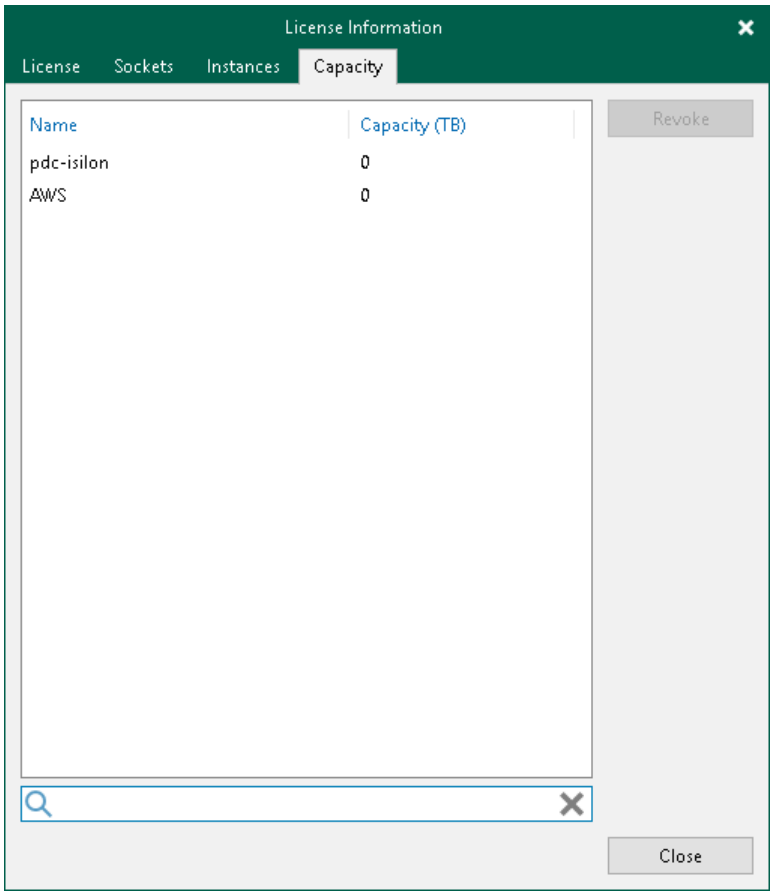


# Creating Current License Consumption Report

If you have an instance license and at least one instance is already consumed by Veeam Backup & Replication, you can create a **Current License Consumption** report. To run the report, click **Create Report** in the **License Information** window.

The report contains the following information about instance license consumption:

- License information

  This section contains the general information about the instance license:

  - License edition (Community, Standard, Enterprise, Enterprise Plus)
  - License expiration date
  - Name of a person or organization to which the license was issued
  - Support ID
  - Installation ID

- Summary of license consumption

  The section contains a short overview of instance consumption by different workload types (virtual machines, physical machines, unstructured data backup, and so on).

- License consumption by workloads

  This section contains detailed information about consumed instances, which are grouped by workload types. Each sub-section lists the following information about the workloads:

  - Workload name
  - Number of the consumed instances
  - Workload sub-type (vSphere, Hyper-V, Proxmox VE VM, and so on)
  - Name of the job that consumed the instances
  - Date of the last job run
  - Notes

# Revoking License

You can revoke licenses from protected workloads or licensed hosts, and re-apply them to other objects that you plan to protect. License revoking can be helpful, for example, if a licensed host goes out of service or you do not want to protect some workloads anymore.

> **NOTE**
>
> If you manually revoke license instances allocated for an unstructured data source, the next run of the unstructured data backup job that protects this data source will trigger the recalculation of the data source protected size and reallocation of license instances that Veeam Backup & Replication will consume. For more information, see the Instance Consumption for Object Storage Backup, File Backup and File to Tape Jobs section.

To revoke a license, do the following:

1. From the main menu, select **License**.
2. In the **License Information** window:
   - For protected workloads, open the **Instances** tab and click **Manage**.
   - For licensed hosts, open the **Sockets** tab and click **Manage**.
3. In the displayed window, select a protected workload or a licensed host and click **Revoke**.
   Veeam Backup & Replication will revoke the license from the selected object, and the license will be freed for other objects in the backup infrastructure.

The steps to revoke licenses for Veeam Agent machines are slightly different. For more information, see the user guide for the Veeam Agent depending on the operating system of the protected computer. For example, if you work with Veeam Agent for Microsoft Windows, see the Viewing Licensed Veeam Agents and Revoking License section in the Veeam Agent for Microsoft Windows User Guide.

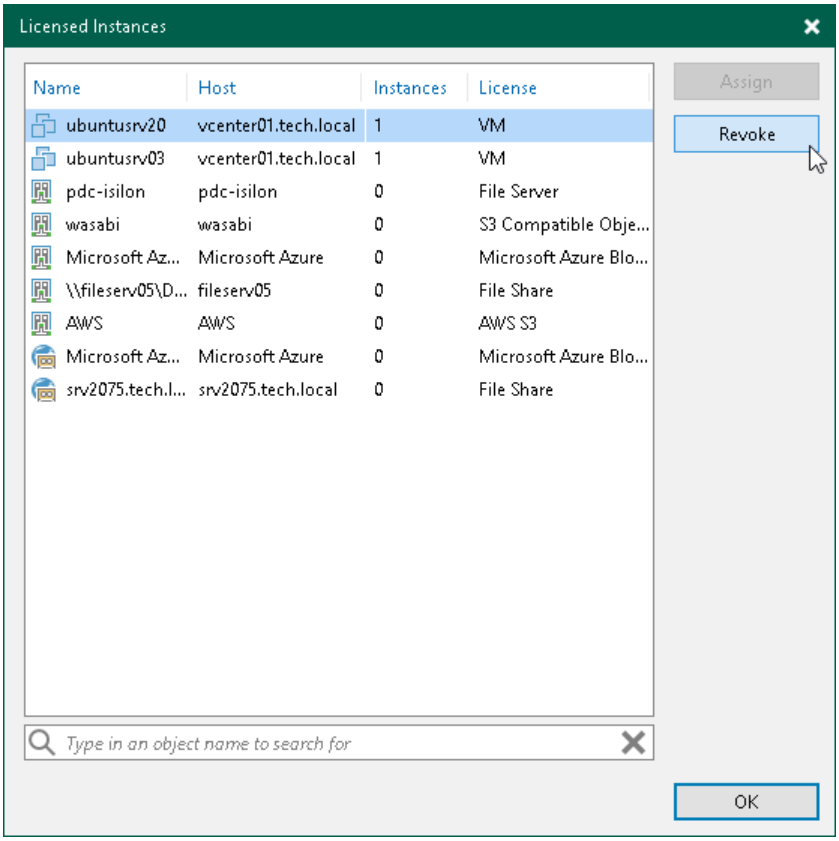

# Removing License

You can remove the installed license. When you remove a license, Veeam Backup & Replication will switch to the Veeam Backup & Replication Community Edition. For more information, see Veeam Backup & Replication Community Edition.

You can also remove a part of merged license. If you do so, Veeam Backup & Replication will operate under the other part of the merged license. For more information, see Merging Licenses.

To remove a Promo license, remove the license on top of which it was installed.

To remove a license, do the following:

1. From the main menu, select **License**.

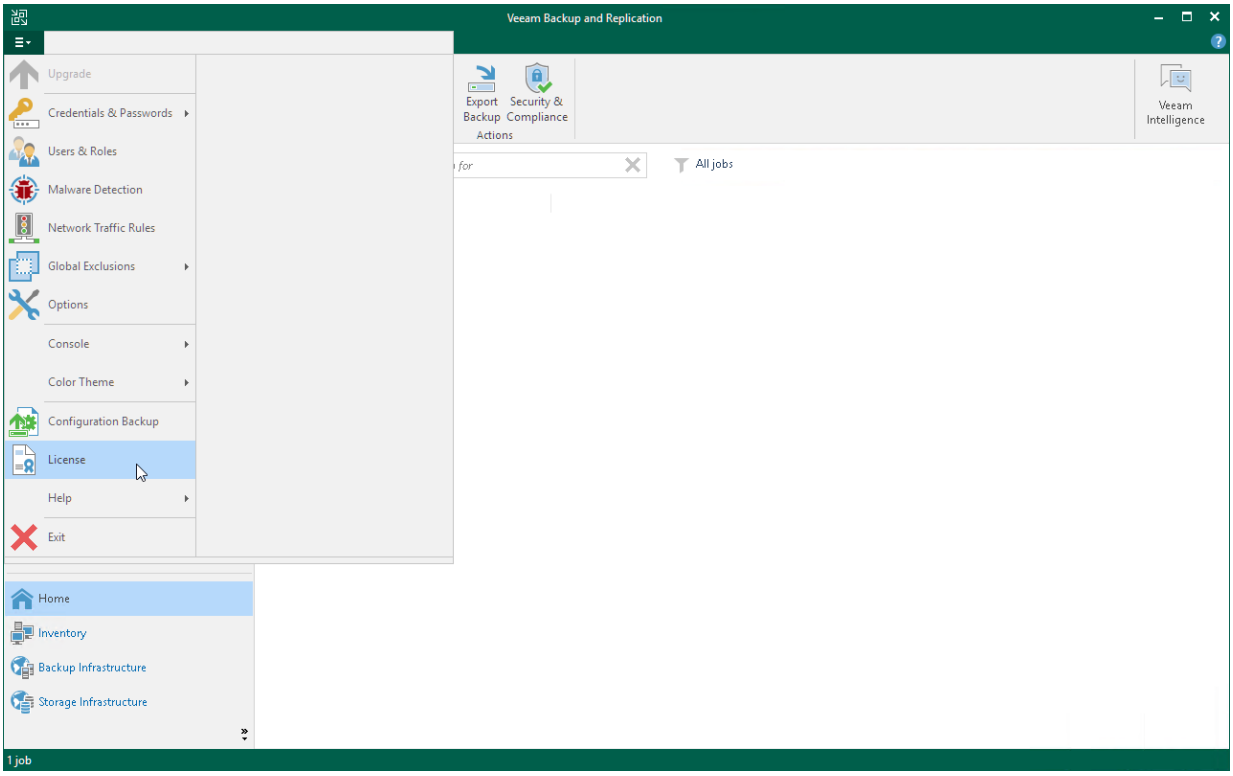

2. In the **License Information** window, click **Remove**.

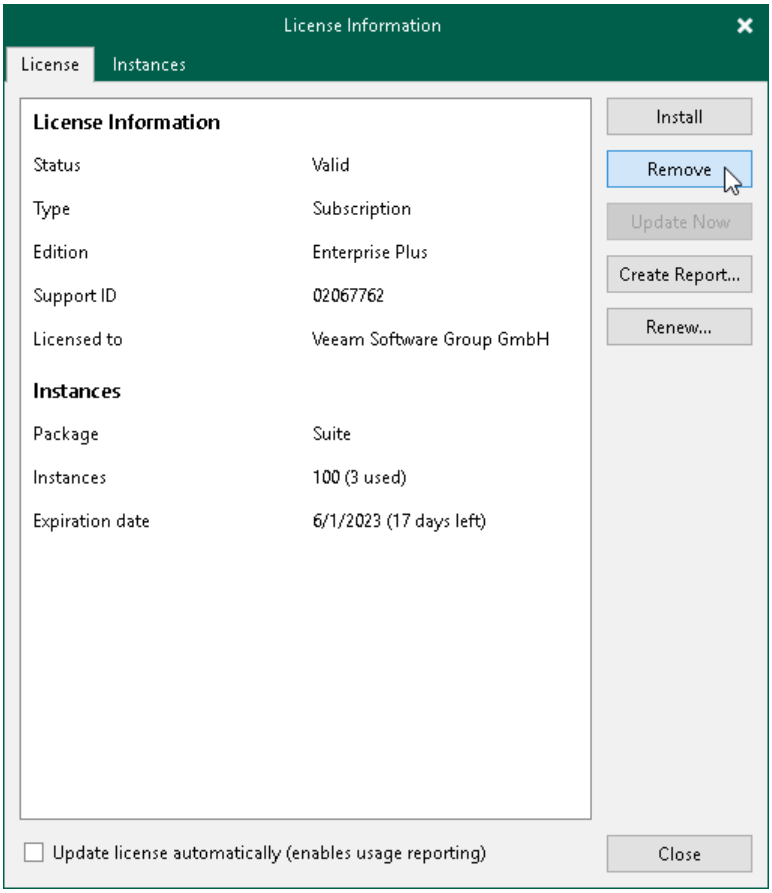


If you have both socket and instance licenses, you will be asked which license you want to remove.

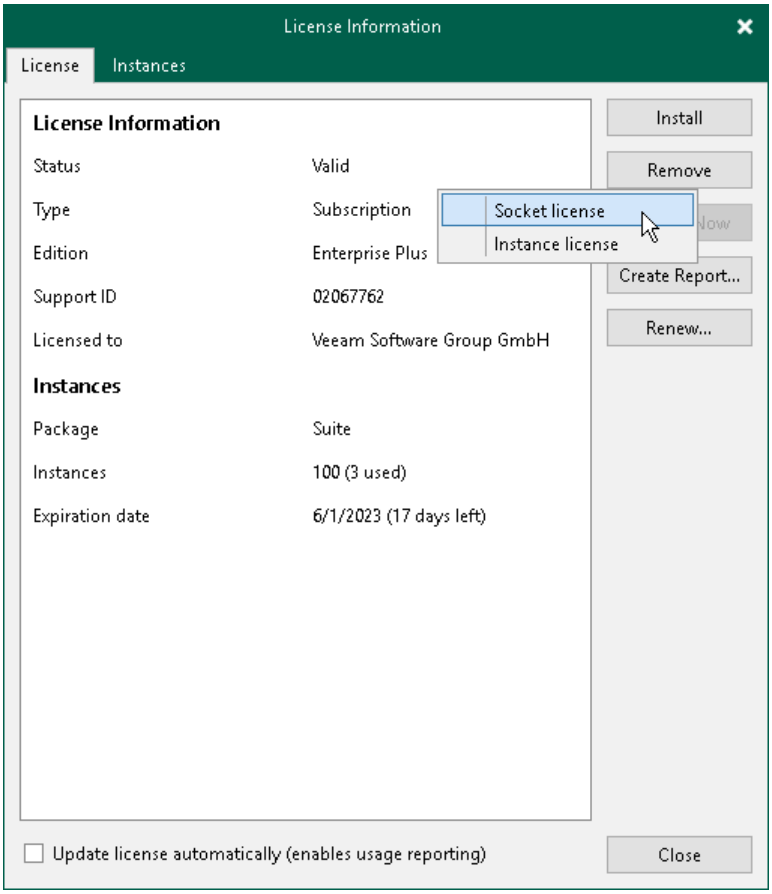

3. Select the license you want to remove and confirm the removal.

# Exceeding License Limit

For Veeam Universal Licenses, Veeam Backup & Replication allows you to protect more workloads than covered by the number of instances specified in the license. An increase in the number of protected workloads is allowed throughout the duration of the contract (license key).

The license limit can be exceeded by a number of instances, or a percentage of the total number of instances specified in the license (depends on which number is greater). The exceeding limit varies according to the license type.

| License Type | Exceeding Limit | Description |
|---|---|---|
| **Socket license** | Not allowed | Workloads that are exceeding the license limit are not processed. |
| **Veeam Universal License (VUL)**<br>**License autoupdate ENABLED** | Less than 10 instances (or 10% of the total instance count) | All protected workloads are processed normally, Veeam Backup & Replication does not display a warning message. |
| | 10-20 instances (or 10%-20% of the total instance count) | All protected workloads are processed normally.<br>Once a week when you open the Veeam Backup & Replication console, a warning message is displayed notifying that you are out of compliance with the Veeam Licensing Policy. |
| | More than 20 instances (or 20% of the total instance count) | Workloads that are exceeding the license limit beyond 20 instances (or 20% of the total instance count) are not processed.<br>Every time you open the Veeam Backup & Replication console, a warning message is displayed notifying that you are out of compliance with the Veeam Licensing Policy. |
| **Veeam Universal License (VUL)**<br>**License autoupdate DISABLED** | Less than 5 instances (or 5% of the total instance count) | All protected workloads are processed normally, Veeam Backup & Replication does not display a warning message. |
| | 5-10 instances (or 5%-10% of the total instance count) | All protected workloads are processed normally.<br>Once a week when you open the Veeam Backup & Replication console, a warning message is displayed notifying that you are out of compliance with the Veeam Licensing Policy. |

| License Type | Exceeding Limit | Description |
|---|---|---|
| | More than 10 instances (or 10% of the total instance count) | Workloads that are exceeding the license limit beyond 10 instances (or 10% of the total instance count) are not processed.<br><br>Every time you open the Veeam Backup & Replication console, a warning message is displayed notifying that you are out of compliance with the Veeam Licensing Policy. |
| **Rental license** | See the Rental License section in the Veeam Cloud Connect Guide. | |

For example, you have a VUL license subscription with 500 instances to protect your workloads. The license autoupdate feature is enabled. According to the table above, you are allowed to use up to 20 instances or 20% of the total instance count (whichever number is greater) over the license limit. As the number of instances in your license is 500, you are allowed to use additional 100 instances (100 makes *20%* of 500, and 100 is greater than *20*). Consider the following:

- Until the license limit is not exceeded by more than 10% of the total instance count (up to 50 instances), Veeam Backup & Replication processes all protected workloads with no restrictions.
- When the license limit is exceeded by 10%-20% (50 to 100 instances), Veeam Backup & Replication processes protected workloads, and displays a warning message once a week when you open the Veeam Backup & Replication console. In the message, Veeam Backup & Replication provides information on the number of exceeded instances and the number of instances by which the license can be further exceeded.
- If the license limit is exceeded by more than 20% (100 instances and more), Veeam Backup & Replication does not process the workloads exceeding the limit, and displays a warning message every time you open the Veeam Backup & Replication console. In the message, Veeam Backup & Replication provides information on the number of instances by which the license is exceeded.

When the license limit is exceeded, the logs will include the number of instances necessary to finish the job successfully.

> **NOTE**
>
> Capacity-based (per-TB) license limit for unstructured data backup cannot be exceeded.

# License Expiration

When the license expires, Veeam Backup & Replication behaves in the following way depending on the license type:

- Evaluation and NFR licenses: Veeam Backup & Replication will stop processing workloads.
- Paid licenses: Veeam Backup & Replication will switch to the grace period.
- Promo license: Veeam Backup & Replication will remove the granted instances and stop processing workloads for them. Promo license does not have a grace period. Upon expiration of the primary license, the promo license will also expire, regardless of its own expiration date.

Perpetual Socket and Perpetual Instance licenses do not expire. However, such licenses have support expiration date. Veeam Backup & Replication will inform you about the support expiration date.

## Grace Period

To ensure a smooth license update and provide sufficient time to install a new license file, Veeam Backup & Replication offers a grace period. Grace period is available for paid licenses.

During the grace period, you can perform all types of data protection and disaster recovery operations. However, Veeam Backup & Replication will inform you about the license expiration when you open the Veeam Backup & Replication console. The license status in the **License Information** window will appear as *Expired*.

You must update your license before the end of the grace period. If you do not update the license, Veeam Backup & Replication stops processing workloads. All existing jobs fail with the *Error* status. However, you will be able to restore machine data from existing backups.

## Grace Period Duration

Before the license expires, Veeam Backup & Replication notifies you about soon license expiration.

The number of days for notification and grace period depends on the type of license:

| License Type | License Expiration Notification | Grace Period |
|---|---|---|
| **Subscription** | 30 days | 30 days |
| **Perpetual Instance** | 14 days before Support expiration date | n/a |
| **Perpetual Socket** | 14 days before Support expiration date | n/a |
| **Rental** | 7 days | 60 days |
| **Evaluation** | 30 days | 0 days |

| License Type | License Expiration Notification | Grace Period |
|---|---|---|
| **NFR** | 30 days | 0 days |
| **Promo** | 7 days | n/a |

# Switching to Veeam Backup & Replication Community Edition

If you do not want to renew your license, you can switch to the free product version named Veeam Backup & Replication Community (free) Edition. For more information, see Veeam Backup & Replication Community Edition.

To do so, remove your license. For more information, see Removing License.

# Expiration of Merged Licenses

When the merged license expires, Veeam Backup & Replication stops processing workloads after the grace period.

If you merged licenses with different expiration dates, the merged license will expire on the date that is closer. For example, if you merged a Perpetual license and a Subscription license, the expiration date will be inherited from the Subscription license.

In such case, you can update your Subscription license or continue using the Perpetual license. To continue using the Perpetual license, remove the Subscription license. For more information, see Removing License.

# Updating License

To be able to use all data protection and disaster recovery features, you must update your license upon expiry. There are two methods to update the license in Veeam Backup & Replication:

- Update the license manually.
- Update the license automatically.

> **NOTE**
>
> When updating the license, Veeam Backup & Replication requires internet access to connect to the Veeam License Update Server. If your network is not connected to the internet, instead of update you can download a new license file from my.veeam.com and then install it. For more information on license installation, see Installing License.

# Updating License Manually

You can update the license manually on demand. When you update the license manually, Veeam Backup & Replication connects to the Veeam License Update Server, downloads a new license from it (if the license is available) and installs it on the backup server.

> **IMPORTANT**
>
> To update the license, the backup server in your Veeam Backup & Replication installation must have access to the Veeam License Update Server (vbr.butler.veeam.com, autolk.veeam.com) using TCP on port 443.

The new license key differs from the previously installed license key in the license expiration date or support expiration date. If you have obtained a license for a greater number of instances, counters in the new license also display the new number of license instances.

To update the license:

1. From the main menu, select **License**.
2. In the **License Information** window, click **Update Now**.

Statistics on the manual license update process is available under the **System** node in the **History** view. You can double-click the **License key auto-update** job to examine session details for the license update operation.

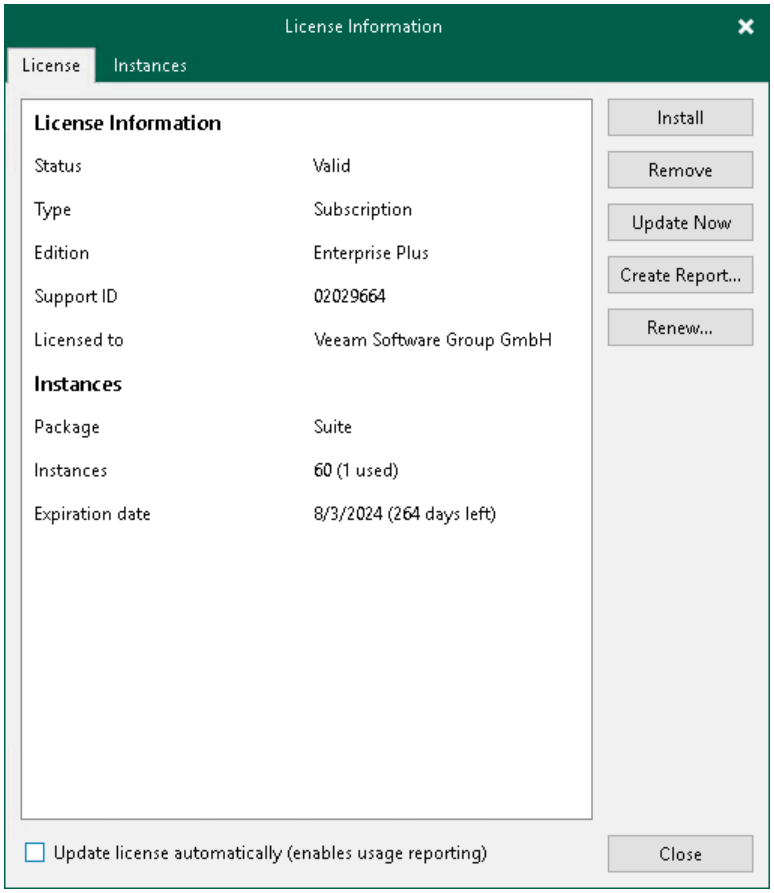

# Manual License Update Results

Manual license update can complete with the following results:

- **Operation is successful**. A new license key is successfully generated, downloaded and installed on the backup server or Veeam Backup Enterprise Manager server.
- **A new license is not required**. The currently installed license key does not need to be updated.
- **The Veeam License Update Server has failed to generate a new license**. Such situation can occur due to some error on the Veeam License Update Server side.
- **Veeam Backup & Replication has received an invalid answer**. Such situation can occur due to connectivity issues between the Veeam License Update Server and Veeam Backup & Replication.
- **Licensing by the contract has been terminated**. In such situation, Veeam Backup & Replication automatically disables automatic license update on the backup server or Veeam Backup Enterprise Manager server.

# Updating License Automatically

You can instruct Veeam Backup & Replication to automatically update the license installed on the backup server or Veeam Backup Enterprise Manager server. With automatic license update, you do not need to download and install the license manually each time when you purchase the license extension. If the automatic update option is enabled, Veeam Backup & Replication proactively communicates with the Veeam License Update Server to obtain and install a new license before the current license expires.

## Requirements and Limitations for Automatic License Update

- Automatic license update is not available in the Veeam Backup & Replication Community Edition.
- Only licenses that contain a real contract number in the Support ID can be updated with the **Update license key automatically** option.
- If you are managing backup servers with Veeam Backup Enterprise Manager, all license management tasks must be performed in the Veeam Backup Enterprise Manager console. Automatic update settings configured in Veeam Backup Enterprise Manager override automatic update settings configured in Veeam Backup & Replication.

  For example, if the automatic update option is enabled in Veeam Backup Enterprise Manager but disabled in Veeam Backup & Replication, automatic update will be performed anyway. For more information, see the Veeam Backup Enterprise Manager User Guide.

> **NOTE**
>
> Veeam Backup & Replication does not automatically update an existing per-VM or socket license that was obtained for an earlier version of the product to a new instance license.
>
> To overcome this issue, after you upgrade to Veeam Backup & Replication 11, you must obtain in the Veeam Customer Support Portal a new instance license and install it on the backup server manually.

## How Automated License Update Works

To update installed licenses automatically, Veeam Backup & Replication performs the following actions:

1. After you enable automatic license update, Veeam Backup & Replication starts sending requests to the Veeam License Update Server on the web (autolk.veeam.com) and checks if a new license key is available. Veeam Backup & Replication sends requests once a week. Communication with the Veeam License Update Server is performed over the HTTPS protocol.
2. Seven days before the expiration date of the current license, Veeam Backup & Replication starts sending requests once a day.
3. When a new license key becomes available, Veeam Backup & Replication automatically downloads it and installs on the backup server or Veeam Backup Enterprise Manager server.

The new license key differs from the previously installed license key in the license expiration date and support expiration date. If you have obtained a license for a greater number of instances, counters in the new license also display the new number of license instances.

# Automatic License Update Results

Automatic license update can complete with the following results:

- **Operation is successful**. A new license key is successfully generated, downloaded and installed on the backup server or Veeam Backup Enterprise Manager server.
- **A new license is not required**. The currently installed license key does not need to be updated.
- **The Veeam License Update Server has failed to generate a new license**. Such situation can occur due to some error on the Veeam License Update Server side.
- **Veeam Backup & Replication has received an invalid answer**. Such situation can occur due to connectivity issues between the Veeam License Update Server and Veeam Backup & Replication.
- **Licensing by the contract has been terminated**. In such situation, Veeam Backup & Replication automatically disables automatic license update on the backup server or Veeam Backup Enterprise Manager server.

# Automatic Update Retries

If Veeam Backup & Replication fails to update the license, it displays a notification in the session report and sends an email notification to users specified in the global email settings (if global email settings are configured on the backup server). You can resolve the issue, while Veeam Backup & Replication will keep retrying to update the license.

Veeam Backup & Replication retries to update the license key in the following way:

- If Veeam Backup & Replication fails to establish a connection to the Veeam License Update Server, retry takes place every 60 min.
- If Veeam Backup & Replication establishes a connection but you are receiving the *"General license key generation error has occurred"* message, the retry takes place every 24 hours.

The retry period ends one month after the license expiration date or the support expiration date (whichever is earlier). The retry period is equal to the number of days in the month of license expiration. For example, if the license expires in January, the retry period will be 31 day; if the license expires in April, the retry period will be 30 days.

If the retry period is over but the new license has not been installed, the automatic update feature is automatically disabled.

For more information about error cases, see the License Update Session Data section in the Veeam Backup Enterprise Manager Guide.

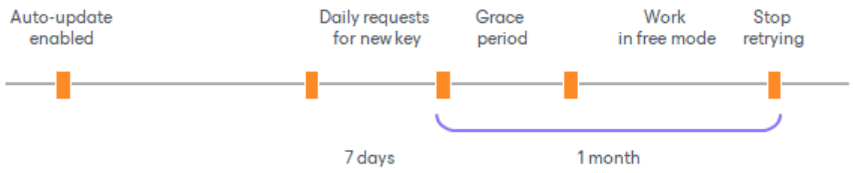


# Enabling Automatic License Update

By default, automatic license update is disabled. To enable automatic license update:

1. From the main menu, select **License**.

2. In the **License Information** window, select the **Update license automatically (enabled usage reporting)** check box.

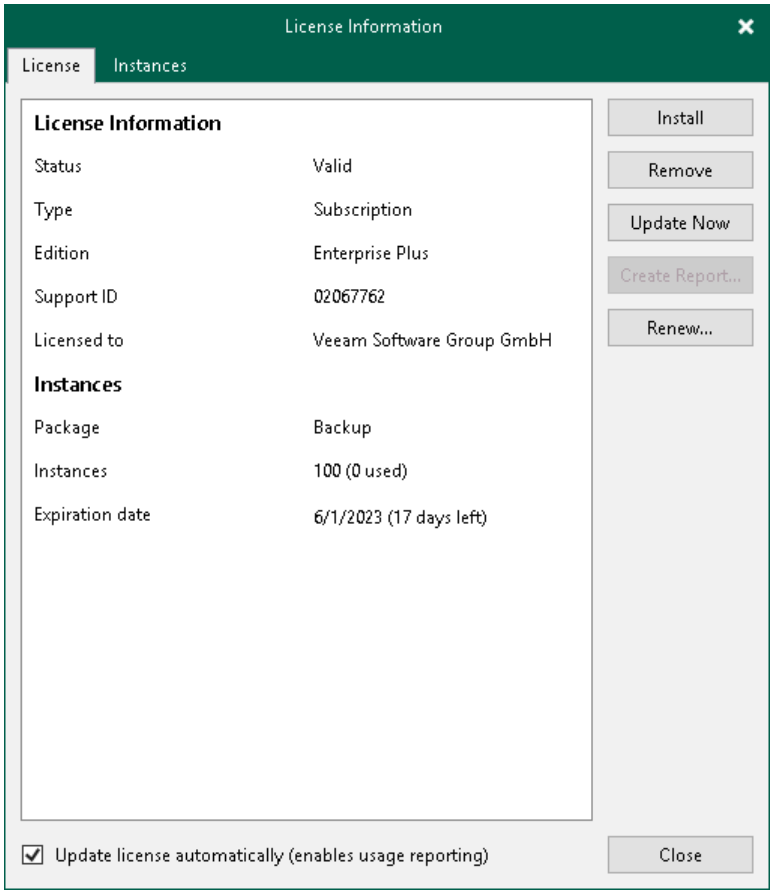


During the installation of a Subscription or Rental license, you will see a dialog box with a suggestion to enable automatic license update.

Statistics on the automatic license update process is available under the **System** node in the **History** view. You can double-click the **License key auto-update** job to examine session details for the scheduled or ad-hoc automatic license update.

> **NOTE**
>
> [For Rental, Subscription, Perpetual licenses] Enabling license auto update activates Automatic License Usage Reporting.

# Automatic License Usage Reporting

When license auto update is enabled for Rental, Subscription, Perpetual licenses, Veeam Backup & Replication performs automatic license usage reporting.

As part of reporting, Veeam Backup & Replication collects statistics on the current license usage and sends it periodically to the Veeam License Update Server. The report provides information about the contract ID, product installation ID, and the maximum number of licensed objects that were managed by Veeam Backup & Replication over the past week (high watermark). The reporting process runs in the background mode, once a week at a random time and day.

The type of reported objects is defined by the product and the installed license. The report can include information about VMs, workstations or servers protected with Veeam backup agents, and so on.

The collected data does not include information on the usage of Veeam Backup & Replication by any individual person identifiable for Veeam, or any data protected by Veeam Backup & Replication.

The collected data allows our back-end system to automatically approve your monthly usage reports as long as they do not deviate from the high watermark value significantly. This helps to keep our report processing costs low, thus allowing us to maintain low rental prices for our solution. Veeam may also use collected data for any other internal business purposes it deems appropriate, including (but not limited to) evaluation, improvement and optimization of Veeam licensing models.

By enabling license auto update you agree with collection, transmission and use of the reporting data. You must not enable license auto update in case you do not agree with such collection, transmission and use.

# Deployment

This section describes how to install, upgrade, update and uninstall Veeam Backup & Replication and the Veeam Backup & Replication console, both through UI and using commands for silent deployment.

# Installation

To start working with Veeam Backup & Replication, you must configure a backup server — install Veeam Backup & Replication on a machine that meets the system requirements. To do this, you can use the setup wizard or install the product in the unattended mode.

When you install Veeam Backup & Replication, the Veeam Backup & Replication console is automatically installed on the backup server. If you want to access Veeam Backup & Replication remotely, you can install the Veeam Backup & Replication console on a dedicated machine.

# Installing Veeam Backup & Replication

Before you install Veeam Backup & Replication, check prerequisites. Then use the **Veeam Backup & Replication** setup wizard to install the product.

## Before You Begin

Before you install Veeam Backup & Replication, check the following prerequisites:

- A machine on which you plan to install Veeam Backup & Replication must meet the system requirements. For more information, see System Requirements.
- A user account that you plan to use for installation must have sufficient permissions. For more information, see Permissions.
- Backup infrastructure components communicate with each other over specific ports. These ports must be open. For more information, see Ports.
- You must remove Veeam Backup & Replication components of versions that are not supported by the upgrade procedure from the target machine. You may also need to remove earlier versions of other Veeam products and components.
- Before deploying Veeam Backup & Replication, define where the Veeam Backup & Replication server will be located. Depending on what kind of protection are you planning to use, the Veeam Backup & Replication server should be located on the source site or the Disaster Recovery site.
  - **When replication or CDP is used**: If you plan to use replication or Continuous Data Protection (CDP), the Veeam Backup & Replication server should be deployed in the disaster recovery site. In this case, if the production host crashes, Veeam Backup & Replication will automatically fail over to the replica without any manual operations. The source backup infrastructure still can be managed with the same Veeam Backup & Replication server with the help of backup proxies deployed in the source site.
  - **When only backup features are used**: If you plan to use Veeam Backup & Replication for backup jobs only, the backup server should be placed in the production site.
- We strongly recommend that no highly-transactional and business-critical software is deployed on the same machine as the Veeam backup server. This could be (but not limited to) software such as Active Directory, Exchange Server or other intensive production databases on the SQL server instance. We recommend that only Veeam Backup & Replication runs on the backup server.

## Configuration Database

Before deploying Veeam Backup & Replication, decide which database engine and version you need to use:

- If you do not prepare a database engine in advance, Veeam Backup & Replication will automatically install PostgreSQL locally on the backup server.
- Decide if you need the database engine installed on the same server as Veeam Backup & Replication or on a remote server. Veeam Backup & Replication requires a database engine deployed either locally on the backup server or remotely.

  It is recommended to run the database engine instance locally to eliminate latency issues. However, in some scenarios a remote instance can be the better choice:

  - High Availability. SQL Clustering and Always On Availability Group on external SQL Servers can be used for high availability of the configuration database. To learn about the configuration details, see this Veeam KB article.

    - Licensing. Some enterprises have dedicated virtual clusters for SQL Servers due to licensing constraints. In such cases, you can place the Veeam configuration database on an existing instance to lower the total cost of ownership.

- If you prefer to use an Express Edition of Microsoft SQL Server as your database engine, note that its usage is limited by 10 GB of configuration database. The Express Edition is enough for evaluation purposes and not very large environments (<500 VMs). If your infrastructure is large (more than 500 VMs), you may consider to install a Microsoft SQL Server in advance.
- If Microsoft SQL Server is installed by the previous product version, Veeam Backup & Replication will connect to the existing configuration database, upgrade it (if necessary) and use it for work.
- If you already have an installed instance of PostgreSQL and want to use it for the configuration database, ensure that the LocalSystem account is added to your PostgreSQL configuration to successfully run the installation of Veeam Backup & Replication.

# Configuring PostgreSQL Instance

The default PostgreSQL instance is configured to consume a minimum amount of resources, which may not be enough for Veeam Backup & Replication performance.

When installing Veeam Backup & Replication, you can choose what PostgreSQL instance to use for the Veeam Backup & Replication configuration database. You can use an already installed PostgreSQL instance or install a new one.

- If you let the setup install a new PostgreSQL instance, it will be configured automatically.
- If you want to use an already installed PostgreSQL instance, make sure that the instance configuration is sufficient for the Veeam Backup & Replication performance.

To adjust the configuration of an existing PostgreSQL instance, see Adjusting PostgreSQL Instance Configuration.

# Related Topics

Planning and Preparation

# Step 1. Start Setup Wizard

To start the setup wizard:

1. Download the latest version of the Veeam Backup & Replication installation image from the Download Veeam Products page.
2. Mount the installation image to the machine on which you plan to install Veeam Backup & Replication or burn the image file to a flash drive or other removable storage device. If you plan to install Veeam Backup & Replication on a VM, use built-in tools of the virtualization management software to mount the installation image to the VM.

   To extract the content of the ISO, you can also use the latest versions of utilities that can properly extract data from ISOs of large size and can properly work with long file paths.
3. After you mount the image or insert the disk, Autorun will open a splash screen. If Autorun is not available or disabled, run the `Setup.exe` file from the image or disk.
4. In the splash screen, click **Install**.

**IMPORTANT**

It is strongly recommended that you install Veeam Backup & Replication using Autorun or the Setup.exe file. If you run other installation files from the ISO folders, you may miss some components that need to be installed, and Veeam Backup & Replication may not work as expected.

# Step 2. Select Component

At the **Select Veeam Backup & Replication Component** step of the wizard, select **Install Veeam Backup & Replication**.

To open Veeam Help Center from the setup wizard, click **View Documentation**.

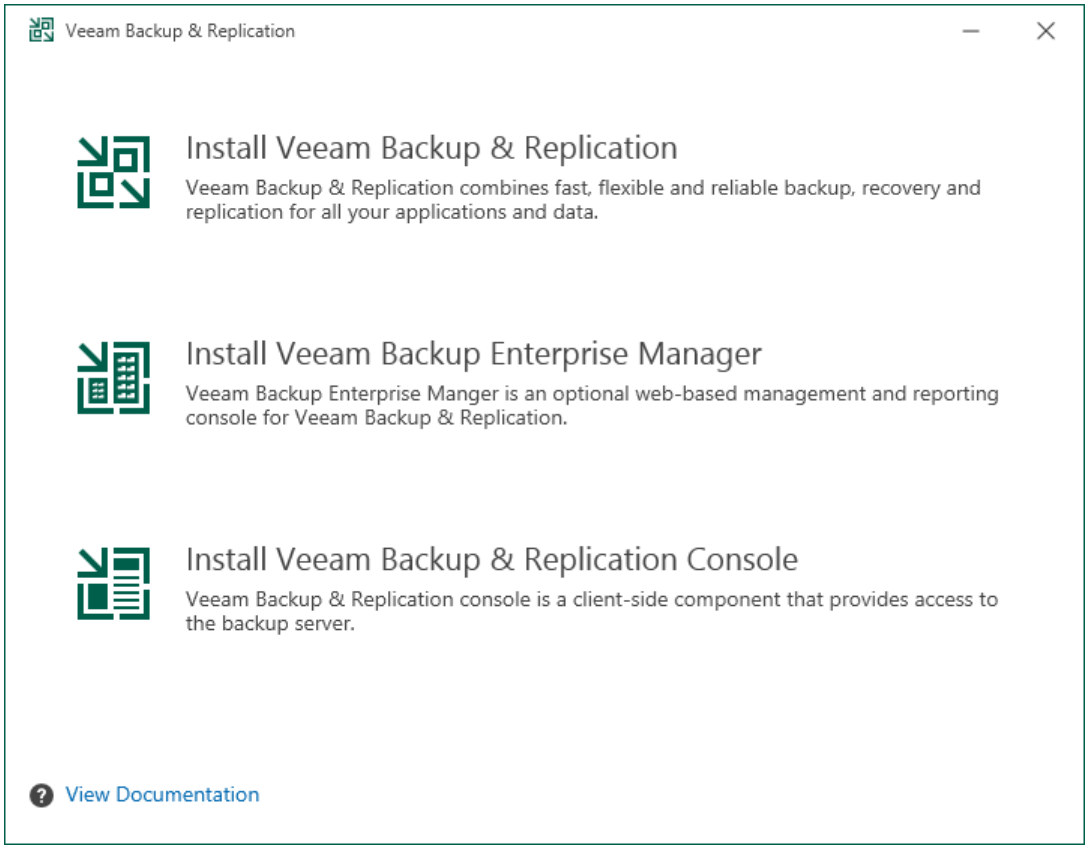

# Step 3. Read and Accept License Agreement

At the **License Agreement** step of the wizard, read Veeam License Agreement and licensing policy as well as license agreements of 3rd party components that Veeam incorporates and license agreements of required software. To accept the license agreements and continue installing Veeam Backup & Replication, click **I Accept**.

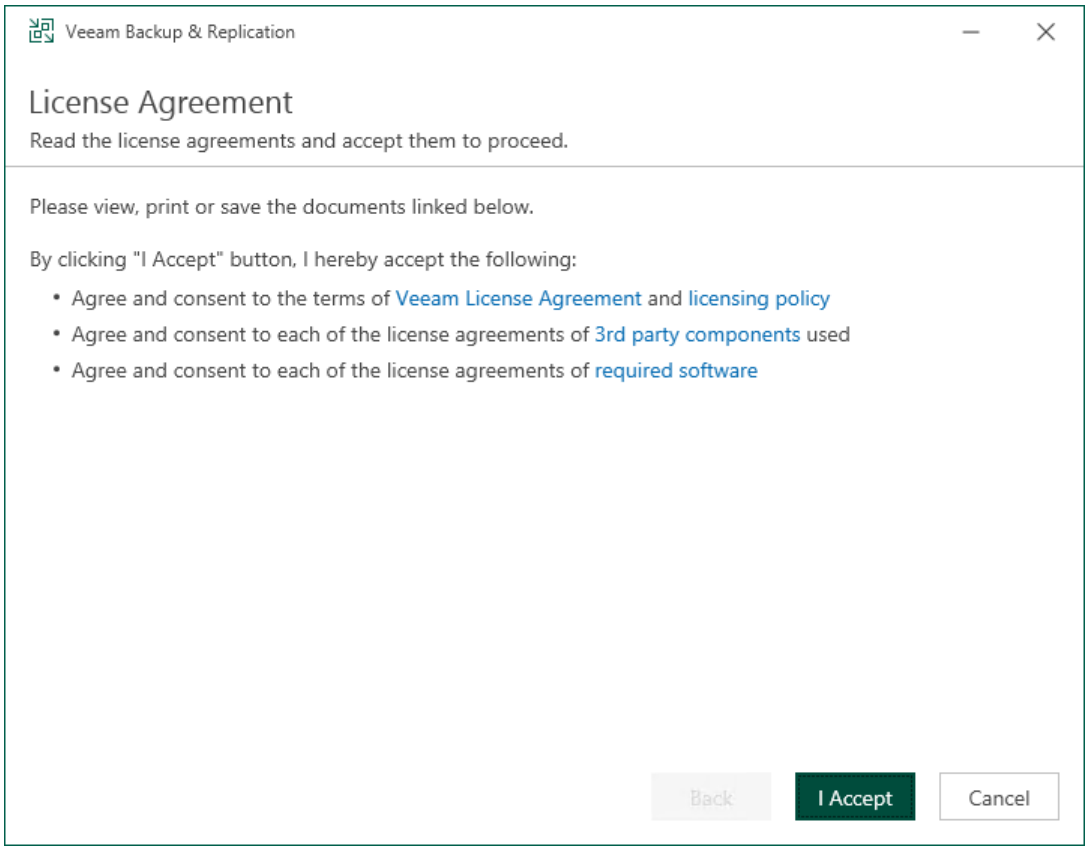

# Step 4. Provide License File

At the **License** step of the wizard, specify what license you want to install for Veeam Backup & Replication. For more information, see Licensing.

If a valid license is already installed on the machine, the setup wizard will inform you about it. In this case, you can skip the **Provide License** step and move to the next step of the wizard.

If you do not install a license, the product will operate in the Veeam Backup & Replication Community (free) Edition. For more information, see Veeam Backup & Replication Community Edition.

To install a license, you have 2 options to choose from:

- Browse your local server or network locations for a license file:
  a. Click **Browse license file**.
  b. Select a valid license file for Veeam Backup & Replication.
- Select a license from your account at the Veeam website:
  a. Click **Sign in with Veeam**.
  b. Enter your credentials for accessing the Veeam website and click **Sign in**.
  c. Select one of the available licenses and click **Install selected license**.

To install new licenses automatically when you renew or expand your contract, select the **Update license automatically** check box. If you enable the automatic license update, and therefore enable usage reporting, you will double the number of workloads by which you can exceed your installed license. For more information, see Exceeding License Limit.

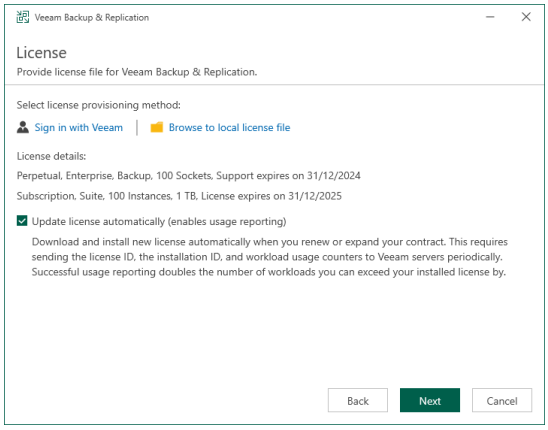

# Step 5. Install Missing Software

At the **System Configuration Check** step of the wizard, the setup wizard checks if the required software is installed on the machine. If some of the required components are missing, the setup will try to install them automatically. After the components are successfully installed, reboot is required. To reboot the machine, click **Reboot**.

If the setup wizard cannot install some of the required software components automatically, install them manually and click **Retry**.

> **NOTE**
>
> If all required software is already installed on the machine, the **System Configuration Check** step will be skipped. For more information on the necessary software, see System Requirements.

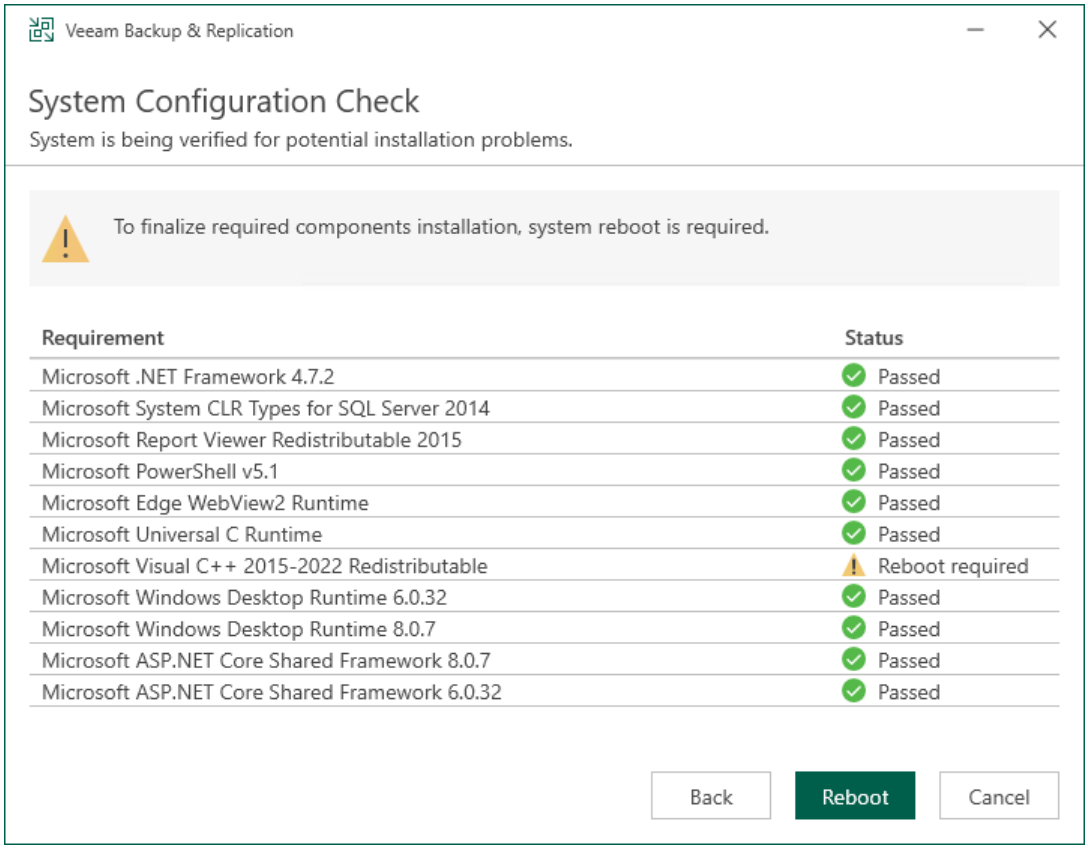

# Step 6. Review Default Installation Settings

At the **Ready to Install** step of the wizard, you can select to install Veeam Backup & Replication with default installation settings or specify custom installation settings.

- To use the default installation settings, click **Install**.
- To use custom installation settings, click **Customize Settings**. The setup wizard will include additional steps that will let you configure installation settings.

The following table lists the default installation settings.

| Setting | Default Value | Description |
|---|---|---|
| **Installation folder** | *C:\Program Files\Veeam\Backup and Replication* | Folder where Veeam Backup & Replication will be installed. |
| **vPower cache folder** | *C:\ProgramData\Veeam\Backup\IRCache* | The `IRCache` folder on a volume with the maximum amount of free space. The IR cache folder stores the write cache for machines that are started from backups during recovery verification or restore operations. Make sure that you have at least 10 GB of free disk space to store the write cache. |
| **Guest catalog folder** | *C:\VBRCatalog* | The `VBRCatalog` folder on a volume with the maximum amount of free space.<br>The guest catalog folder stores indexing data for VM guest OS files. Indexing data is required for browsing and searching for VM guest OS files inside backups and performing 1-click restore. |
| **Service account** | *LOCAL SYSTEM* | Account under which the Veeam Backup Service runs. |
| **Database engine** | *PostgreSQL* | The setup wizard installs PostgreSQL as a database engine locally on the Veeam Backup & Replication server. |
| **SQL server** | *<host_name>:5432* | The local host name and port number to be used by SQL server. |
| **Database name** | *VeeamBackup* | The setup deploys the Veeam Backup & Replication configuration database on the locally installed instance of PostgreSQL. |

| Setting | Default Value | Description |
|---|---|---|
| **Catalog service port** | *9393* | The catalog service port is used by the Veeam Guest Catalog Service to replicate catalog data from backup servers to Veeam Backup Enterprise Manager. |
| **Service port** | *9392* | The service port is used by Veeam Backup Enterprise Manager to collect data from backup servers. In addition to it, the Veeam Backup & Replication console uses this service port to connect to the backup server. |
| **Secure connections port** | *9401* | The secure connections port is used by the mount server to communicate with the backup server. |
| **REST API service port** | *9419* | This service port is used to access the Veeam Backup & Replication REST API. |
| **Check for updates** | *Automatically* | Veeam Backup & Replication will check for product updates weekly. When a new product build is published on the Veeam update server, a notification is displayed in the Windows Action Center. |

The following Veeam services and components are also deployed when installing Veeam Backup & Replication. They have predefined installation locations that cannot be changed during the Veeam Backup & Replication installation:

| Veeam Component | Default Installation Path |
|---|---|
| **AWS Plug-in for Veeam Backup & Replication** | *%ProgramFiles%\Veeam\Plugins\AWS* |
| **Google Cloud Plug-in for Veeam Backup & Replication** | *%ProgramFiles%\Veeam\Plugins\GCP* |
| **Veeam Kasten Plug-in for Veeam Backup & Replication** | *%ProgramFiles%\Veeam\Plugins\Kasten* |
| **Microsoft Azure Plug-in for Veeam Backup & Replication** | *%ProgramFiles%\Veeam\Plugins\Microsoft Azure* |
| **Nutanix AHV Plug-in for Veeam Backup & Replication** | *%ProgramFiles%\Veeam\Plugins\Nutanix AHV* |

| Veeam Component | Default Installation Path |
| --- | --- |
| oVirt KVM Plug-in for Veeam Backup & Replication | *%ProgramFiles%\Veeam\Plugins\RHV* |
| Proxmox Virtual Environment Plug-in for Veeam Backup & Replication | *%ProgramFiles%\Veeam\Plugins\PVE* |
| Veeam Agent for Linux Redistributable | *%ProgramData%\Veeam\Agents* |
| Veeam Agent for Mac Redistributable | *%ProgramData%\Veeam\Agents* |
| Veeam Agent for Microsoft Windows Redistributable | *%ProgramData%\Veeam\Agents* |
| Veeam Agent for Unix Redistributable | *%ProgramData%\Veeam\Agents* |
| Veeam Backup Transport | *%ProgramFiles(x86)%\Veeam\Backup Transport* |
| Veeam Backup vPowerNFS | *%ProgramFiles(x86)%\Veeam\vPowerNFS* |
| Veeam Backup VSS Integration | *%ProgramFiles%\Veeam\Backup File System VSS Integration* |
| Veeam Distribution Service | *%ProgramFiles%\Veeam\Veeam Distribution Service* |
| Veeam Installer Service | *%WinDir%\Veeam\Backup* |
| Veeam Mount Server | *%ProgramFiles%\Common Files\Veeam\Backup and Replication\Mount Service* |
| Veeam Plug-ins for Enterprise Applications Redistributable | *%ProgramData%\Veeam\Plugins* |
| VMware VDDK | *%ProgramFiles(x86)%\Veeam\Backup Transport* |

For more information, see Veeam Backup & Replication Services.

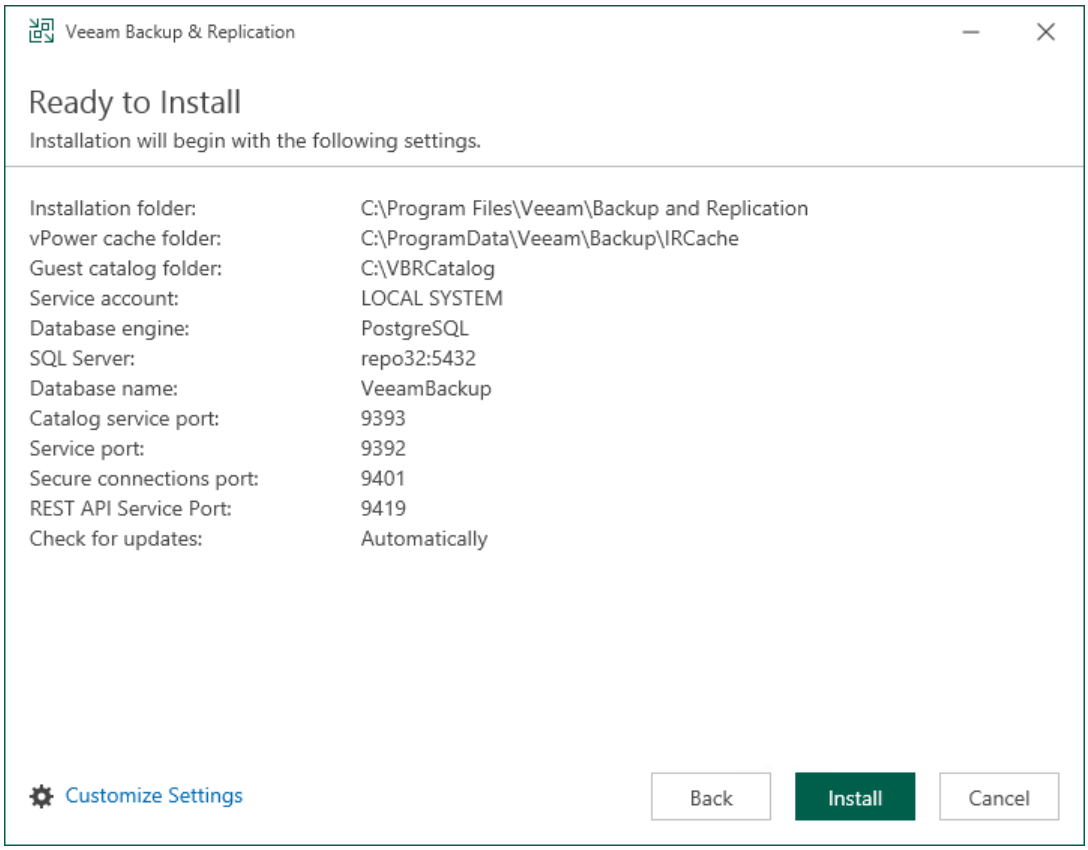

# Step 7. Specify Service Account Settings

The **Service Account** step of the wizard is available if you have selected to configure installation settings manually.

You can select an account under which you want to run the Veeam Backup Service:

- LOCAL SYSTEM account (recommended, used by default)
- Custom user account

The user name of the custom account must be specified in the *DOMAIN|USERNAME* format and have the following rights and permissions:

- The account must be a member of the *Administrators* group on the machine where Veeam Backup & Replication is installed.
- The account must have *db_owner* rights for the configuration database.

Veeam Backup & Replication automatically grants the *Log on as service* right to the specified user account.

> **NOTE**
>
> You cannot use a gMSA for running the Veeam Backup Service.

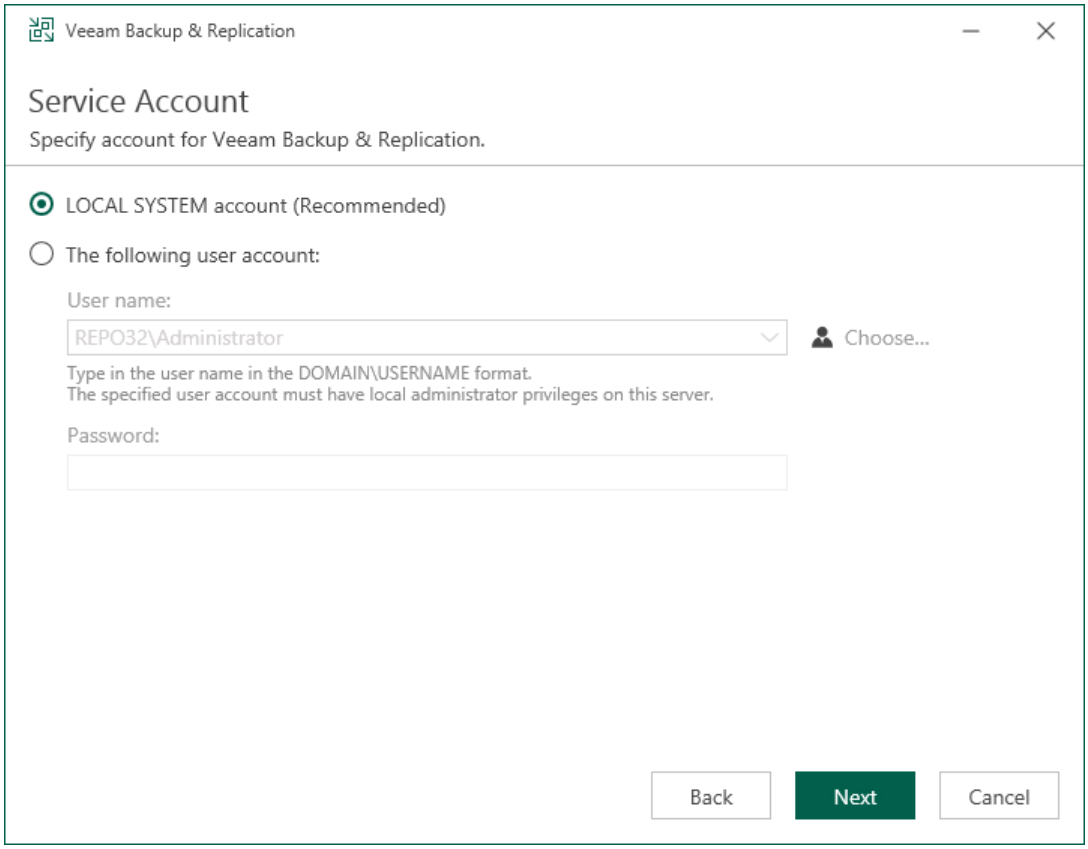

# Step 8. Specify Database Engine and Instance

The **Database** step of the wizard is available if you have selected to configure installation settings manually.

At this step of the wizard, you can choose a database engine (PostgreSQL or Microsoft SQL Server) for the Veeam Backup & Replication configuration database, specify a new or existing instance where you want to deploy the configuration database, and specify the authentication mode.

> **NOTE**
>
> Consider limitations and considerations in Before You Begin.

1. Select one of the following database engines that you want to use for the configuration database:
   - PostgreSQL
   - Microsoft SQL Server
2. Specify instance settings:
   - [For PostgreSQL] You can use an already installed PostgreSQL instance or install a new one.
     - To install a new PostgreSQL instance, select the **Install new instance** option. The setup will install PostgreSQL on the Veeam Backup & Replication server and create a database with the *VeeamBackup* name.
     - To use an already installed PostgreSQL instance, select the **Use existing instance** option. Enter the instance name in the *HOSTNAME:PORT* format. In the **Database name** field, specify a name for the Veeam Backup & Replication configuration database.

     > **IMPORTANT**
     >
     > If you use the already installed PostgreSQL instance or make any changes in the machine hardware, perform the additional configuration of the PostgreSQL instance, as described in section Before You Begin.

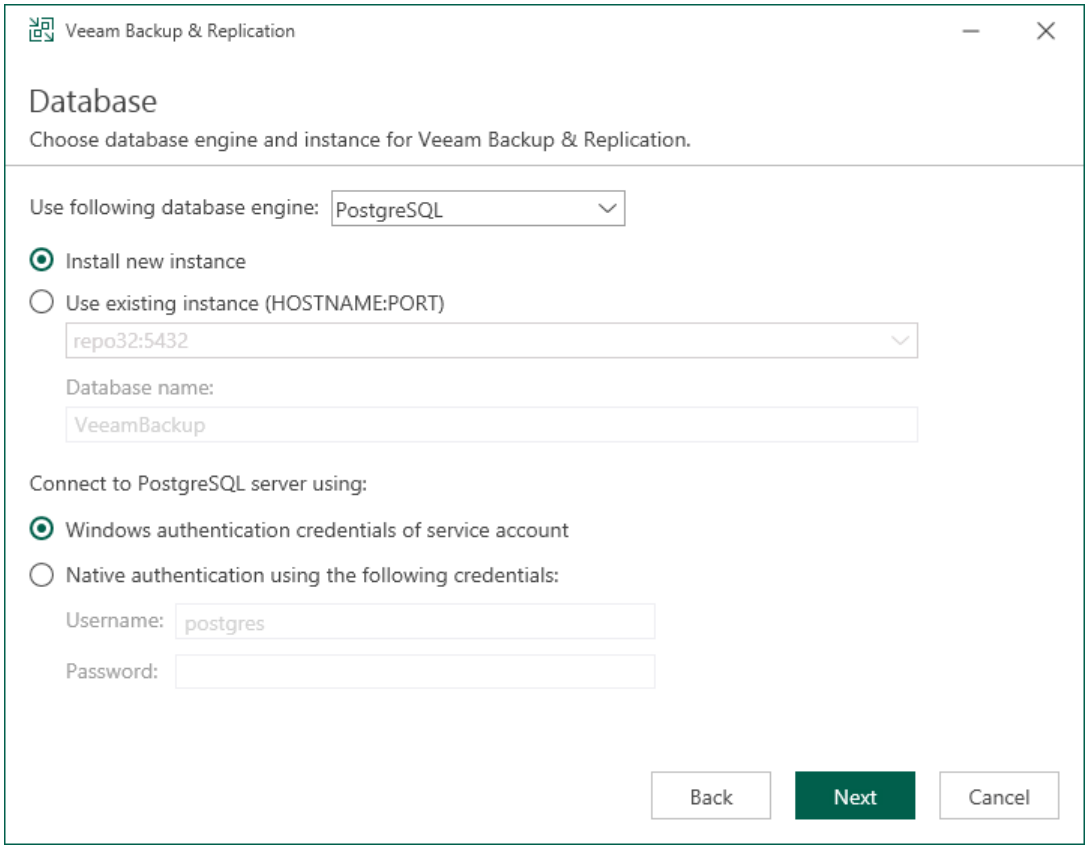


- [For Microsoft SQL Server] You can use an already installed Microsoft SQL Server database only.

  i. In the **SQL Server instance** field, enter the instance name in the *HOSTNAME\INSTANCE* format or select an instance from the drop-down list. You can also click **Browse** to choose a Microsoft SQL Server on a remote machine.

ii. In the **Database name** field, specify a name for the Veeam Backup & Replication configuration database.

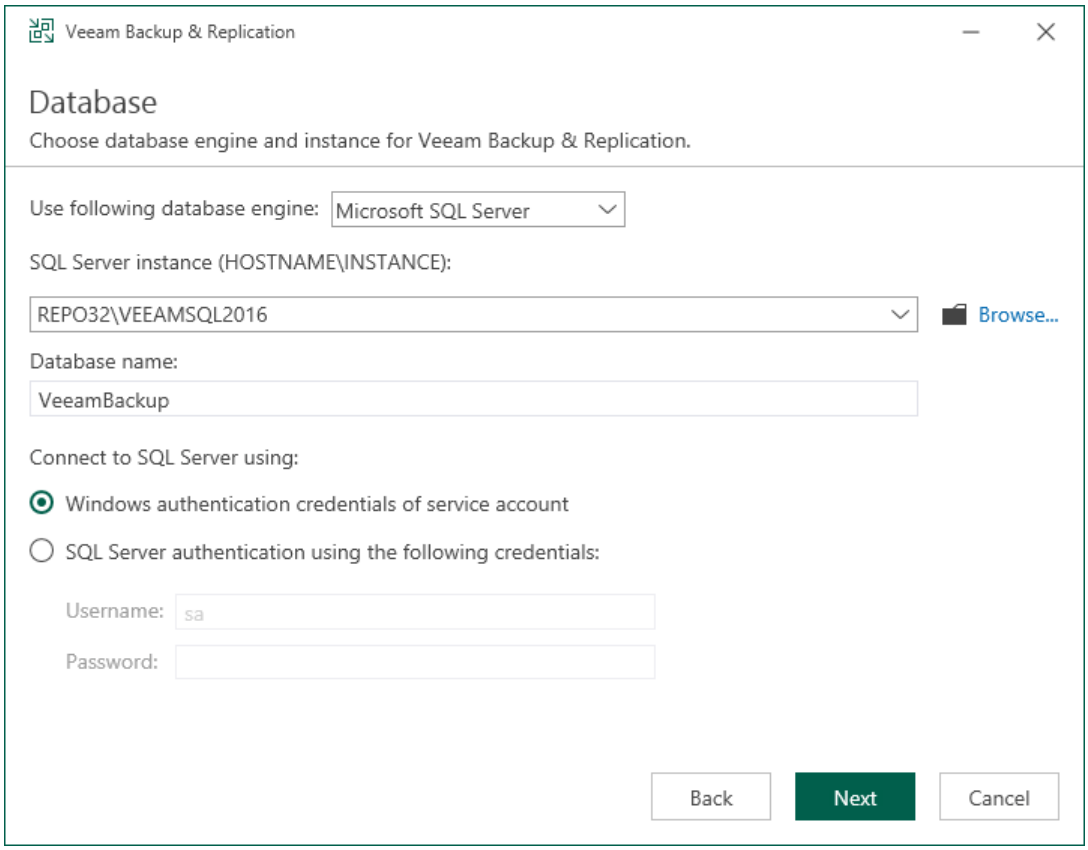


3. Select an authentication mode to connect to the database server instance: Microsoft Windows authentication or native database server authentication. If you select the native authentication, enter credentials of the database account.

If a configuration database with the specified name already exists (for example, it was created by a previous installation of Veeam Backup & Replication), the setup wizard will notify about it. To connect to the detected database, click **Yes**. If necessary, Veeam Backup & Replication will automatically upgrade the database to the latest version.

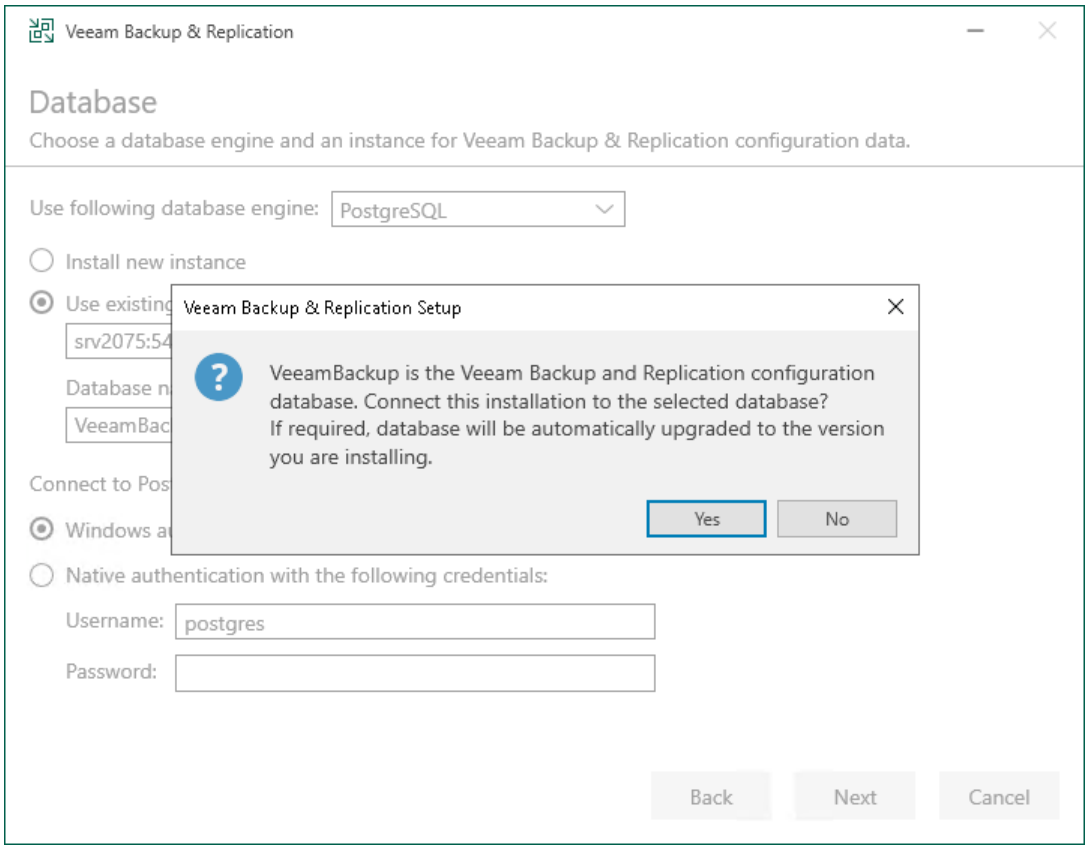

# Step 9. Perform Configuration Check

At the **Configuration Check** step of the wizard, the setup checks the Veeam Backup & Replication configuration.

If the check returns errors, solve their causes before continuing the installation. After you solve them, click **Retry** to check if there are any issues left.

If the check returns warning or information messages, you can continue the installation and address them later. However, we recommend that you closely investigate warning and information messages: if not properly addressed, their causes may lead to serious problems with further system operation.

To view the details of a certain message, point the cursor to the line with the message. The dialog box will display the detailed description.

To copy a list of detected issues with detailed descriptions for further investigation, click **Copy list to clipboard**.

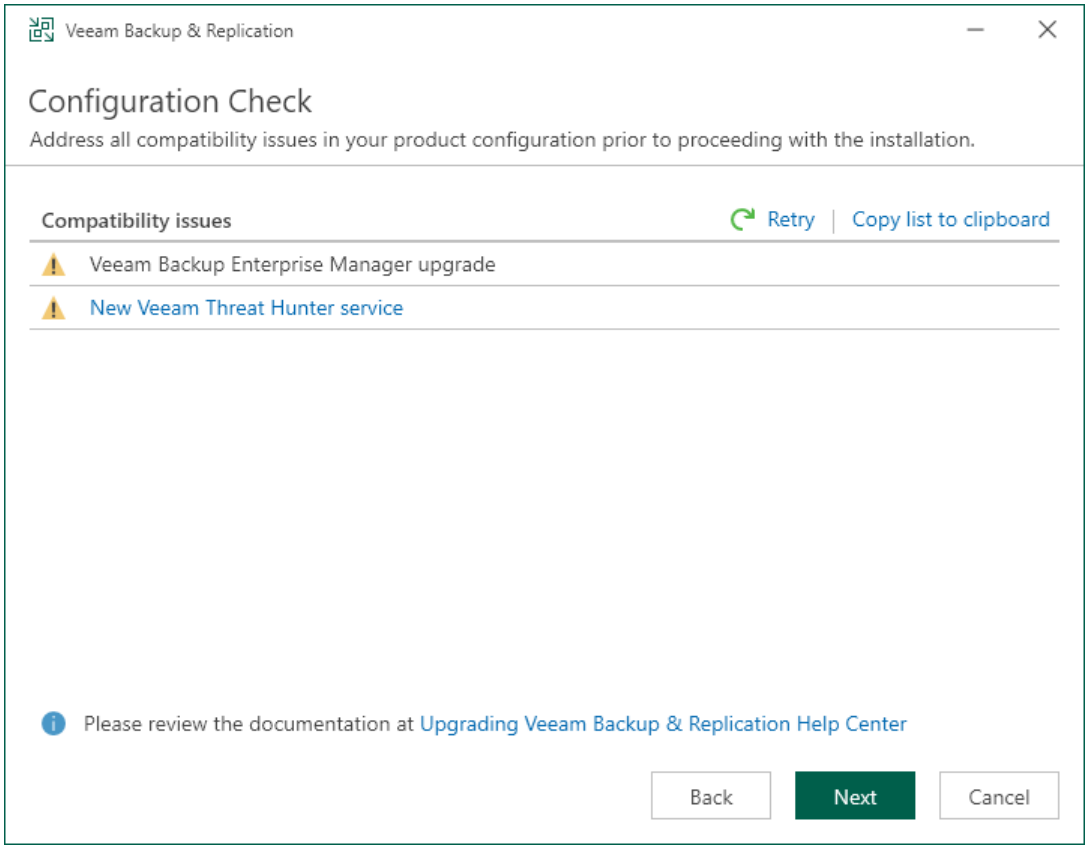

# Step 10. Specify Data Locations

The **Data Locations** step is available if you have selected to configure installation settings manually and to install a new instance of the database server.

At this step of the wizard, you can specify the installation folder and where the write cache and indexing data must be stored.

1. To change the default installation folder, click **Browse** next to the **Installation path** field.

   By default, the setup wizard uses the following installation folder: `C:\Program Files\Veeam\Backup and Replication`.

   Veeam Backup & Replication setup wizard calculates the space available on the selected disk and displays this information for your convenience.

2. To change the path to the folder where index files will be stored, click **Browse** next to the **Guest file system catalog** field.

   By default, the setup wizard creates the *VBRCatalog* folder on a volume with the maximum amount of free space, for example: `C:\VBRCatalog`.

3. [For VMware environments] The instant recovery cache folder stores the write cache for machines that are started from backups during recovery verification or restore operations. To change the path to the IR cache folder, click **Browse** next to the **Instant recovery write cache** field. Make sure that you have at least 100 GB of free disk space to store the write cache.

   By default, the setup wizard creates the IR cache folder on a volume with the maximum amount of free space, for example: `C:\ProgramData\Veeam\Backup\IRCache`.

   You do not need to configure this data location for Microsoft Hyper-V environments.

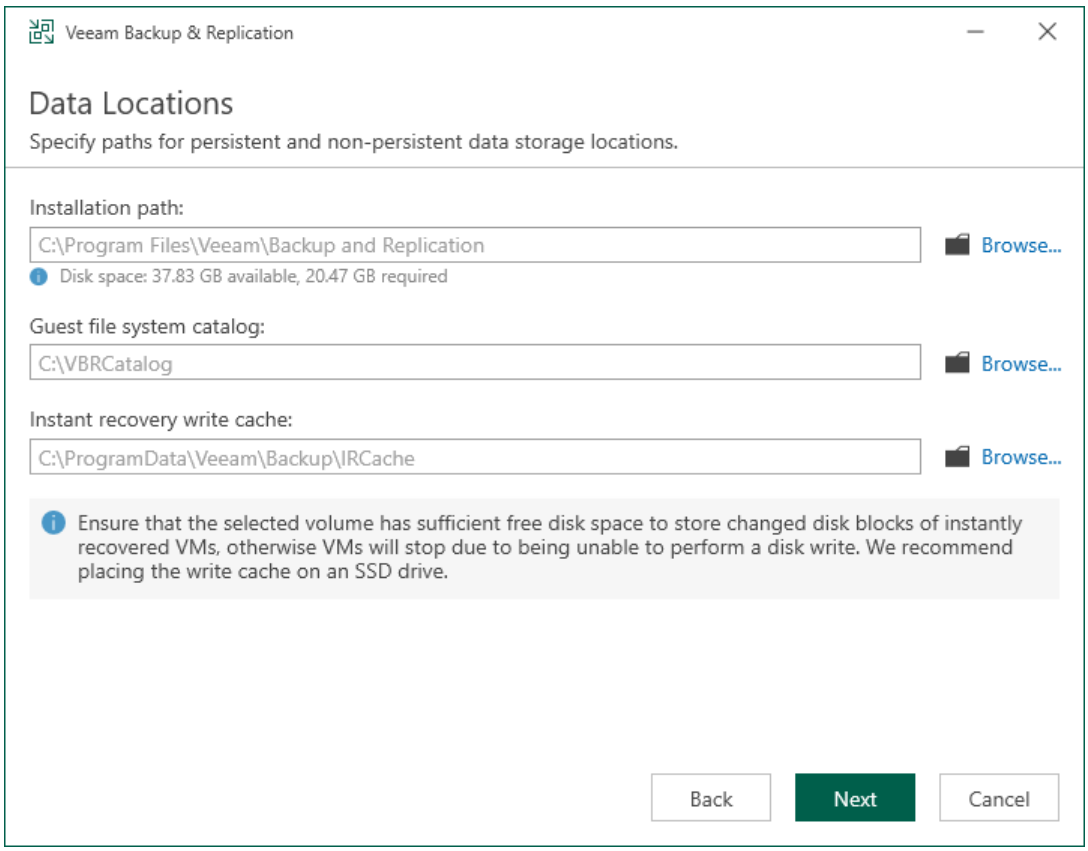

# Step 11. Specify Service Ports

The **Port Configuration** step of the wizard is available if you have selected to configure installation settings manually.

At this step of the wizard, you can customize port number values that will be used for communication between backup infrastructure components. For more information about Veeam Backup & Replication ports, see Ports.

- **Catalog service port**. This port is used by the Veeam Guest Catalog Service to replicate catalog data from backup servers to Veeam Backup Enterprise Manager. By default, port 9393 is used.
- **Veeam Backup service port**. This port is used by Veeam Backup Enterprise Manager to collect data from backup servers. In addition to it, the Veeam Backup & Replication console uses this service port to connect to the backup server. By default, port 9392 is used.
- **Secure connections port**. This port is used by the mount server to communicate with the backup server. By default, port 9401 is used.
- **REST API service port**. This port is used to communicate with the Veeam Backup & Replication REST API. By default, port 9419 is used.

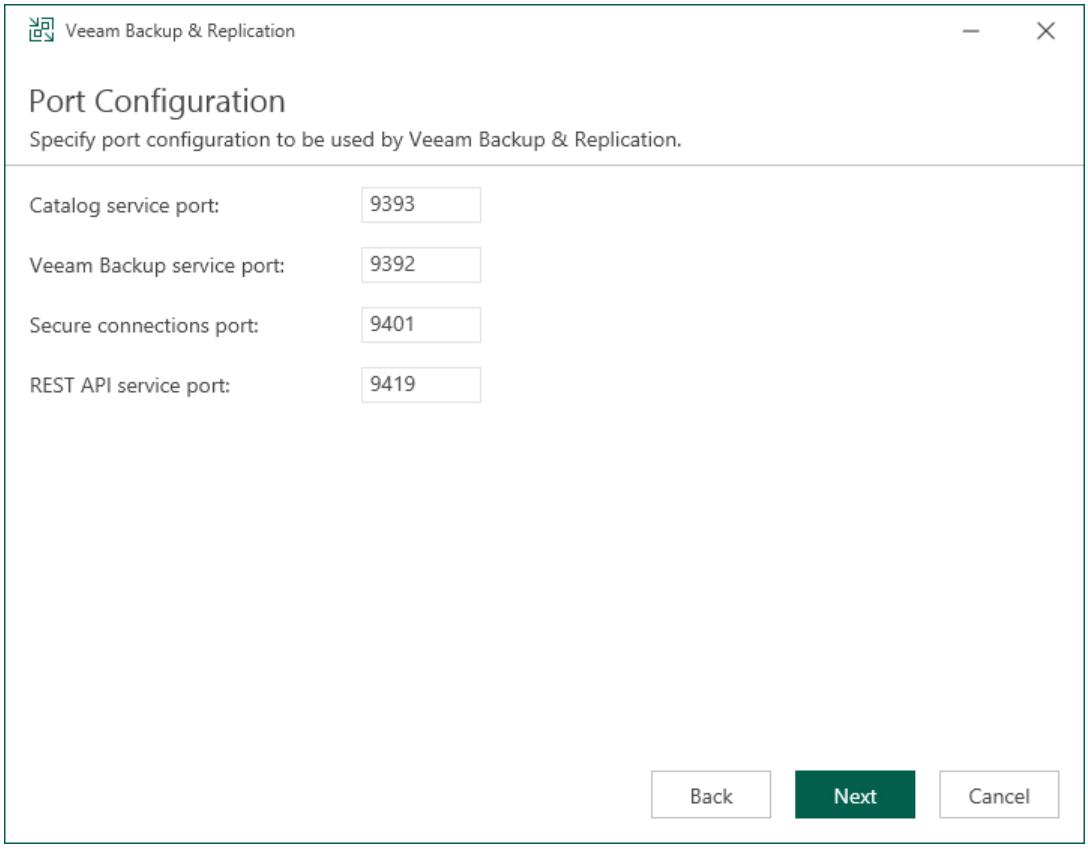

# Step 12. Begin Installation

The **Ready to Install** step of the wizard is available if you have selected to configure installation settings manually.

At this step of the wizard, you can review the Veeam Backup & Replication installation settings and start the installation process:

1. If you want Veeam Backup & Replication to check for product updates weekly, select the **Check for product updates automatically** check box. When a new product build is published on the Veeam update server, a notification will be displayed in the Windows Action Center.
2. Click **Install** to begin the installation.
3. Wait for the installation process to complete and click **Finish** to exit the setup wizard.

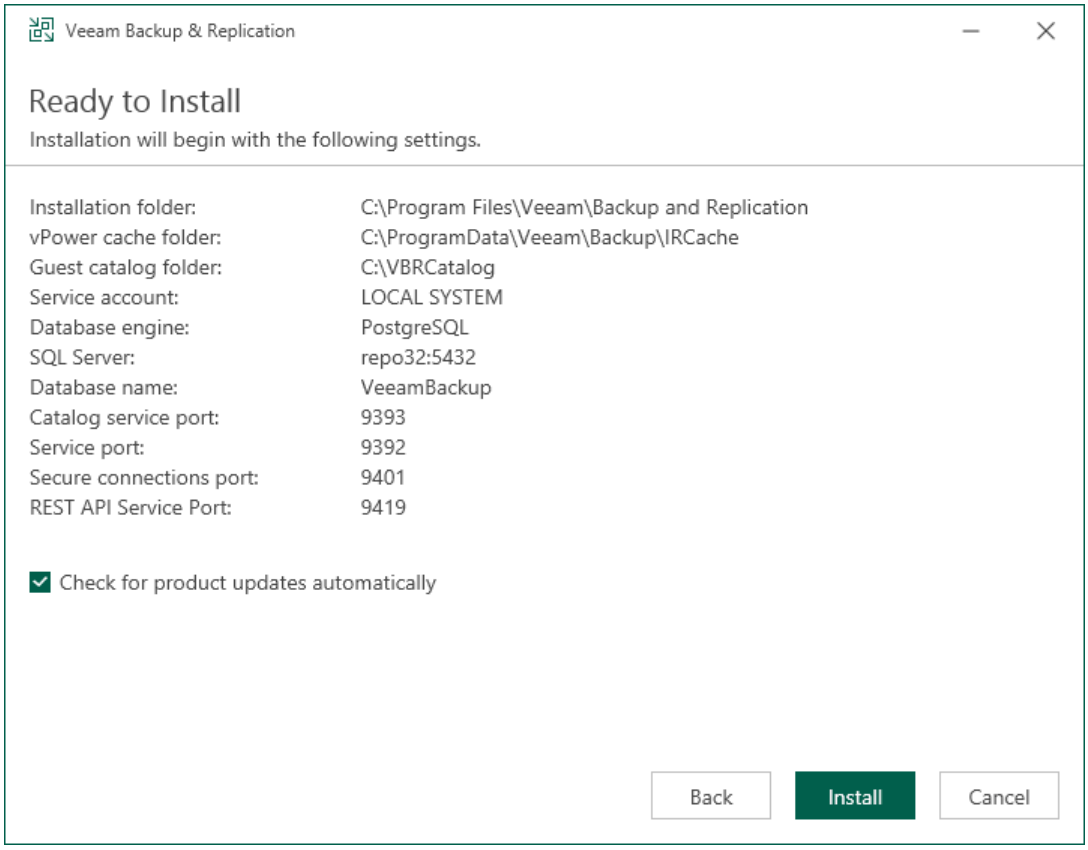


# Adjusting PostgreSQL Instance Configuration

If you selected to use an already installed PostgreSQL instance at the Specify Database Engine and Instance step of the wizard, make sure that the instance configuration is sufficient for the Veeam Backup & Replication performance.

To adjust the configuration of an existing PostgreSQL instance, take the following steps after you install Veeam Backup & Replication:

1. On a backup server, run the Set-VBRPSQLDatabaseServerLimits cmdlet. The cmdlet generates the necessary PostgreSQL configuration and saves it to a dump SQL file.

```
Set-VBRPSQLDatabaseServerLimits -OSType <String> -CPUCount <number of CPU
cores> -RamGb <RAM in GB> -DumpToFile <file path>
```

For example:

```
Set-VBRPSQLDatabaseServerLimits -OSType Windows -CPUCount 16 -RamGb 32 -Du
mpToFile "C:\config.sql"
```

2. On the machine with the PostgreSQL instance where you want to deploy the Veeam Backup & Replication configuration database, use the `psql` tool to apply the configuration from the dump file.

   The tool is located in the PostgreSQL installation folder.

```
psql -U <user> -f <file path>
```

   For example:

```
psql -U postgres -f "C:\config.sql"
```

   After you apply the configuration from the dump file, all changes will be written into the `postgressql.auto.conf` file located in the PostgreSQL installation folder. This file is loaded when the service starts and takes precedence over the default PostgreSQL configuration file.

3. Include the pg_stat_statements library to the PostgreSQL configuration. To add the library, you can manually edit the `shared_preload_libraries` option in the `postgresql.conf` file.

   Alternatively, you can do it by by executing the SQL code:

   a. Check the content of the `shared_preload_libraries` variable.

```
SELECT * FROM pg_settings
WHERE name = 'shared_preload_libraries';
```

   b. Add the `pg_stat_statements` library to the shared preloaded libraries.

   - If the `shared_preload_libraries` value is empty, assign `pg_stat_statements` to the `shared_preload_libraries` variable.

```
ALTER SYSTEM SET shared_preload_libraries = pg_stat_statements;
```

   - If the `shared_preload_libraries` value is not empty, add `pg_stat_statements` to the current value separated by comma.

```
ALTER SYSTEM SET shared_preload_libraries = <existing libraries>, pg
_stat_statements;
```

4. Restart the PostgreSQL service for the new configuration to take effect.
5. Install the `pg_stat_statements` extension. The extension is used to analyze the PostgreSQL performance.

```
CREATE EXTENSION IF NOT EXISTS "pg_stat_statements";
```

# Installing Veeam Backup & Replication Console

By default, the Veeam Backup & Replication console is installed on the backup server automatically when you install Veeam Backup & Replication. You do not need to install the console manually.

However, in addition to the default console, you can install the Veeam Backup & Replication console on a dedicated machine to access the backup server remotely. You can install as many remote consoles as you need. For more information, see Backup & Replication Console.

Before you install the Veeam Backup & Replication console, check prerequisites. Then use the **Veeam Backup & Replication Console Setup** wizard to install the console.

## Before You Begin

Before you install the Veeam Backup & Replication console, consider the following:

- The Veeam Backup & Replication console must be of the same version as Veeam Backup & Replication installed on the backup server.
- A machine on which you plan to install the Veeam Backup & Replication console must meet the system requirements. For more information, see System Requirements.
- A user account that you plan to use for installation must have sufficient permissions. For more information, see Permissions.
- Backup infrastructure components communicate with each other over specific ports. These ports must be open. For more information, see Ports.
- We do not recommend installing the Veeam Backup & Replication console on the machine that is used in the role of a different backup infrastructure component, for example, as a backup repository. Different components have different services installed on the server. Upon upgrading the console to the new version, your backup infrastructure may get out-of-order.

# Step 1. Start Setup Wizard

To start the setup wizard:

1. Download the latest version of the Veeam Backup & Replication installation image from the Download Veeam Products page.
2. Mount the installation image to the machine on which you plan to install the Veeam Backup & Replication console or burn the image file to a flash drive or other removable storage device. If you plan to install the Veeam Backup & Replication console on a VM, use built-in tools of the virtualization management software to mount the installation image to the VM.

   To extract the content of the ISO, you can also use the latest versions of utilities that can properly extract data from ISOs of large size and can properly work with long file paths.
3. After you mount the image or insert the disk, Autorun will open a splash screen. If Autorun is not available or disabled, run the `Setup.exe` file from the image or disk.
4. In the splash screen, click **Install**.

> **IMPORTANT**
>
> It is strongly recommended that you install the Veeam Backup & Replication console using Autorun or the **Setup.exe** file. If you run other installation files from the ISO folders, you may miss some components that need to be installed, and the Veeam Backup & Replication console may not work as expected.

# Step 2. Select Component

At the **Select Veeam Backup & Replication Component** step of the wizard, select **Install Veeam Backup & Replication Console**.

To open Veeam Help Center from the setup wizard, click **View Documentation**.

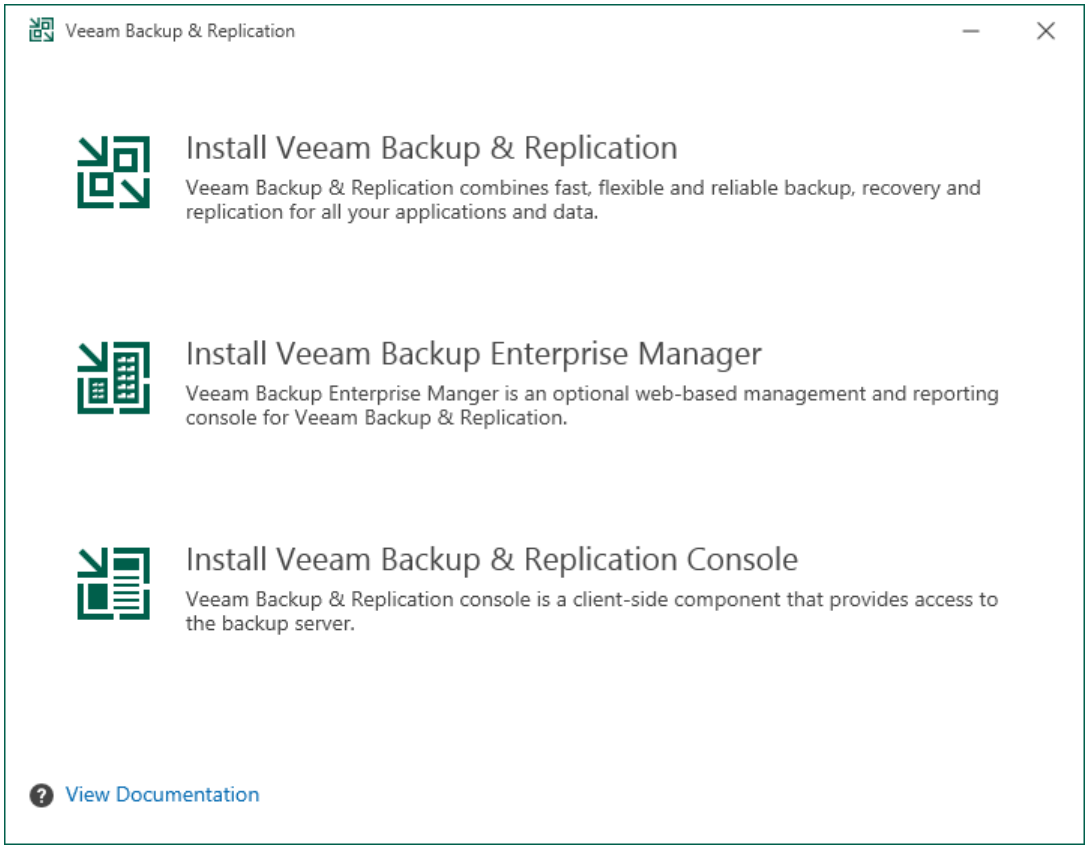

# Step 3. Read and Accept License Agreement

At the **License Agreement** step of the wizard, read Veeam License Agreement and licensing policy as well as license agreements of 3rd party components that Veeam incorporates and license agreements of required software. To accept the license agreements and continue installing the Veeam Backup & Replication console, click **I Accept**.

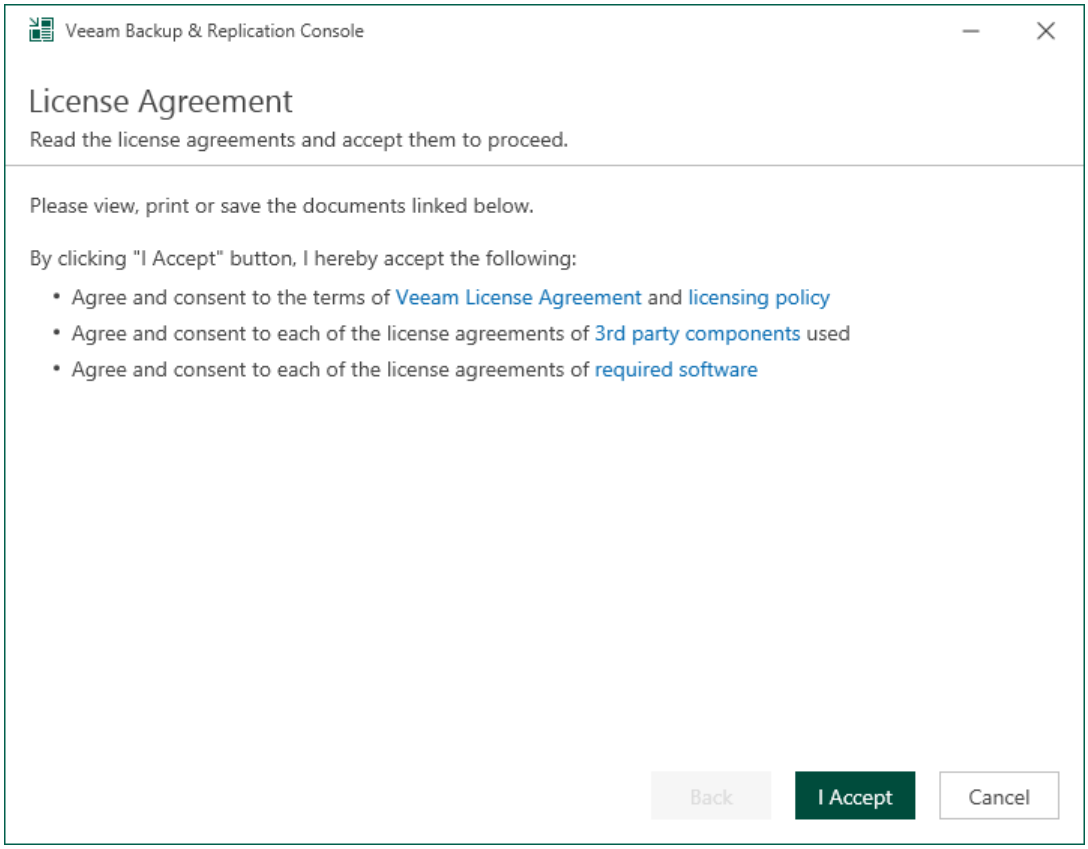

# Step 4. Install Missing Software

At the **System Configuration Check** step of the wizard, the setup wizard checks if the required software is installed on the machine. If some of the required components are missing, the setup will try to install them automatically. After the components are successfully installed, reboot is required. To reboot the machine, click **Reboot**.

If the setup wizard cannot install some of the required software components automatically, install them manually and click **Retry**.

> **NOTE**
>
> If all required software is already installed on the machine, the **System Configuration Check** step will be skipped. For more information on the necessary software, see System Requirements.

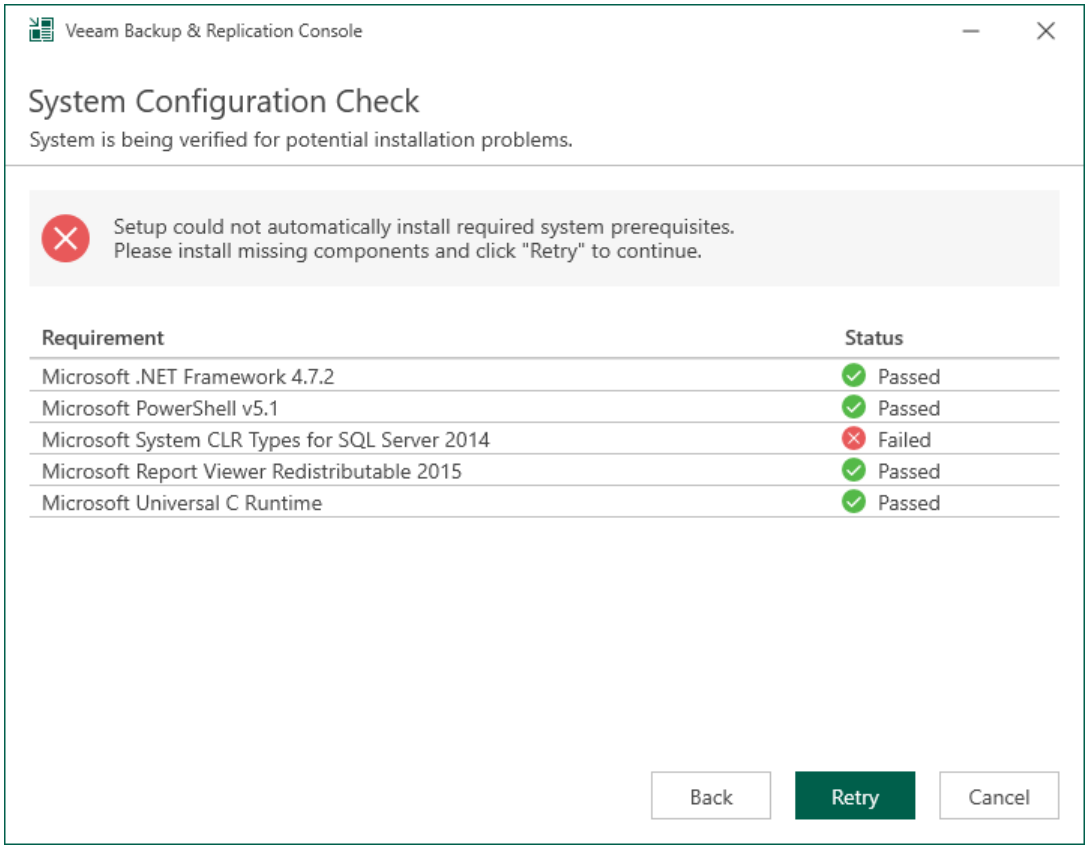

# Step 5. Review Default Installation Settings

At the **Ready to Install** step of the wizard, you can select to install the Veeam Backup & Replication console with default installation settings or specify custom installation settings.

- To use the default installation settings, click **Install**.
- To use custom installation settings, click **Customize Settings**. The setup wizard will include additional steps that will let you configure installation settings.

The following table lists the default installation settings.

| Setting | Default Value | Description |
|---|---|---|
| **Installation folder** | *C:\Program Files\Veeam\Backup and Replication* | Folder where the Veeam Backup & Replication console will be installed. |

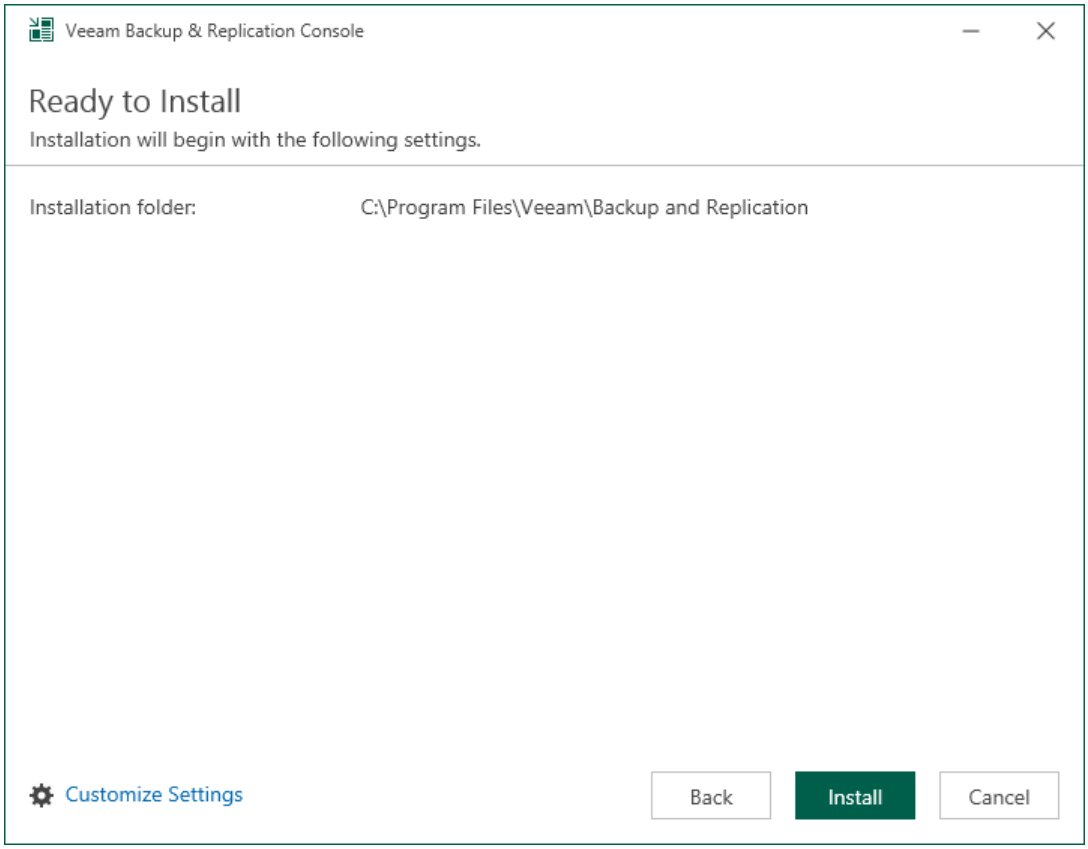

# Step 6. Specify Installation Path

The **Data Locations** step is available if you have selected to configure installation settings manually.

At this step of the wizard, you can choose the installation folder for the Veeam Backup & Replication console.

1. On the right of the **Installation path** field, click **Browse**.
2. In the **Select Folder** window, select the installation folder for the product. The default folder is `C:\Program Files\Veeam\Backup and Replication\`.

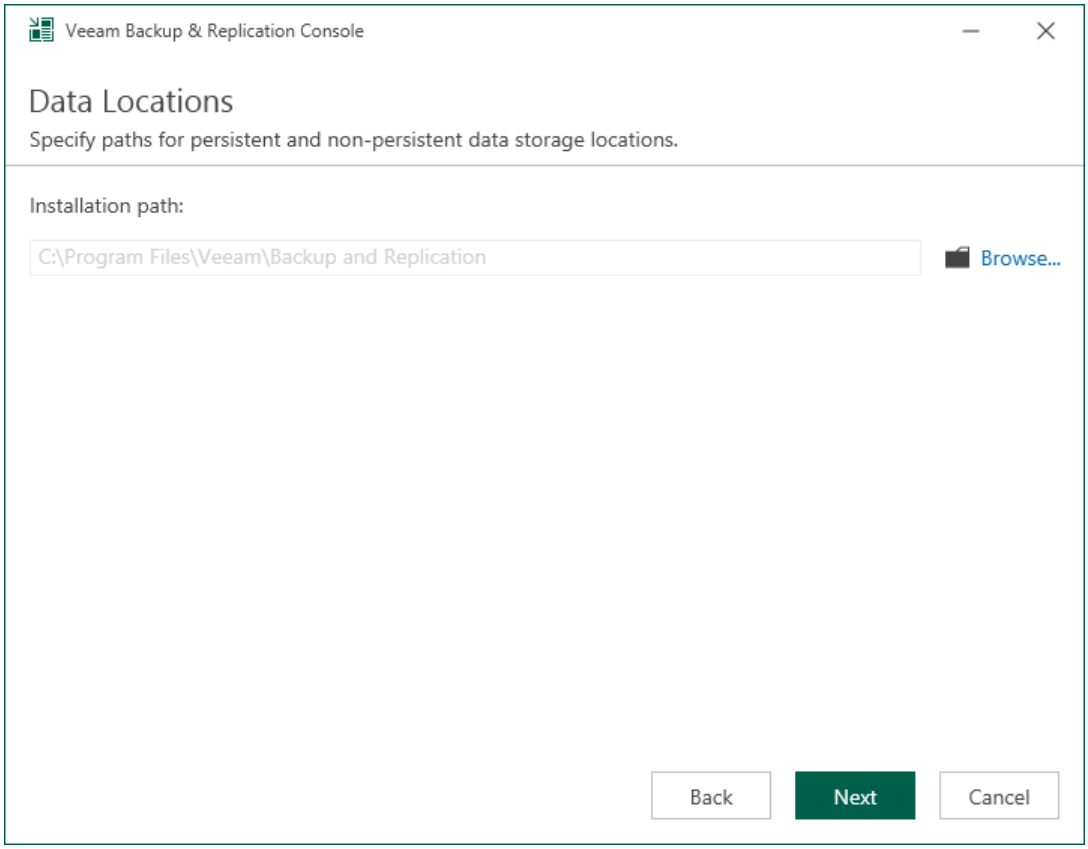

# Step 7. Begin Installation

The **Ready to Install** step of the wizard is available if you have selected to configure installation settings manually.

At this step of the wizard, you can review the Veeam Backup & Replication console installation settings and start the installation process

1. Click **Install** to begin the installation.
2. Wait for the installation process to complete and click **Finish** to exit the setup wizard.

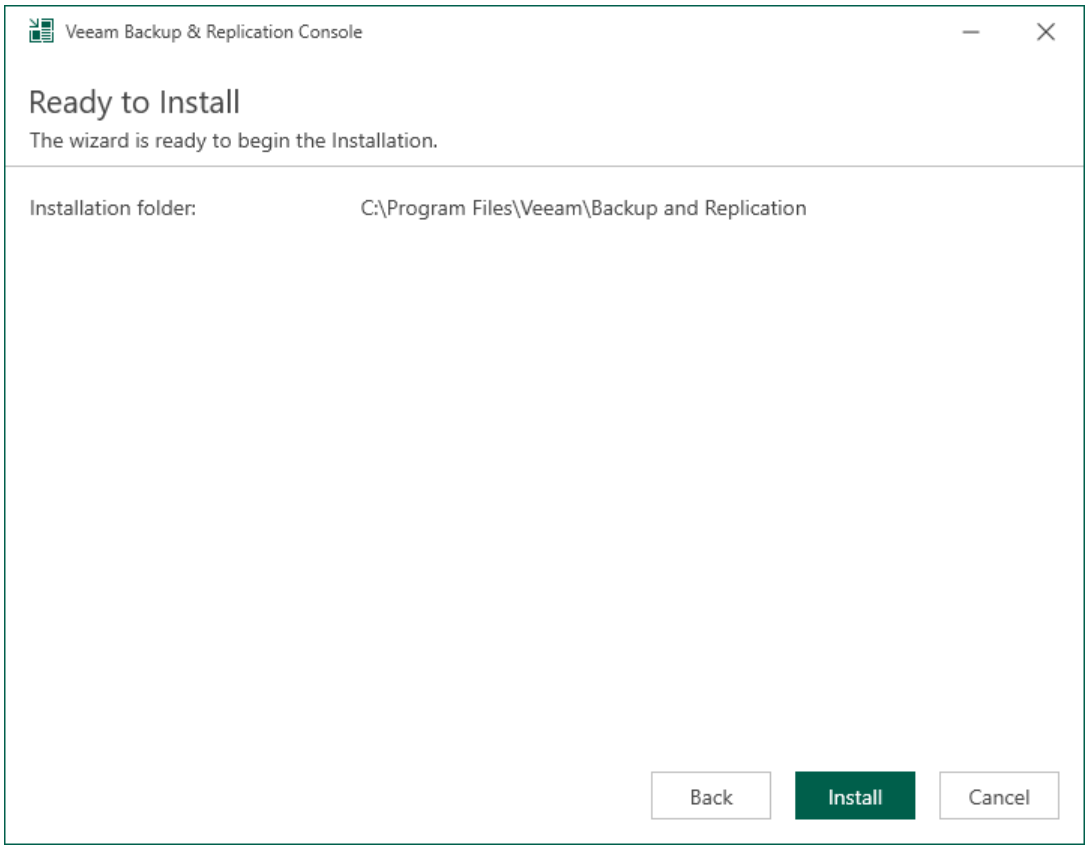

# Uninstalling Veeam Backup & Replication

To uninstall Veeam Backup & Replication:

1. From the **Start** menu, select **Control Panel > Programs and Features**.
2. In the programs list, right-click **Veeam Backup & Replication** and select **Uninstall**. If you have Veeam Backup Enterprise Manager installed on this machine, Veeam Backup & Replication will uninstall both components. Wait for the process to complete.
3. If the program list contains additional Veeam Backup & Replication components, right-click the remaining components and select **Uninstall**.

The Veeam Backup & Replication configuration database is not removed during the uninstall process. All configuration data stored in the database remains as well.

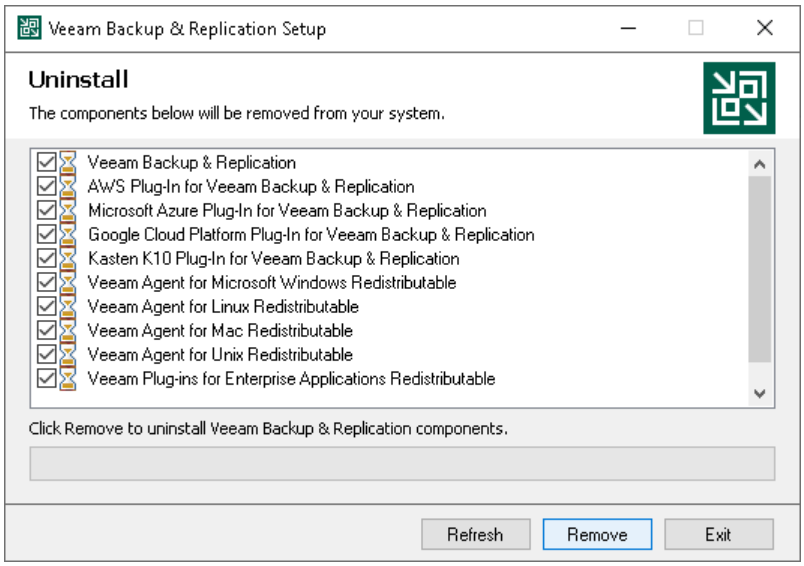

# Upgrade and Update

This section describes how to update and upgrade Veeam Backup & Replication, Veeam Backup & Replication Console, and other backup infrastructure components.

# Upgrading to Veeam Backup & Replication 12.3

To perform upgrade of Veeam Backup & Replication to version 12.3, you must be running version 11a (build 11.0.1.1261) or later on the supported operating system (refer to the System Requirements section of this document). For information on upgrade from earlier versions, see this Veeam KB article.

Before you upgrade Veeam Backup & Replication, check prerequisites. Then use the Veeam Backup & Replication upgrade wizard to install the product.

After you upgrade Veeam Backup & Replication, perform the finalizing steps.

## Upgrade Checklist

Use the following checklist to ensure your infrastructure is ready for the Veeam Backup & Replication upgrade. The built-in configuration check mechanism of the Veeam Backup & Replication Upgrade wizard performs some of the checks. Still, you can control them manually before starting the upgrade procedure.

## Licensing

1. Veeam Backup & Replication 12.3 uses the same license file format introduced with version 10, so you can use your existing version 10 or 11 license file to install version 12.3. Your support contract must be active as of the date when the product version you are installing was built.
2. Your support contract must be active as of the date when the product build you are installing was built. This is determined by the support expiration date in the installed license. If required, you can install a new license during the upgrade procedure.
3. Are you using Veeam Backup Starter? This edition has been discontinued, so Veeam Backup & Replication 12.3 will not accept such a license file. Download a replacement license file from the Customer Portal before upgrading.

## System Requirements

1. Check if the backup server to be upgraded is installed on the supported operating system version according to the System Requirements section. If it is not, create a configuration backup, install Veeam Backup & Replication 12.3 on the supported OS first, then restore the configuration backup created earlier. For information on how to perform the migration, see the Migrating Veeam Backup & Replication to Another Backup Server section.
2. Ensure that the backup server has sufficient disk space. The minimum disk space is calculated on-the-fly after the system configuration check during the upgrade procedure. It is based on the list of required packages to be installed on the machine and usually is about 9 GB. We recommend providing the minimum disk space calculated during the system configuration check, but not less than 41.5 GB of disk space: 3 x ISO size (36.5 GB) in the selected installation path (for example, `D:\VBR`) plus 5 GB for the database operations on the system volume (for example, `C:`).
3. Make sure that other servers that you plan to use as backup infrastructure components meet the system requirements listed in the System Requirements section of this document. In particular, ensure all backup infrastructure servers are based on 64-bit operating systems.

   Note that starting from Veeam Backup & Replication 12 (build 12.0.0.1420), Microsoft Windows Server Semi-Annual Channel versions are not supported. If you are upgrading from earlier Veeam Backup & Replication versions, make sure you do not have such servers in your backup infrastructure.

4. Make sure that the environment you are going to protect with Veeam Backup & Replication meets the requirements listed in the Supported Platforms and Applications section of this document. In particular:
   - Make sure that VMware ESXi and VMware vCenter server are upgraded to the minimum supported version 6.0 or remove these servers from the backup server configuration to continue.
   - Make sure that VMware Cloud Director is upgraded to the minimum supported version 10.1 or remove the hosts from the backup server configuration to continue.
5. Backup infrastructure components communicate with each other over specific ports. These ports must be open. For more information, see the Ports section.
6. Make sure that all necessary permissions are granted. For information on permissions, see the Permissions for VMware vSphere section.
7. Are you using a hardened repository? Consider the following:
   - Make sure that any Linux server associated with the hardened repository is configured to use a non-root account. For more information, see this Veeam KB article.
   - Veeam Backup & Replication does not support symlinks in the path to the hardened repository. If necessary, you can re-map backups to the new paths by using the steps described in the Switching from Linux Repository to Hardened Repository sub-section. If you have any problems, contact Veeam Customer Support.
   - Make sure that the Linux server used as a hardened repository does not have Veeam Agent for Linux installed. Update of hardened repository components will fail if Veeam Backup & Replication detects Veeam Agent for Linux - managed or standalone - on the server. For more information on uninstalling Veeam Agent for Linux, see Uninstalling Veeam Agent and Other Veeam Components if Veeam Agent for Linux is managed by Veeam Backup & Replication or Uninstalling Veeam Agent for Linux if Veeam Agent is operating in the standalone mode.
8. Do you have any jobs using the **Transform previous backup chains into rollbacks** option? This option has been removed from the product and such jobs are no longer supported. For more information, see this Veeam KB article.
9. Are you using file to tape jobs? Consider the following:
   - File to tape jobs have been re-engineered for scalability and are no longer free starting from Veeam Backup & Replication 12.0 (build 12.0.0.1402). You will have a grace period of 3 months following the upgrade to Veeam Backup & Replication 12 from earlier versions during which your existing jobs will not consume a license. For more information, see the Instance Consumption for Object Storage Backup, File Backup and File to Tape Jobs section. Note that backup to tape jobs do not consume licenses.
   - File to tape jobs can now process Distributed File Systems (DFS) data.
10. Are you using installations of Veeam Backup & Replication and Veeam Backup for Microsoft 365 on the same machine? First upgrade Veeam Backup for Microsoft 365, second upgrade Veeam Backup & Replication.
11. Azure compute accounts based on Azure AD user credentials (created with the **Use the existing account** option) are obsolete. You need to replace these accounts with new ones to restore workloads to Microsoft Azure, use the Microsoft Azure archive storage or Microsoft Azure Plug-in for Veeam Backup & Replication appliance.
12. Are you using integration with Veeam Backup for Microsoft Azure? If yes, after you upgrade to Veeam Backup & Replication 12.3 and replace the obsolete accounts from p.11, select the existing Microsoft Azure Compute account in the Manage Cloud Credentials, click Edit, and go through the Microsoft Azure Compute Account to update account permissions. Otherwise, you can face problems when adding an external repository with backups created by Veeam Backup for Microsoft Azure 6.0.

13. Are you using Server 2019 based ReFS backup repositories? If yes, avoid upgrading them to Server 2022 and mounting ReFS volumes from Server 2019 to new Server 2022 installations until you read this thread on Veeam R&D forums. Microsoft has addressed the known regression in the ReFS format upgrade code, and the fix is now publicly available.
14. Are you using Scale-Out Backup Repositories with immutable performance tier extents? Make sure that all extents have the same immutability settings.
15. Are you using a customized `AntivirusInfos.xml` file? During the upgrade, Veeam Backup & Replication will replace it with the default file. Make sure that you save your customized file at another path and after the upgrade make necessary changes to the default file.
16. Are you using Azure Local (formerly known as Azure Stack HCI)? You may need to apply an updated WDAC policy before upgrading if you have WDAC enabled on your Azure Local nodes. For more information, see this Veeam KB article.

# Integration with Veeam Management and Monitoring Products

1. Are you using **Veeam ONE** to monitor your backup infrastructure? If yes, upgrade it first. Veeam ONE supports monitoring of backup servers version 11a or later.
2. Are you using **Veeam Backup Enterprise Manager**? If yes, consider the following:
   - **Important!** Starting with Veeam Backup Enterprise Manager 12, a new port (port 9405) is used for certificate communication between Enterprise Manager and Veeam Backup & Replication. Ensure that your firewalls are configured to take into account this new port to avoid communication issues between Enterprise Manager and Veeam Backup & Replication.
   - Start the upgrade procedure with this component. Veeam Backup & Replication should be upgraded after that. If you have a backup server installed on the same machine, upgrade it immediately after completing upgrade of the Veeam Backup Enterprise Manager server. Otherwise, the Veeam Configuration Database Connection Utility (DBConfig) utility will not work properly for Veeam Backup & Replication.
   - From Veeam Backup Enterprise Manager, you cannot edit jobs that are managed by backup servers of earlier versions as well as Veeam Agent backup jobs, file share backup jobs, and backup copy jobs. To edit settings of such jobs, use the Veeam Backup & Replication console.
3. Are you using **Veeam Backup Enterprise Manager** server added to **Veeam ONE**? If yes, first upgrade Veeam ONE, second upgrade Veeam Backup Enterprise Manager, and third upgrade Veeam Backup & Replication.
4. Are you using **Cloud Connect**? If yes, consider the following:
   - Check with your Cloud Connect service provider if they have already upgraded their system to at least the version you are upgrading to.
   - Ensure your Cloud Connect tenants use the supported Veeam product versions. The minimal supported tenant versions are: Veeam Backup & Replication 12.0, Veeam Agent for Microsoft Windows 6.0, Veeam Agent for Linux 6.0, Veeam Agent for Mac 2.0.
5. Are you using **Veeam Recovery Orchestrator**? If yes, note that Veeam Disaster Recovery Orchestrator 6.0 is compatible with Veeam Backup & Replication 11a and 12, Veeam Recovery Orchestrator 7.0 — with Veeam Backup & Replication 12, 12.1 and 12.2, Veeam Recovery Orchestrator 7.1 — with Veeam Backup & Replication 12, 12.1, 12.2 and 12.3. If necessary, upgrade Veeam Recovery Orchestrator before upgrading to Veeam Backup & Replication 12.3.

# Integration with Veeam Backup for Public Clouds

Are you using Veeam Backup & Replication integrated with Veeam Backup for Public Cloud solutions? If yes, first upgrade Veeam Backup & Replication to version 12.3. Second upgrade plug-ins for Veeam Backup for AWS, Veeam Backup for Microsoft Azure and Veeam Backup for Google Cloud. Third upgrade connected appliances to the most recent version.

# Integration with Veeam Backup for Hypervisors

Are you using Veeam Backup & Replication integrated with Veeam Backup for Hypervisors?

- **Veeam Backup for Nutanix AHV**: plugin for this product is included into the Veeam Backup & Replication 12 package. During the upgrade to version 12.3, the plugin for this product will be automatically upgraded to the required version.
- **Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization**: plugin for this product is included into the Veeam Backup & Replication 12 package. During the upgrade to version 12.3, the plugin for this product will be automatically upgraded to the required version.
- **Veeam Backup for Proxmox Virtual Environment**: plugin for this product is included into the Veeam Backup & Replication 12 package. During the upgrade to version 12.3, the plugin for this product will be automatically upgraded to the required version.

# Integration with Veeam Backup Agents and Enterprise Plug-Ins

1. Are you using **Veeam Agents** managed through Veeam Backup & Replication?
    - If you use **Veeam Agent for Microsoft Windows** or **Veeam Agent for Linux** 4.0, they will stop working after upgrading to Veeam Backup & Replication 12.3. In this case, we recommend immediately upgrading Veeam Agent for Microsoft Windows or Veeam Agent for Linux to 6.3. If you use Veeam Agent for Microsoft Windows or Veeam Agent for Linux 5.0 or later, they will continue working after upgrading to Veeam Backup & Replication 12.3, but new features implemented in Veeam Backup & Replication 12.3 will not be supported. In this case, you can upgrade Veeam Agent for Microsoft Windows or Veeam Agent for Linux to 6.3 later if the support of new features is not critical for you.
    - If you use **Veeam Agent for Mac** 1.0, it will stop working after upgrading to Veeam Backup & Replication 12.3. In this case, we recommend immediately upgrading Veeam Agent for Mac to 2.3. If you use Veeam Agent for Mac 1.0.1, it will continue working after upgrading to Veeam Backup & Replication 12.3, but new features implemented in Veeam Backup & Replication 12.3 will not be supported. In this case, you can upgrade Veeam Agent for Mac to 2.3 later if the support of new features is not critical for you.
    - If you use **Veeam Agent for IBM AIX** or **Veeam Agent for Oracle Solaris** 3.0 or later, they will continue working after upgrading to Veeam Backup & Replication 12.3, but new features implemented in Veeam Backup & Replication 12.3 will not be supported. In this case, you can upgrade Veeam Agent for IBM AIX or Veeam Agent for Oracle Solaris to 4.6 later if the support of new features is not critical for you.

2. Are you using **Veeam Plug-in for Oracle RMAN**, **Veeam Plug-in for SAP HANA**, **Veeam Plug-in for SAP on Oracle**, **Veeam Plug-in for IBM Db2** or **Veeam Plug-in for Microsoft SQL Server**? If yes, you upgrade Veeam Backup & Replication first, then you can upgrade Veeam Plug-ins (Veeam Plug-in for Oracle RMAN, Veeam Plug-in for SAP HANA, Veeam Plug-in for SAP on Oracle, Veeam Plug-in for IBM Db2, Veeam Plug-in for Microsoft SQL Server).

   Note that Veeam Backup & Replication supports only two latest minor updates of Veeam Plug-ins. For example, Veeam Backup & Replication 12.3 supports Veeam Plug-ins of versions 12.3 and 12.2.

# Storage System Snapshot Integration

1. **IBM FlashSystem**: In Veeam Backup & Replication 12.3, the storage snapshot integration functionality was re-implemented as an independent plug-in. After the upgrade, download and install the IBM FlashSystem Plug-in for Veeam Backup & Replication from the Veeam Download page to continue using jobs utilizing storage snapshots on IBM FlashSystem arrays.
2. Make sure your storage systems work on a supported operating system:
   - **Huawei**: Storage arrays of this vendor are not supported for Veeam storage snapshot integration. Remove the storage arrays from the backup server configuration and use other backup transport modes instead.
   - **Cisco HyperFlex**: The minimum supported operating system version is v4.0(2x). Upgrade to it or remove the storage arrays from the backup server configuration.
   - **HPE 3PAR WSAPI**: The minimum supported WSAPI version is 1.5. Upgrade to it or remove the storage arrays from the backup server configuration.
   - **HPE Nimble**: The minimum supported operating system version is 5.0. Upgrade to it or remove the storage arrays from the backup server configuration.
   - **Dell Data Domain**: The supported operating system version is 7.3 to 8.1. Upgrade to it or the backup jobs pointed to this repository will fail to start.
   - **HPE StoreOnce**: The minimum supported operating system version is 3.18.18 for Gen3 and 4.2.3 for Gen4. Upgrade to it or the backup jobs pointed to this repository will fail to start.
3. **IBM HyperSwap**: If you upgrade Veeam Backup & Replication to version 12.3 from the previous versions, select the secondary destination for IBM HyperSwap configurations explicitly in the backup job.

# Other Changes

1. Veeam Backup & Replication 12.3 supports protecting unstructured data backups with backup to tape jobs. If you have entire repositories added as sources for backup to tape jobs, make sure these repositories contain only backups that you want to protect with backup to tape jobs.
2. If you use persistent agents for guest OS processing in a Kerberos-only environment, perform the steps listed in this Veeam KB article.
3. If you use Kerberos authentication for Guest OS processing with persistent guest agent components, create several Service Principal Names in Active Directory before upgrading to Veeam Backup & Replication 12.3. For more information, see this Veeam KB article.
4. If you have HPE StoreOnce backup copy jobs where the source backup repository has immutability enabled and the target backup repository has immutability disabled, backups copied by backup copy jobs will not be immutable. For immutability to work, make sure both HPE StoreOnce repositories have immutability enabled.

5. Note that background retention is now applied to daily backups belonging to disabled backup jobs, as well as to orphaned backups. Background retention is still applied to GFS backups belonging to disabled backup jobs, as well as to orphaned backups.

# Upgrade Process

1. Make sure the latest run for all existing jobs has completed successfully. Rerun the failed jobs.
2. Ensure there are no running jobs, restore sessions, Instant Recovery sessions, and SureBackup jobs. We recommend that you do not stop running jobs and let them complete successfully.
3. Disable any periodic and backup copy jobs temporarily to prevent them from starting during the upgrade.
4. Disable CDP policies. Otherwise the CDP filter will not be upgraded.
5. Ensure there are no active tasks from standalone (unmounted) agents.
6. Ensure there are no active Veeam Recovery Orchestrator tasks.
7. Perform the configuration backup, as described in the Running Configuration Backups Manually section.

   Ensure you have configuration backup encryption enabled, otherwise stored credentials will not be included in it. For more information, see the Creating Encrypted Configuration Backups section.

# Step 1. Start Upgrade Wizard

To start the upgrade wizard, take the following steps:

1. Download the latest version of the Veeam Backup & Replication installation image from the Veeam Product Downloads page.
2. Mount the installation image to the machine where Veeam Backup & Replication is installed, or burn the image file to a flash drive or other removable storage device. If you plan to upgrade Veeam Backup & Replication on a VM, use built-in tools of the virtualization management software to mount the image to the VM.

   To extract the content of the ISO, you can also use the latest versions of utilities that can properly extract data from ISO files of large size and can properly work with long file paths.
3. After you mount the image or insert the disk, Autorun opens a splash screen. If Autorun is not available or disabled, run the `Setup.exe` file from the image or disk.
4. Click **Install**.

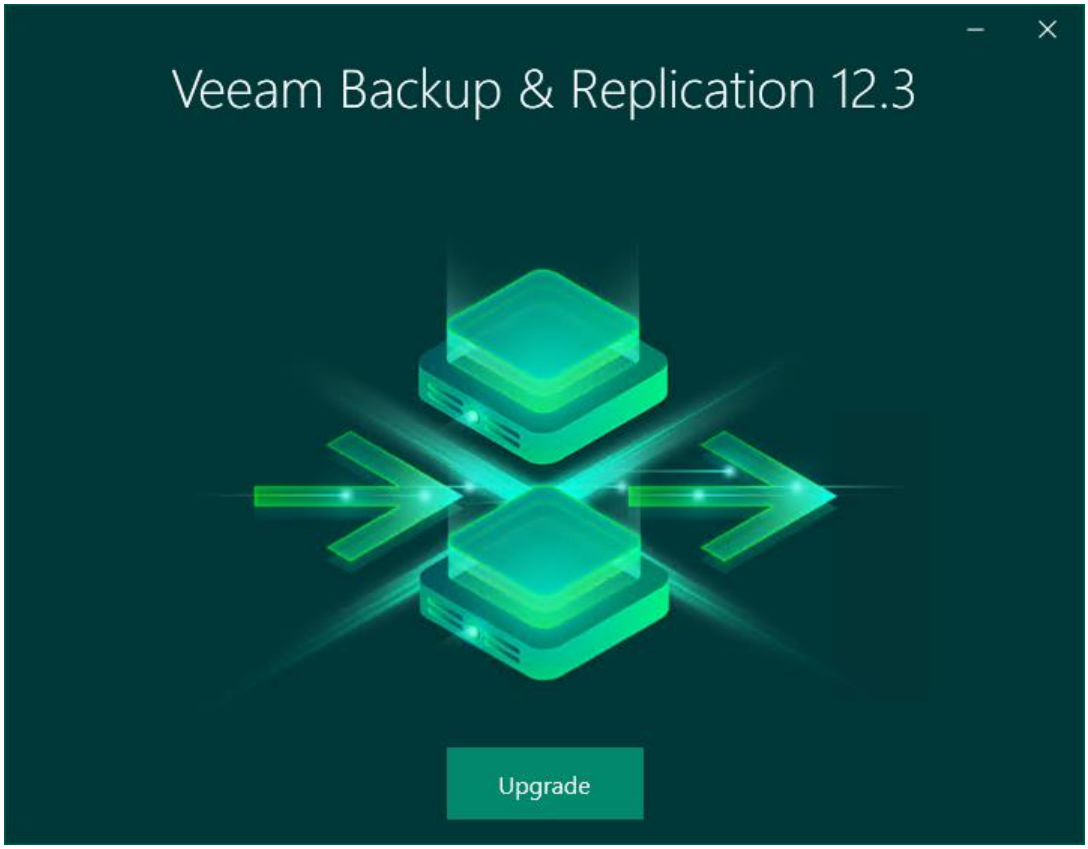

# Step 2. Select Component

At this step of the wizard, select **Upgrade Veeam Backup & Replication**.

To open Veeam Help Center from the upgrade wizard, click **View Documentation**.

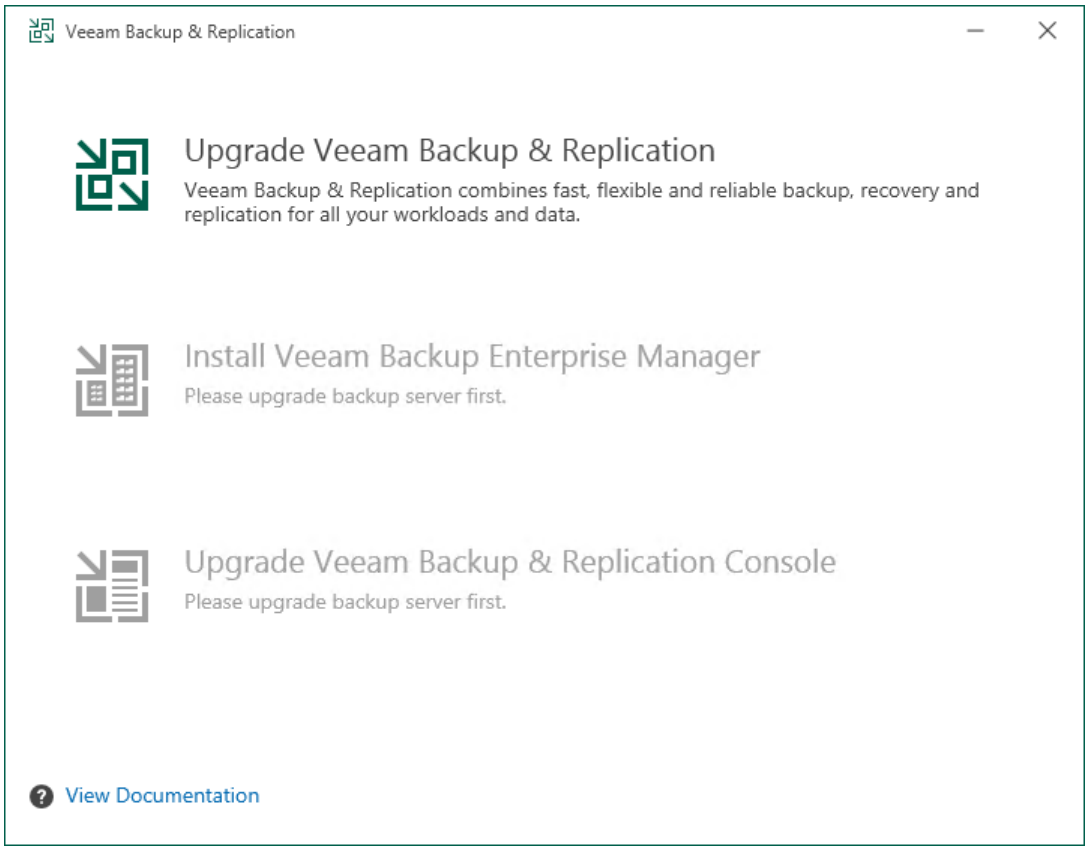

# Step 3. Read and Accept License Agreement

At the **License Agreement** step of the wizard, read Veeam License Agreement and licensing policy as well as license agreements of 3rd party components that Veeam incorporates and license agreements of required software. To accept the license agreements and continue installing Veeam Backup & Replication, click **I Accept**.

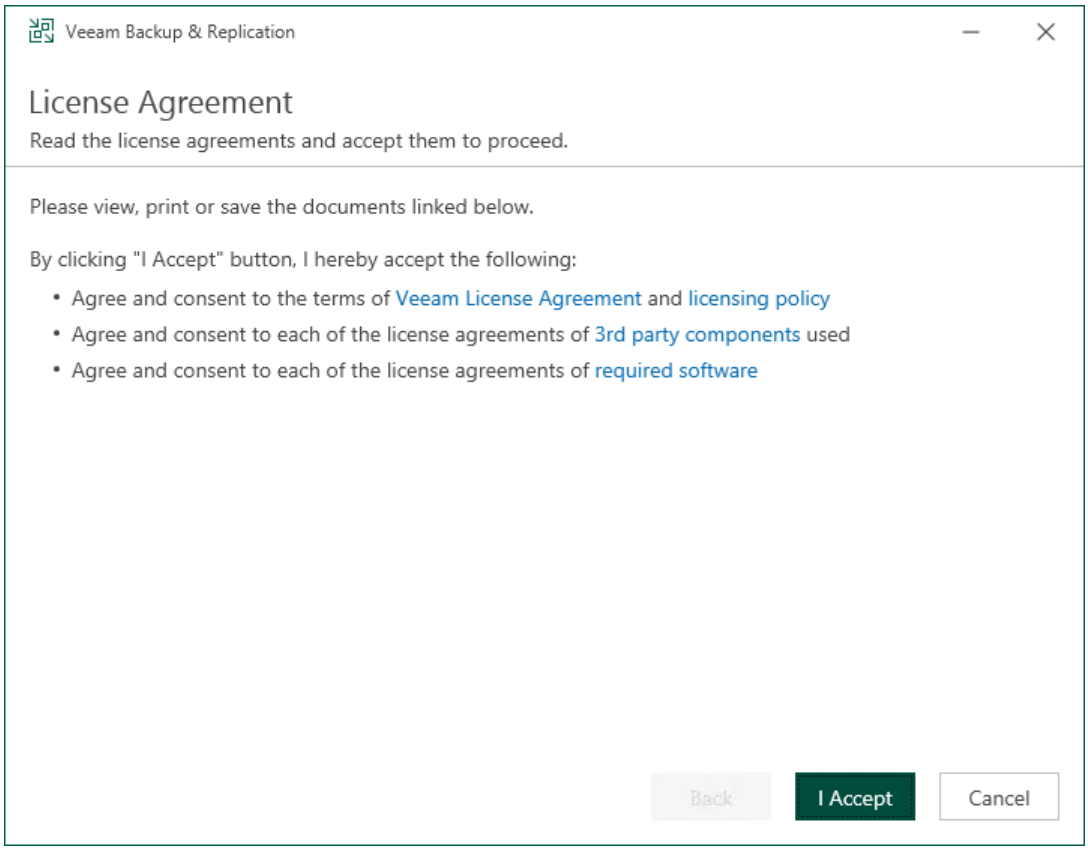

# Step 4. Provide License File

At the **License** step of the wizard, specify what license you want to install for Veeam Backup & Replication. You can leave the license file used in the previous version of Veeam Backup & Replication or install a new one.

To install a license, you have 2 options to choose from:

- Browse your local server or network locations for a license file:
  a. Click **Browse license file**.
  b. Select a valid license file for Veeam Backup & Replication.
- Select a license from your account at the Veeam website:
  a. Click **Sign in with Veeam**.
  b. Enter your credentials for accessing the Veeam website and click **Sign in**.
  c. Select one of the available licenses and click **Install selected license**.

To install new licenses automatically when you renew or expand your contract, select the **Update license automatically** check box. If you enable the automatic license update, and therefore enable usage reporting, you will double the number of workloads by which you can exceed your installed license. For more information, see Exceeding License Limit.

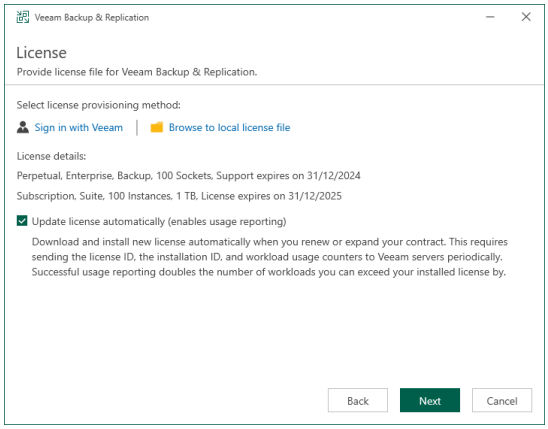

# Step 5. Install Missing Software

At the **System Configuration Check** step of the wizard, the setup wizard checks if the required software is installed on the machine. If some of the required components are missing, the setup will try to install them automatically. After the components are successfully installed, reboot is required. To reboot the machine, click **Reboot**.

If the setup wizard cannot install some of the required software components automatically, install them manually and click **Retry**.

> **NOTE**
>
> If all required software is already installed on the machine, the **System Configuration Check** step will be skipped. For more information on the necessary software, see System Requirements.

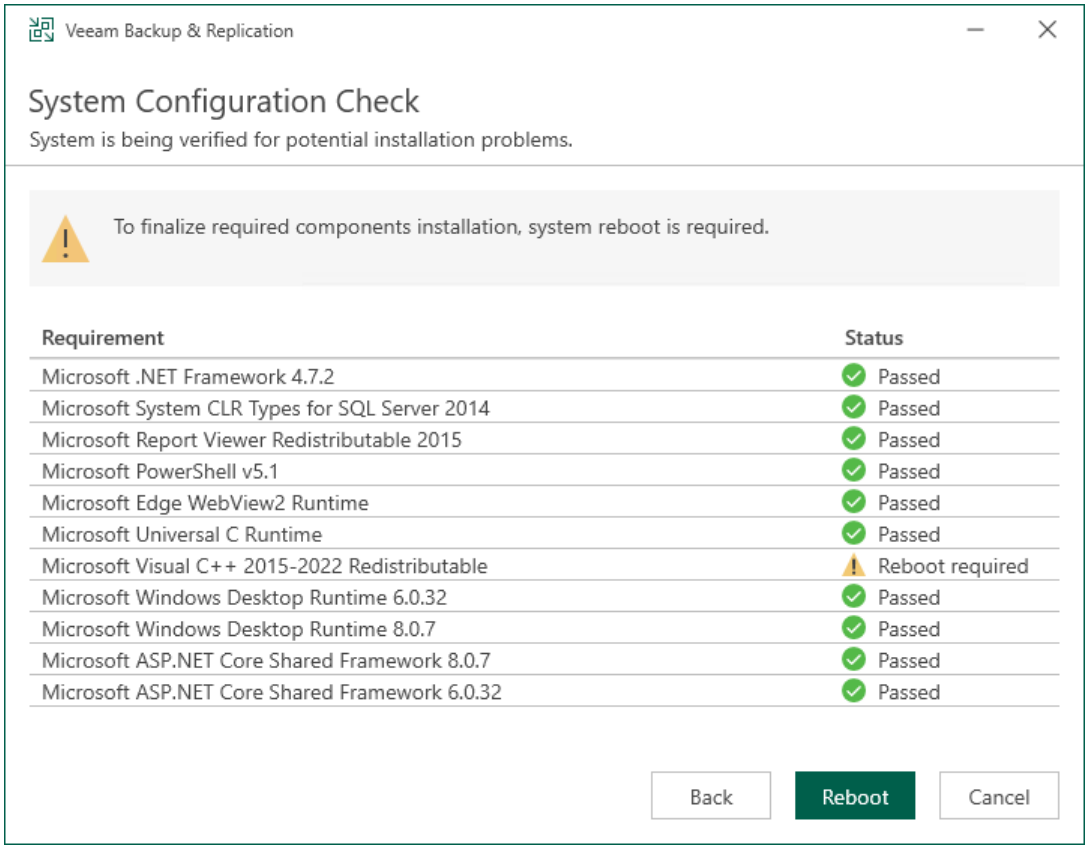

# Step 6. Perform Configuration Check

At the **Configuration Check** step of the wizard, the setup checks the Veeam Backup & Replication configuration.

If the check returns errors, solve their causes before continuing the upgrade. After you solve them, click **Retry** to check if there are any issues left.

If the check returns warning or information messages, you can continue the upgrade and address them later. However, we recommend that you closely investigate warning and information messages: if not properly addressed, their causes may lead to serious problems with further system operation.

To view the details of a certain message, point the cursor to the line with the message. The dialog box will display the detailed description.

To copy a list of detected issues with detailed descriptions for further investigation, click **Copy list to clipboard**.

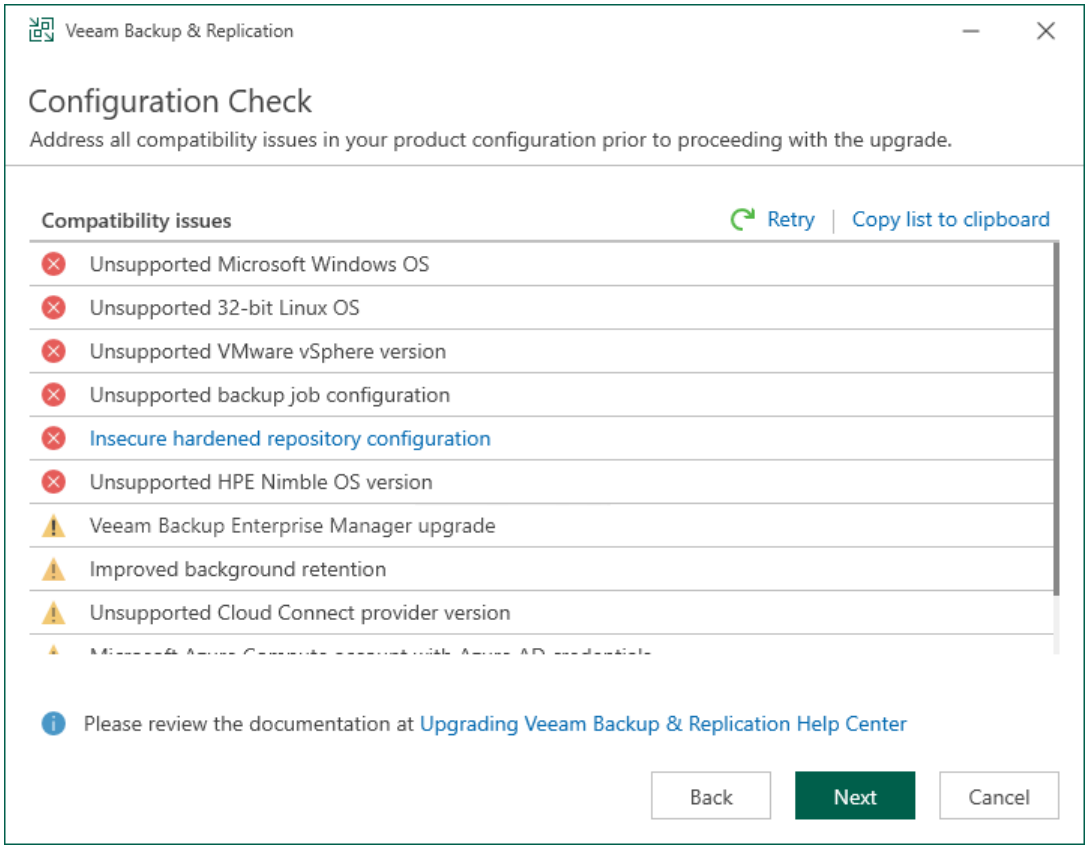

# Step 7. Review Components and Begin Upgrade

At the **Upgrade** step of the wizard, you can review the components that will be upgraded.

To upgrade the remote backup infrastructure components and required Veeam services after the Veeam Backup & Replication server is upgraded, select the **Update remote components automatically** check box. Otherwise, the backup server will prompt you to upgrade them during the first run of the backup server after the upgrade.

Click **Next** to begin the upgrade process.

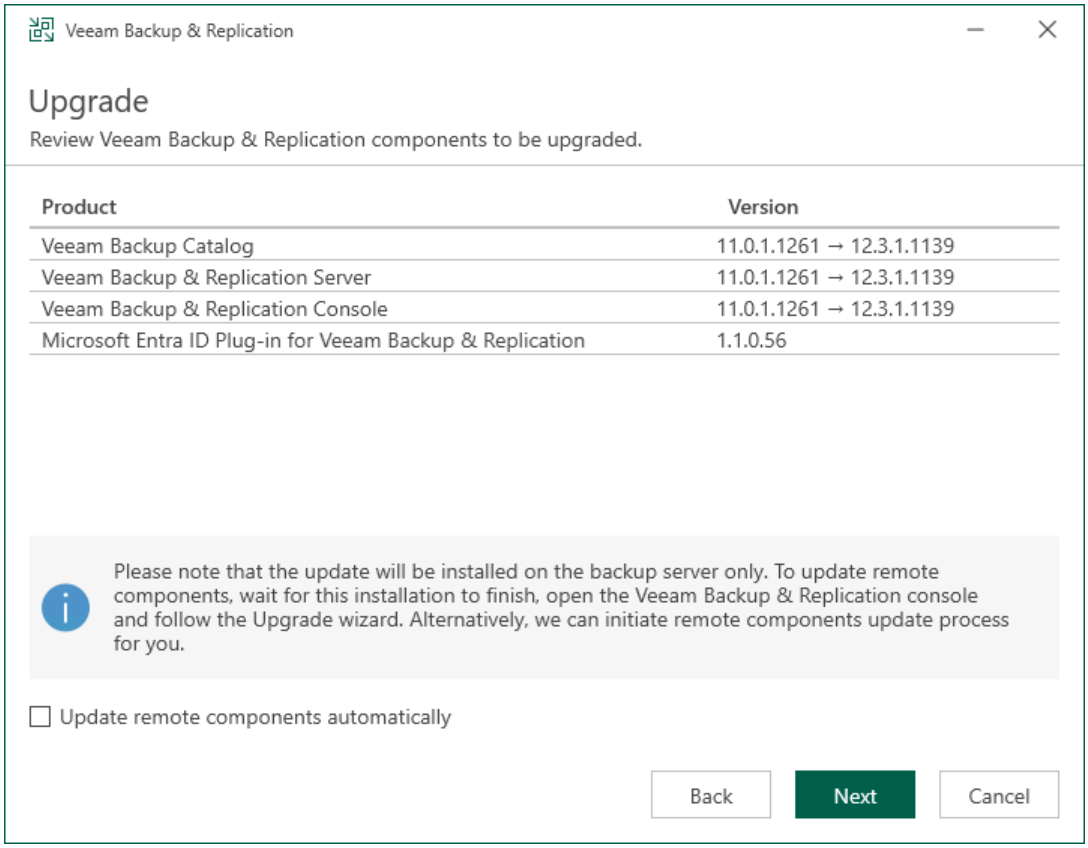


# After Upgrade

After you upgrade Veeam Backup & Replication to version 12.3, perform the following steps:

1. If you use remote backup consoles, upgrade them to version 12.3. If you use remote backup consoles of version 11a, you can upgrade them to version 12.3 automatically when connecting to backup server version 12.3.
2. Download and install the latest available update from the Veeam Updates page.
3. Open the Veeam Backup & Replication console. If necessary, the automated upgrade wizard will automatically appear, prompting you to upgrade the product components running on remote servers. Follow the wizard to complete the upgrade process.
4. If some remote servers are unavailable at the time of upgrade, you can run the upgrade wizard at any time later from the main product menu, or by closing and re-opening the Veeam Backup & Replication console. Note that the out-of-date product components cannot be used by jobs until they are updated to the backup server version.

5. Azure compute accounts based on Microsoft Entra ID (formerly Azure Active Directory) user credentials (created with the **Use existing account option** selected) are obsolete. Replace these accounts with new ones to restore workloads to Microsoft Azure and use Azure archive storage or the Microsoft Azure Plug-In for Veeam Backup & Replication appliance.
6. If you use the Virtual Labs functionality, open settings of each Virtual Lab, and click through the wizard to redeploy each virtual lab with the new proxy appliance version.
7. If you are using Linux servers for your backup infrastructure components, the process of upgrade to version 12.3 will automatically deploy the new persistent data mover only to Linux servers with the VMware Backup Proxy role. To deploy it on other Linux servers, click through the Linux server properties, or use Set-VBRLinux PowerShell cmdlet to mass-deploy. Until you do this, those Linux servers will continue using the legacy run-time data mover to avoid issues with backup repository not meeting the persistent data mover requirements.
8. Enable any scheduled jobs that you have disabled before the upgrade.

Note that immediately after the upgrade, the backup server performance may decrease. This happens due to the maintenance job that optimizes the configuration database. The process may take up to an hour depending on the database size.

> **IMPORTANT**
>
> You must upgrade Veeam components on all remote servers with which the backup server communicates during data protection and disaster recovery tasks. If you do not upgrade components on remote servers, Veeam Backup & Replication jobs will fail. For more information, see Server Components Upgrade.

# Updating Veeam Backup & Replication

Apart from major version releases of Veeam Backup & Replication (for example, *12.3*), Veeam Software provides updates (for example, update *12.3.1.1139* for v12.3). Updates contain bug fixes, performance enhancements and introduce new features.

> **NOTE**
>
> During the update procedure, Veeam Backup & Replication may recreate Windows Firewall rules with default settings. If you have manually modified Windows Firewall rules before installing a new cumulative patch, reapply those modifications after the update.

Before you update Veeam Backup & Replication, check prerequisites. Then use the Veeam Backup & Replication update wizard to run the product update.

After you upgrade Veeam Backup & Replication, perform the finalizing steps.

For information on how to update Veeam Backup & Replication in the unattended mode, see Upgrading Veeam Backup & Replication in Silent Mode.

## Before You Begin

Before you install an update for Veeam Backup & Replication 12.3, check the following prerequisites:

- Make sure you have Veeam Backup & Replication 12.3 (build 12.3.0.310) installed.

  For information on how to upgrade from product version 11a (build 11.0.1.1261) or later, see Upgrading to Veeam Backup & Replication 12.3.

- If you use Veeam Backup Enterprise Manager, start the update procedure with this component. For details, see Updating Veeam Backup Enterprise Manager. Veeam Backup & Replication should be updated after that.

- Stop all restore processes in the Veeam Backup & Replication Console.

- Stop all jobs and disable them in the Veeam Backup & Replication Console.

- Close the Veeam Backup & Replication Console for all users.

# Step 1. Start Update Wizard

To start the update wizard, take the following steps:

1. Download the Updated ISO from this Veeam KB article.
2. Mount the installation image to the machine where Veeam Backup & Replication is installed, or burn the image file to a flash drive or other removable storage device. If you plan to upgrade Veeam Backup & Replication on a VM, use built-in tools of the virtualization management software to mount the image to the VM.

   To extract the content of the ISO, you can also use the latest versions of utilities that can properly extract data from ISO files of large size and can properly work with long file paths.
3. After you mount the image or insert the disk, Autorun opens a splash screen. If Autorun is not available or disabled, run the `Setup.exe` file from the image or disk.
4. Click **Install**.

# Step 2. Select Component

At this step of the wizard, select **Update Veeam Backup & Replication**.

To open Veeam Help Center from the update wizard, click **View Documentation**.

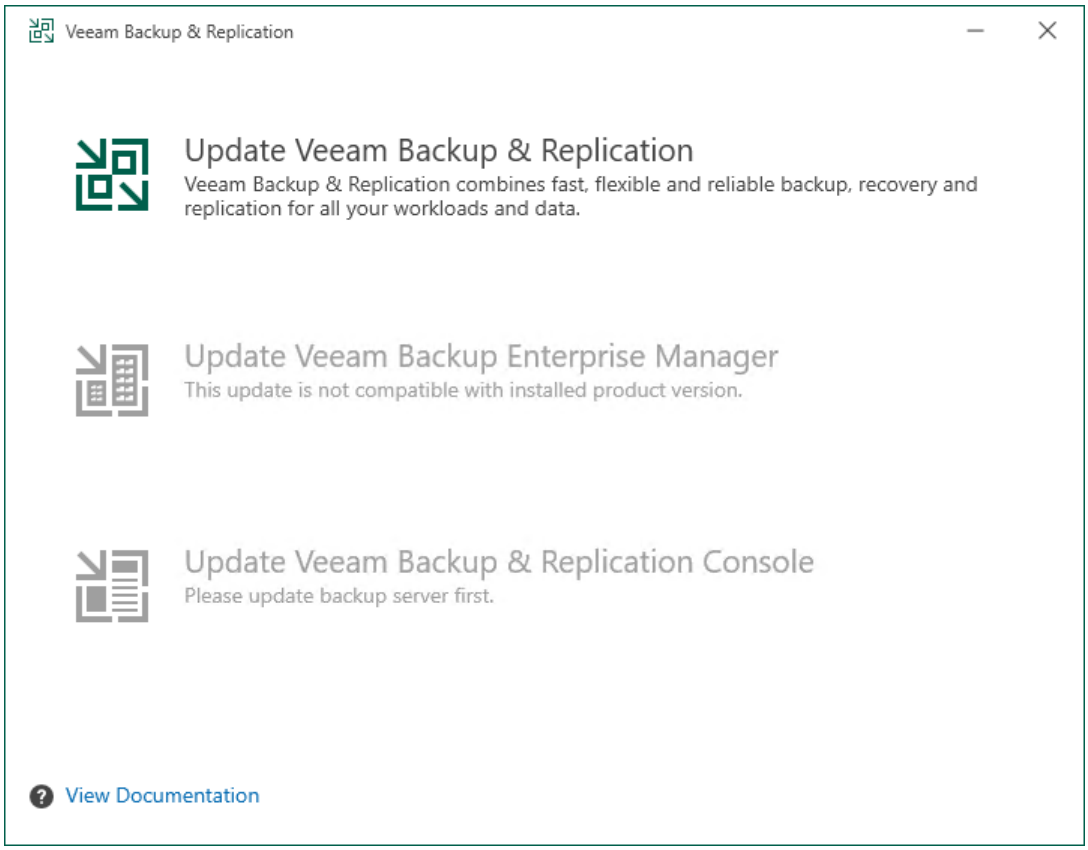

# Step 3. Read and Accept License Agreement

At the **License Agreement** step of the wizard, read Veeam License Agreement and licensing policy as well as license agreements of 3rd party components that Veeam incorporates and license agreements of required software. To accept the license agreements and continue installing Veeam Backup & Replication, click **I Accept**.

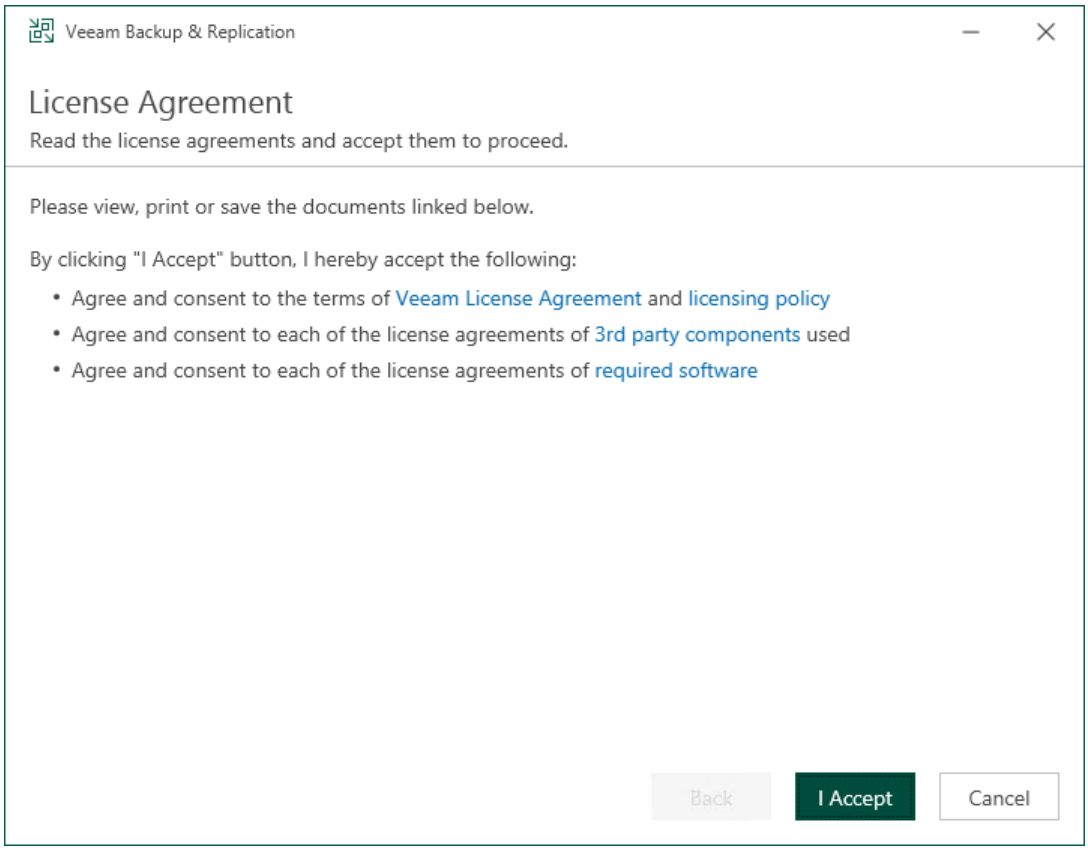

# Step 4. Perform Configuration Check

At the **Configuration Check** step of the wizard, the setup checks the Veeam Backup & Replication configuration.

If the check returns errors, solve their causes before continuing the update. After you solve them, click **Retry** to check if there are any issues left.

If the check returns warning or information messages, you can continue the update and address them later. However, we recommend that you closely investigate warning and information messages: if not properly addressed, their causes may lead to serious problems with further system operation.

To view the details of a certain message, point the cursor to the line with the message. The dialog box will display the detailed description.

To copy a list of detected issues with detailed descriptions for further investigation, click **Copy list to clipboard**.

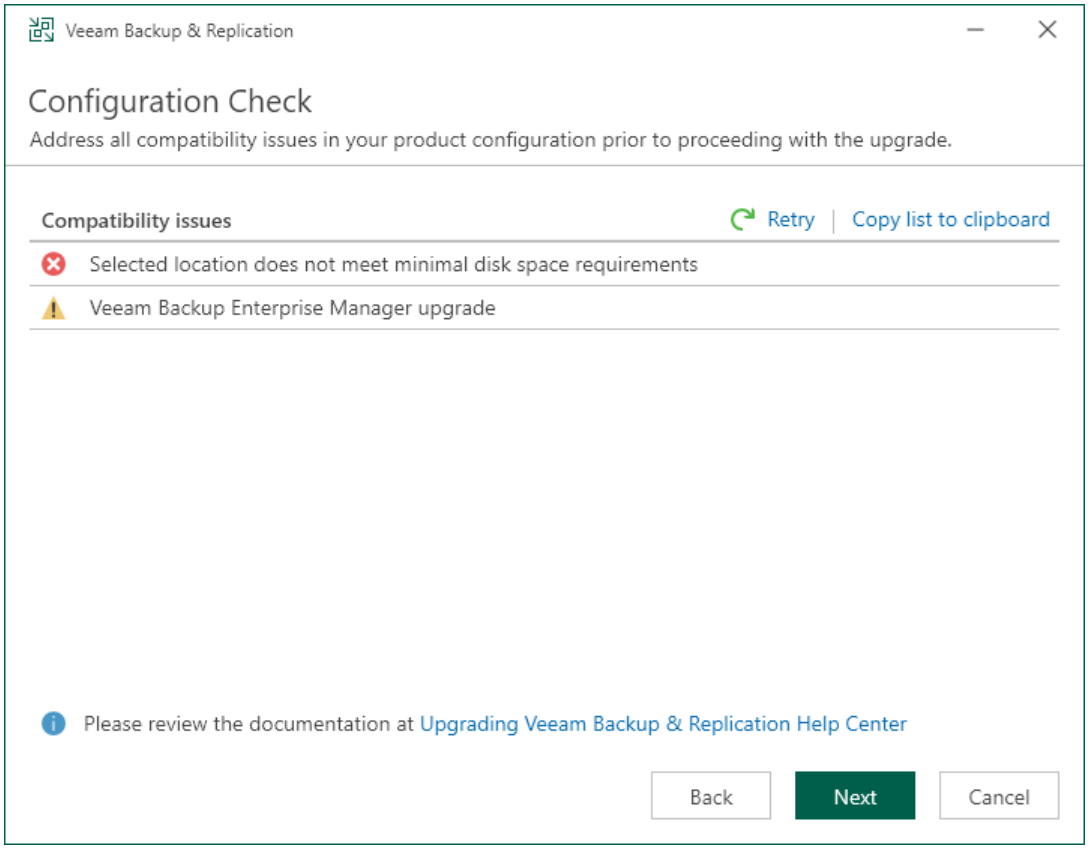

# Step 5. Review Components and Begin Update

At the **Upgrade** step of the wizard, you can review the components that will be upgraded.

To upgrade the remote backup infrastructure components and required Veeam services after the Veeam Backup & Replication server is upgraded, select the **Update remote components automatically** check box. Otherwise, the backup server will prompt you to upgrade them during the first run of the backup server after the update.

Click **Next** to begin the update process.

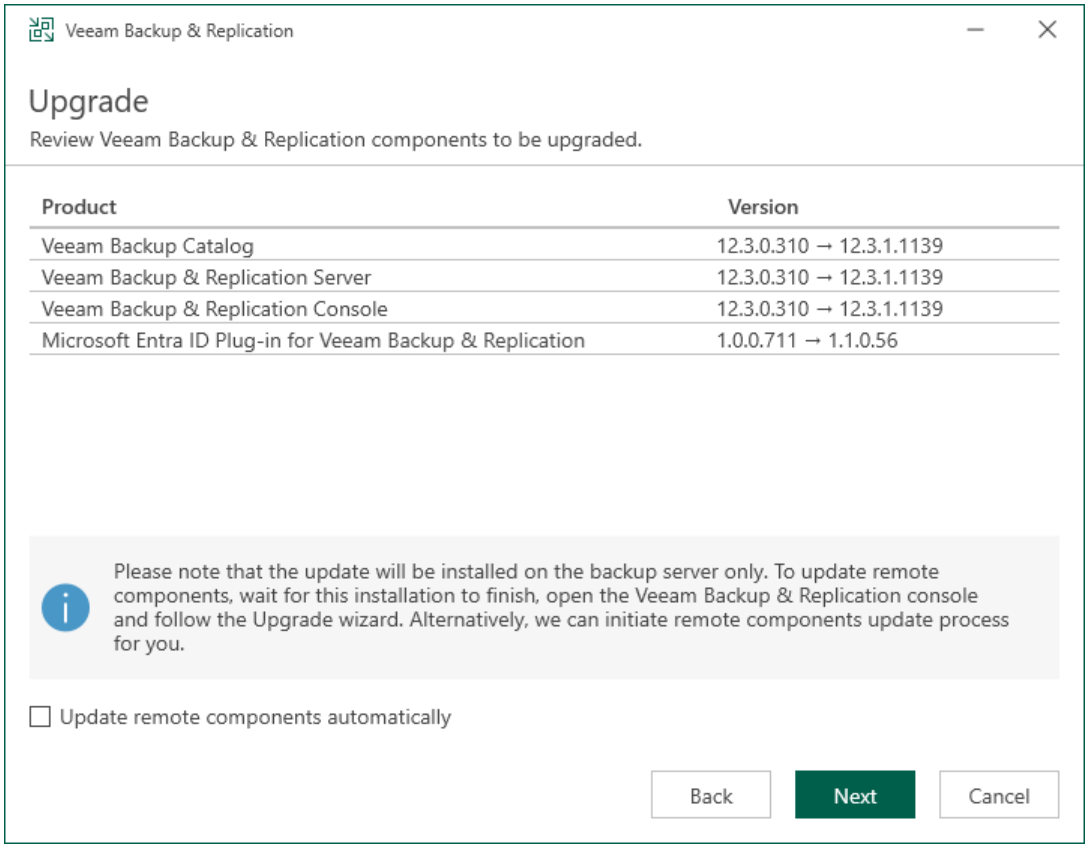


# After Update

After you install updates on the backup server, consider the following:

1. The first connection to the backup server from a remote Veeam Backup & Replication console may require running it once as administrator. Elevated privileges are required to update the Veeam Backup & Replication console binaries.
2. Enable any scheduled jobs that you have disabled before the update.

# Veeam Backup & Replication Update Notifications

Veeam Backup & Replication may automatically notify you about updates that must or can be installed to enhance your work experience with the product. Update notifications eliminate the risk of using out-of-date components in the backup infrastructure or missing critical updates that can have a negative impact on data protection and disaster recovery tasks.

After a new build of Veeam Backup & Replication is published on the Veeam update server, the backup console will display a notification in the Windows Action Center (or an icon in the system tray for earlier Windows versions). If the update is not installed, this notification will keep appearing once a week as a reminder.

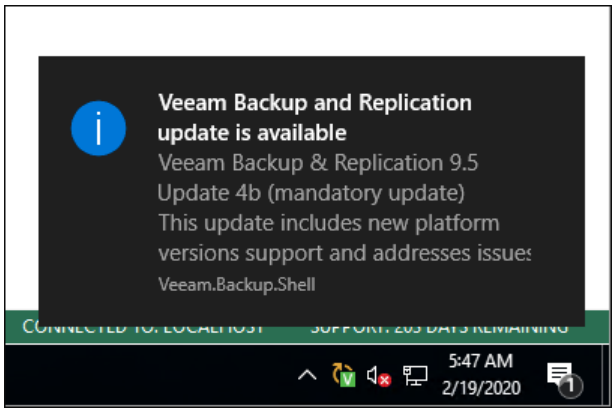


You can also see available updates in the **Managed Servers** > **Missing Updates** node in the **Backup Infrastructure** view.

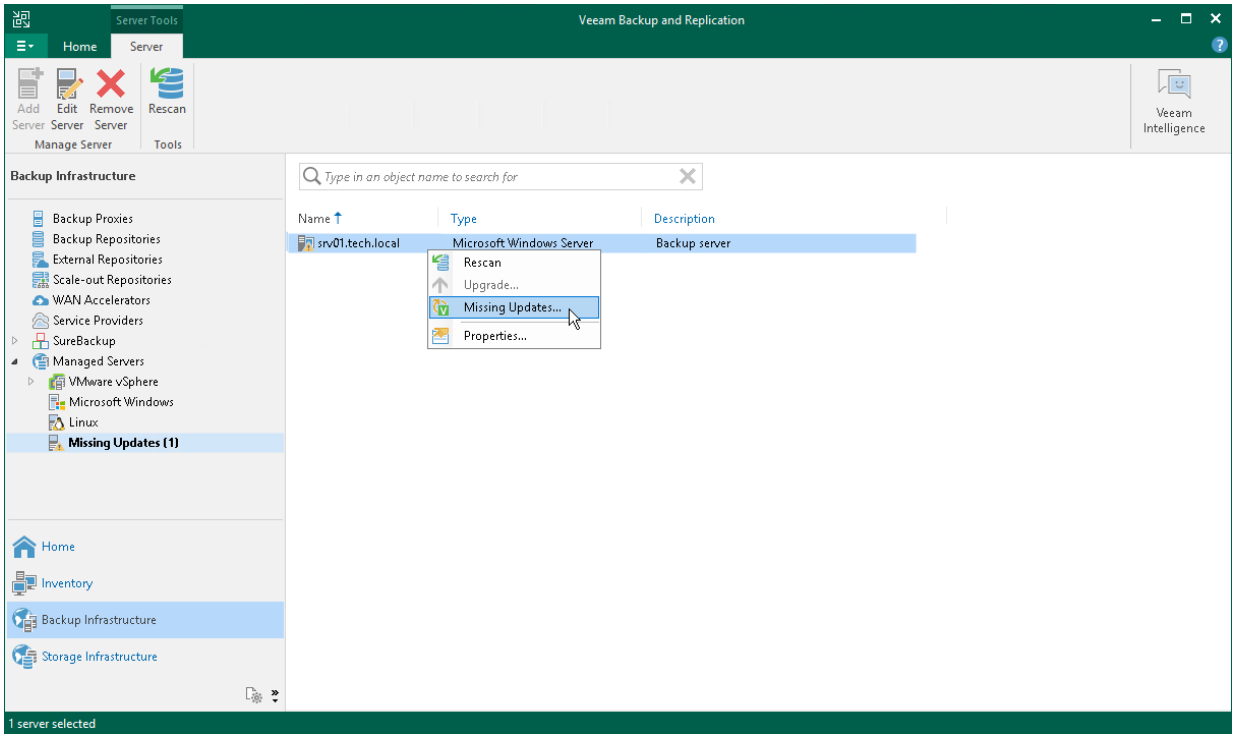


The update notifications are enabled by default. You can disable notifications by clearing the **Check for product and hypervisor updates periodically** check box on the **Options** > **Notifications** tab, as described in section Specifying Other Notification Settings. However, it is recommended that you leave update notifications enabled not to miss critical updates.

Veeam Backup & Replication notifies about new product versions and new product updates.

# How Update Notification Works

To check for updates, Veeam Backup & Replication uses a special XML file on the Veeam Update Notification Server (dev.veeam.com). The XML file contains information about the most up-to-date product version and updates.

Veeam Backup & Replication downloads an XML file from the Veeam Update Notification Server once a week. It also collects information about the installed product. The collected information is compared with the information in the downloaded file. If new product versions and updates are available, Veeam Backup & Replication informs you about them.

> **NOTE**
>
> Make sure that the backup server is connected to the internet and update notification is enabled in Veeam Backup & Replication options. In the opposite case, update notification will not function.

# Installing Updates

To install a product update, double-click the Veeam Backup & Replication notification in the Windows Action Center (or an icon in the system tray for earlier Windows versions). Veeam Backup & Replication will open a KB webpage with the update description and links to the installation archive of the new product version or new update.

# Upgrading Veeam Backup & Replication Console

To perform upgrade of Veeam Backup & Replication console to version 12.3, you must be running version 11a (build 11.0.1.1261) or later on the supported operating system (refer to the System Requirements section of this document). For information on upgrade from earlier versions, see this Veeam KB article.

Use the **Veeam Backup & Replication Console Upgrade** wizard to install the product.

There is an alternative method to upgrade the Veeam Backup & Replication console. If you connect the remote console to the backup server of a later version, the console requests the upgrade. Click **Yes** to upgrade it to the version of the backup server.

> **NOTE**
>
> You cannot downgrade the Veeam Backup & Replication console by connecting to a backup server of an earlier version. In this case, you have to either upgrade the backup server to the version of the console or reinstall the console of an earlier version.

You can also use this approach to install minor updates on the Veeam Backup & Replication console.

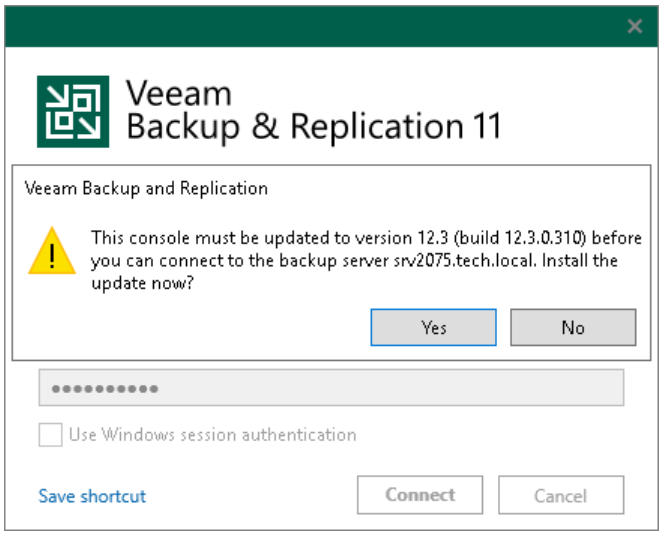

# Step 1. Start Upgrade Wizard

To start the upgrade wizard, take the following steps:

1. Download the latest version of the Veeam Backup & Replication installation image from the Veeam Product Downloads page.
2. Mount the installation image to the machine where the Veeam Backup & Replication console is installed, or burn the image file to a flash drive or other removable storage device. If you plan to upgrade the Veeam Backup & Replication console on a VM, use built-in tools of the virtualization management software to mount the image to the VM.

   To extract the content of the ISO, you can also use the latest versions of utilities that can properly extract data from ISO files of large size and can properly work with long file paths.
3. After you mount the image or insert the disk, Autorun opens a splash screen. If Autorun is not available or disabled, run the `Setup.exe` file from the image or disk.
4. Click **Upgrade**.

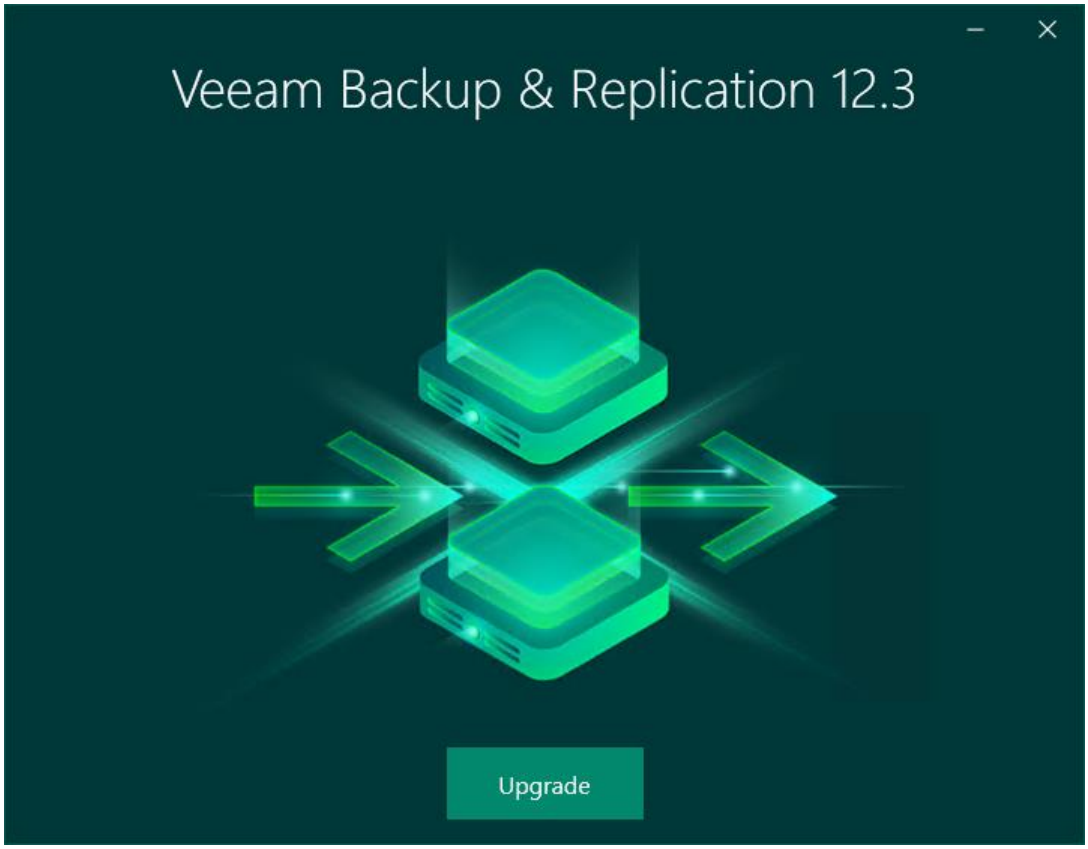

# Step 2. Select Component

At this step of the wizard, select **Upgrade Veeam Backup & Replication Console**.

To open Veeam Help Center from the upgrade wizard, click **View Documentation**.

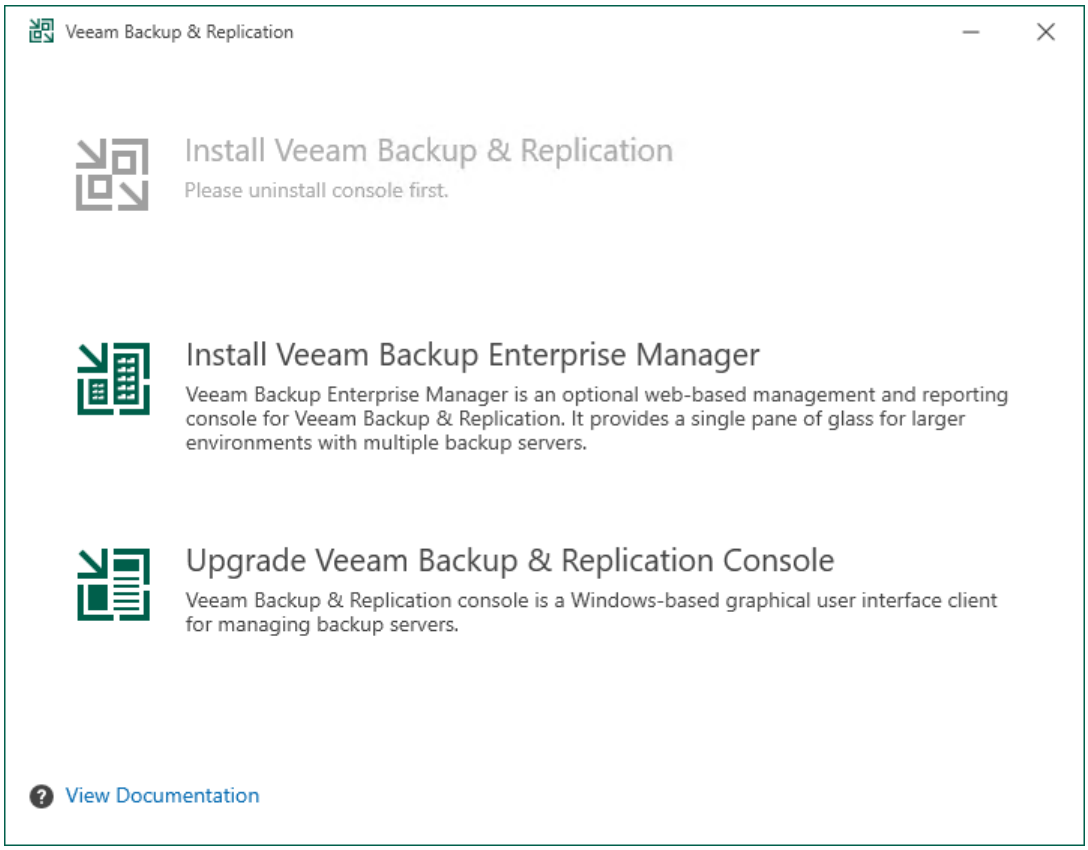

# Step 3. Read and Accept License Agreement

At the **License Agreement** step of the wizard, read Veeam License Agreement and licensing policy as well as license agreements of 3rd party components that Veeam incorporates and license agreements of required software. To accept the license agreements and continue installing the Veeam Backup & Replication console, click **I Accept**.

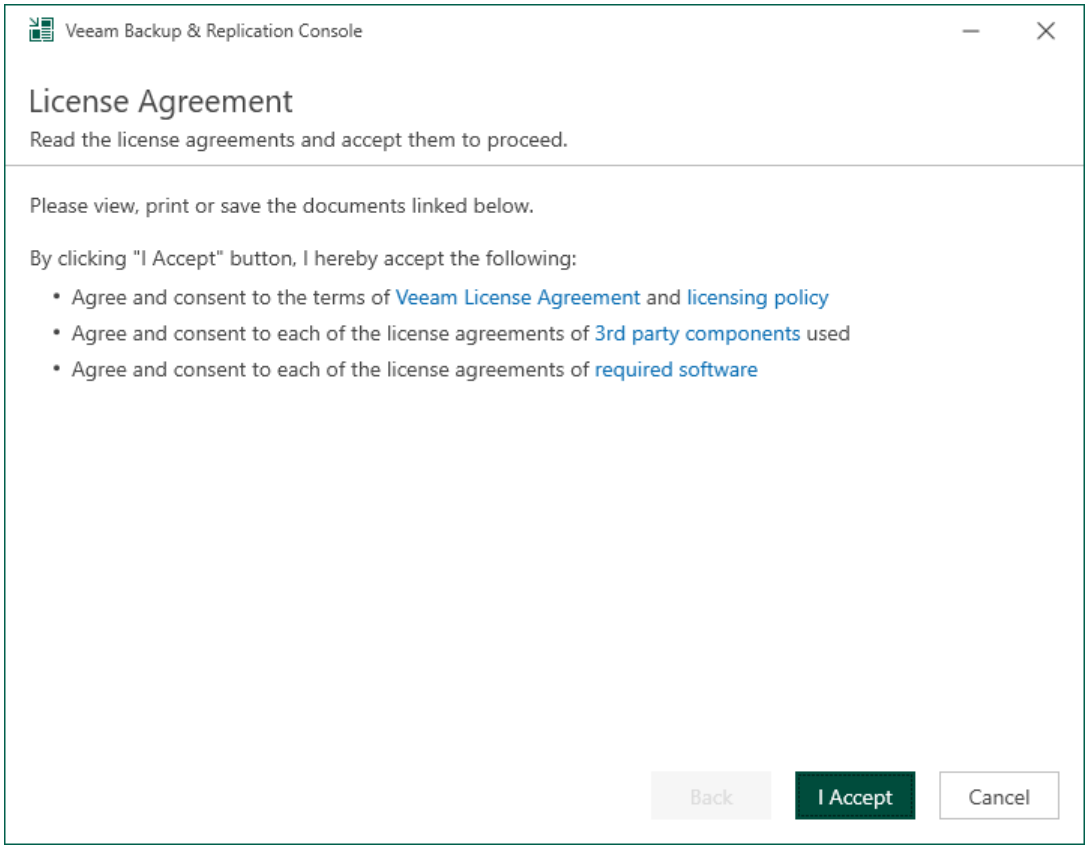

# Step 4. Install Missing Software

At the **System Configuration Check** step of the wizard, the setup wizard checks if the required software is installed on the machine. If some of the required components are missing, the setup will try to install them automatically. After the components are successfully installed, reboot is required. To reboot the machine, click **Reboot**.

If the setup wizard cannot install some of the required software components automatically, install them manually and click **Retry**.

> **NOTE**
>
> If all required software is already installed on the machine, the **System Configuration Check** step will be skipped. For more information on the necessary software, see System Requirements.

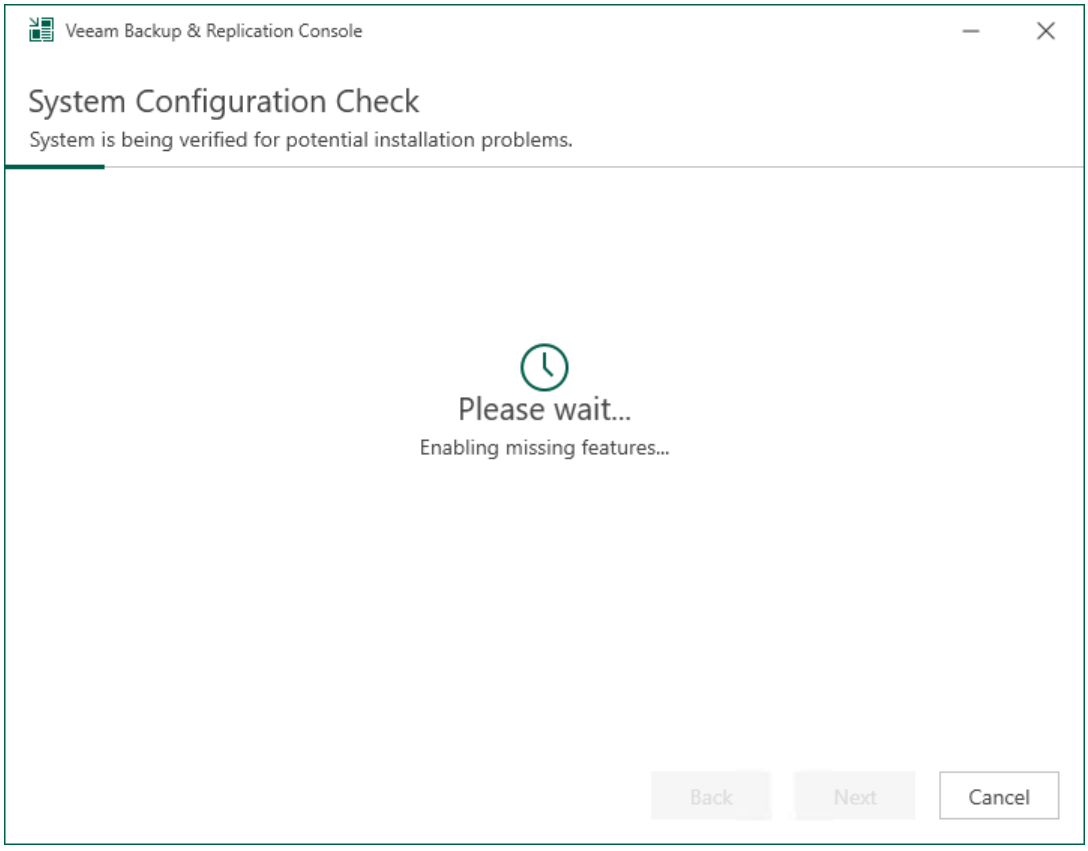

# Step 5. Review Components and Begin Upgrade

At the **Upgrade** step of the wizard, you can review the components that will be upgraded and installed.

Click **Upgrade** to begin the upgrade process.

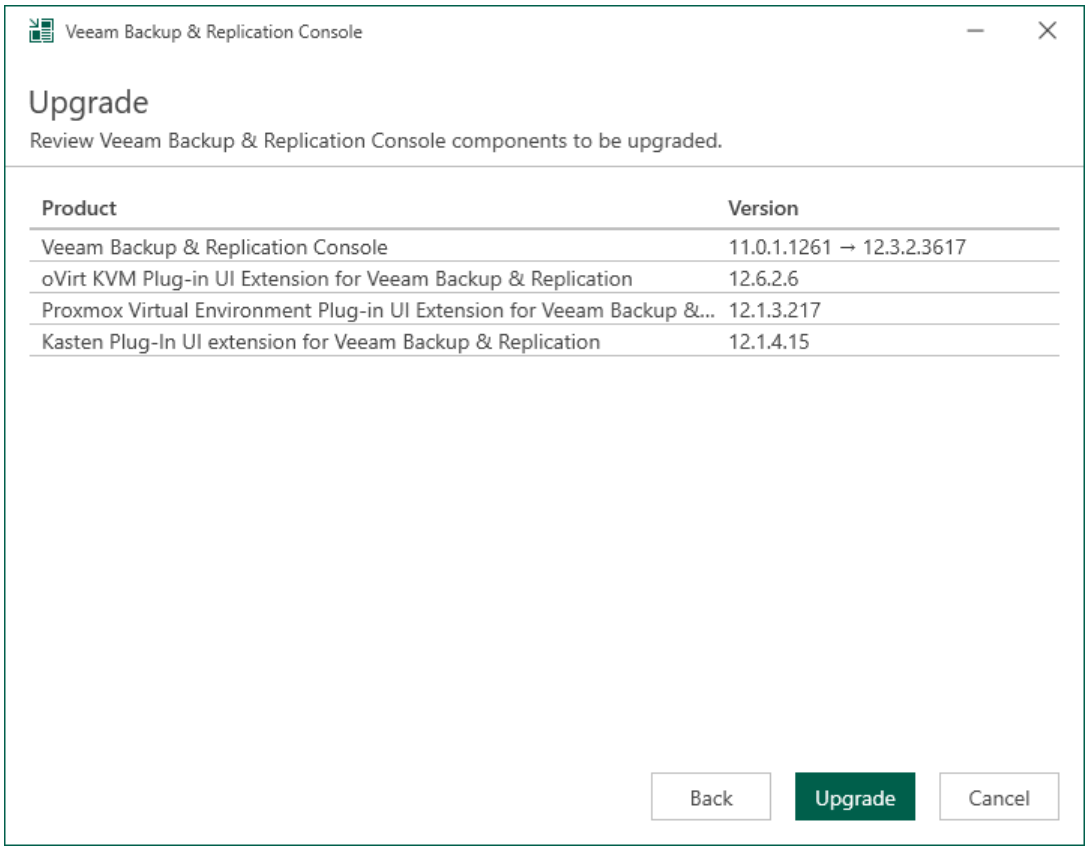

# Upgrading Infrastructure Components

Every time you launch the Veeam Backup & Replication console, Veeam Backup & Replication automatically checks if Veeam Backup & Replication components installed on managed servers are up to date. If a later version of components is available, the system displays the **Components Update** window once a week for all users (once per user) who access the backup server. Components upgrade may be necessary, for example, after you have upgraded Veeam Backup & Replication.

You can manually check if components upgrade is required. To do this, select **Upgrade** from the main menu. If components on all managed servers are up to date, the menu item will be disabled.

> **NOTE**
>
> When you upgrade Veeam Backup & Replication to the current version, you can postpone the upgrade of the I/O filter on the vCenter Servers to a later time. Veeam Backup & Replication supports the postponed upgrade for the minor versions within the 12 version. You must postpone the upgrade or upgrade all clusters on the vCenter Servers. Partially upgraded vCenter Servers or clusters have limited functionality. You cannot add VMs from such vCenter Servers or clusters to CDP policies, commit failback and perform some other operations.

## Bulk Upgrade

To upgrade components on managed servers:

1. In the **Components Update** window, select a server and click **Details**. Veeam Backup & Replication will display the current and latest available versions for installed components.
2. In the **Components Update** window, select check boxes next to servers for which you want to upgrade components and click **Apply**.

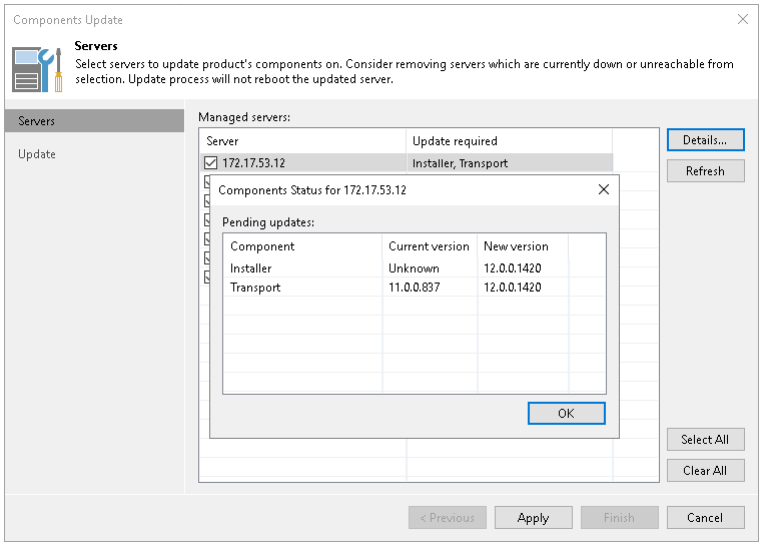


## Individual Upgrade

You can update components on every managed server separately. If components installed on the server require upgrade, Veeam Backup & Replication displays a warning icon next to the server.

To update components for an individual managed server:

1. Open the **Backup Infrastructure** view.

2. In the inventory pane, click **Managed servers**.

3. In the working area, select the server and click **Upgrade** on the ribbon. Alternatively, right-click the selected server and select **Upgrade**.

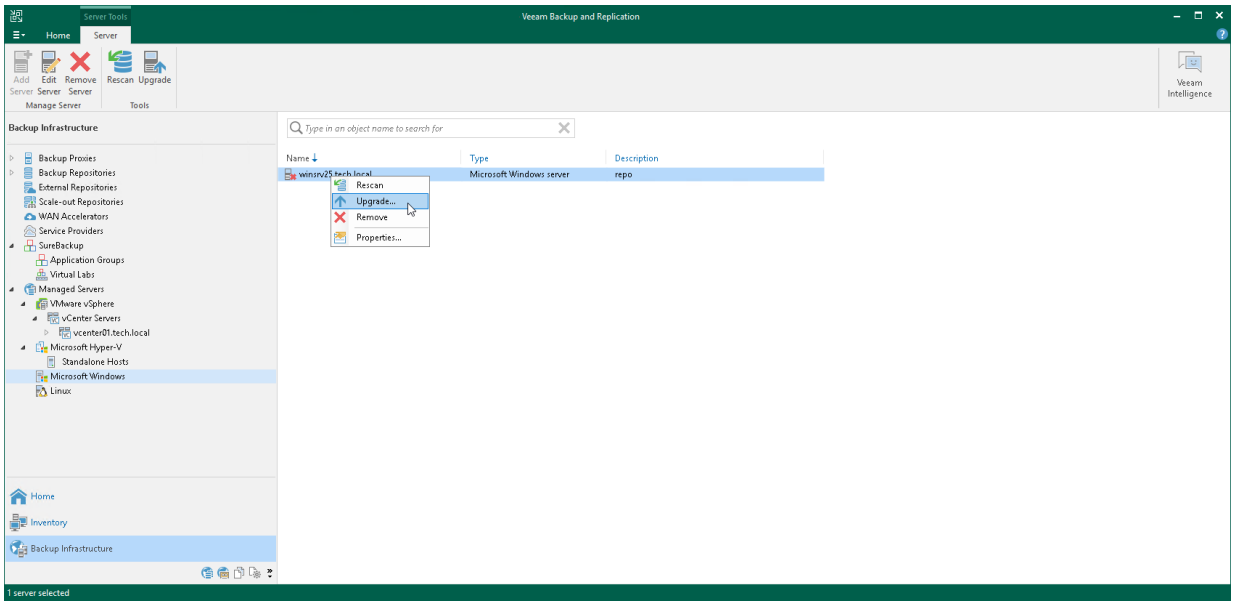

# Silent Deployment

You can install, upgrade and uninstall Veeam Backup & Replication and the Veeam Backup & Replication console in the unattended automated mode with a special XML answer file by using the command line interface. Answer files contain all the necessary installation, upgrade or uninstallation settings in the proper order and their thorough description.

You can find the template answer files on the Veeam Backup & Replication installation disk in the `\Setup\Silent\AnswerFiles\VBR` and `\Setup\Silent\AnswerFiles\VBRConsole` folder. This folder contains the following templates of answer files used for installing, uninstalling, and upgrading Veeam Backup & Replication:

- `VbrAnswerFile_install.xml` — for installing Veeam Backup & Replication
- `VbrAnswerFile_uninstall.xml` — for uninstalling Veeam Backup & Replication
- `VbrAnswerFile_upgrade.xml` — for upgrading Veeam Backup & Replication
- `VbrConsoleAnswerFile_install.xml` — for installing the Veeam Backup & Replication console
- `VbrConsoleAnswerFile_uninstall.xml` — for uninstalling the Veeam Backup & Replication console
- `VbrConsoleAnswerFile_upgrade.xml` — for upgrading the Veeam Backup & Replication console

# Installing Veeam Backup & Replication in Silent Mode

You can install Veeam Backup & Replication in the silent automated mode with a special XML answer file by using the command line interface. The answer file contains all the necessary installation settings in the proper order and their thorough description.

## Before You Begin

Before starting the installation of Veeam Backup & Replication in the unattended mode with the answer file, consider the following:

- The user account that you use to run the silent installation must be in the local Administrators group on the machine where the silent installation will run. The silent installation cannot be run under the LocalSystem and NetworkService accounts.
- If the user account that you use to run the unattended installation is logged on the machine using the network logon method, the unattended installation will fail. To avoid this, use an additional `/SkipNetworkLogonErrors` command line key. For example, it is required when the unattended installation is started within a remote PowerShell session.
- When configuring the answer file, remove or comment out unused `[Optional]` parameters. Otherwise, the installation session will fail.

## Installing Veeam Backup & Replication

To install Veeam Backup & Replication in the silent mode with the answer file, take the following steps:

1. Copy the `VbrAnswerFile_install.xml` file to your local drive.

   You can find the template answer file on the Veeam Backup & Replication installation disk in the `\Setup\Silent\AnswerFiles\VBR` folder. This folder contains the following templates of answer files used for installing, uninstalling, and upgrading Veeam Backup & Replication:

   - `VbrAnswerFile_install.xml` — for installing Veeam Backup & Replication
   - `VbrAnswerFile_uninstall.xml` — for uninstalling Veeam Backup & Replication
   - `VbrAnswerFile_upgrade.xml` — for upgrading Veeam Backup & Replication

2. Configure installation parameters according to your needs. For details, see Configuration Parameters.

   Check that the answer file has the correct bundle (`Vbr`) and mode (`install`) specified in line 2:

   ```
   <unattendedInstallationConfiguration bundle="Vbr" mode="install" version="
   1.0">
   ```

3. After you make all the necessary changes in your answer file, start the installation by running the `Veeam.Silent.Install.exe` file located on the Veeam Backup & Replication installation disk in the `\Setup\Silent` folder. Use the following command line keys in your installation command:

```
D:\Setup\Silent\Veeam.Silent.Install.exe /AnswerFile E:\MyAnswerFileVBRIns
tall.xml /SkipNetworkLogonErrors
```

where:

- `/AnswerFile` — required key for specifying the path to your custom answer file, for example: `E:\MyAnswerFileVBRInstall.xml`.
- `/SkipNetworkLogonErrors` — optional key that allows skipping additional pre-installation validations that do not work under the network logon, which will block the silent installation from running.
- `/LogFolder` — optional key for specifying the path where the setup should save log files if it is different from the default path. The default path is: `C:\ProgramData\Veeam\Setup\Temp`.

# Configuration Parameters

The configuration file contains the following parameters:

| Parameter | Required? | Default | Description |
|---|---|---|---|
| **ACCEPT_EULA** | Yes | | Specify `1` to accept the Veeam license agreement. |
| **ACCEPT_LICENSING_POLICY** | Yes | | Specify `1` to accept the Veeam licensing policy. |
| **ACCEPT_THIRDPARTY_LICENSES** | Yes | | Specify `1` to accept the license agreement for 3rd party components that Veeam incorporates. |
| **ACCEPT_REQUIRED_SOFTWARE** | Yes | | Specify `1` to accept all required software license agreements. |
| **VBR_LICENSE_FILE** | No | | Specify the path to the license file on the machine where you want to install Veeam Backup & Replication. If you do not specify this parameter (or leave it empty value), Veeam Backup & Replication will be installed using the current license file. To install the Community Edition, set the parameter to `0`. |

| Parameter | Required? | Default | Description |
|---|---|---|---|
| VBR_LICENSE_AUTOUPDATE | No | 1 | Specify `1` to enable automatic license update and usage reporting. Specify `0` if you want to update the license manually. For Community Edition, NFR and Evaluation licenses, specify `1`. For licenses without ID information, specify `0`. |
| VBR_SERVICE_USER | No | | Specify the user account under which the Veeam Backup Service will run. If you do not specify this parameter, the service will run under the LocalSystem account. |
| VBR_SERVICE_PASSWORD | No | | Specify the password for the account under which the Veeam Backup Service will run. |
| VBR_SQLSERVER_INSTALL | Yes | 1 | Specify `0` to use an existing SQL server instance or specify `1` to create a new SQL server instance. |
| VBR_ENTRAID_DATABASE_INSTALL | No | 1 | Specify `1` to install a bundled PostgreSQL server for the Microsoft EntraID backup functionality. If set to `0`, PostgreSQL database will not be installed for Microsoft EntraID backup. |
| VBR_SQLSERVER_ENGINE | No | 1 | Specify `1` to use PostgreSQL engine or specify `0` to use Microsoft SQL engine. Note that if you want to create a new SQL server instance, you can only choose the PostgreSQL engine, that is, set the parameter to 1. |

| Parameter | Required? | Default | Description |
|---|---|---|---|
| **VBR_SQLSERVER_SERVER** | No | localhost:5432 | Specify the SQL server and instance (for Microsoft SQL) or the SQL server and port (for PostgreSQL) to be used for deploying the configuration database. Note that if you want to create a new SQL instance, you can only connect to a local host. |
| **VBR_SQLSERVER_DATABASE** | No | VeeamBackup | Specify a configuration database name. If you do not specify this parameter, the default `VeeamBackup` name is used. |
| **VBR_SQLSERVER_AUTHENTICATION** | No | 0 | Specify the authentication mode to connect to the SQL Server where the Veeam Backup & Replication configuration database is deployed. Specify `1` to use the SQL Server authentication mode or specify `0` to use the Microsoft Windows authentication mode. |
| **VBR_SQLSERVER_USERNAME** | No | | Specify the LoginID to connect to the SQL Server in the SQL Server authentication mode. |
| **VBR_SQLSERVER_PASSWORD** | No | | Specify the password to connect to the SQL Server in the SQL Server authentication mode. |
| **VBRC_SERVICE_PORT** | No | 9393 | Specify the TCP port to be used by the Veeam Guest Catalog Service. If you do not specify this parameter, the default `9393` port is used. |

| Parameter | Required? | Default | Description |
|---|---|---|---|
| VBR_SERVICE_PORT | No | 9392 | Specify the TCP port to be used by the Veeam Backup Service. If you do not specify this parameter, the default `9392` port is used. If the specified port number is already occupied, the setup will assign the next available port number to the component. |
| VBR_SECURE_CONNECTIONS_PORT | No | 9401 | Specify the TCP port to be used for communication between the mount server and the backup server. If you do not specify this parameter, the default `9401` port is used. |
| VBR_RESTSERVICE_PORT | No | 9419 | Specify the TCP port to be used for communication with REST API service. If you do not specify this parameter, the default `9419` port is used. |
| INSTALLDIR | No | %ProgramFiles%\Veeam\Backup and Replication | Specify the path to the directory where Veeam Backup & Replication will be installed. If you do not specify this parameter, the default `%ProgramFiles%\Veeam\Backup and Replication` installation path is used. |
| VM_CATALOGPATH | No | C:\VBRCatalog | Specify the path to the catalog folder where index files will be stored. If you do not specify this parameter, a path is selected based on the free space across all available disks. |

| Parameter | Required? | Default | Description |
|---|---|---|---|
| **VBR_IRCACHE** | No | C:\ProgramData\Veeam\Backup\IRCache | Specify the path to the folder where the instant recovery cache will be stored. If you do not specify this parameter, a path is selected based on the free space across all available disks. |
| **VBR_CHECK_UPDATES** | No | 1 | Specify `1` to automatically check for new product versions and updates. Specify `0` if you do not want Veeam Backup & Replication to check for updates automatically. |
| **AHV_INSTALL** | No | 1 | Specify `1` if you want to install Nutanix AHV Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |
| **KVM_INSTALL** | No | 1 | Specify `1` if you want to install oVirt KVM Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |
| **PVE_INSTALL** | No | 1 | Specify `1` if you want to install Proxmox Virtual Environment Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |
| **AWS_INSTALL** | No | 1 | Specify `1` if you want to install AWS Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |

| Parameter | Required? | Default | Description |
|---|---|---|---|
| AZURE_INSTALL | No | 1 | Specify 1 if you want to install Microsoft Azure Plug-in for Veeam Backup & Replication. Specify 0 if you do not want to install the plugin. |
| GCP_INSTALL | No | 1 | Specify 1 if you want to install Google Cloud Plug-in for Veeam Backup & Replication. Specify 0 if you do not want to install the plugin. |
| KASTEN_INSTALL | No | 1 | Specify 1 if you want to install Veeam Kasten Plug-in for Veeam Backup & Replication. Specify 0 if you do not want to install the plugin. |
| REBOOT_IF_REQUIRED | No | 0 | Specify 1 if you want to reboot the machine where you install Veeam Backup & Replication after the installation finishes. Specify 0 if you do not want to reboot the machine. |

Note that you must specify "1" in ACCEPT_EULA, ACCEPT_LICENSING_POLICY, ACCEPT_THIRDPARTY_LICENSES and ACCEPT_REQUIRED_SOFTWARE parameters to proceed with the installation.

# Installation Result Codes

The installation result is written into the installation log file located at your selected log folder. It may show one of the following result codes:

| Result Code | Result |
|---|---|
| 0 | success |
| 1603 | install failure |
| 3010 | reboot required |

| Result Code | Result |
|---|---|
| 3011 | logoff required |

# Installation Error Codes

The installation error codes accompanied by their detailed description are displayed in the command line dialog. They can also be found in the `UnattendedInstallationResult_%DATE%_%TIME%.xml` file in the log folder (by default, `C:\ProgramData\Veeam\Setup\Temp`). You can use such an XML file for retrieving installation results from the scripts or utilities that are used to run the installation. The error message may show one of the following error codes:

| Error Code | Description |
|---|---|
| 0 | Installation has been completed successfully. |
| 1 | Product is already installed. |
| 2 | Uninstallation has been completed successfully. |
| 11 | Unable to start the setup program, because machine reboot is pending. |
| 12 | Reboot is required to finalize prerequisites installation. |
| 13 | Reboot is required to finalize the product installation. |
| 14 | Logoff is required to finalize the product installation. |
| 101 | Failed to start the installer. |
| 102 | Invalid answer file provided. |
| 103 | Invalid launch conditions. |
| 104 | Failed to initialize setup properties. |
| 105 | Failed to validate setup properties. |
| 106 | System configuration check detected some issues. |

| Error Code | Description |
| --- | --- |
| 107 | Failed to install prerequisites. |
| 108 | Failed to install a database server. |
| 109 | Failed to install the product. |
| 110 | Failed to update the product. |
| 111 | Failed to change a service status. |
| 112 | Failed to uninstall the product. |
| 113 | Unexpected error occurred. |

# Installing Veeam Backup & Replication Console in Silent Mode

You can install the Veeam Backup & Replication console in the silent automated mode with a special XML answer file by using the command line interface. The answer file contains all the necessary installation settings in the proper order and their thorough description.

## Before You Begin

Before starting the installation of the Veeam Backup & Replication console in the unattended mode with the answer file, consider the following:

- The user account that you use to run the silent installation must be in the local Administrators group on the machine where the silent installation will run. The silent installation cannot be run under the LocalSystem and NetworkService accounts.
- If the user account that you use to run the unattended installation is logged on the machine using the network logon method, the unattended installation will fail. To avoid this, use an additional `/SkipNetworkLogonErrors` command line key. For example, it is required when the unattended installation is started within a remote PowerShell session.
- When configuring the answer file, remove or comment out unused `[Optional]` parameters. Otherwise, the installation session will fail.

## Installing Veeam Backup & Replication Console

To install the Veeam Backup & Replication console in the silent mode with the answer file, take the following steps:

1. Copy the `VbrConsoleAnswerFile_install.xml` file to your local drive.

   You can find the template answer file on the Veeam Backup & Replication installation disk in the `\Setup\Silent\AnswerFiles\VBRConsole` folder. This folder contains the following templates of answer files used for installing, uninstalling, and upgrading the Veeam Backup & Replication console:

   - `VbrConsoleAnswerFile_install.xml` — for installing the Veeam Backup & Replication console
   - `VbrConsoleAnswerFile_uninstall.xml` — for uninstalling the Veeam Backup & Replication console
   - `VbrConsoleAnswerFile_upgrade.xml` — for upgrading the Veeam Backup & Replication console

2. Configure installation parameters according to your needs. For details, see Configuration Parameters.

   Check that the answer file has the correct bundle (`VbrConsole`) and mode (`install`) specified in line 2:

   ```
   <unattendedInstallationConfiguration bundle="VbrConsole" mode="install" ve
   rsion="1.0">
   ```

3. After you make all the necessary changes in your answer file, start the installation by running the `Veeam.Silent.Install.exe` file located on the Veeam Backup & Replication installation disk in the `\Setup\Silent` folder. Use the following command line keys in your installation command:

```
D:\Setup\Silent\Veeam.Silent.Install.exe /AnswerFile E:\MyAnswerFileConsol
e.xml /SkipNetworkLogonErrors
```

where:

- `/AnswerFile` — required key for specifying the path to your custom answer file, for example: `E:\MyAnswerFileConsole.xml`.
- `/SkipNetworkLogonErrors` — optional key that allows skipping additional pre-installation validations that do not work under the network logon, which will block the silent installation from running.
- `/LogFolder` — optional key for specifying the path where the setup should save log files if it is different from the default path. The default path is: `C:\ProgramData\Veeam\Setup\Temp`.

# Configuration Parameters

The configuration file contains the following parameters:

| Parameter | Required? | Default | Description |
|---|---|---|---|
| **ACCEPT_EULA** | Yes | | Specify `1` to accept the Veeam license agreement. |
| **ACCEPT_LICENSING_POLICY** | Yes | | Specify `1` to accept the Veeam licensing policy. |
| **ACCEPT_THIRDPARTY_LICENSES** | Yes | | Specify `1` to accept the license agreement for 3rd party components that Veeam incorporates. |
| **ACCEPT_REQUIRED_SOFTWARE** | Yes | | Specify `1` to accept all required software license agreements. |
| **INSTALLDIR** | No | %ProgramFiles%\Veeam\Backup and Replication | Specify the path to the directory where the Veeam Backup & Replication console will be installed. If you do not specify this parameter, the default `%ProgramFiles%\Veeam\Backup and Replication` installation path is used. |

| Parameter | Required? | Default | Description |
|---|---|---|---|
| **AHV_INSTALL** | No | 1 | Specify `1` if you want to install Nutanix AHV Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |
| **KVM_INSTALL** | No | 1 | Specify `1` if you want to install oVirt KVM Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |
| **PVE_INSTALL** | No | 1 | Specify `1` if you want to install Proxmox Virtual Environment Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |
| **AWS_INSTALL** | No | 1 | Specify `1` if you want to install AWS Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |
| **AZURE_INSTALL** | No | 1 | Specify `1` if you want to install Microsoft Azure Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |
| **GCP_INSTALL** | No | 1 | Specify `1` if you want to install Google Cloud Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |
| **KASTEN_INSTALL** | No | 1 | Specify `1` if you want to install Veeam Kasten Plug-in for Veeam Backup & Replication. Specify `0` if you do not want to install the plugin. |

| Parameter | Required? | Default | Description |
|---|---|---|---|
| REBOOT_IF_REQUIRED | No | 0 | Specify `1` if you want to reboot the machine where you install Veeam Backup & Replication after the installation finishes. Specify `0` if you do not want to reboot the machine. |

Note that you must specify "1" in ACCEPT_EULA, ACCEPT_LICENSING_POLICY, ACCEPT_THIRDPARTY_LICENSES and ACCEPT_REQUIRED_SOFTWARE parameters to proceed with the installation.

# Installation Result Codes

The installation result is written into the installation log file located at your selected log folder. It may show one of the following result codes:

| Result Code | Result |
|---|---|
| 0 | success |
| 1603 | install failure |
| 3010 | reboot required |
| 3011 | logoff required |

# Installation Error Codes

The installation error codes accompanied by their detailed description are displayed in the command line dialog. They can also be found in the `UnattendedInstallationResult_%DATE%_%TIME%.xml` file in the log folder (by default, `C:\ProgramData\Veeam\Setup\Temp`). You can use such an XML file for retrieving installation results from the scripts or utilities that are used to run the installation. The error message may show one of the following error codes:

| Error Code | Description |
|---|---|
| 0 | Installation has been completed successfully. |
| 1 | Product is already installed. |
| 2 | Uninstallation has been completed successfully. |

| Error Code | Description |
|---|---|
| 11 | Unable to start the setup program, because machine reboot is pending. |
| 12 | Reboot is required to finalize prerequisites installation. |
| 13 | Reboot is required to finalize the product installation. |
| 14 | Logoff is required to finalize the product installation. |
| 101 | Failed to start the installer. |
| 102 | Invalid answer file provided. |
| 103 | Invalid launch conditions. |
| 104 | Failed to initialize setup properties. |
| 105 | Failed to validate setup properties. |
| 106 | System configuration check detected some issues. |
| 107 | Failed to install prerequisites. |
| 108 | Failed to install a database server. |
| 109 | Failed to install the product. |
| 110 | Failed to update the product. |
| 111 | Failed to change a service status. |
| 112 | Failed to uninstall the product. |
| 113 | Unexpected error occurred. |

# Upgrading Veeam Backup & Replication in Silent Mode

You can upgrade Veeam Backup & Replication in the silent automated mode with a special XML answer file by using the command line interface. The answer file contains all the necessary upgrade settings in the proper order and their thorough description.

## Before You Begin

Before starting the upgrade of Veeam Backup & Replication in the unattended mode with the answer file, consider the following:

- The user account that you use to run the silent installation must be in the local Administrators group on the machine where the silent installation will run. The silent installation cannot be run under the LocalSystem and NetworkService accounts.
- If the user account that you use to run the unattended upgrade is logged on the machine using the network logon method, the unattended upgrade will fail. To avoid this, use an additional `/SkipNetworkLogonErrors` command line key. For example, it is required when the unattended upgrade is started within a remote PowerShell session.
- When configuring the answer file, remove or comment out unused `[Optional]` parameters. Otherwise, the upgrade session will fail.

## Upgrading Veeam Backup & Replication

To upgrade Veeam Backup & Replication in the silent mode with the answer file, take the following steps:

1. Copy the `VbrAnswerFile_upgrade.xml` file to your local drive.

   You can find the template answer file on the Veeam Backup & Replication installation disk in the `\Setup\Silent\AnswerFiles\VBR` folder. This folder contains the following templates of answer files used for installing, uninstalling, and upgrading Veeam Backup & Replication:

   - `VbrAnswerFile_install.xml` — for installing Veeam Backup & Replication
   - `VbrAnswerFile_uninstall.xml` — for uninstalling Veeam Backup & Replication
   - `VbrAnswerFile_upgrade.xml` — for upgrading Veeam Backup & Replication

2. Configure upgrade parameters according to your needs. For details, see Configuration Parameters.

   Check that the answer file has the correct bundle (`Vbr`) and mode (`upgrade`) specified in line 2:

   ```
   <unattendedInstallationConfiguration bundle="Vbr" mode="upgrade" version="
   1.0">
   ```

3. After you make all the necessary changes in your answer file, start the upgrade by running the `Veeam.Silent.Install.exe` file located on the Veeam Backup & Replication installation disk in the `\Setup\Silent` folder. Use the following command line keys in your installation command:

```
D:\Setup\Silent\Veeam.Silent.Install.exe /AnswerFile E:\MyAnswerFileVBRUpg
rade.xml /SkipNetworkLogonErrors
```

where:

- `/AnswerFile` — required key for specifying the path to your custom answer file, for example: `E:\MyAnswerFileVBRUpgrade.xml`.
- `/SkipNetworkLogonErrors` — optional key that allows skipping additional pre-upgrade validations that do not work under the network logon, which will block the silent upgrade from running.
- `/LogFolder` — optional key for specifying the path where the setup should save log files if it is different from the default path. The default path is: `C:\ProgramData\Veeam\Setup\Temp`.

# Configuration Parameters

The configuration file contains the following parameters:

| Parameter | Required? | Default | Description |
|---|---|---|---|
| **ACCEPT_EULA** | Yes | | Specify `1` to accept the Veeam license agreement. |
| **ACCEPT_LICENSING_POLICY** | Yes | | Specify `1` to accept the Veeam licensing policy. |
| **ACCEPT_THIRDPARTY_LICENSES** | Yes | | Specify `1` to accept the license agreement for 3rd party components that Veeam incorporates. |
| **ACCEPT_REQUIRED_SOFTWARE** | Yes | | Specify `1` to accept all required software license agreements. |
| **VBR_LICENSE_FILE** | No | | Specify the path to the license file on the machine where you want to install Veeam Backup & Replication. If you do not specify this parameter (or leave it empty value), Veeam Backup & Replication will be installed using the current license file. To install the Community Edition, set the parameter to `0`. |
| **VBR_LICENSE_AUTOUPDATE** | No | 1 | Specify `1` to enable automatic license update and usage reporting. Specify `0` if you want to update the license manually. For Community Edition, NFR and Evaluation licenses, specify `1`. For licenses without ID information, specify `0`. |

| Parameter | Required? | Default | Description |
|---|---|---|---|
| **VBR_SERVICE_PASSWORD** | No | | Specify the password for the account under which the Veeam Backup Service will run. |
| **VBR_ENTRAID_DATABASE_INSTALL** | No | 1 | Specify `1` to install a bundled PostgreSQL server for the Microsoft EntraID backup functionality. If set to `0`, PostgreSQL database will not be installed for Microsoft EntraID backup. |
| **VBR_SQLSERVER_PASSWORD** | No | | Specify the password to connect to the SQL Server in the SQL Server authentication mode. |
| **VBR_AUTO_UPGRADE** | No | 0 | Specify `1` to automatically upgrade existing components in the backup infrastructure. Specify `0` if you do not want Veeam Backup & Replication to upgrade the existing components automatically. |
| **REBOOT_IF_REQUIRED** | No | 0 | Specify `1` if you want to reboot the machine where you install Veeam Backup & Replication after the installation finishes. Specify `0` if you do not want to reboot the machine. |

Note that you must specify "1" in ACCEPT_EULA, ACCEPT_LICENSING_POLICY, ACCEPT_THIRDPARTY_LICENSES and ACCEPT_REQUIRED_SOFTWARE parameters to proceed with the upgrade.

# Upgrade Result Codes

The upgrade result is written into the upgrade log file located at your selected log folder. It may show one of the following result codes:

| Result Code | Result |
|---|---|
| **0** | success |
| **1603** | install failure |
| **3010** | reboot required |
| **3011** | logoff required |

# Upgrade Error Codes

The upgrade error codes accompanied by their detailed description are displayed in the command line dialog. They can also be found in the `UnattendedInstallationResult_%DATE%_%TIME%.xml` file in the log folder (by default, `C:\ProgramData\Veeam\Setup\Temp`). You can use such an XML file for retrieving upgrade results from the scripts or utilities that are used to run the upgrade. The error message may show one of the following error codes:

| Error Code | Description |
|---|---|
| 0 | Installation has been completed successfully. |
| 1 | Product is already installed. |
| 2 | Uninstallation has been completed successfully. |
| 11 | Unable to start the setup program, because machine reboot is pending. |
| 12 | Reboot is required to finalize prerequisites installation. |
| 13 | Reboot is required to finalize the product installation. |
| 14 | Logoff is required to finalize the product installation. |
| 101 | Failed to start the installer. |
| 102 | Invalid answer file provided. |
| 103 | Invalid launch conditions. |
| 104 | Failed to initialize setup properties. |
| 105 | Failed to validate setup properties. |
| 106 | System configuration check detected some issues. |
| 107 | Failed to install prerequisites. |
| 108 | Failed to install a database server. |
| 109 | Failed to install the product. |

| Error Code | Description |
| --- | --- |
| 110 | Failed to update the product. |
| 111 | Failed to change a service status. |
| 112 | Failed to uninstall the product. |
| 113 | Unexpected error occurred. |

# Upgrading Veeam Backup & Replication Console in Silent Mode

You can upgrade the Veeam Backup & Replication console in the silent automated mode with a special XML answer file by using the command line interface. The answer file contains all the necessary upgrade settings in the proper order and their thorough description.

## Before You Begin

Before starting the upgrade of the Veeam Backup & Replication console in the unattended mode with the answer file, consider the following:

- The user account that you use to run the silent installation must be in the local Administrators group on the machine where the silent installation will run. The silent installation cannot be run under the LocalSystem and NetworkService accounts.
- If the user account that you use to run the unattended upgrade is logged on the machine using the network logon method, the unattended upgrade will fail. To avoid this, use an additional `/SkipNetworkLogonErrors` command line key. For example, it is required when the unattended upgrade is started within a remote PowerShell session.
- When configuring the answer file, remove or comment out unused `[Optional]` parameters. Otherwise, the upgrade session will fail.

## Upgrading Veeam Backup & Replication Console

To upgrade the Veeam Backup & Replication console in the silent mode with the answer file, take the following steps:

1. Copy the `VbrConsoleAnswerFile_upgrade.xml` file to your local drive.

   You can find the template answer file on the Veeam Backup & Replication installation disk in the `\Setup\Silent\AnswerFiles\VBRConsole` folder. This folder contains the following templates of answer files used for installing, uninstalling, and upgrading the Veeam Backup & Replication console:

   - `VbrConsoleAnswerFile_install.xml` — for installing the Veeam Backup & Replication console
   - `VbrConsoleAnswerFile_uninstall.xml` — for uninstalling the Veeam Backup & Replication console
   - `VbrConsoleAnswerFile_upgrade.xml` — for upgrading the Veeam Backup & Replication console

2. Configure upgrade parameters according to your needs. For details, see Configuration Parameters.

   Check that the answer file has the correct bundle (`VbrConsole`) and mode (`upgrade`) specified in line 2:

   ```
   <unattendedInstallationConfiguration bundle="VbrConsole" mode="upgrade" ve
   rsion="1.0">
   ```

3. After you make all the necessary changes in your answer file, start the upgrade by running the `Veeam.Silent.Install.exe` file located on the Veeam Backup & Replication installation disk in the `\Setup\Silent` folder. Use the following command line keys in your upgrade command:

```
D:\Setup\Silent\Veeam.Silent.Install.exe /AnswerFile E:\MyAnswerFileConsol
eUpgrade.xml /SkipNetworkLogonErrors
```

where:

- `/AnswerFile` — required key for specifying the path to your custom answer file, for example: `E:\MyAnswerFileConsoleUpgrade.xml`.
- `/SkipNetworkLogonErrors` — optional key that allows skipping additional pre-upgraded validations that do not work under the network logon, which will block the silent upgrade from running.
- `/LogFolder` — optional key for specifying the path where the setup should save log files if it is different from the default path. The default path is: `C:\ProgramData\Veeam\Setup\Temp`.

# Configuration Parameters

The configuration file contains the following parameters:

| Parameter | Required? | Default | Description |
|---|---|---|---|
| **ACCEPT_EULA** | Yes | | Specify `1` to accept the Veeam license agreement. |
| **ACCEPT_LICENSING_POLICY** | Yes | | Specify `1` to accept the Veeam licensing policy. |
| **ACCEPT_THIRDPARTY_LICENSES** | Yes | | Specify `1` to accept the license agreement for 3rd party components that Veeam incorporates. |
| **ACCEPT_REQUIRED_SOFTWARE** | Yes | | Specify `1` to accept all required software license agreements. |
| **REBOOT_IF_REQUIRED** | No | 0 | Specify `1` if you want to reboot the machine where you upgrade the Veeam Backup & Replication console after the upgrade finishes. Specify `0` if you do not want to reboot the machine. |

Note that you must specify "1" in ACCEPT_EULA, ACCEPT_LICENSING_POLICY, ACCEPT_THIRDPARTY_LICENSES and ACCEPT_REQUIRED_SOFTWARE parameters to proceed with the installation.

# Upgrade Result Codes

The upgrade result is written into the installation log file located at your selected log folder. It may show one of the following result codes:

| Result Code | Result |
|---|---|
| 0 | success |
| 1603 | install failure |
| 3010 | reboot required |
| 3011 | logoff required |

# Upgrade Error Codes

The upgrade error codes accompanied by their detailed description are displayed in the command line dialog. They can also be found in the `UnattendedInstallationResult_%DATE%_%TIME%.xml` file in the log folder (by default, `C:\ProgramData\Veeam\Setup\Temp`). You can use such an XML file for retrieving upgrade results from the scripts or utilities that are used to run the upgrade. The error message may show one of the following error codes:

| Error Code | Description |
|---|---|
| 0 | Installation has been completed successfully. |
| 1 | Product is already installed. |
| 2 | Uninstallation has been completed successfully. |
| 11 | Unable to start the setup program, because machine reboot is pending. |
| 12 | Reboot is required to finalize prerequisites installation. |
| 13 | Reboot is required to finalize the product installation. |
| 14 | Logoff is required to finalize the product installation. |
| 101 | Failed to start the installer. |

| Error Code | Description |
|---|---|
| 102 | Invalid answer file provided. |
| 103 | Invalid launch conditions. |
| 104 | Failed to initialize setup properties. |
| 105 | Failed to validate setup properties. |
| 106 | System configuration check detected some issues. |
| 107 | Failed to install prerequisites. |
| 108 | Failed to install a database server. |
| 109 | Failed to install the product. |
| 110 | Failed to update the product. |
| 111 | Failed to change a service status. |
| 112 | Failed to uninstall the product. |
| 113 | Unexpected error occurred. |

# Uninstalling Veeam Backup & Replication in Silent Mode

You can uninstall Veeam Backup & Replication in the silent automated mode with a special XML answer file by using the command line interface. The answer file contains all the necessary uninstallation settings in the proper order and their thorough description.

## Before You Begin

Before starting the uninstallation of Veeam Backup & Replication in the unattended mode with the answer file, consider the following:

- The user account that you use to run the silent installation must be in the local Administrators group on the machine where the silent installation will run. The silent installation cannot be run under the LocalSystem and NetworkService accounts.
- If the user account that you use to run the unattended uninstallation is logged on the machine using the network logon method, the unattended uninstallation will fail. To avoid this, use an additional `/SkipNetworkLogonErrors` command line key. For example, it is required when the unattended uninstallation is started within a remote PowerShell session.
- When configuring the answer file, remove or comment out the unused `[Optional]` parameter. Otherwise, the uninstallation session will fail.

## Uninstalling Veeam Backup & Replication

To uninstall Veeam Backup & Replication in the silent mode with the answer file, take the following steps:

1. Copy the `VbrAnswerFile_uninstall.xml` file to your local drive.

   You can find the template answer file on the Veeam Backup & Replication installation disk in the `\Setup\Silent\AnswerFiles\VBR` folder. This folder contains the following templates of answer files used for installing, uninstalling, and upgrading Veeam Backup & Replication:

   - `VbrAnswerFile_install.xml` — for installing Veeam Backup & Replication
   - `VbrAnswerFile_uninstall.xml` — for uninstalling Veeam Backup & Replication
   - `VbrAnswerFile_upgrade.xml` — for upgrading Veeam Backup & Replication

2. Configure uninstallation parameters according to your needs. For details, see Configuration Parameters.

   Check that the answer file has the correct bundle (`Vbr`) and mode (`uninstall`) specified in line 2:

   ```
   <unattendedInstallationConfiguration bundle="Vbr" mode="uninstall" version
   ="1.0">
   ```

3. After you make all the necessary changes in your answer file, start the uninstallation by running the `Veeam.Silent.Install.exe` file located on the Veeam Backup & Replication installation disk in the `\Setup\Silent` folder. Use the following command line keys in your uninstallation command:

```
D:\Setup\Silent\Veeam.Silent.Install.exe /AnswerFile E:\MyAnswerFileVBRUni
nstall.xml /SkipNetworkLogonErrors
```

where:

- `/AnswerFile` — required key for specifying the path to your custom answer file, for example: `E:\MyAnswerFileVBRUninstall.xml`.
- `/SkipNetworkLogonErrors` — optional key that allows skipping additional pre-uninstallation validations that do not work under the network logon, which will block the silent uninstallation from running.
- `/LogFolder` — optional key for specifying the path where the setup should save log files if it is different from the default path. The default path is: `C:\ProgramData\Veeam\Setup\Temp`.

# Configuration Parameters

The configuration file contains only the following parameter:

| Parameter | Required? | Default | Description |
|---|---|---|---|
| **REBOOT_IF_REQUIRED** | No | 0 | Specify `1` if you want to reboot the machine where you uninstall Veeam Backup & Replication after the uninstallation finishes. Specify `0` if you do not want to reboot the machine. |

# Uninstallation Result Codes

The uninstallation result is written into the uninstallation log file located at your selected log folder. It may show one of the following result codes:

| Result Code | Result |
|---|---|
| **0** | success |
| **1603** | install failure |
| **3010** | reboot required |
| **3011** | logoff required |

# Uninstallation Error Codes

The uninstallation error codes accompanied by their detailed description are displayed in the command line dialog. They can also be found in the `UnattendedInstallationResult_%DATE%_%TIME%.xml` file in the log folder (by default, `C:\ProgramData\Veeam\Setup\Temp`). You can use such an XML file for retrieving uninstallation results from the scripts or utilities that are used to run the uninstallation. The error message may show one of the following error codes:

| Error Code | Description |
|---|---|
| 2 | Uninstallation has been completed successfully. |
| 11 | Unable to start the setup program, because machine reboot is pending. |
| 101 | Failed to start the installer. |
| 102 | Invalid answer file provided. |
| 103 | Invalid launch conditions. |
| 112 | Failed to uninstall the product. |
| 113 | Unexpected error occurred. |

# Uninstalling Veeam Backup & Replication Console in Silent Mode

You can uninstall the Veeam Backup & Replication console in the silent automated mode with a special XML answer file by using the command line interface. The answer file contains all the necessary uninstallation settings in the proper order and their thorough description.

## Before You Begin

Before starting the uninstallation of the Veeam Backup & Replication console in the unattended mode with the answer file, consider the following:

- The user account that you use to run the silent installation must be in the local Administrators group on the machine where the silent installation will run. The silent installation cannot be run under the LocalSystem and NetworkService accounts.
- If the user account that you use to run the unattended uninstallation is logged on the machine using the network logon method, the unattended uninstallation will fail. To avoid this, use an additional `/SkipNetworkLogonErrors` command line key. For example, it is required when the unattended uninstallation is started within a remote PowerShell session.
- When configuring the answer file, remove or comment out the unused `[Optional]` parameter. Otherwise, the uninstallation session will fail.

## Uninstalling Veeam Backup & Replication Console

To uninstall the Veeam Backup & Replication console in the silent mode with the answer file, take the following steps:

1. Copy the `VbrAnswerFile_uninstall.xml` file to your local drive.

   You can find the template answer file on the Veeam Backup & Replication installation disk in the `\Setup\Silent\AnswerFiles\VBRConsole` folder. This folder contains the following templates of answer files used for installing, uninstalling, and upgrading Veeam Backup & Replication:

   - `VbrConsoleAnswerFile_install.xml` — for installing the Veeam Backup & Replication console
   - `VbrConsoleAnswerFile_uninstall.xml` — for uninstalling the Veeam Backup & Replication console
   - `VbrConsoleAnswerFile_upgrade.xml` — for upgrading the Veeam Backup & Replication console

2. Configure uninstallation parameters according to your needs. For details, see Configuration Parameters.

   Check that the answer file has the correct bundle (`VbrConsole`) and mode (`uninstall`) specified in line 2:

   ```
   <unattendedInstallationConfiguration bundle="VbrConsole" mode="uninstall"
   version="1.0">
   ```

3. After you make all the necessary changes in your answer file, start the uninstallation by running the `Veeam.Silent.Install.exe` file located on the Veeam Backup & Replication installation disk in the `\Setup\Silent` folder. Use the following command line keys in your uninstallation command:

```
D:\Setup\Silent\Veeam.Silent.Install.exe /AnswerFile E:\MyAnswerFileConsol
eUninstall.xml /SkipNetworkLogonErrors
```

where:

- `/AnswerFile` — required key for specifying the path to your custom answer file, for example: `E:\MyAnswerFileConsoleUninstall.xml`.
- `/SkipNetworkLogonErrors` — optional key that allows skipping additional pre-uninstallation validations that do not work under the network logon, which will block the silent uninstallation from running.
- `/LogFolder` — optional key for specifying the path where the setup should save log files if it is different from the default path. The default path is: `C:\ProgramData\Veeam\Setup\Temp`.

# Configuration Parameters

The configuration file contains only the following parameter:

| Parameter | Required? | Default | Description |
|---|---|---|---|
| **REBOOT_IF_REQUIRED** | No | 0 | Specify `1` if you want to reboot the machine where you uninstall the Veeam Backup & Replication console after the uninstallation finishes. Specify `0` if you do not want to reboot the machine. |

# Uninstallation Result Codes

The uninstallation result is written into the uninstallation log file located at your selected log folder. It may show one of the following result codes:

| Result Code | Result |
|---|---|
| **0** | success |
| **1603** | install failure |
| **3010** | reboot required |
| **3011** | logoff required |

# Uninstallation Error Codes

The uninstallation error codes accompanied by their detailed description are displayed in the command line dialog. They can also be found in the `UnattendedInstallationResult_%DATE%_%TIME%.xml` file in the log folder (by default, `C:\ProgramData\Veeam\Setup\Temp`). You can use such an XML file for retrieving uninstallation results from the scripts or utilities that are used to run the uninstallation. The error message may show one of the following error codes:

| Error Code | Description |
|---|---|
| 2 | Uninstallation has been completed successfully. |
| 11 | Unable to start the setup program, because machine reboot is pending. |
| 101 | Failed to start the installer. |
| 102 | Invalid answer file provided. |
| 103 | Invalid launch conditions. |
| 112 | Failed to uninstall the product. |
| 113 | Unexpected error occurred. |

# Getting Started with Veeam Backup & Replication

This section describes Veeam Backup & Replication UI and basic concepts.

# Logging in to Veeam Backup & Replication

To log in to Veeam Backup & Replication, you must open the Veeam Backup & Replication console and specify connection settings to access the backup server.

1. To open the Veeam Backup & Replication console, do one of the following:
   - Double-click the console icon on the desktop.
   - From the Microsoft Windows **Start** menu, select **All Programs** > **Veeam** > **Veeam Backup & Replication Console**.
   - Use the Microsoft Windows search to find the **Veeam Backup & Replication Console** program on the computer.
2. In the **Server** field, type the name or IP address of the backup server or select it from the list of recent connections. By default, the console connects to the backup server installed locally — localhost.
3. In the **Port** field, enter the port over which you want to connect to the backup server. The port number is set at the **Port Configuration** step of the setup wizard for Veeam Backup & Replication. By default, port 9392 is used.
4. In the **Username and Password** fields, enter credentials of the user account that you want to use to connect to the backup server. The user account must be added to the Local Users group on the backup server or a group of domain users who have access to the backup server.

   You can also select the **Use Windows session authentication** check box. In this case, you will log in to Veeam Backup & Replication using the account under which you are currently logged in to Microsoft Windows.
5. Click **Connect**.

   If multi-factor authentication (MFA) is enabled, after user clicks **Connect**, they will get the instruction how to set up MFA or have to enter a 6-digit confirmation code generated in the mobile authenticator application. For more information, see Multi-Factor Authentication.

To create a shortcut for the connection, click **Save shortcut**. You can create as many shortcuts as you need.

> **NOTE**
>
> If you create a shortcut for a connection, the credentials for this connection will be stored in the Windows Credentials Manager. The credentials are saved after the first successful login.

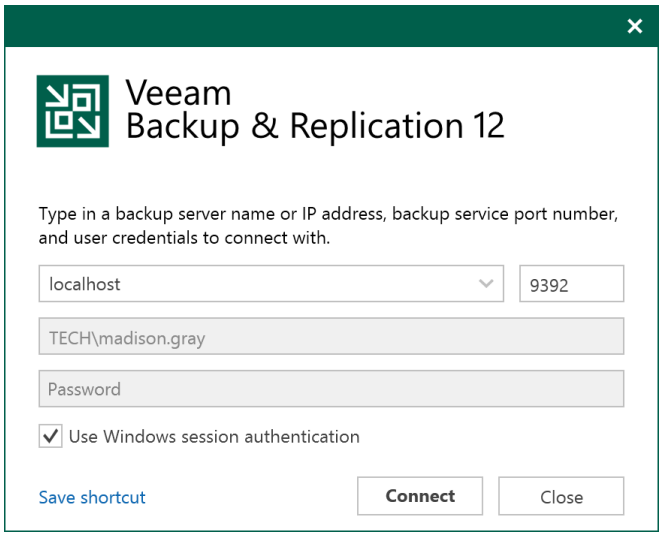

# Veeam Backup & Replication UI

The user interface of Veeam Backup & Replication is designed to let you quickly find commands that you need and perform data protection and disaster recovery tasks.

- Main Menu
- Navigation Pane
- Ribbon and Tabs
- Views
- Working Area
- Job Filter
- Footer
- Changing Color Theme
- Infrastructure Icons

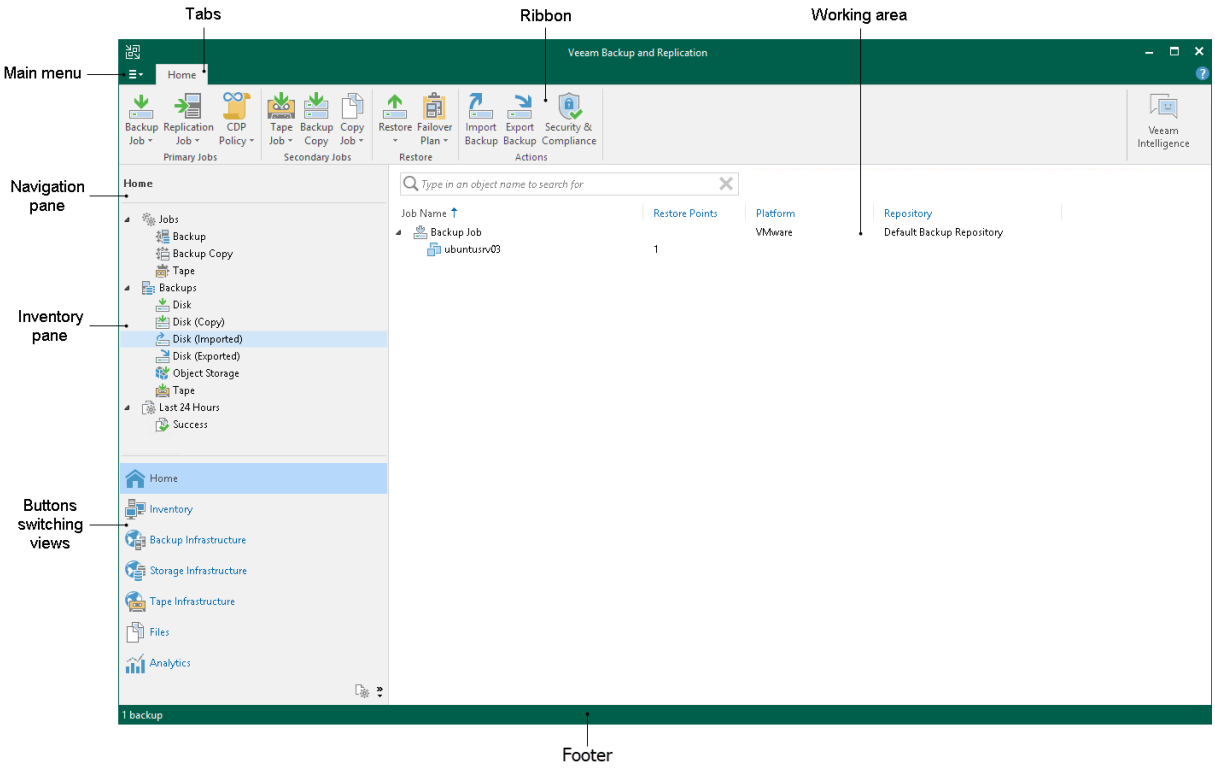


## TIP

To open online help, press [F1] in any Veeam Backup & Replication wizard or window. You will be redirected to the required section of the user guide.

# Main Menu

The main menu in Veeam Backup & Replication contains commands related to general application settings. You can perform the following operations using the main menu:

- Upgrade backup infrastructure components.
- Manage credentials.
- Manage cloud credentials.
- Manage passwords.
- Configure application settings.
- Set up user roles.
- Configure malware detection settings.
- Exclude VMs globally.
- Configure network traffic rules.
- Perform configuration backup and restore.
- Start PuTTy and Microsoft PowerShell consoles, and open a remote desktop connection to the backup server.
- Change the color theme.
- Work with licenses.
- View Veeam Backup & Replication help and export program logs.
- Exit Veeam Backup & Replication.

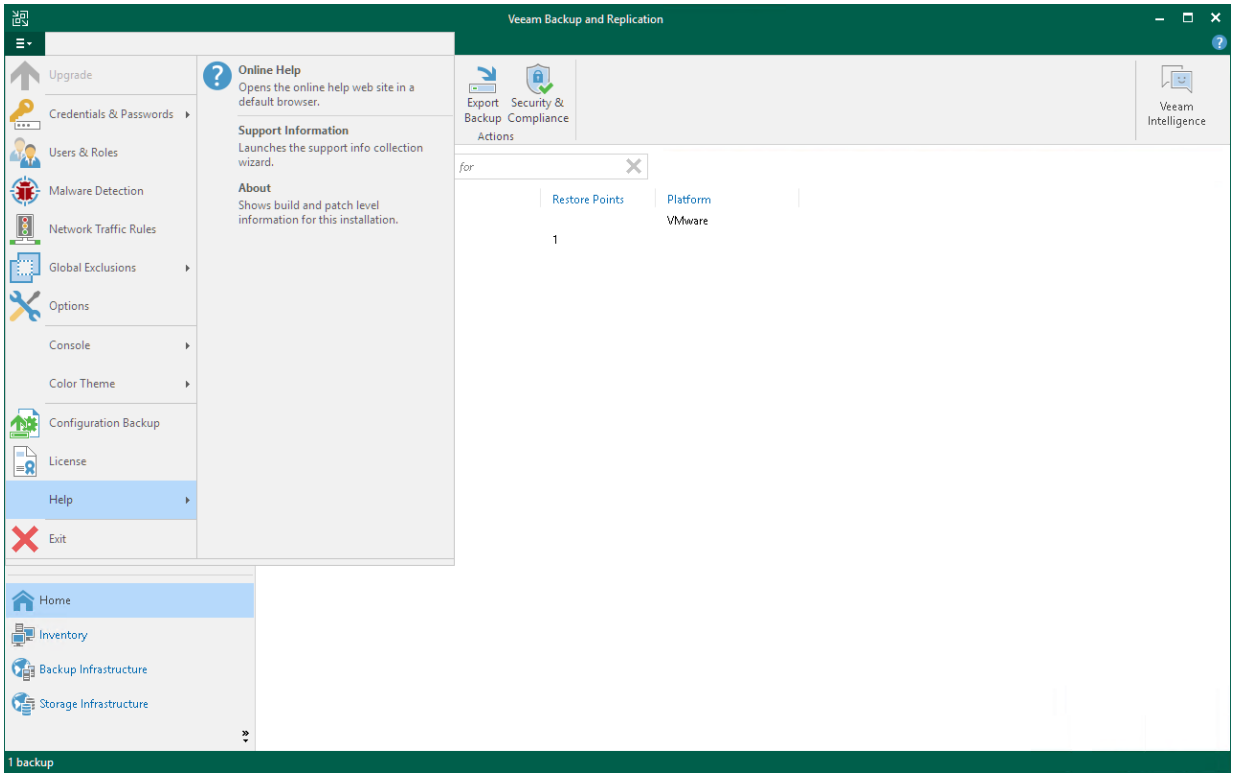

# Navigation Pane

The navigation pane, located on the left of the window, provides centralized navigation and lets you easily access Veeam Backup & Replication items organized in views.

The navigation pane consists of two areas:

- The upper pane, or the inventory pane, displays a hierarchy or list of items relevant for a specific view.

  Items displayed in the inventory pane differ depending on the active view. For example, in the **Backup Infrastructure** view, the inventory pane displays a list of backup infrastructure components — virtualization servers, backup proxies, backup repositories and so on. In the **Inventory** view, the inventory pane displays a list of servers added to the backup infrastructure.

- The lower pane contains a set of buttons that let you switch between views. For more information on views and how to show/hide a view button, see Views.

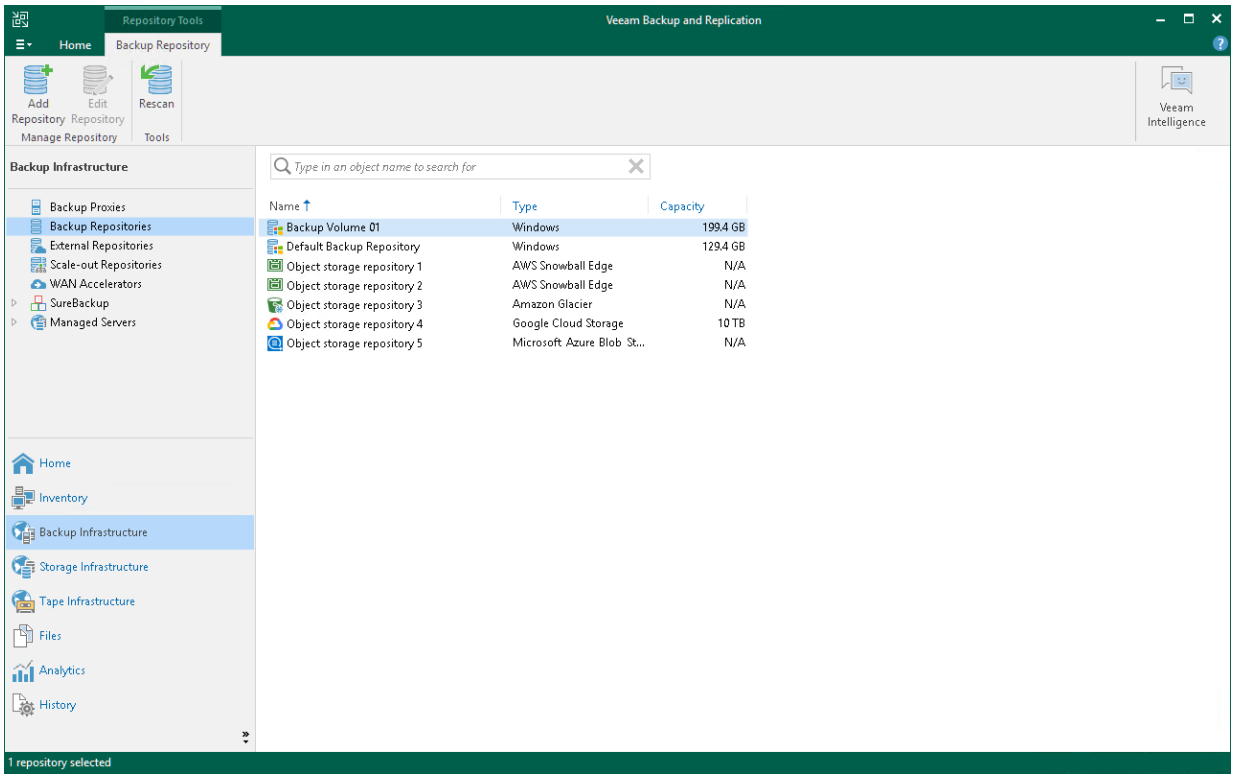

# Ribbon and Tabs

Operation commands in Veeam Backup & Replication are organized in logical groups and displayed under tabs on the ribbon. The ribbon is displayed at the top of the main application window.

On the ribbon, the following tabs are displayed:

- The **Home** tab provides quick access to the most common operations. It lets you configure different types of jobs, perform restore and import operations. This tab is always available, no matter which view is currently active.
- Other tabs contain commands specific for certain items and appear when these items are selected. For example, if you open the **Home** view and select a backup job in the working area, the **Job** tab containing buttons for operations with jobs will appear on the ribbon. If you open the **Files** view and select a file or folder, the **File Tools** tab containing buttons for operations with files will appear on the ribbon.

> **TIP**
>
> Commands for operations with items in Veeam Backup & Replication are also available from the shortcut menu.

You can minimize the ribbon. To do that, right-click the area with buttons on the ribbon and select **Minimize the Ribbon**. To restore the ribbon, right-click on the minimized ribbon and clear the **Minimize the Ribbon** option.

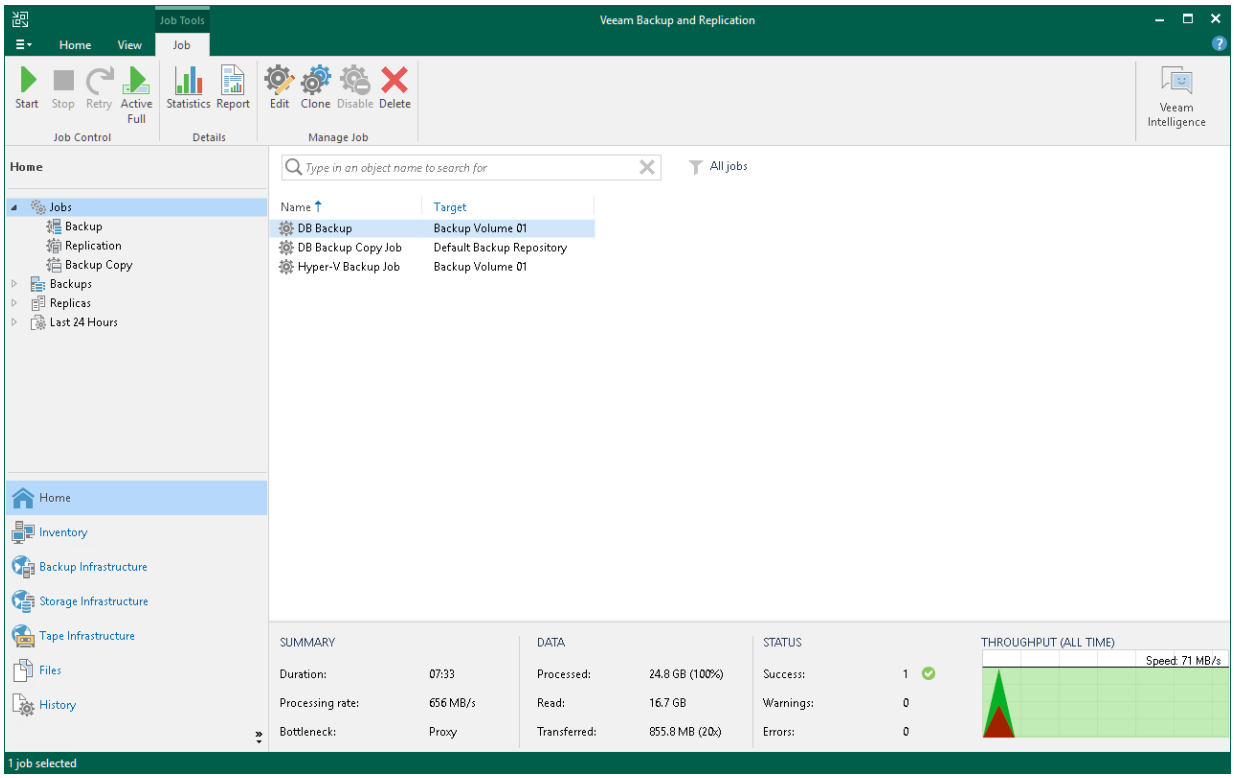

# Views

Veeam Backup & Replication displays its items in views. When you click the button of a specific view in the navigation pane, the view content is displayed in the working area of Veeam Backup & Replication.

Veeam Backup & Replication offers the following views:

- The **Home** view is intended for work with jobs. It also displays a list of created backups and replicas that can be used for various restore operations, and provides statistics for recently performed jobs. For more information about job statistics, see Reporting.
- The **Inventory** view displays the inventory of the virtual infrastructure. The inventory can be presented from different perspectives: **Computer**, **Storage**, **VM Folders, VM Tags** and **vCloud**. You can use this view to work with VMs, and VM containers or groups.
- The **Backup Infrastructure** view displays a list of backup infrastructure components: servers, hosts, backup proxies, backup repositories and so on. You can use this view for backup infrastructure setup — here you can configure backup infrastructure components that will be used for data protection and disaster recovery tasks.
- The **Storage Infrastructure** view displays a list of storage systems, volumes and snapshots. You can use this view to restore VM data from storage snapshots.
- The **Tape Infrastructure** view displays a hierarchy of tape libraries connected to the tape server. You can use this view to archive data to tapes and restore data from tapes.
- The **Cloud Connect Infrastructure** view displays components of the Veeam Cloud Connect infrastructure. This view can be used by SP to manage TLS certificates, configure cloud gateways and create accounts for users who plan to work with cloud resources.
- The **Files** view displays a file tree of servers added to the backup infrastructure. You can use this view for file copying operations.
- The **Analytics** view displays the following dashboards configured in Veeam ONE for monitoring the backup infrastructure and data protection operations in the virtual environment: Veeam Threat Center, Veeam Backup & Replication Overview, and Backup Heatmap. Dashboards in the view are available if you have Veeam ONE integrated with your Veeam Backup & Replication installation. For more information, see the Configuring Analytics View section.
- The **History** view displays statistics on operations performed with Veeam Backup & Replication. For more information, see the Viewing History Statistics section.

In some situations, some views may not be displayed. Consider the following:

- Right after installation, Veeam Backup & Replication displays only **Backup Infrastructure** and **History** views. To display other views, you must add at least one server or virtualization host to the backup infrastructure.
- Right after installation, Veeam Backup & Replication does not save changes that you make to the navigation pane or views: for example, if you resize panes, display or hide specific views. After you restart the Veeam Backup & Replication console, the main window settings are back to default ones. To save these settings, you must add at least one server or virtualization host to the backup infrastructure.
- Views can be shown as icons if they do not fit into the pane. To show the views in the full size, drag and drop the upper border of the pane.
- To display the **Cloud Connect Infrastructure** view, you must install a valid license that supports the Veeam Cloud Connect functionality.

- The **Analytics** view is available if you have the Microsoft Edge WebView2 Runtime component installed. The component is not installed for Microsoft Windows Server 2012 and 2012 R2 due to the version incompatibility, so the **Analytics** view is not available for the backup server running these Microsoft Windows versions.

You can hide views that you do not plan to use. For example, if you do not use tapes for data archiving, you can hide the **Tape Infrastructure** view.

To hide a view:

1. Click the arrow icon (») at the bottom of the navigation pane.
2. Click the view in the list.

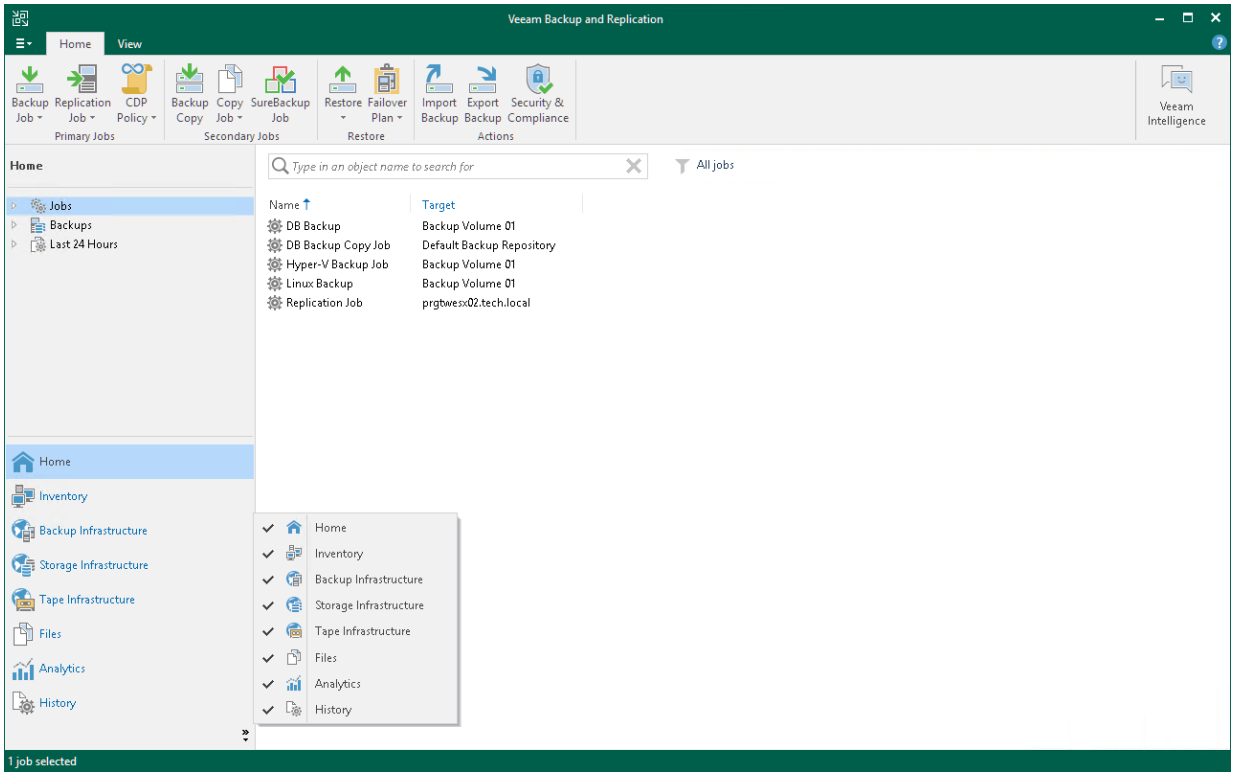

# Working Area

The working area of Veeam Backup & Replication displays a list of items relating to a specific view.

The working area looks different depending on the view that is currently active. For example, if you open the **History** view, the working area will display a list of job sessions and restore tasks performed with Veeam Backup & Replication. If you open the **Inventory** view, the working area will display a list of VMs that reside on servers connected to Veeam Backup & Replication.

Every item is described with a set of properties that are presented as column headers. You can click column headers to sort items by a specific property. For example, to sort VMs by the amount of provisioned storage space, click the **Provisioned Size** header.

To hide or display properties, right-click a column header and, in the opened menu, clear or select check boxes near property names.

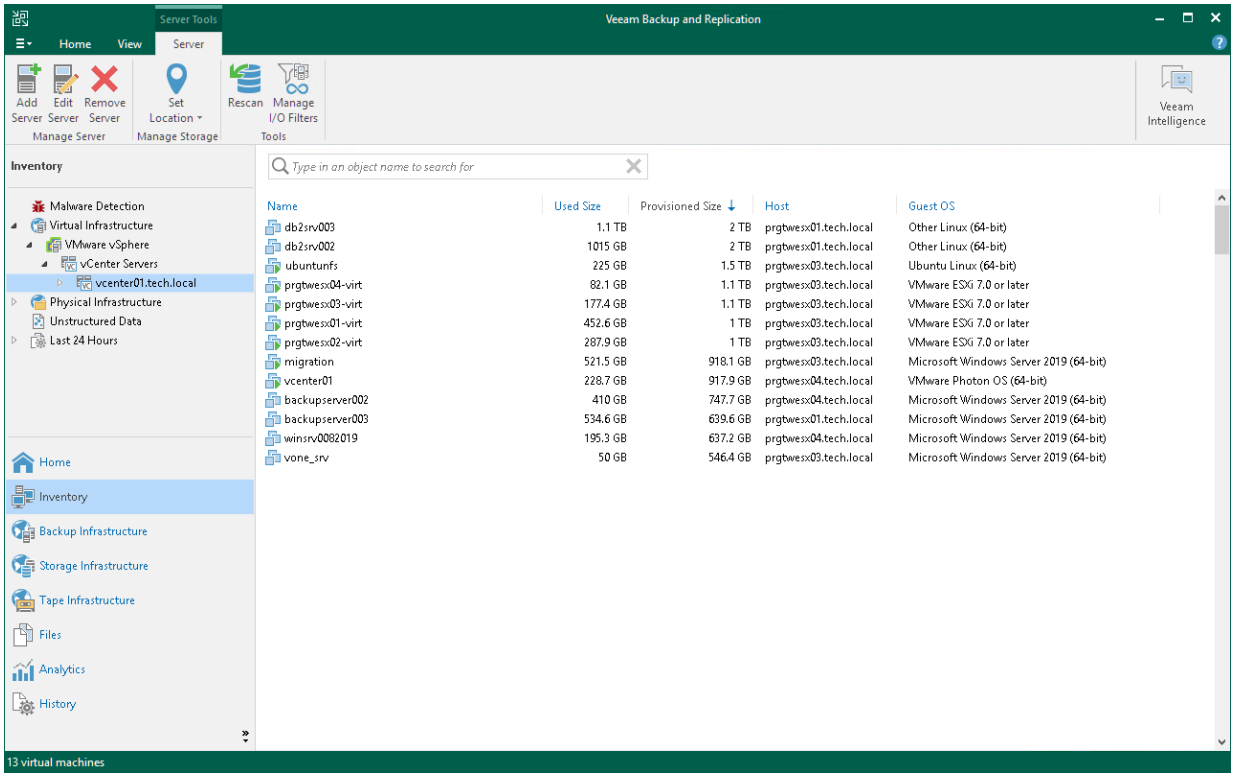

# Job Filter

A job filter allows you to filter jobs by different parameters. For example, you can create a filter that will show only VM backup copy jobs.

## Creating Job Filters

To create a filter, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Jobs** node. At the top of the working area, click the filter icon and then **Add new**. Alternatively, right-click the **Jobs** node and select **Add view**.
3. In the **Add View** window, specify a filter name. The name can be up to 200 characters long.
4. If you want to show this filter as a subnode of the **Jobs** node in the inventory pane, select the **Add this view to the job management tree** check box.
5. In the **Workload** section, select workloads which jobs must process:
   - **Application** – Veeam Backup & Replication will show Veeam Plug-in jobs and jobs in which SQL transaction log backup, Oracle archive log backup or PostreSQL WAL files backup is enabled.
   - **Cloud VM** – Veeam Backup & Replication will show jobs that process VMs stored in clouds.
   - **File** – Veeam Backup & Replication will show jobs that process files.
   - **Server** – Veeam Backup & Replication will show Veeam Agent jobs that process servers. For more information on processed computer types, see Selecting Protected Computer Type.
   - **Workstation** – Veeam Backup & Replication will show Veeam Agent jobs that process workstations. For more information on processed computer types, see Selecting Protected Computer Type.
   - **VM** – Veeam Backup & Replication will show jobs that process VMs.
6. In the **Job type** section, select job types:
   - **Backup** – backup jobs. For example, VM backup jobs or file backup jobs.
   - **Backup copy** – backup copy jobs. For example, VM backup copy jobs or file backup copy jobs.
   - **Copy** – copy jobs. For example, file copy jobs or VM copy jobs.
   - **Replication** – replication jobs.
   - **SureBackup** – SureBackup jobs.
   - **Tape** – tape jobs. For example, backup to tape jobs or file to tape jobs. For more information, see the Creating Backup to Tape Jobs and Creating File to Tape Jobs sections in the Veeam Backup & Replication User Guide.
7. In the **Last result** section, select statuses with which jobs must finish: *Success*, *Warning* or *Failed*.

   Jobs that that have never started are considered as *Failed*.
8. In the **Status** section, select states of jobs: *Running*, *Inactive* or *Disabled*.

   Backup copy jobs in the *Idle* state are considered Inactive.

9. If you want to show jobs with manually set **High priority** flag, select the **High priority job** check box. For more information on job priorities, see Job Priorities.
10. If you want to show jobs that were inactive for some period of time, select the **Last run is over N days ago** check box and specify the period in days.
11. If you want to show jobs whose names include specific keywords, select the **Job name includes** check box and enter keywords.

    To show jobs that include any of the specified keywords, separate these keywords by a semicolon without a space. For example, if you enter *"Backup Job;Daily"*, Veeam Backup & Replication will show all jobs that include *"Backup Job"* or *"Daily"* keywords in their names.

> **NOTE**
>
> Only the user who creates filters can access them — that is, other users cannot use these filters.

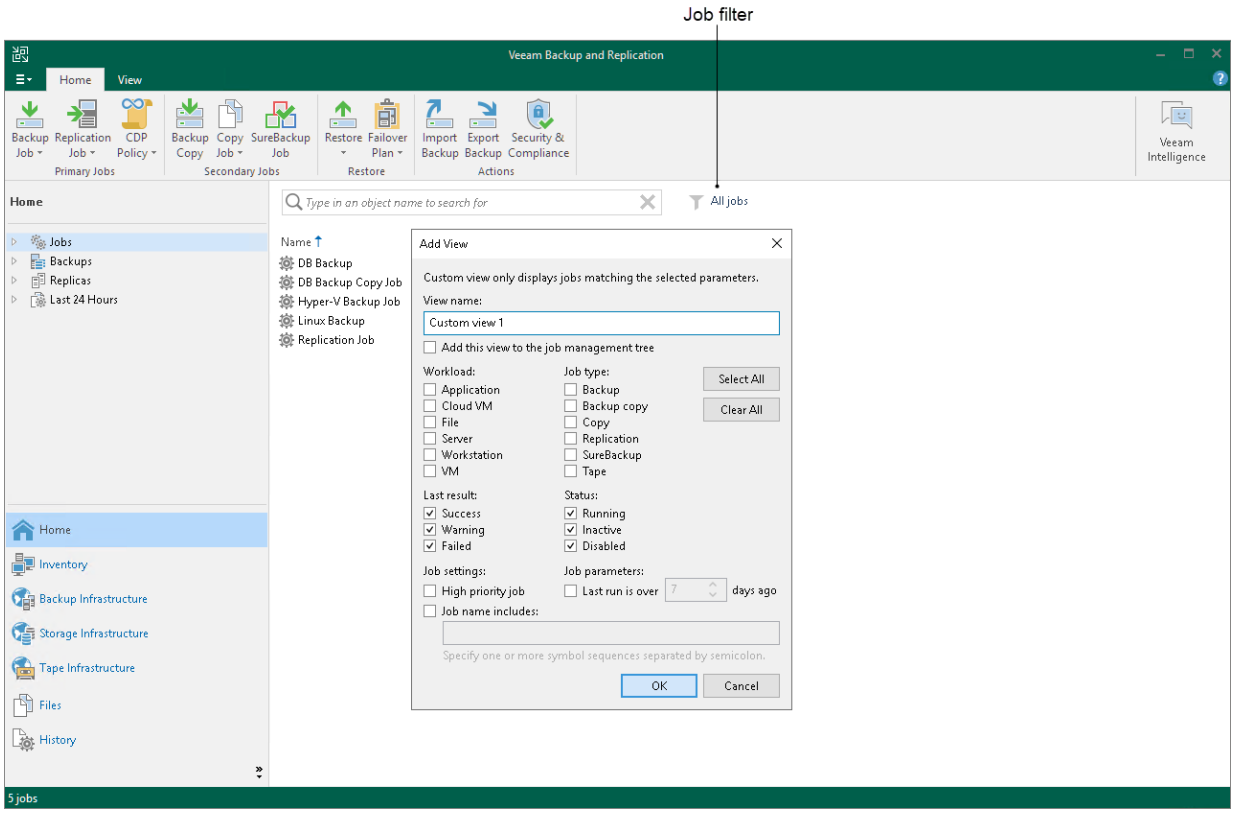


# Editing and Deleting Job Filters

To edit or delete a job filter, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Jobs** node.
3. At the top of the working area, click the filter icon .
4. Hover the mouse over a filter that you want to edit or delete.
5. Click the ellipsis button.
6. Select **Edit view** or **Delete view**.

If you have added the filter as a subnode of the **Jobs** node in the inventory pane, you can right-click the subnode and then click **Edit view** or **Delete view**.

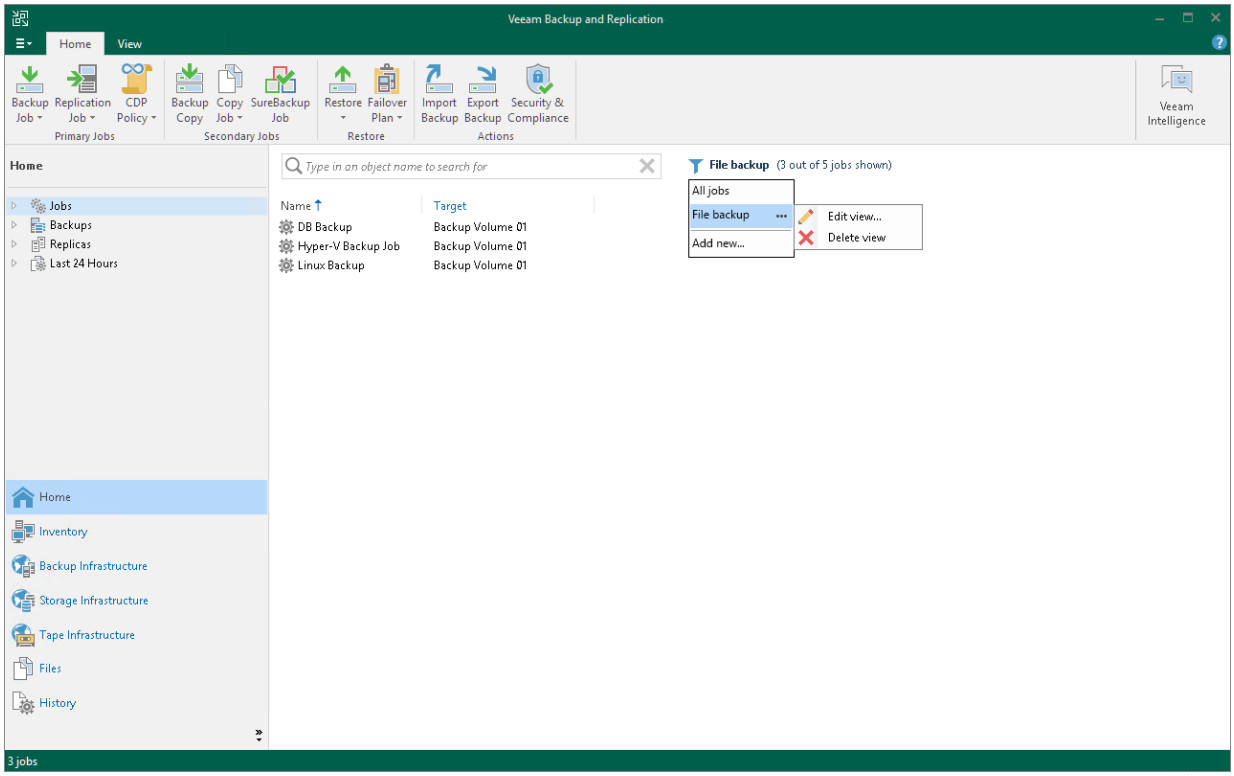

# Footer

The footer contains the following information, starting from left to right:

- The number of items you selected for further actions or the total number of items in the working area if no items are selected.
- The backup server which this Veeam Backup & Replication console is connected to. By default, the console connects to the backup server installed locally — localhost.
- The backup server build number.

  > **TIP**
  >
  > You can also check this installation build number in **Main Menu** > **Help** > **About**. Our build numbers follow a dotted x.x.x.x system. For more information on how to update your installation of Veeam Backup & Replication to the latest version, see Updating Veeam Backup & Replication 12.

- The information about your license edition and its expiration date.

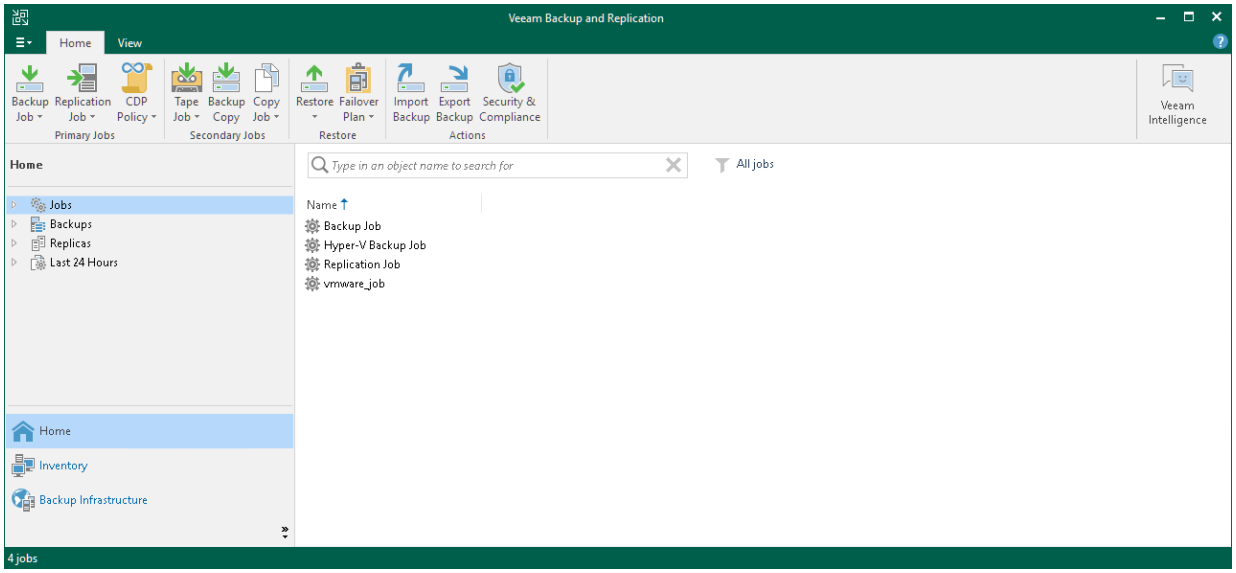

# Changing Color Theme

By default, Veeam Backup & Replication uses a '*Veeam Green*' color theme for the UI. If necessary, you can change the color theme. Changing the color theme can be helpful, for example, if you connect to different backup servers from one remote machine on which the Veeam Backup & Replication console is installed. In this case, you will be able to easily differentiate with which backup server you are currently working.

To change the color theme for Veeam Backup & Replication:

1. From the main menu, select **Color Theme**.
2. Choose one of color themes: *Grass Green*, *Veeam Green*, *Marine Blue*, *Ocean Graphite*.

> **NOTE**
>
> Color theme settings are applicable for a specific combination of a backup server and user account. For example, the color theme is initially set to the default one. You log on to the Veeam Backup & Replication console under some user account and change the color theme to **Marine Blue**. If you log on to the same backup server under the same account next time, the color theme will be set to **Marine Blue**. If you log on to the same backup server under another account, Veeam Backup & Replication will use the color theme that was previously set for this account — that is, the default color theme.

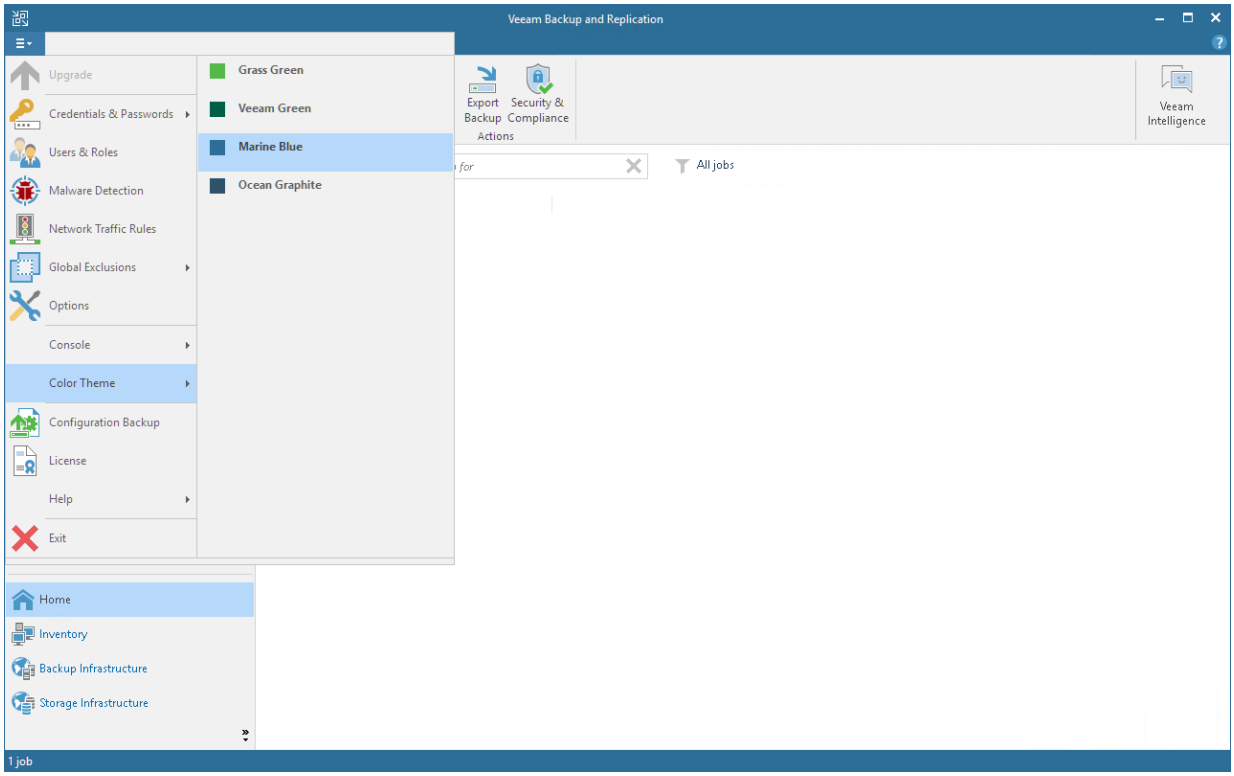

# Infrastructure Icons

The user interface icons display Veeam Backup & Replication infrastructure objects, jobs and their current state.

## Jobs

The following icons represent the jobs configured in Veeam Backup & Replication.

| Icon | Description |
|---|---|
| | Job |
| | High priority job |
| | Job in the disabled state |
| | Job in the running state |

## Job Nodes

The following icons represent the job nodes displayed in the inventory pane of the **Home** view.

| Icon | Description |
|---|---|
| | Backup jobs |
| | Replication jobs |
| | Backup copy jobs |
| | SureBackup jobs |
| | Tape jobs |
| | File copy jobs |
| | VM copy jobs |
| | Agent backup jobs |
| | Agent backup copy jobs |

# Backups

The following icons represent the backups created with Veeam Backup & Replication.

| Icon | Description |
|---|---|
| | Backup |
| | Encrypted backup |
| | Backup to tape |
| | Incomplete backup |

# Backup Nodes

The following icons represent the backup nodes displayed in the inventory pane of the **Home** view.

| Icon | Description |
|---|---|
| | Backups on disk |
| | Backup copies on disk |
| | Imported backups on disk |
| | Exported backups on disk |
| | Orphaned backups on disks |
| | Encrypted backups on disk |
| | Backups in external repository |
| | Encrypted backups in external repository |
| | Backups in cloud repository |
| | Orphaned backups in cloud repository |
| | Backups on tape |

| Icon | Description |
|---|---|
| | Encrypted backups on tape |
| | Backups in object storage |
| | Backup copies in object storage |
| | Encrypted backups in object storage |
| | Imported backups in object storage |
| | Orphaned backups in object storage |
| | VeeamZIP backups in object storage |
| | Backups in object storage with data archiving |
| | Backups in capacity tier |
| | Imported backups in capacity tier |
| | Orphaned backups in capacity tier |
| | VeeamZIP backups in capacity tier |
| | Imported backups in archive tier |
| | Orphaned backups in archive tier |
| | VeeamZIP backups in archive tier |

# Restore Points

The following icons indicate the states of the restore points in backups of all types.

| Icon | Description |
|---|---|
| | Full restore point |

| Icon | Description |
|---|---|
| | Incremental restore point |
| | Reverse incremental restore point |
| | Missing full restore point |
| | Missing incremental restore point |
| | Missing reverse incremental restore point |
| | Infected restore point |

# Scale-Out Backup Repository Restore Points

The following icons indicate the states of the restore points in the scale-out backup repository.

| Icon | Description |
|---|---|
| | Full restore point; on performance tier only |
| | Full restore point; on performance tier and offloaded to capacity tier |
| | Full restore point; on capacity tier only |
| | Full restore point; on capacity tier and partially downloaded to performance tier<br>(This can happen when download to the performance tier fails) |
| | Incremental restore point; on performance tier only |
| | Incremental restore point; on performance tier and offloaded to capacity tier |
| | Incremental restore point; on capacity tier only |
| | Incremental restore point; on capacity tier and partially downloaded to performance tier<br>(This can happen when download to the performance tier fails) |
| | Rollback restore point; on performance tier only |

| Icon | Description |
|---|---|
| | Rollback restore point; on performance tier and offloaded to capacity tier |
| | Rollback restore point; on capacity tier only |
| | Rollback restore point; on capacity tier and partially downloaded to performance tier<br>(This can happen when download to the performance tier fails) |
| | Rollback restore point; on archive tier only |

# Replicas

The following icons represent the replicas configured in Veeam Backup & Replication and displayed in the inventory pane of the **Home** view.

| Icon | Description |
|---|---|
| | Replicas in the active state |
| | Replicas in the ready state |
| | Failover plan |
| | Cloud failover plan |

# Microsoft Entra ID Backups

The following icons represent Microsoft Entra ID tenants and backups created by Veeam Backup for Microsoft Entra ID.

| Icon | Description |
|---|---|
| | Microsoft Entra ID tenant / Microsoft Entra ID tenant backup |
| | Microsoft Entra ID audit and sign-in log backup |
| | Microsoft Entra ID tenant backup is corrupted |

# VMware vSphere Virtual Infrastructure

The following icons represent the hypervisor objects of VMware vSphere workloads and servers added to the backup infrastructure.

| Icon | Description |
|---|---|
| | vCenter |
| | vCenter with no license |
| | VM |
| | Deleted VM |
| | VM in the running state |
| | VM in the paused state |
| | VM in the disabled state |
| | VM in the failback state |
| | VM is corrupted |
| | VM is not consistent |
| | Infected VM |
| | VM in the upgrade state |
| | VM with the vSphere Fault Tolerance feature enabled |
| | Cluster |
| | Cluster with no license |
| | Datacenter |
| | Datastore |

| Icon | Description |
|---|---|
| | Datastore cluster |
| | ESXi host |
| | ESXi host with no license<br>(This can happen when you have Per-Socket licensing) |
| | ESXi host in the standby state |
| | Provider VDC |
| | Resource pool |

# VMware Cloud Director

The following icons represent the VMware Cloud Director objects added to Veeam Backup & Replication.

| Icon | Description |
|---|---|
| | VMware Cloud Director server |
| | Cloud Director organization |
| | Cloud Director organization VDC |
| | Cloud Director vApp |
| | Cloud Director vApp in the running state |
| | Cloud Director vApp in the paused state |
| | Cloud Director vApp is not consistent |
| | Cloud Director vApp in the mounting state |
| | Cloud Director vApp in the failback state |
| | Cloud Director vApp is infected |

| Icon | Description |
|---|---|
| | Cloud Director vApp is corrupted |
| | Cloud Director vApp has an error |

# Physical Infrastructure

The following icons represent the physical infrastructure objects added to Veeam Backup & Replication.

| Icon | Description |
|---|---|
| | Protection group |
| | Protection group in the disabled state |
| | Protection group is outdated |
| | Protection group is unavailable |
| | Unmanaged protection group |
| | Veeam Agent computer |
| | Veeam Agent computer is corrupted |
| | Encrypted Veeam Agent computer |
| | Infected Veeam Agent computer |
| | Microsoft Windows-based Veeam Agent computer |
| | Microsoft Windows-based Veeam Agent computer is corrupted |
| | Linux-based Veeam Agent computer |
| | Linux-based Veeam Agent computer is corrupted |
| | macOS-based Veeam Agent computer |

| Icon | Description |
|---|---|
| | macOS-based Veeam Agent computer is corrupted |
| | IBM AIX-based Veeam Agent computer |
| | IBM AIX-based Veeam Agent computer is corrupted |
| | Oracle Solaris-based Veeam Agent computer |
| | Oracle Solaris-based Veeam Agent computer is corrupted |

# Unstructured Data

The following icons represent the unstructured data sources added to Veeam Backup & Replication.

| Icon | Description |
|---|---|
| | File server |
| | NAS filer |
| | File share |
| | SMB file server |
| | NFS file server |
| | Amazon S3 object storage |
| | Microsoft Azure Blob storage |
| | Microsoft Azure Data Lake storage |
| | S3 compatible object storage |

# Backup Proxies

The following icons represent the backup proxies added to Veeam Backup & Replication.

| Icon | Description |
|---|---|
| | VMware backup proxy |
| | VMware backup proxy is disabled |
| | Hyper-V backup proxy |
| | Hyper-V backup proxy is disabled |
| | Backup proxy is busy |
| | Backup proxy is disabled |
| | Backup proxy is outdated |
| | Backup proxy is unavailable |
| | General-purpose backup proxy (for file shares) |
| | General-purpose backup proxy is disabled |
| | General-purpose backup proxy is outdated |
| | General-purpose backup proxy is unavailable |
| | VMware CDP proxy |
| | VMware CDP proxy is disabled |
| | VMware CDP proxy is outdated |
| | VMware CDP proxy is unavailable |

# Backup Repositories and Scale-Out Backup Repositories

The following icons represent the backup repositories and scale-out backup repositories added to Veeam Backup & Replication.

| Icon | Description |
|---|---|
| | Backup repositories |
| | External repositories |
| | Cloud repositories |
| | Scale-out backup repository |
| | Repository is busy |
| | Repository is full |
| | Repository is outdated |
| | Repository is unavailable |
| | Windows server repository |
| | Windows server repository in the maintenance mode |
| | Windows server repository in the sealed mode |
| | Linux server repository |
| | Linux server repository in the maintenance mode |
| | Linux server repository in the sealed mode |
| | Hardened repository |
| | Hardened repository in the maintenance mode |
| | Hardened repository in the sealed mode |

| Icon | Description |
| --- | --- |
| | SMB share repository |
| | SMB share repository in the maintenance mode |
| | SMB share repository in the sealed mode |
| | NFS share repository |
| | NFS share repository in the maintenance mode |
| | NFS share repository in the sealed mode |
| | Veeam Data Cloud Vault |
| | Veeam Data Cloud Vault in the maintenance mode |
| | Veeam Data Cloud Vault in the sealed mode |
| | S3 compatible repository |
| | S3 compatible repository in the maintenance mode |
| | S3 compatible repository in the sealed mode |
| | S3 compatible repository with data archiving repository |
| | S3 compatible repository with data archiving repository in the maintenance mode |
| | S3 compatible repository with data archiving repository in the sealed mode |
| | Amazon S3 repository |
| | Amazon S3 repository in the maintenance mode |
| | Amazon S3 repository in the sealed mode |
| | Amazon Glacier repository |

| Icon | Description |
| --- | --- |
| | Amazon Glacier repository in the maintenance mode |
| | Amazon Glacier repository in the sealed mode |
| | AWS Snowball Edge Storage repository |
| | AWS Snowball Edge Storage repository in the maintenance mode |
| | AWS Snowball Edge Storage repository in the sealed mode |
| | IBM Cloud repository |
| | IBM Cloud repository in the maintenance mode |
| | IBM Cloud repository in the sealed mode |
| | Google Cloud Storage repository |
| | Google Cloud Storage repository in the maintenance mode |
| | Google Cloud Storage repository in the sealed mode |
| | Microsoft Azure Blob Storage repository |
| | Microsoft Azure Blob Storage repository in the maintenance mode |
| | Microsoft Azure Blob Storage repository in the sealed mode |
| | Microsoft Azure Archive Storage repository |
| | Microsoft Azure Archive Storage repository in the maintenance mode |
| | Microsoft Azure Archive Storage repository in the sealed mode |
| | Microsoft Azure Data Box Storage repository |
| | Microsoft Azure Data Box Storage repository in the maintenance mode |

| Icon | Description |
| --- | --- |
| | Microsoft Azure Data Box Storage repository in the sealed mode |
| | 11:11 Cloud Object Storage repository |
| | 11:11 Cloud Object Storage repository in the maintenance mode |
| | 11:11 Cloud Object Storage repository in the sealed mode |
| | Wasabi repository |
| | Wasabi repository in the maintenance mode |
| | Wasabi repository in the sealed mode |
| | Infinidat InfiniGuard repository |
| | Infinidat InfiniGuard repository in the maintenance mode |
| | Infinidat InfiniGuard repository in the sealed mode |
| | Quantum DXi repository |
| | Quantum DXi repository in the maintenance mode |
| | Quantum DXi repository in the sealed mode |
| | ExaGrid repository |
| | ExaGrid repository in the maintenance mode |
| | ExaGrid repository in the sealed mode |
| | Fujitsu ETERNUS CS800 repository |
| | Fujitsu ETERNUS CS800 repository in the maintenance mode |
| | Fujitsu ETERNUS CS800 repository in the sealed mode |

| Icon | Description |
|---|---|
| | HPE StoreOnce repository |
| | HPE StoreOnce repository in the maintenance mode |
| | Dell Data Domain repository |
| | Dell Data Domain repository in the maintenance mode |
| | Dell Data Domain repository in the sealed mode |

# External Repositories

The following icons represent the external repositories added to Veeam Backup & Replication.

| Icon | Description |
|---|---|
| | Amazon S3 external repository |
| | Microsoft Azure Blob Storage external repository |
| | Google Cloud Storage external repository |
| | External repository is unavailable |

# WAN Accelerators

The following icons represent the WAN accelerators added to Veeam Backup & Replication.

| Icon | Description |
|---|---|
| | WAN accelerator |
| | WAN accelerator is busy |
| | WAN accelerator has an error |
| | WAN accelerator is outdated |

# SureBackup

The following icons represent the SureBackup objects.

| Icon | Description |
|---|---|
| | Application group |
| | Virtual lab |

# Managed Servers

For more information about the icons that represent the hypervisor objects of VMware vSphere workloads and servers, see the VMware vSphere Virtual Infrastructure section.

| Icon | Description |
|---|---|
| | Microsoft Windows server |
| | Linux host |
| | Server is unavailable |
| | Server is outdated |
| | Backup server |
| | Domain |

# Tape Infrastructure

For more information about the icons that represent the tape infrastructure components, see Tape Infrastructure Icons.

# Veeam Backup & Replication Services

Veeam Backup & Replication uses the following services:

- **Veeam AHV Service** (VeeamAHVSvc) enables interaction between Veeam Backup & Replication and Nutanix AHV infrastructure.
- **Veeam AWS Service** (VeeamAWSSvc) enables interaction between Veeam Backup & Replication and AWS infrastructure.
- **Veeam Azure Service** (VeeamAzureSvc) enables interaction between Veeam Backup & Replication and Microsoft Azure infrastructure.
- **Veeam Backup Server RESTful API Service** (VeeamBackupRESTSvc) provides access to Veeam Backup & Replication by using the web API. For more information, see the REST API Reference.
- **Veeam Backup Service** (VeeamBackupSvc) is a Microsoft Windows service that coordinates all operations performed by Veeam Backup & Replication such as backup, replication, recovery verification and restore tasks. This service runs on the backup server.
- **Veeam Backup Update Service** (VeeamBackupUpdateSvc) enables the update of backup servers. The service runs locally and does not connect to the internet. It only requires a connection to the configuration database server of the backup server.
- **Veeam Backup VSS Integration Service** (VeeamFilesysVssSvc) manages Microsoft VSS snapshots used for NAS, file share and Veeam Agent backups.
- **Veeam Broker Service** (VeeamBrokerSvc) interacts with the virtual infrastructure to collect and cache the virtual infrastructure topology. Jobs and tasks query information about the virtual infrastructure topology from the broker service, which accelerates job and task performance.
- **Veeam CDP Coordinator Service** (VeeamBackupCdpSvc) communicates with VMware vCenter, assigns continuous data protection (CDP) tasks and manages the infrastructure components involved in CDP. For more information on CDP, see Continuous Data Protection (CDP).
- **Veeam CDP Proxy Service** (VeeamCdpProxySvc) sends and receives virtual machine data during CDP.
- **Veeam Cloud Connect Service** (VeeamCloudSvc) provides transparent connection to cloud resources over the secure SSL/TLS channel.
- **Veeam Cloud Gateway Service** (VeeamGateSvc) manages data exchange with Veeam Cloud Connect users. For more information on Veeam Cloud Connect, see the Veeam Cloud Connect Guide.
- **Veeam Data Analyzer Service** (VeeamDataAnalyzerSvc) analyzes the malware detection metadata for potential threats.
- **Veeam Data Mover Service** (VeeamTransportSvc) sends and receives protected data during backup, replication and restore processes.
- **Veeam Distribution Service** (VeeamDistributionSvc) distributes the Veeam Agent setup packages to the protected computers. For more information, see the Veeam Agent Management Guide.
- **Veeam Deployment Service** (VeeamMBPDeploymentService) enables installing, updating and configuring Veeam Service Provider Console components.
- **Veeam Explorers Recovery Service** (VeeamExplorersRecoverySvc) executes data recovery workflows on behalf of some Veeam Explorers. This service runs on the mount server associated with the backup repository. For more details on which Explorers use this service and how they use it, see the Veeam Explorers User Guide.

- **Veeam GCP Service** (VeeamGCPSvc) enables interaction between Veeam Backup & Replication and Google Cloud infrastructure.
- **Veeam Guest Agent** (VeeamGuestHelper) performs application-aware processing, file system indexing and enables guest OS file restore to the original location.
- **Veeam Guest Catalog Service** (VeeamCatalogSvc) manages guest OS file system indexing for VMs and replicates system index data files to enable search through guest OS files. Index data is stored in the Veeam Backup Catalog — a folder on the backup server. The Veeam Guest Catalog Service running on the backup server works in conjunction with search components installed on Veeam Backup Enterprise Manager and (optionally) a dedicated Microsoft Search Server.
- **Veeam Installer Service** (VeeamDeploySvc) installs, updates or removes Veeam services when you add, update or remove backup infrastructure components. For more details, see Veeam Installer Service.
- **Veeam Kubernetes Service** (VeeamKastenSvc) enables integration between Veeam Backup & Replication and Kubernetes Infrastructure via Kasten K10.
- **Veeam KVM Service** (VeeamRHVSvc) enables interaction between Veeam Backup & Replication and oVirt KVM infrastructure.
- **Veeam Log Shipping Service** (VeeamLogShipper) copies Microsoft SQL Server, Oracle and PostgreSQL transaction logs to the backup repository. In addition, Veeam Explorer for Microsoft SQL Server uses this service to restore transaction logs if the administrative share (for example, `\\myserver\ADMIN$`) is unavailable.
- **Veeam Management Agent Service** (VeeamManagementAgentSvc) collects data from Veeam products discovered on the machine.
- **Veeam Mount Service** (VeeamMountSvc) mounts backups and replicas for file-level access, browsing the VM guest file system and restoring VM guest OS files and application items to the original location.
- **Veeam ONE Agent** (VeeamOneAgentSvc) enables remediation actions and communication between Veeam ONE and monitored Veeam Backup & Replication servers.
- **Veeam Oracle Restore Service** (Veeam.Oracle.Service) enables item-level recovery of Oracle data for Windows machines. This runtime component is deployed on the target machine or staging server.
- **Veeam PVE Service** (VeeamPVESvc) enables interaction between Veeam Backup & Replication and Proxmox Virtual Environment.
- **Veeam SQL Agent Service** (Veeam.SQL.Agent.Service) enables item-level recovery of Microsoft SQL Server data. This persistent component is deployed on the target machine or staging server if the administrative share (for example, `\\myserver\ADMIN$`) is unavailable. For more information, see the How Data Recovery Works section of the Veeam Explorers User Guide.
- **Veeam SQL Restore Service** (Veeam.SQL.Service) enables item-level recovery of Microsoft SQL Server data. This runtime component is deployed on the target machine or staging server if the administrative share (for example, `\\myserver\ADMIN$`) is available. For more information, see the How Data Recovery Works section of the Veeam Explorers User Guide.
- **Veeam Threat Hunter Service** (VeeamThreatHunterSvc) performs signature-based malware detection scans. For more information, see Veeam Threat Hunter for Scan Backup.
- **Veeam Tape Access Service** (VeeamTapeSvc) provides access to tape devices connected to this tape server for Veeam Backup & Replication tape jobs.
- **Veeam vPower NFS Service** (VeeamNFSSvc) implements vPower NFS server that enables running virtual machines directly from backup files.

- **Veeam VSS Hardware Provider Service** (VeeamVssProviderSvc) extends the Microsoft VSS to enable Veeam Agent backup from storage snapshots.
- **Veeam WAN Accelerator Service** (VeeamWANSvc) optimizes the network traffic consumption by identifying and caching repeatedly transferred data.

# Veeam Installer Service

The Veeam Installer Service is a Windows or Linux service that manages Veeam components and services in the following cases:

- When you add, update or remove a physical or virtual machine as a managed server in the Veeam Backup & Replication console. Depending on the role selected for the server, the Veeam Installer Service deploys, updates, and removes required services such as the Veeam Data Mover Service, the Veeam Mount Service, and so on.

  For hardened repository, the Veeam Installer Service deploys the Veeam Data Mover Service and Veeam Immutability Service. For more information about these services, see How Immutability Works.

- When you discover protected machines or perform guest processing tasks for them. The Veeam Installer Service deploys, updates, and removes persistent agent components for guest processing tasks. For more information, see Persistent Agent Components.

By default, the service uses port 6160. For more information, see Ports.

> **NOTE**
>
> The Veeam Installer Service also manages public keys of self-signed TLS certificates generated by Veeam Backup & Replication. These keys verify installation packages to prevent managed servers and protected machines from being compromised. When Veeam Backup & Replication sends request to deploy or update components and services, the Veeam Installer Service validates the public certificate, the checksum and the digital signature of the installation package.

## Veeam Installer Service for Windows

The service can be installed in one of the following ways:

- Automatic installation by Veeam Backup & Replication when you add a Microsoft Windows machine as a managed server in the Veeam Backup & Replication console.
- Manual installation by running the `VeeamInstallerSvc.msi` file located at `C:\Program Files\Veeam\Backup and Replication\Backup\Packages`.
- Automatic installation by using a logon script.
- Automatic installation by using domain group policies.
- Installation by using the standard third-party software distribution tool.

## Veeam Installer Service for Linux

The service is installed in the following situations:

- Automatic installation by Veeam Backup & Replication when you add a Linux machine as a managed server in the Veeam Backup & Replication console.
- Automatic installation by Veeam Backup & Replication when you create a protection group with Linux machines. For more information, see Protection Groups in the Veeam Agent Management Guide.
- When you manually install persistent agents components on Linux machines. For more information, see Installing Persistent Agent Components on Linux VMs.

To deploy the Veeam Installer Service on Linux machines, Veeam Backup & Replication uses the SSH connection. After deployment, the Veeam Installer Service will communicate with backup infrastructure components and get updates from the backup server without the SSH connection.

The Veeam Installer Service for Linux (`veeamdeploymentsvc`) runs the following processes:

- Deployer — deploys, updates, and removes required services and components. Runs with root permissions.
- Web listener — listens to port 6160 for new request from Veeam Backup & Replication to deploy or update components and services. Also, validates the public certificate, the checksum and the digital signature of the installation package. By default, runs with root permissions as a child deployer process. For the hardened repository, runs with reduced permissions.
- Watchdog — monitors deployer and web listener processes and restarts them if required. Runs with root permissions as a child deployer process.

# Installing Veeam Installer Service with GPO

This topic describes how you can deploy the Veeam Installer service setup file to remote computers using GPO. You must create an MST file with custom configuration parameters and use this MST file to deploy the service on remote computers.

# Step 1. Unpack Service Setup File

Unpack the content of the service setup file:

1. Obtain the necessary version of the service setup file.

   The service setup files reside in the `C:\Program Files\Veeam\Backup and Replication\Backup\Packages` folder on the machine where Veeam Backup & Replication is installed.

2. Perform installation of the service in the administrator mode to unpack the content of the setup file:

   a. In the command prompt, run a command to start the wizard.

      For example, for the VeeamInstaller service, the command is:

      ```
      msiexec /a <...>/VeeamInstallerSvc.msi
      ```

   b. At the **Network Location** step of the wizard, specify a directory to which setup files must be unpacked.

   c. Click **Install**.

3. Check the output directory and make sure it includes the `.msi` file and the `Windows` folder.

4. Copy the unpacked files to a network share.

   The network share must be accessible from all machines on which you want to deploy the service. Make sure you set at least `Read` permissions on the files.

# Step 2. Create MST Configuration File

Create an MST configuration file:

1. In the output directory, open the service setup file and edit it with Orca.

   For details on Orca, see Windows Dev Center.

2. In the menu, choose **Transform** > **New Transform**.

3. In the **Tables** pane, click **Property**.
4. Add the following properties to the table and set their values to 1 to accept the terms and proceed with installation:
   - `ACCEPT_THIRDPARTY_LICENSES` — specifies if you want to accept the terms of the license agreement for the 3rd party components.
   - `ACCEPT_REQUIRED_SOFTWARE` — specifies if you want to accept all required software license agreements.
   - `ACCEPT_LICENSING_POLICY` — specifies if you want to accept the Veeam licensing policy.
   - `ACCEPT_EULA` — specifies if you want to accept the terms of the Veeam license agreement.
5. In the menu, choose **Transform** > **Generate Transform**.
6. Save the MST file with configuration details.
7. Close Orca.
8. Copy the MST to a network share.

   The network share must be accessible from all machines on which you want to deploy the service. Make sure you set at least `Read` permissions on the files.

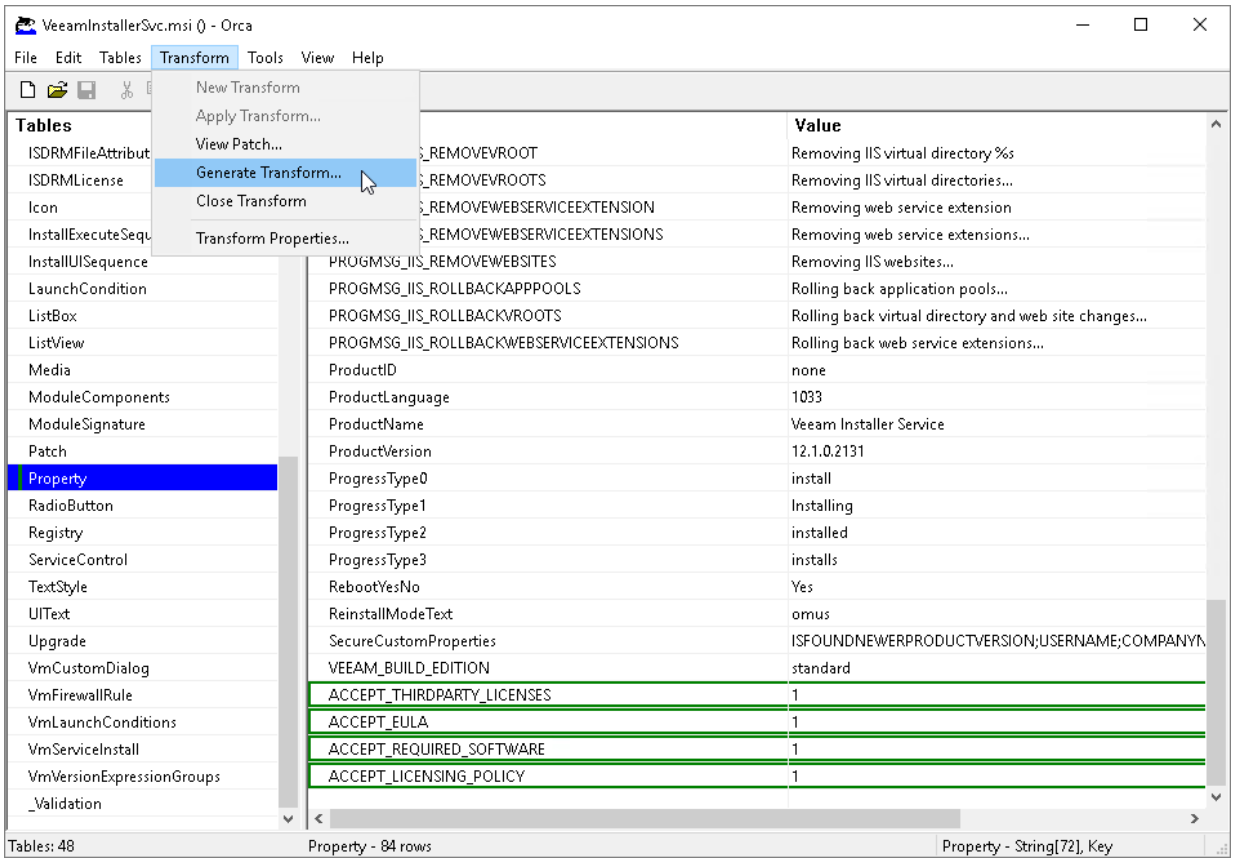


# Step 3. Create Group Policies

Create a Group Policy that will install and configure the service on the machines:

1. Log on to a domain controller.
2. Open the Group Policy Management Console.

3. Right-click the organizational unit which includes computers on which the service must be deployed, and choose to create a new Group Policy Object.
4. Right-click the Group Policy Object and choose **Edit**.
5. In the left pane of the Group Policy Management Editor, expand **Computer Configuration** > **Policies** > **Software Settings**.
6. Right-click **Software Installation** and select **New** > **Package**.
7. In the **Open** window, point to the service setup file located on the network share.
8. In the **Deploy Software** window, choose the **Advanced** deployment method.
9. Open the **Modifications** tab, click **Add** and choose the `MST` file located on the network share.
10. Click **OK**.
11. In the left pane of the Group Policy Management Editor, expand **Computer Configuration** > **Policies** > **Administrative Templates** > **System** > **Logon**.
12. Right-click the **Always wait for the network at computer startup and logon** policy setting and choose **Edit**.
13. In the policy setting window, select **Enabled** and click **OK**.
14. Close the Group Policy Management Editor.

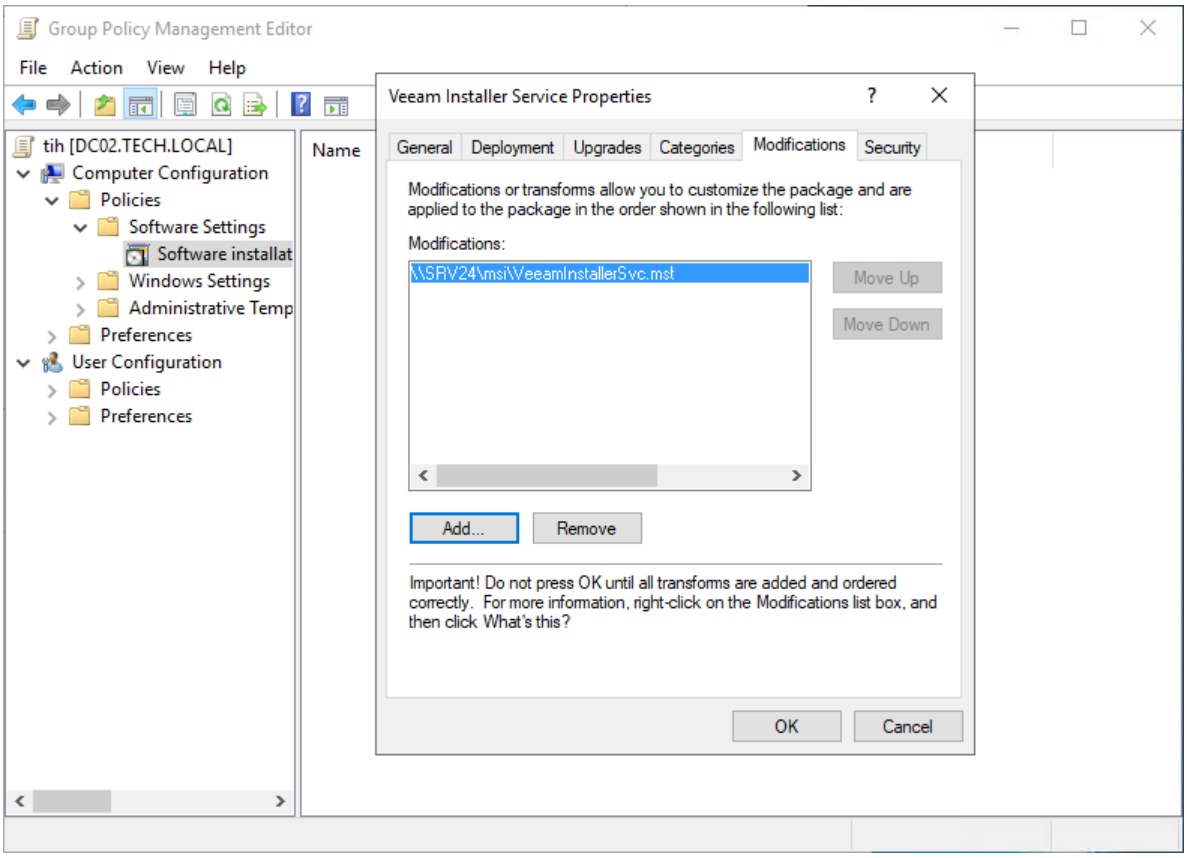


# Step 4. Apply Group Policies to Client Computers

Apply the created Group Policy to client computers.

# Veeam Intelligence

Veeam Intelligence is a chatbot that helps with common issues and questions related to Veeam products. It is trained on Veeam technical documentation to provide accurate answers. You can communicate with Veeam chatbot in any language you want and create both simple and complex inquiries.

> **IMPORTANT**
>
> Consider the following:
>
> - You must have a paid license and support contract or Evaluation license to use Veeam Intelligence. If you have a Community (free) Edition or NFR license, Veeam Intelligence will not be available.
> - Only currently installed license is used for authentication. If the license is changed, you must close the chatbot window and open it again.
> - We do not recommend to share any confidential information when using Veeam Intelligence as the queries are sent outside your organization. Veeam takes no responsibility for the accuracy of the information that chatbot provides.
> - Veeam Intelligence is under constant development. For more information on Veeam Intelligence updates, see this Veeam KB article.

## Limitations of Veeam Intelligence

Veeam Intelligence has the following limitations:

- There is a limit of up to 30 questions per 24 hours for each license. If this limit is reached, Veeam Intelligence will stop processing further questions. You can continue using it after the 24-hour timeout expires.
- Veeam Intelligence is not available on Windows 2012. It requires Windows 2012 R2 or later.
- To use Veeam Intelligence, you must have an available internet connection on the server that runs Veeam Backup & Replication Console.

## Using Veeam Intelligence

To start a new conversation:

1. Click **Veeam Intelligence** button located in the upper right corner of the Veeam Backup & Replication window.
2. In the chat window, type the question. Consider the following examples:
   - How to create a Replication Job?
   - Where are Veeam Backup & Replication logs stored?
   - How to change license?
   - How can I use Veeam Backup & Replication to make sure my VMs are regularly backed up, easily recovered, and securely stored?

   You can move the chat window and position it anywhere within the UI. The window remains open alongside any wizard.
3. Click **Send** button or press **Enter** to send your question.

4. When you do not need AI Assistant, you can do the following:
    - Click **Minimize** to hide the chatbot window and preserve the conversation in the current session.
    - Click **Plus** or **Close** button to erase the conversation and its context.

The replies from the Veeam Intelligence are standardized text blocks in Markdown format. At the end of every answer, the chatbot adds links to relevant Veeam documentation and KB articles. Screenshots and images are not supported.

If you find the answer to be insufficient, you can add more details. The bot retains the conversation context and previous questions within a current session, so you do not have to repeat anything.

> **NOTE**
>
> If the Veeam Intelligence service becomes overloaded, the question may get canceled.

# Disabling Veeam Intelligence

If, due to company policy, security concerns, or any other reasons, you are not allowed to leverage the Veeam Intelligence in your environment, you can disable it completely with the following registry key created on the backup server:

```
HKLM\Software\Veeam\Veeam Backup and Replication\

DWORD, AIAssistantDisabled = 1
```

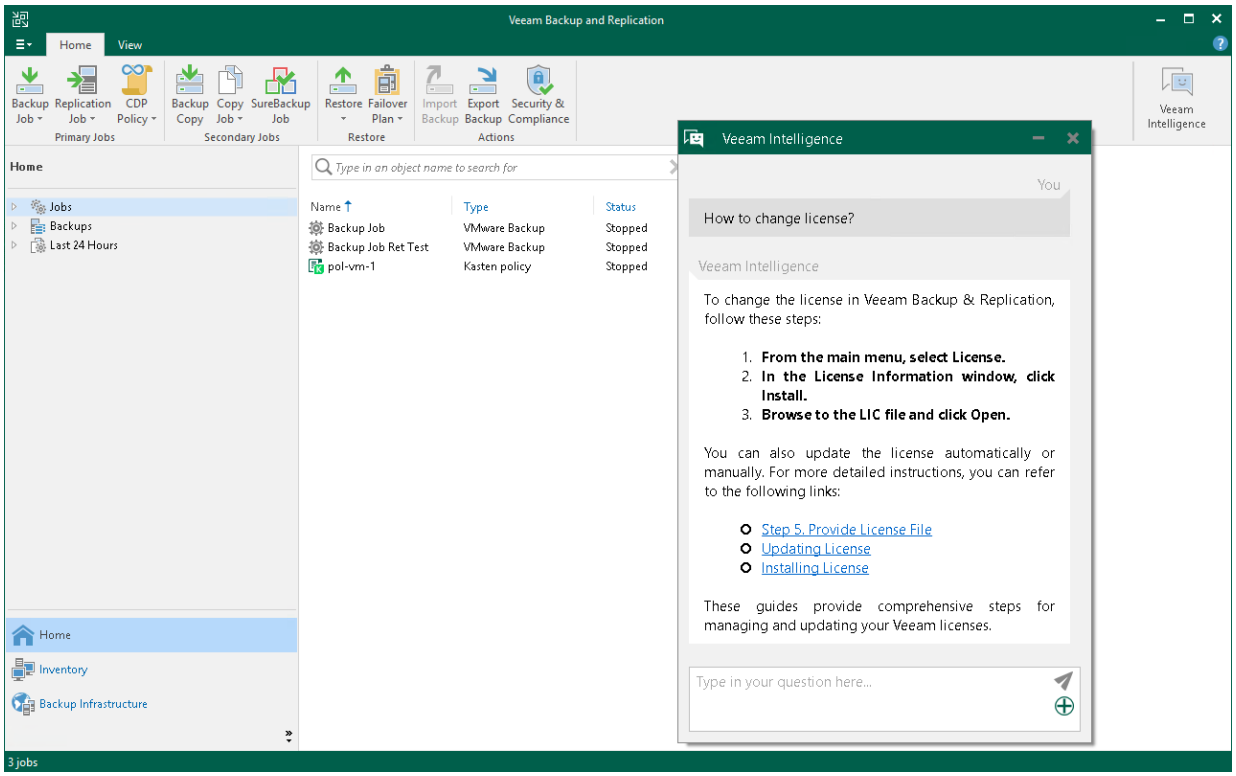

# Resource Scheduling

Veeam Backup & Replication has the built-in mechanism of resource scheduling. Resource scheduling lets Veeam Backup & Replication automatically define what backup infrastructure resources are required for data protection and disaster recovery jobs and tasks, select optimal resources and assign them for the jobs and tasks.

Resource scheduling is performed by the Veeam Backup Service running on the backup server. When a job or task starts, it communicates with the service and informs it about the resources it needs. The service analyzes job settings, parameters specified for backup infrastructure components, current load on the components, and automatically allocates optimal resources to the job.

For resource scheduling, Veeam Backup Service uses the following settings and features:

- Limitation of Concurrent Tasks
- Limitation of Read and Write Data Rates for Backup Repositories
- Performance Bottlenecks

# Limitation of Concurrent Tasks

When you start a data protection or disaster recovery job, Veeam Backup & Replication analyzes the list of VMs added to the job and creates a list of tasks to be processed. Veeam Backup & Replication then defines what backup infrastructure components must be used for the job, checks what backup infrastructure components are currently available, and assigns necessary components to process the created job tasks. The task concept differs depending on the type of the operation and the type of the backup chain being processed.

Backup infrastructure components typically process several tasks at the same time. You can limit the number of tasks that backup infrastructure components must process concurrently. Task limitation helps you balance the workload across the backup infrastructure and avoid performance bottlenecks.

Veeam Backup & Replication lets you limit the number of concurrent tasks for the following backup infrastructure components:

- Backup proxies
- Backup repositories

> **NOTE**
>
> Task limits set for backup infrastructure components influence the job performance. For example, you add a VM with 6 disks to a job and assign a VMware backup proxy that can process maximum 4 tasks concurrently for the job. In this case, Veeam Backup & Replication will create 6 tasks (1 task per each VM disk) and start processing 4 tasks in parallel. The other 2 tasks will be pending.

## How Task Limitation Works

Task limiting is performed by the Veeam Backup Service. The Veeam Backup Service is aware of all backup proxies and backup repositories in the backup infrastructure, and task limitation settings configured for them.

When a job starts, it informs the Veeam Backup Service about the list of tasks created for the job, and backup infrastructure resources that must be used for the job. The Veeam Backup Service detects the number of tasks that required backup infrastructure components are currently processing, and analyzes the number of allowed tasks for these components. If the number of currently processed tasks has reached the allowed limit, the backup infrastructure component will not start processing a new task until one of the currently running tasks finishes.

## Task Limitation for Backup Proxies

During data protection or disaster recovery jobs, Veeam Backup & Replication creates a separate task per each disk of every VM added to the job.

We recommend that you define the maximum number of concurrent tasks depending on the number of CPU cores available on the VMware backup proxy. To calculate the optimum limit, we recommend that you follow the rule: not more than 2 tasks per 1 CPU core.

To limit the number of concurrent tasks on a VMware backup proxy, you must define the **Max concurrent tasks** setting for the backup proxy.

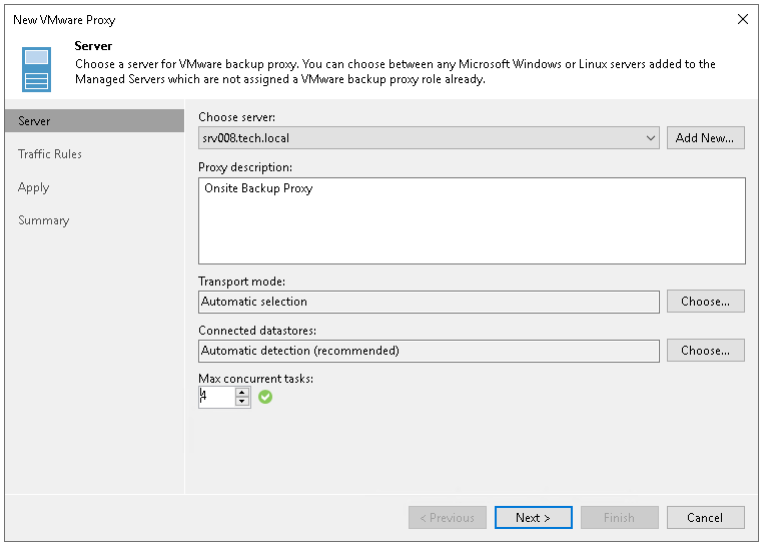


# Task Limitation for Backup Repositories

The number of tasks that Veeam Backup & Replication creates during data protection or disaster recovery jobs depends on the type of backup chains stored on the backup repository:

- For regular backup chains, Veeam Backup & Replication creates 1 task per job.
- For per-machine backup chains, Veeam Backup & Replication creates 1 task per every VM disk (that is, a disk of a VM added to the job).

Synthetic operations performed in the backup repository (such as synthetic full backup, backup files merge and transformation) are also regarded as tasks. The number of tasks performed during these operations also depends on the type of backup chains stored on the backup repository:

- For regular backup chains, Veeam Backup & Replication creates 1 task per job.
- For per-machine backup chains, Veeam Backup & Replication creates 1 task per every VM chain (that is, every VM added to the job).

We recommend that you define the maximum number of concurrent tasks based on the number of CPU cores and RAM available on the backup repository. To calculate the optimum limit, follow the rule: not more than 2 tasks per 1 CPU core, not less than 1 GB RAM for each concurrently processed machine disk and not less than 4 GB RAM for each concurrently processed unstructured data source (in case of deduplicating storage appliances, up to 8 GB RAM). In case of shared folder backup repositories, the same amount of resources is required for gateway servers.

To limit the number of concurrent tasks in a backup repository, you must enable the **Limit maximum concurrent tasks to <N>** option on the backup repository and define the necessary task limit.

If you use backup repositories for backup copy jobs, you must also consider tasks for read operations.

> **NOTE**
>
> When you limit the number of tasks for the backup repository, consider the storage throughput. If the storage system is not able to keep up with the number of tasks that you have assigned, it will be the limiting factor. It is recommended that you test components and resources of the backup infrastructure to define the workload that they can handle.

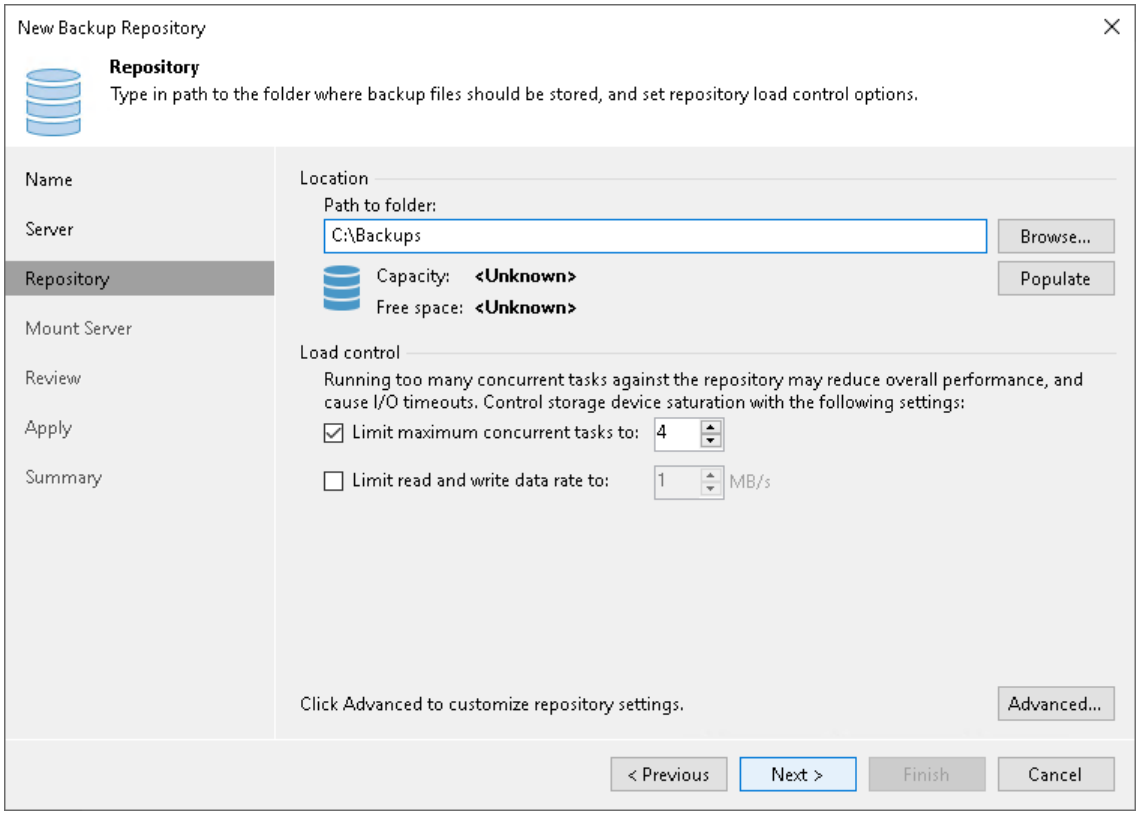


# Task Limitation for Components with Several Roles

One machine can perform several roles. For example, you can assign roles of the VMware backup proxy and backup repository to the same machine, or use a VMware backup proxy as a gateway server for a shared folder backup repository. In such situation, you must make sure that the backup infrastructure component is able to process the cumulative number of tasks specified for different roles.

We do not recommend sharing backup infrastructure components across Veeam installations. However, if you use the same infrastructure component with multiple Veeam Backup & Replication servers, ensure that this component can handle the cumulative workload for all Veeam Backup & Replication servers to which it is added. For more information on sharing backup infrastructure components across Veeam installations, see System Requirements.

# Limitation of Read and Write Data Rates for Backup Repositories

Veeam Backup & Replication can limit the speed with which Veeam Backup & Replication must read and write data to/from the backup repository. The data read and write speed is controlled with the **Limit read and write data rates to <N> MB/s** option that you can enable in backup repository settings.

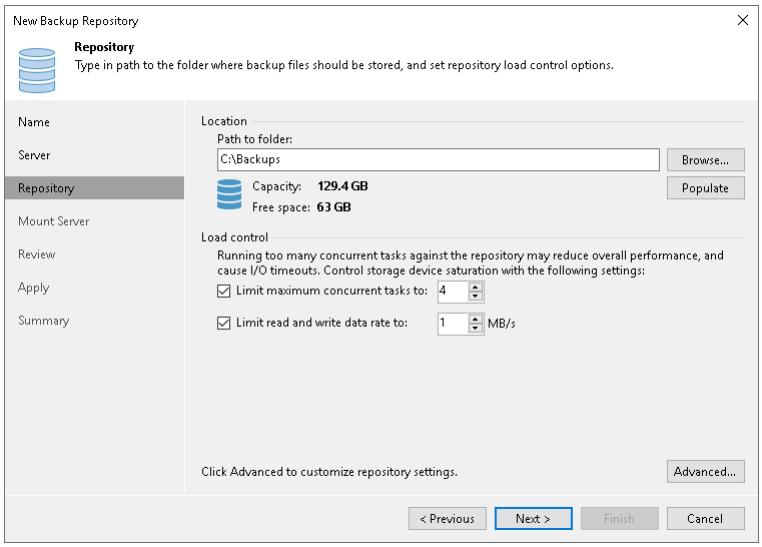


The Veeam Backup Service is aware of read and write data rate settings configured for all backup repositories in the backup infrastructure. When a job targeted at a backup repository starts, the Veeam Backup Service informs the Veeam Data Mover running on this backup repository about the allowed read/write speed set for this repository so that the Veeam Data Mover can limit the read/write speed to the specified value.

If the backup repository is used by a number of tasks simultaneously, Veeam Backup & Replication splits the allowed read/write speed rate between these tasks equally. Note that the specified limit defines the allowed read speed and the allowed write speed at the same time.

For example, you set the **Limit read and write data rates to** option to 8 MB/s and start two backup jobs. Each job processes 1 VM with 1 VM disk. In this case, Veeam Backup & Replication will create 2 tasks and target them at the backup repository. The data write rate will be split between these 2 tasks equally: 4 MB/s for one task and 4 MB/s for the other task.

If at this moment you start some job reading data from the same backup repository, for example, a backup copy job processing 1 VM with 1 disk, Veeam Backup & Replication will assign the read speed rate equal to 8 MB/s to this job. If you start 2 backup copy jobs at the same time (each processing 1 VM with 1 disk), Veeam Backup & Replication will split the read speed rate between these 2 jobs equally: 4 MB/s for one backup copy job and 4 MB/s for the other backup copy job.

# Performance Bottlenecks

As any backup application handles a great amount of data, it is important to make sure the data flow is efficient and all resources engaged in the backup process are optimally used. Veeam Backup & Replication provides advanced statistics about the data flow efficiency and lets you identify bottlenecks in the data transmission process.

Veeam Backup & Replication processes VM data in cycles. Every cycle includes a number of stages:

1. Reading VM data blocks from the source
2. Processing VM data on the VMware backup proxy
3. Transporting data over the network
4. Writing data to the target

When one data processing cycle is over, the next cycle begins. VM data therefore goes over the "data pipe".

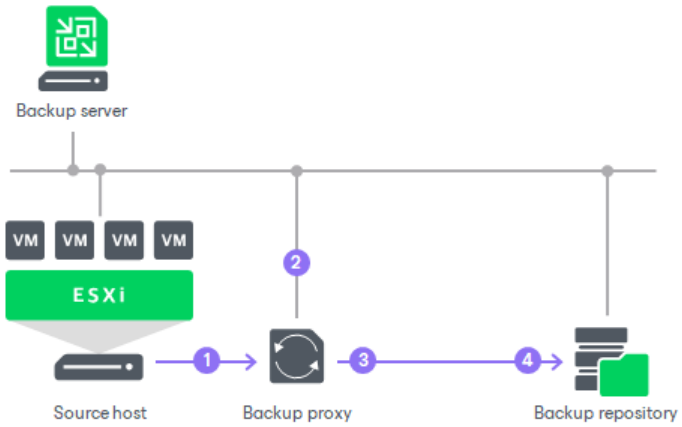


To evaluate the data pipe efficiency, Veeam Backup & Replication analyzes performance of all components in the data flow working as the cohesive system, and evaluates key factors on the source and target sites. Veeam Backup & Replication checks the following points in the data pipe:

1. **Source** – source disk reader component responsible for retrieving data from the source storage.
2. **Proxy** – VMware backup proxy component responsible for processing VM data.
3. **Source WAN accelerator** – WAN accelerator deployed on the source site. Used for backup copy and replication jobs working through WAN accelerators.
4. **Network** – network queue writer component responsible for getting processed VM data from the VMware backup proxy and sending it over the network to the backup repository or another VMware backup proxy.
5. **Target WAN Accelerator** – WAN accelerator deployed on the target site. Used for backup copy and replication jobs working through WAN accelerators.
6. **Target** – target disk writer component (backup storage or replica datastore).

The resource usage level for these points is evaluated in percent. This percent rate defines the amount of time for which components are busy during the job. An efficient data flow assumes that there is no latency at any point of the data pipe, and all its components work for approximately equal amount of time.

If any of the components operates inefficiently, there may appear a bottleneck in the data path. The insufficient component will work 100% of time while the others will be idling, waiting for data to be transferred. As a result, the whole data flow will slow down to the level of the slowest point in the data path, and the overall time of data processing will increase.

To identify a bottleneck in the data path, Veeam Backup & Replication detects the component with the maximum workload: that is, the component that works for the most time of the job. For example, you use a low-speed storage device as the backup repository. Even if VM data is retrieved from the SAN storage on the source site and transported over a high-speed link, VM data flow will still be impaired at the backup repository. The backup repository will be trying to consume transferred data at the rate that exceeds its capacity, and the other components will stay idle. As a result, the backup repository will be working 100% of job time, while other components may be employed, for example, for 60% only. In terms of Veeam Backup & Replication, such data path will be considered insufficient.

The bottleneck statistics for a job is displayed in the job session data. The bottleneck statistics does not necessarily mean that you have a problem in your backup infrastructure. It informs you about the weakest component in the data path. However, if you feel that the job performance is low, you may try taking some measures to get rid of the bottleneck. For example, in the case described above, you can limit the number of concurrent tasks for the backup repository.

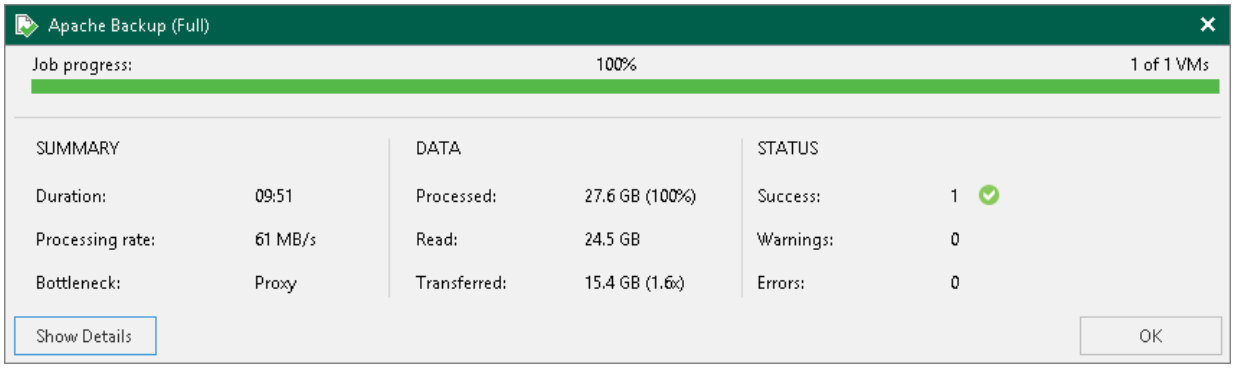


# Throttling as Bottleneck

In addition to main points in the data pipe, Veeam Backup & Replication may report throttling as a bottleneck. This can happen in the following cases:

- If you limit the read and write data rates for a backup repository, a backup repository may become a bottleneck. Veeam Backup & Replication will report *Throttling* in the bottleneck statistics.
- If you set up network throttling rules, network may become a bottleneck. Veeam Backup & Replication will report *Throttling* in the bottleneck statistics.

# Job Priorities

Resources in the backup infrastructure are limited. To make sure that the most crucial jobs are the first to get free resources to provide the reliable data protection, Veeam Backup & Replication uses the system of priorities to allocate resources to different jobs.

The resource scheduler within Veeam Backup & Replication uses several stages to prioritize jobs and provide free resources to them:

1. **Type** – at the first stage, the resource scheduler identifies the priority of the jobs awaiting free resources based on their type:
   a. **Backup restore jobs** – these jobs have the highest priority (**800**) and are the first to get free system resources.
   b. **Continuous data protection jobs** – these jobs have priority (**700**).
   c. **SnapshotDeleter jobs** – these jobs have priority (**600**).
   d. **Quick backup jobs** – these jobs have priority (**500**).
   e. **High priority jobs** – jobs with the enabled **High priority** option have priority (**400**). You can enable the **High priority** option for the following jobs: backup jobs, replication jobs, agent jobs managed by backup server, file backup jobs.

      > **TIP**
      >
      > In the list of jobs in the Veeam Backup & Replication console, jobs with the **High priority** option enabled are marked with a red flag ( ).

   f. **Regular backup and replication jobs** – these jobs have priority (**300**).
   g. **Backup copy jobs** – these jobs have priority (**200**). Immediate backup copy jobs have priority (**400**) when the related source backup job is a high priority job.
   h. **Archive jobs** – these jobs have the lowest priority (**100**) and are the last to get free system resources. These are the jobs that move backups to capacity and archive tiers.
2. **Priority** – at the second stage, the resource scheduler identifies the priority of the jobs within each type group from the first stage based on their startup type:
   a. **Scheduled VSS proxy jobs** – the jobs with the configured job schedule and using a VSS proxy have the highest priority (**40**) within the group and are the first to get free system resources.
   b. **Scheduled jobs** – the jobs with the configured job schedule have priority (**30**) within the group.
   c. **Manually started VSS proxy jobs** – the manually started jobs using a VSS proxy have priority (**20**) within the group.
   d. **Manually started jobs** – the manually started jobs have the lowest priority (**10**) within the group and are the last to get free system resources.
3. **Start time** – if the jobs have the same type and priority, resources are first allocated to jobs that were started earlier with an exception of tape jobs.

> **NOTE**
>
> The resource scheduler does not take into account the start time for tape jobs. The source job may start when the tape job is still running. For more information on the priority settings for tape jobs, see Schedule for Backup to Tape Job.

You can check the job type and priority of a certain job in service logs. For more information on logs, see Managing Logs.

# Configuring Veeam Backup & Replication

This section describes basic operations related to Veeam Backup & Replication configuration management.

# Configuring General Settings

You can set up general settings for Veeam Backup & Replication. General settings are applied to all jobs, backup infrastructure components and other objects managed by the backup server.

# Specifying I/O Settings

You can specify data processing settings.

Consider the following:

- The **Enable storage latency control** option is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.
- The **Set custom thresholds on individual datastores** option is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise Plus edition is required.
- The **Enable storage latency control** option is not supported for vVols/vSAN storage.

To specify data processing settings:

1. From the main menu, select **Options**.
2. Click the **I/O Control** tab.
3. To control the I/O load on the production storage where VMs reside, select the **Enable storage latency control** check box. When you enable storage latency control, Veeam Backup & Replication monitors storage read latency on production datastores during data protection and disaster recovery activities. To monitor the storage latency, Veeam Backup & Replication uses real-time metrics from the hypervisor where VMs reside. By default, metrics from the hypervisor are collected every 20 seconds. These settings are inherited from VMware vSphere.

   Specify two thresholds:

   a. In the **Stop assigning new tasks to datastore** at field, specify the I/O latency limit at which Veeam Backup & Replication must not assign new tasks targeted at the datastore.

   b. In the **Throttle I/O of existing tasks at** field, specify the I/O latency limit at which Veeam Backup & Replication must decrease the speed of data retrieval or writing to/from the datastore. When the I/O latency for this datastore reaches this value, the Veeam Data Mover working with this datastore will slow down data retrieval or writing.

   The value in the **Stop assigning new tasks to datastore at** field cannot be greater than the value in the **Throttle I/O of existing tasks at** field.

> **NOTE**
>
> If you enable the storage latency control option, Veeam Backup & Replication starts processing VM disks residing on the same datastore with a 40-60 second time offset. This offset helps Veeam Backup & Replication evaluate the current I/O load on the datastore. For example, if you launch a job processing a VM with two disks, Veeam Backup & Replication will start processing the first VM disk, wait for 40-60 seconds to evaluate the I/O workload on the datastore, and then start processing the second VM disk.
>
> Consider this behavior. If you schedule jobs that process multiple VM disks residing on the same datastore to start at the same time, the jobs performance will degrade.

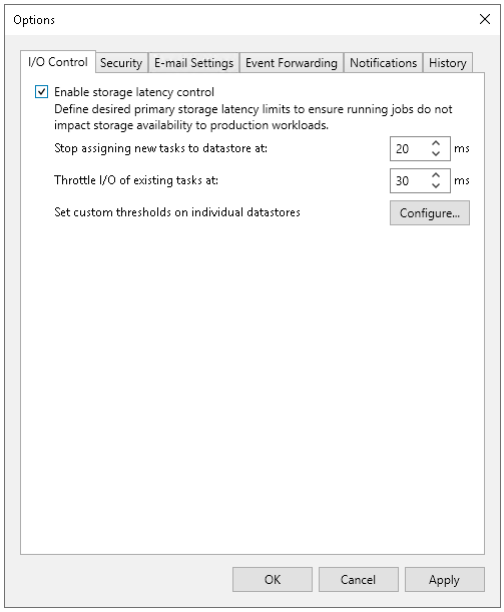


You can set the I/O latency limit for every storage in the virtual infrastructure separately.

To set the I/O latency limit for every storage separately:

1. From the main menu, select **Options**.
2. Click the **I/O Control** tab.
3. Click **Configure**.
4. Click **Add** > **Datastore**, select the necessary datastore and click **OK** to add it to the storage list.
5. Select the added datastores in the list and click **Edit**.
6. Specify the I/O thresholds for the datastores.

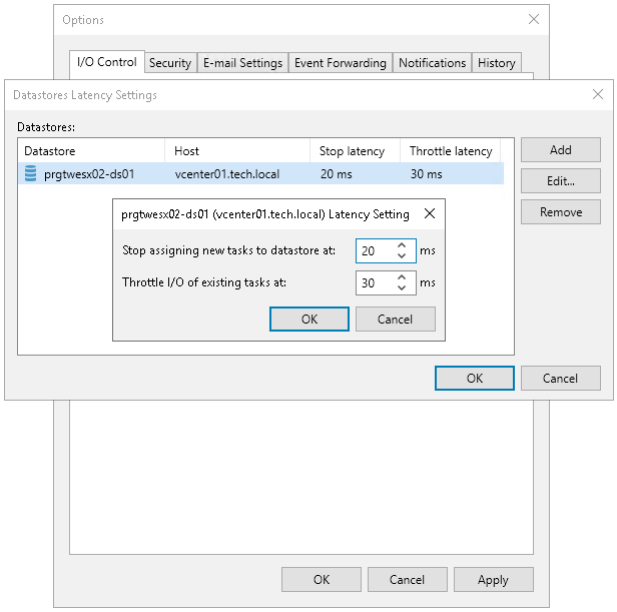

# Configuring Security Settings

In the **Security** tab, you can configure the following:

- Backup Server Certificate

  Configure a TLS certificate to establish secure communication from backup infrastructure components to the backup server.

- Linux Hosts Authentication

  Enable the fingerprint check for Linux machines to protect connection from man-in-the-middle attacks.

- Cloud Connect

  Enable access to the cloud gateway for the Remote Access Console connected to an external network.

- FIPS Compliance

  Enable FIPS-compliant operation mode.

- Audit Logs Location

  Select a folder for storing audit logs.

## Backup Server Certificate

When you configure the Veeam Backup & Replication infrastructure, you can specify what TLS certificate must be used to establish a secure connection from backup infrastructure components to the backup server. Veeam Backup & Replication offers the following options for TLS certificates:

- Keep the default self-signed TLS certificate generated by Veeam Backup & Replication at the process of upgrading to a new version of Veeam Backup & Replication.
- Use Veeam Backup & Replication to generate a new self-signed TLS certificate. To learn more, see Generating Self-Signed Certificate.
- Select an existing TLS certificate from the certificate store. To learn more, see Importing Certificate from Certificate Store.
- Import a TLS certificate from a file in the PFX format. To learn more, see Importing Certificate from PFX Files.

If you plan to use a certificate issued by your own Certificate Authority (CA), make sure that the certificate meets the requirements. For more information, see Using Certificate Signed by Internal CA.

**IMPORTANT**

If you update the TLS certificate used on the backup server, you must also do the following:

- If multi-factor authentication is enabled, any Veeam Backup & Replication consoles connected to the backup server must be restarted to avoid connection issues.
- For AHV Backup proxies, pass through the **Edit Nutanix Proxy** wizard. To do this, in the **Backup Infrastructure** view, right-click a proxy and select **Properties**. In the wizard, click **Finish**. Also, restart the Veeam AHV Service.
- For RHV Backup proxies, pass through the **Edit Red Hat Virtualization Proxy** wizard. To do this, in the **Backup Infrastructure** view, right-click a proxy and select **Properties**. In the wizard, click **Finish**.
- For VMware clusters, pass through the **I/O filter Management** wizard as described in section Installing I/O Filter.
- For VMware CDP proxies, pass through the **Edit VMware CDP Proxy** wizard. To do this, in the **Backup Infrastructure** view, right-click a proxy and select **Properties**. In the wizard, click **Finish**.

If you remove the old certificate from the Microsoft Windows certificate store, you must also reconfigure Veeam Agents added to the **Computers with pre-installed agents** protection group. To do this, repeat the configuration step of the Veeam Agent deployment scenario as described in the subsections of the Deploying Veeam Agents Using Generated Setup Files section. Other protection groups will be automatically reconfigured during the next rescan operation.

If you do not remove the old certificate from the Microsoft Windows certificate store, all protection groups will be automatically reconfigured the next time Veeam Agents connect to the backup server.

## Generating Self-Signed Certificate

You can use Veeam Backup & Replication to generate a self-signed certificate for authenticating parties in the Veeam Backup & Replication infrastructure.

To generate TLS certificates, Veeam Backup & Replication employs the RSA Full cryptographic service provider by Microsoft Windows installed on the backup server. The created TLS certificate is saved to the *Shared* certificate store. The following types of users can access the generated TLS certificate:

- User who created the TLS certificate
- LocalSystem user account
- Local Administrators group

If you use a self-signed TLS certificate generated by Veeam Backup & Replication, you do not need to take additional actions to deploy the TLS certificate on a protected computer. When Veeam Backup & Replication discovers a protected computer, a matching TLS certificate with a public key is installed on the protected computer automatically. During discovery, Veeam Installer Service deployed on the protected computer retrieves the TLS certificate with a public key from the backup server and installs a TLS certificate with a public key on the protected computer.

**NOTE**

When you generate a self-signed TLS certificate with Veeam Backup & Replication, you cannot include several aliases to the certificate and specify a custom value in the *Subject* field. The *Subject* field value is taken from the Veeam Backup & Replication license installed on the Veeam backup server.

> **IMPORTANT**
>
> If you update the TLS certificate used on the backup server, you must also update info about the certificate on the specific backup infrastructure components as described in section Backup Server Certificate.

To generate a self-signed TLS certificate, do the following:

1. From the main menu, select **Options**.
2. Click the **Security** tab.
3. In the **Security** tab, click **Install**.
4. At the **Certificate Type** step of the wizard, select **Generate new certificate**.

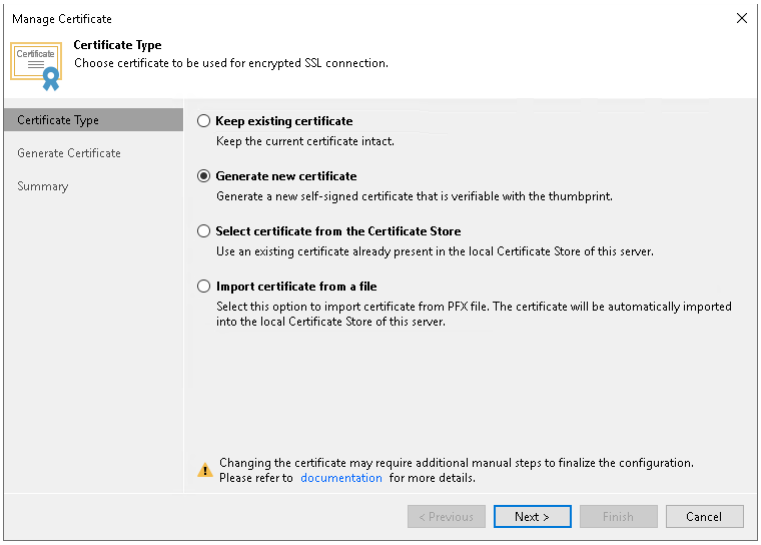


5. At the **Generate Certificate** step of the wizard, specify a friendly name for the created self-signed TLS certificate.

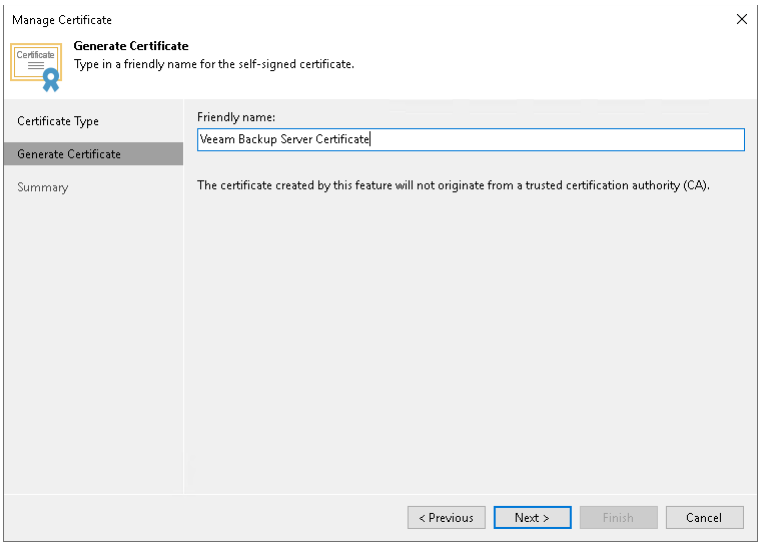


6. At the **Summary** step of the wizard, review the certificate properties. Use the **Copy to clipboard** link to copy and save information about the generated TLS certificate. You will be able to use the copied information to verify the TLS certificate with the certificate thumbprint.

7. Click **Finish**. Veeam Backup & Replication will save the generated certificate in the *Shared* certificate store on the Veeam backup server.

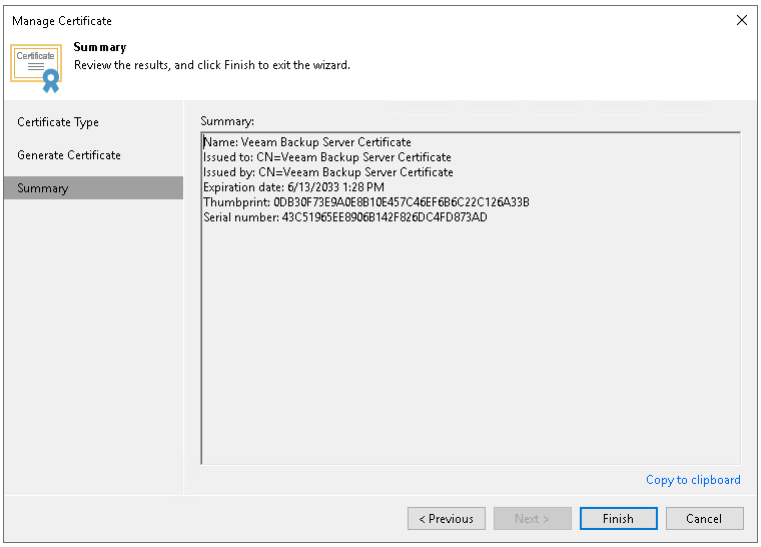


# Importing Certificate from Certificate Store

If the Veeam backup server has been issued a TLS certificate signed by a CA and the TLS certificate is located in the Microsoft Windows certificate store, you can use this certificate for authenticating parties in the Veeam Backup & Replication infrastructure.

> **IMPORTANT**
>
> If you update the TLS certificate used on the backup server, you must also update info about the certificate on the specific backup infrastructure components as described in section Backup Server Certificate.

To select a certificate from the Microsoft Windows certificate store, do the following:

1. From the main menu, select **Options**.
2. Click the **Security** tab.
3. In the **Security** tab, click **Install**.
4. At the **Certificate Type** step of the wizard, choose **Select certificate from the Certificate Store**.

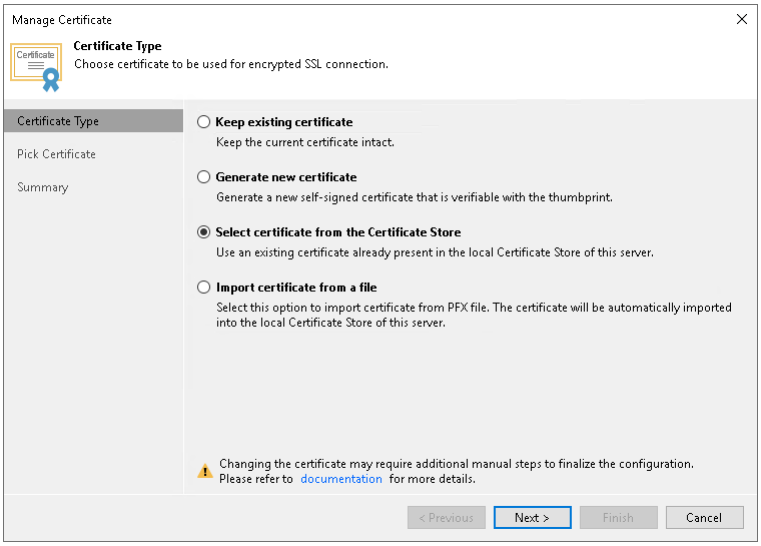

5. At the **Pick Certificate** step of the wizard, select a TLS certificate that you want to use. You can select only certificates that contain both a public key and a private key. Certificates without private keys are not displayed in the list.

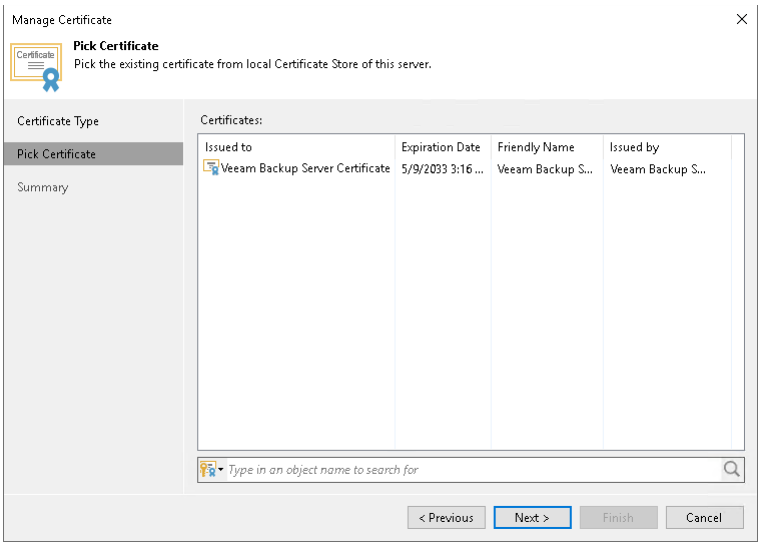


6. At the **Summary** step of the wizard, review the certificate properties.
7. Click **Finish** to apply the certificate.

## Importing Certificate from PFX Files

You can import a TLS certificate in the following situations:

- Your organization uses a TLS certificate signed by a CA and you have a copy of this certificate in a file of PFX format.
- You have generated a self-signed TLS certificate in the PFX format with a third-party tool and you want to import it to Veeam Backup & Replication.

**NOTE**

Consider the following:

- The TLS certificate must pass validation on the Veeam backup server. Otherwise, you will not be able to import the TLS certificate.
- If a PFX file contains a certificate chain, only the end entity certificate will be imported.

**IMPORTANT**

If you update the TLS certificate used on the backup server, you must also update info about the certificate on the specific backup infrastructure components as described in section Backup Server Certificate.

To import a TLS certificate from a PFX file, do the following:

1. From the main menu, select **Options**.
2. Click the **Security** tab.
3. In the **Security** tab, click **Install**.

4. At the **Certificate Type** step of the wizard, choose **Import certificate from a file**.

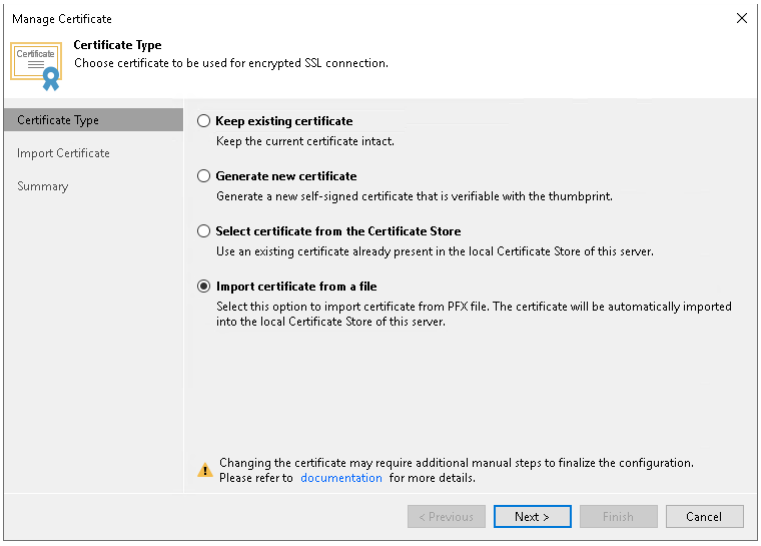


5. At the **Import Certificate** step of the wizard, specify a path to the PFX file.
6. If the PFX file is protected with a password, specify the password in the **Password** field.

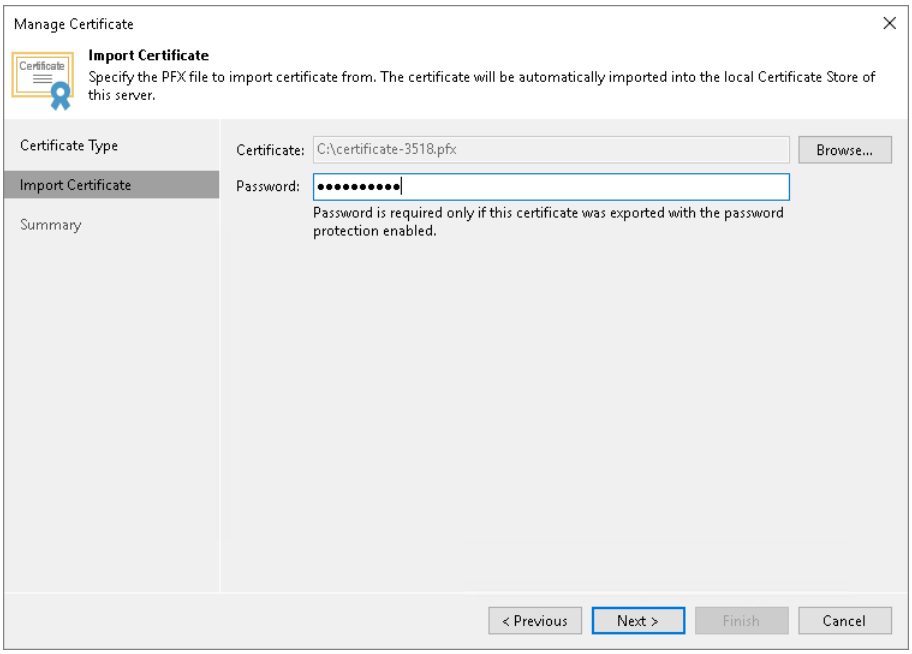


7. At the **Summary** step of the wizard, review the certificate properties. Use the **Copy to clipboard** link to copy and save information about the TLS certificate. You can use the copied information on a protected computer to verify the TLS certificate with the certificate thumbprint.
8. Click **Finish** to apply the certificate.

## Using Certificate Signed by Internal CA

If you want to use a certificate signed by an internal Certificate Authority (CA), consider the following:

- Make sure that Veeam Backup & Replication server trusts the CA. That means that the Certification Authority certificate must be added to the Trusted Root Certification Authority store on the Veeam Backup & Replication server. Also, Certificate Revocation List (CRL) must be accessible from the Veeam Backup & Replication server.

- If you use Windows Server Certification Authority, issue a Veeam Backup & Replication certificate based on the built-in *Subordinate Certification Authority* template or templates similar to it. You can manage templates with the **Certificate Templates** MMC snap-in.

> **IMPORTANT**
>
> The following certificates are not supported:
>
> - Certificates issued by public CAs
> - Elliptic Curve Signature (ECC) certificates
> - Cryptography API: Next Generation (CNG) certificates

A certificate signed by a CA must meet the following requirements:

- The certificate subject is equal to the fully qualified domain name of the Veeam Backup & Replication server. For example: `vbrserver.domain.local`.

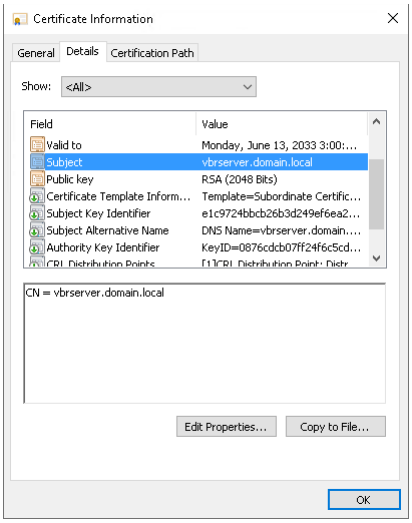


- The **Subject Alternative Name** field contains both the FQDN and the NetBIOS name. You can add multiple DNS entries in the following format: `DNS:vbrserver.domain.local,DNS:vbrserver`.
- The minimum key size is 2048 bits.
- The following key usage extensions are enabled in the certificate:
  - Digital Signature
  - Certificate Signing
  - Off-line CRL Signing

- CRL Signing (06)

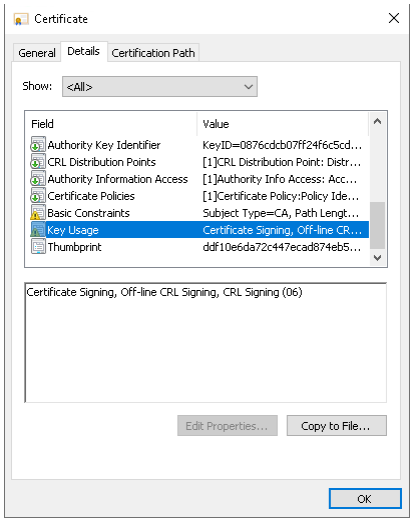


- The **Path Length Constraint** parameter in the **Basic Constraints** extension is set to *0*.
- The **Subject Type** parameter in the **Basic Constraints** extension is set to *CA*.

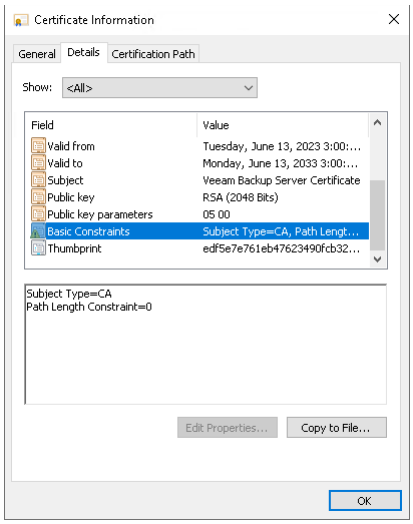


If you use Windows Server Certification Authority, open the **Certificate Templates** MMC snap-in and select the certificate template based on the built-in *Subordinate Certification Authority* template or templates similar to it. On the **Extensions** tab, enable the **Do not allow subject to issue certificates to other CAs** option.

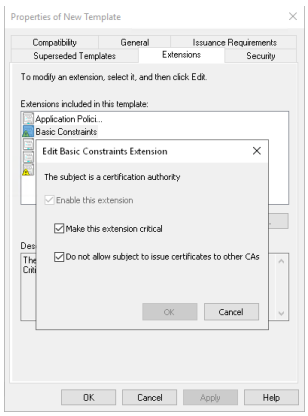


- The key type in the certificate is set to *Exchange*.

To start using the signed certificate, you must select it from the certificates store on the Veeam Backup & Replication server. To learn more, see Importing Certificate from Certificate Store.

# Linux Hosts Authentication

In the **Linux hosts authentication** section of the Veeam Backup & Replication settings, you can specify SSH fingerprint verification settings for protected Linux machines.

> **NOTE**
>
> Veeam Backup & Replication uses the SHA-256 hashing algorithm for verification. If you upgrade from previous versions, SSH fingerprint format will be updated automatically during next rescan or next connection to the Linux machine through SSH.

You can select one of the following options:

- **Add all discovered hosts to the list automatically** — Veeam Backup & Replication allows all Linux servers added to the protection group and all Linux VMs to connect to the backup server. Machine fingerprints are added to the Veeam Backup & Replication database and checked every time when machines establish a connection with the backup server. If SSH fingerprints do not match, the connection fails.
- **Add unknown hosts to the list manually** — this option provides a more secure environment because only trusted Linux servers and Linux VMs can connect to the backup server:
  - Machines that have already established a connection with the backup server and have their fingerprints stored in the Veeam Backup & Replication database. You can export the list of trusted machines to the `known_hosts` file. To do this, click **Export** and specify a path to the folder to save the file.
  - Machines specified in the `known_hosts` file imported to Veeam Backup & Replication. To import the `known_hosts` file, click **Import** and specify a path to the folder where the file resides.

    When you specify a trusted host in the `known_hosts` file, it must follow the same format as the `~/.ssh/known_hosts` file. It must include the network name hash, the type of key, and the public key.

    Example of a trusted host entry:

    ```
    |1|y/XiVUB2z/ZBb3vuOYm0x9RUiQA=|9zTpxEaAKbGPe7JyS/OyIWvsTz8= ssh-ed2551
    9 AAAAC3NzaC1lZDI1NTE5AAAAIHhO7S1tp0EAgainstjkXSAi4a+JIPKnTUpABC8BGyWk9
    ```

  Veeam Backup & Replication displays the number of trusted machines in the **Trusted hosts** field.

  Untrusted Linux VMs are displayed under the **Untrusted** node in the **Inventory** view. Untrusted Linux servers are displayed under the **Unavailable** node in the **Backup Infrastructure** view. These machines cannot connect to the backup server and download Veeam Agent for Linux installation packages during discovery. Also, guest OS processing of untrusted VMs will fail.

To start managing an untrusted Linux machine, you need to validate its fingerprint manually in the Veeam Backup & Replication console. For more details, see Validating SSH Fingerprints.

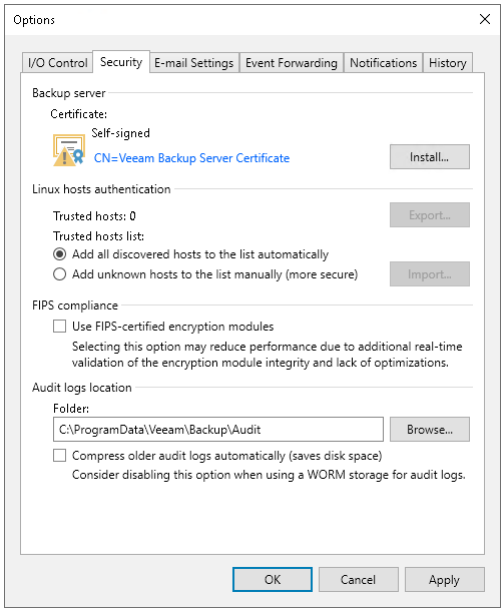


# Validating SSH Fingerprints

When you select the **Add unknown hosts to the list manually** option in Veeam Backup & Replication settings, you need to validate SSH fingerprints of untrusted Linux servers and Linux VMs manually in the Veeam Backup & Replication console. For more information about all SSH fingerprint verification options, see Linux Hosts Authentication.

> **IMPORTANT**
>
> To avoid fingerprint mismatch errors, you must use unique hostnames on your network.

# Validating Linux VMs

During discovery, Veeam Backup & Replication puts untrusted Linux VMs to the **Untrusted** node in the inventory pane. You can validate all untrusted VMs at once or a specific VM:

- To validate all untrusted Linux VMs at once, select the **Untrusted** node and click **Trust All** on the ribbon. Alternatively, you can right-click the **Untrusted** node and select **Trust all**.

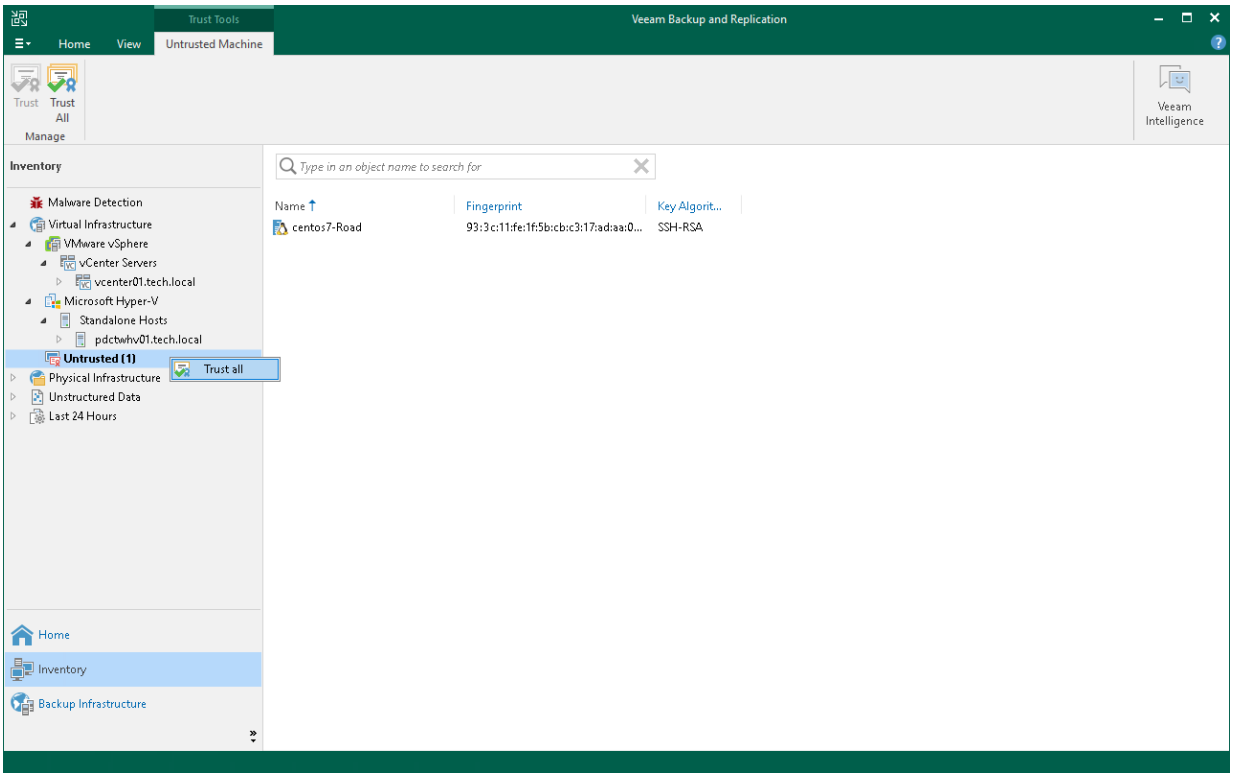

- To validate a specific VM, select it in the working area and click **Trust** on the ribbon. Alternatively, you can right-click the VM and select **Trust**.

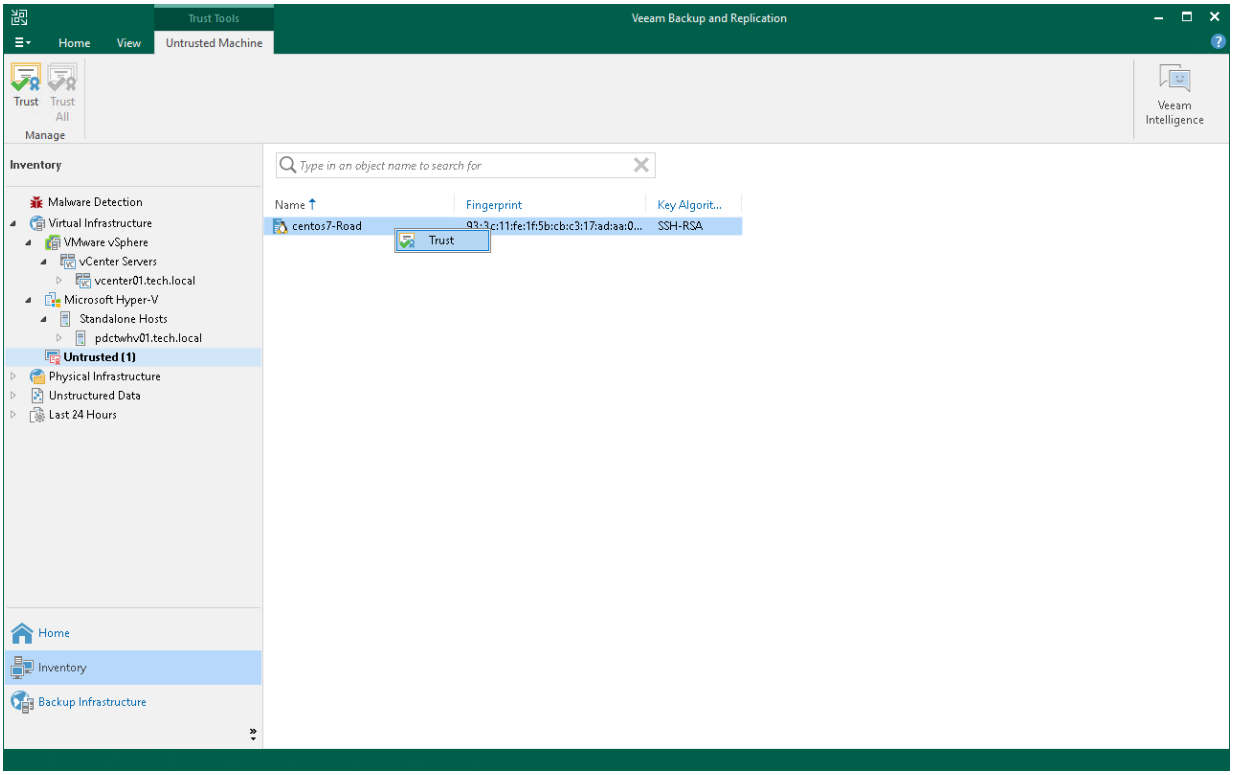


# Validating Linux Servers

If the SSH public key fingerprint of a Linux server is changed, Veeam Backup & Replication puts this machine to the **Unavailable** node in the **Backup Infrastructure** view. To validate the Linux server, do the following:

1. Right-click the Linux server and select **Properties.**
2. In the **SSH Connection** step of the **Edit Linux Server** wizard, click **Apply.**
3. In the dialog box, click **Yes** to confirm that you trust this server.

4. Click **Finish** to close the wizard.

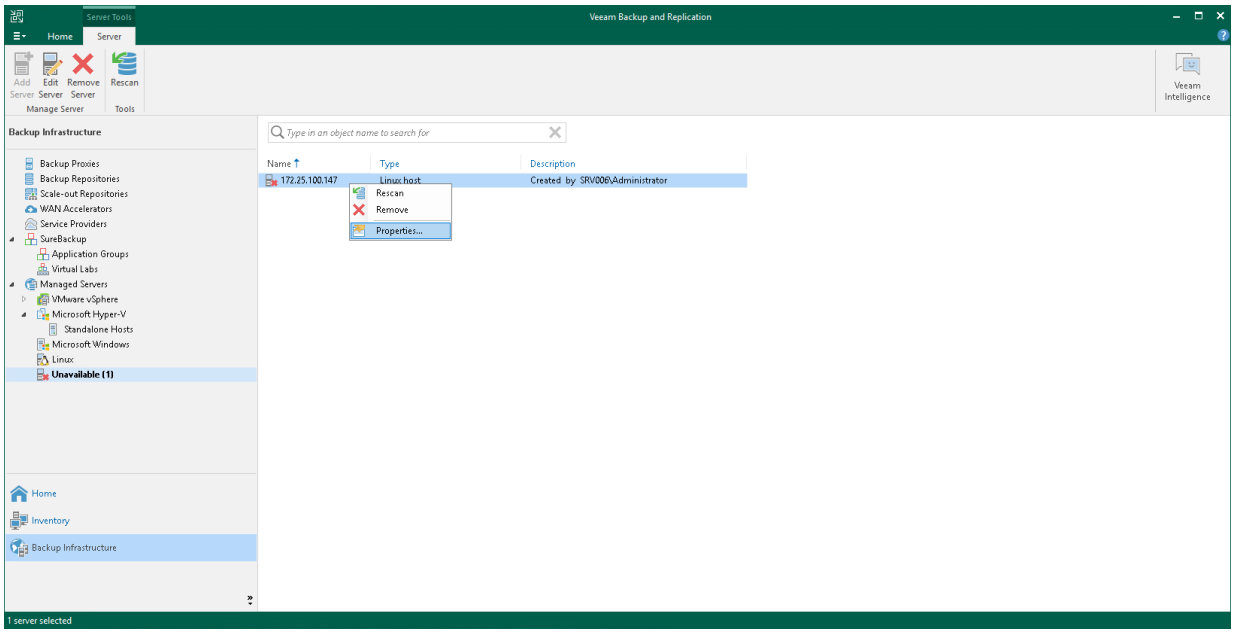


# Cloud Connect Remote Access Console

If you use the Veeam Cloud Connect functionality, you can remotely access the tenant backup server to manage Veeam Backup & Replication deployed on the tenant side. One of the ways to do this is to use the Remote Access Console. In case it is installed on a remote machine connected to an external network, you will need to enable access to the cloud gateway for the Remote Access Console over the internet. For more information, see Veeam Cloud Connect Guide.

To enable access to the cloud gateway for the Remote Access Console:

1. From the main menu, select **Options**.
2. Open the **Security** tab.
3. In the **Cloud Connect** section, select the **Accept remote console connection requests originating from the Internet** check box.

4. Click **OK**.

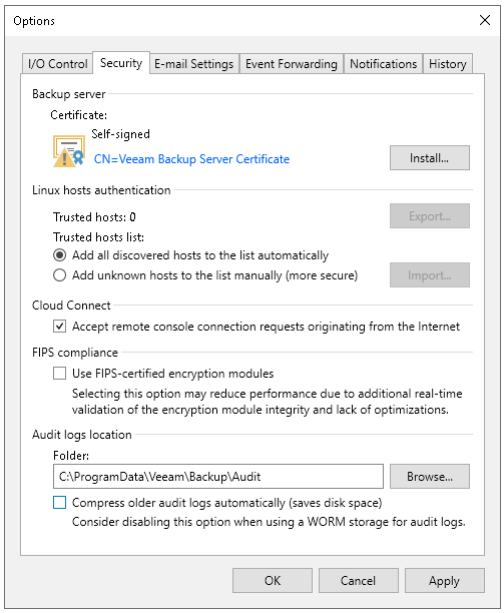


# FIPS Compliance

Veeam Backup & Replication can be configured to run in a FIPS-compliant operation mode. When this mode is enabled, Veeam Backup & Replication uses platform-provided cryptographic APIs and the Veeam Cryptographic Module to meet FIPS-compliance requirements. Additionally, connections cannot be established with components that are not FIPS-compliant.

> **NOTE**
>
> To make your backup infrastructure FIPS-compliant, follow vendor recommendations. For more information on Microsoft Windows Server, see this article.

To enable the FIPS-compliant operation mode:

1. From the main menu on the backup server, select **Options**.
2. Open the **Security** tab.
3. In the **FIPS compliance** section, select the **Use FIPS-certified encryption modules** check box.
4. Click **OK**.

> **NOTE**
>
> If you use Amazon S3 or Amazon S3 Glacier object repositories in your backup infrastructure and enable the FIPS-compliant operation mode, Veeam Backup & Replication checks if these components are FIPS-compliant. If any of them are not, a warning will be displayed.

> **IMPORTANT**
>
> If you have backup infrastructure components based on Linux servers with persistent Veeam Data Movers and select or clear the **Use FIPS-certified encryption modules** check box, you must open the Edit Linux Server wizard for each Linux server with the persistent Veeam Data Mover and proceed to the end of the wizard. This will update server settings. If you do not update the settings, the servers will be unavailable.

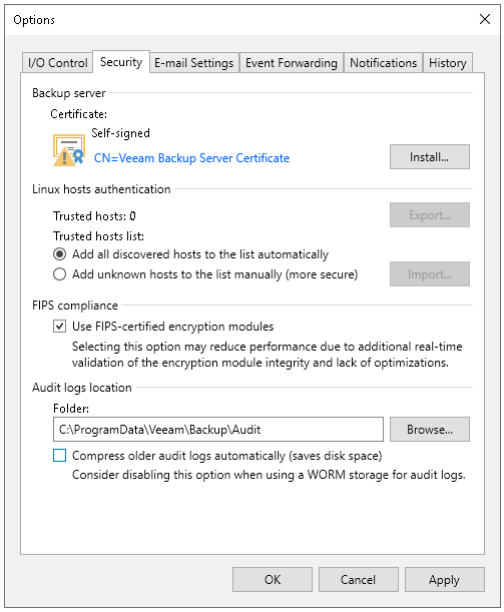


# Audit Logs Location

Veeam Backup & Replication logs activities during **File-Level Restore** sessions involving backups, replicas and unstructured data backups. The results are stored in form of `.csv` files that are called audit logs. For more information about log files, see Managing Logs.

In **Audit Logs Location** field, you can specify folder where the audit logs will be stored. By default, log files are stored in the following folder: `%ProgramData%\Veeam\Backup\Audit`. You can also specify an SMB (CIFS) folder.

> **IMPORTANT**
>
> Storing audit logs on WORM tapes is not supported. Storing audit logs on WORM storage is supported without log compression. This type of storage prevents the data from being deleted or modified. Thus, raw audit logs cannot be deleted after creating compressed files.

If you use an SMB (CIFS) folder, the service account used for Veeam Backup Service on the machine with Veeam Backup & Replication must have access to that SMB (CIFS) folder. By default, this is the *LocalSystem* account, so you will need to grant write access to the *VBR Server Active Directory* computer account.

By default, older audit logs are compressed automatically. To perform this operation, the service account under which Veeam Backup Service runs must have modify permissions on the target folder. If you do not want to compress older audit logs, clear the **Compress older audit logs automatically** check box.

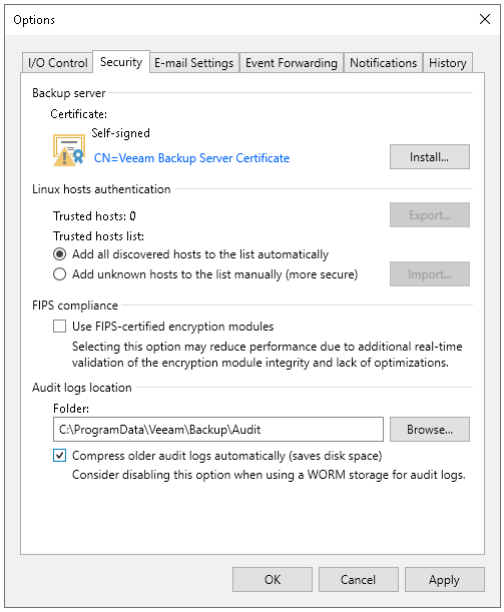

# Specifying Email Notification Settings

You can receive email notifications with results on jobs performed on the backup server.

To receive email notifications, do the following:

- Configure global email notification settings in Veeam Backup & Replication.
- Configure job notification settings.

> **TIP**
>
> To receive email notification about all jobs performed on the backup server in one email, configure email notification settings in Veeam Backup Enterprise Manager.

## Configuring Global Email Notification Settings

To configure global email notification settings:

1. From the main menu, select **Options**.
2. Open the **E-mail Settings** tab and select the **Enable e-mail notifications** check box.
3. Configure mail server.
4. Customize send settings.

## Configuring Mail Server

To configure mail server, perform the following steps:

1. In the **Mail Server** field, specify the authentication method you want to use. Veeam Backup & Replication supports the following methods:
   - SMTP basic authentication
   - Google Gmail OAuth 2.0 authentication
   - Microsoft 365 OAuth 2.0 authentication

   > **NOTE**
   >
   > For more secure environments, it is recommended to use OAuth 2.0 authentication. Also, note that Microsoft and Google consider SMTP basic authentication as an outdated industry standard and plan to disable it. For more information, see this Microsoft article and this Google article.

2. If you want to use SMTP basic authentication, perform the following steps:
   a. In the **Mail server** field, select *SMTP server* from the list.

b. In the **SMTP server** field, enter a full DNS name, or IPv4 or IPv6 address of the SMTP server that will be used for sending email notifications. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.

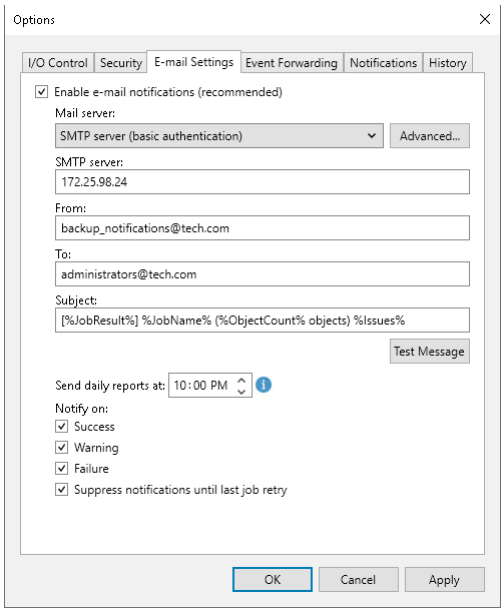


c. To specify user credentials and connection options, click the **Advanced** button:

- Specify the port number and connection timeout for the SMTP server.

> **NOTE**
>
> Sending email notifications using Implicit TLS (over port 465) is not supported. For more information about Implicit TLS, see this RFC section.

- To use a secure connection for email operations, select the **Connect using SSL** check box.
- If you need to connect to the SMTP server using a specific account, select the **This SMTP server requires authentication** check box and select the necessary credentials from the **Log on as** list. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add credentials. For more information, see Credentials Manager.

When you add an SMTP server, Veeam Backup & Replication saves to the configuration database a thumbprint of the TLS certificate. If the certificate is not trusted, Veeam Backup & Replication displays a warning. If you trust the certificate, click **Continue**.

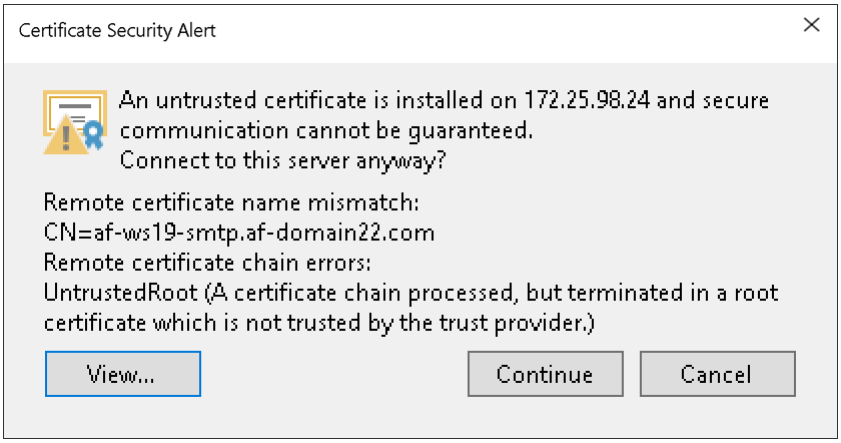

3. If you want to use Google Gmail OAuth 2.0 authentication, perform the following steps:

   a. In the **Mail server** field, select *Google Gmail* from the list and click the **Sign in with Google** button.

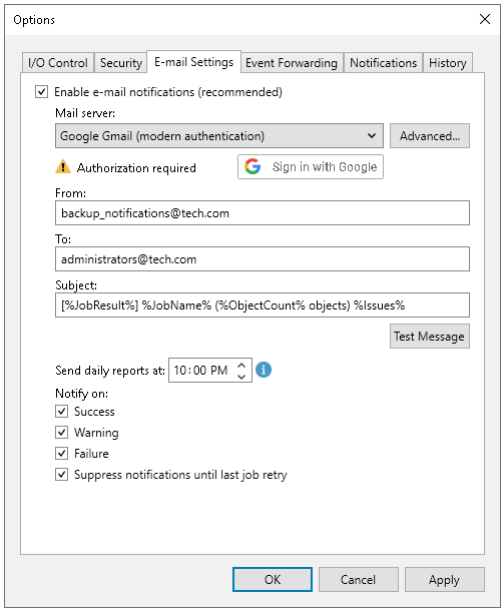


   b. In the opened web browser window, specify the Google account to connect to the Veeam Backup & Replication application. Note that you must select **Send email on your behalf** check box during configuring access options.

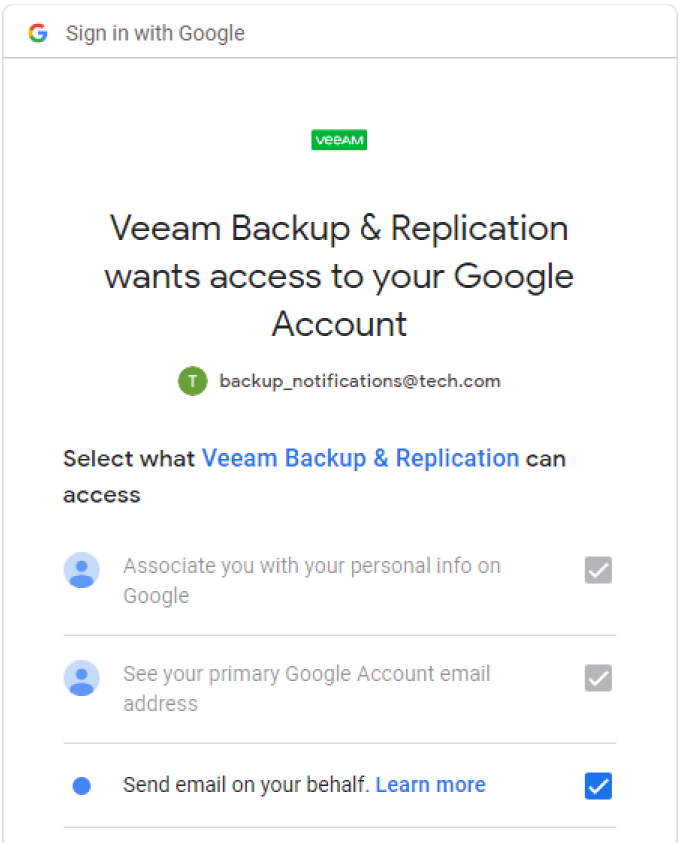

> **NOTE**
>
> Consider the following:
>
> - For security reasons, it is recommended to use a dedicated service account with granular SendMail permissions.
> - To sign in with the Google account, your default web browser must meet Google requirements. For more information, see this article.

If the authentication is successful, the *Token is valid* notice will appear. The token is refreshed automatically. If it was revoked or the Google account password was changed, click the **Re-authorize** link to update configuration.

To specify custom authentication options, click the **Advanced** button:

- Select the **Use custom application registration settings** check box.
- Specify the application client ID and the client secret.

For more information on how to register your custom application, see Registering Application in Google Cloud Console.

4. If you want to use Microsoft 365 OAuth 2.0 authentication, perform the following steps:
   a. In the **Mail server** field, select *Microsoft 365* from the list and click the **Authorize now** link.

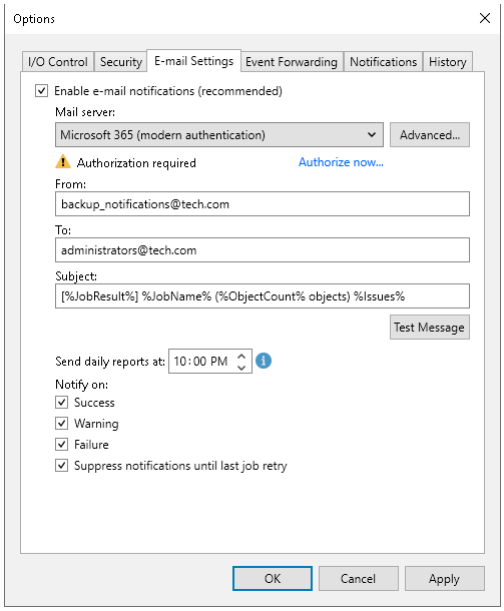


   b. In the opened window, specify your Exchange Online credentials to connect to the Veeam Backup & Replication application.

> **NOTE**
>
> Consider the following:
>
> - For security reasons, it is recommended to use a dedicated service account with granular SendMail permissions.
> - To sign in with Exchange Online credentials, turn off the **Internet Explorer Enhanced Security Configuration** option in Server Manager. For more information, see this article.

If the authentication is successful, the *Token is valid* notice will appear. The token is refreshed automatically. If it was revoked or Exchange Online credentials were changed, click the **Re-authorize** link to update configuration.

To specify custom authentication option, click the **Advanced** button:

- Select the **Use custom application registration settings** check box.
- Specify the application client ID and the tenant ID.

> **NOTE**
>
> For custom applications, note that you must select **Consent on behalf of your organization** check box during configuring access options.

For more information on how to register your custom application, see Registering Application in Microsoft Azure Portal.

# Customizing Send Settings

To customize send settings, perform the following steps:

1. In the **From** field, specify an email from which email notifications must be sent. Note that for OAuth 2.0 authentication, it must be the account you use to connect to the Veeam Backup & Replication application.
2. In the **To** field, specify the recipient addresses. Use a semicolon to separate multiple addresses. Recipients specified in this field will receive notification about every job managed by the backup server. You can leave the field empty if required.

   For every particular job, you can specify additional recipients. For more information, see Configuring Job Notification Settings.

   > **IMPORTANT**
   >
   > If you specify the same email recipient in both job notification and global notification settings, Veeam Backup & Replication will send the job notification only.

3. In the **Subject** field, specify a subject for the sent message. You can use the following variables in the subject:
   a. *%Time%* — completion time
   b. *%JobName%*
   c. *%JobResult%*
   d. *%ObjectCount%* — number of VMs in the job
   e. *%Issues%* — number of VMs in the job that have been processed with the *Warning* or *Failed* status
4. In the **Send daily reports at** field, specify at what time Veeam Backup & Replication will send daily email reports. Daily reports are generated for different purposes throughout Veeam Backup & Replication:
   - Reports about processing results of scale-out repository data. For more information, see Receiving Scale-Out Backup Repository Reports.
   - Reports about processing results of backup copy jobs. For more information, see Notification Settings in the **Creating Backup Copy Jobs for VMs and Physical Machines** section.
   - Reports about processing results of backup copy jobs for transaction log backups. For more information about transaction log backups, see Microsoft SQL Server Logs Backup.
   - Reports about backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux in the **Managed by Agent** mode.

- Reports with statistics for rescan job sessions performed for protection groups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.
- Reports about processing results of backup copy jobs for backups created by Veeam Plug-ins for Enterprise Applications.
- Reports about active Instant Recovery sessions, that is, sessions that were not finalized. For more information about Instant Recovery, see Instant Recovery to VMware vSphere and Instant Recovery to Microsoft Hyper-V.
- Reports about malware detection events that were created in the last 24 hours. For more information, see Notifications in the **Configuring Malware Detection** section.

> **NOTE**
>
> Settings configured for a certain report override global notification settings.

5. In the **Notify on** group, select the **Success**, **Warning** and **Failure** check boxes to receive email notification if a job is run successfully, not successfully or with a warning.
6. Select the **Suppress notifications until the last retry** check box to receive a notification about the final job status. If you do not enable this option, Veeam Backup & Replication will send one notification per every job retry. This option does not apply to immediate backup copy jobs.
7. Veeam Backup & Replication allows sending a test email to check if all settings have been configured correctly. To send a test email, click **Test Message**.

## Registering Application in Google Cloud Console

If you want to use your own web application for email notifications, you need to configure it in the Google Cloud console. To do this, perform the following steps:

1. Log in to the Google Cloud console under a Google account that has permissions to create applications.
2. Create a new project and enable Gmail API for the project. To do this, open **APIs and services** > **Library** > **Gmail API** > **Manage** and click **Enable API**.
3. Create OAuth credentials. To do this, perform the following steps:
   a. Open **APIs and services** > **Credentials**. Click **Create credentials** and select **OAuth client ID**.
   b. In the **Application type** field, select **Desktop app**.
   c. In the **Name** field, specify the name of your OAuth 2.0 client.
   d. Click **Create** to generate the application client ID and the client secret. In the opened window, you can copy credentials or download them in the JSON format. You can also find them later in the **APIs and services** > **Credentials** section when editing your OAuth 2.0 client ID.

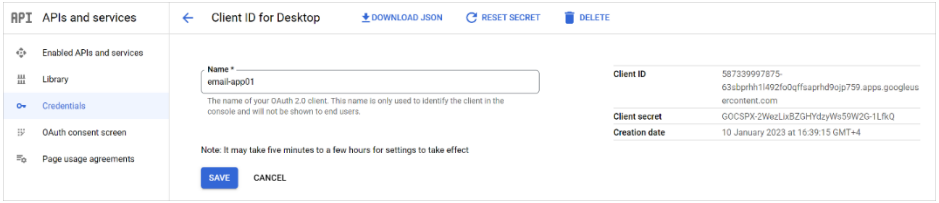


4. Open **APIs and services** > **OAuth consent screen** and click **Edit App**. Specify your application name and the user support email and click **Save and continue**.

5. If your application is in the **Testing** status, you must specify test users. To do this, at the **Test users** step of the **Edit App** wizard, click **Add users**. To apply changes, click **Save and continue**. Note that only test users will have access to the app.

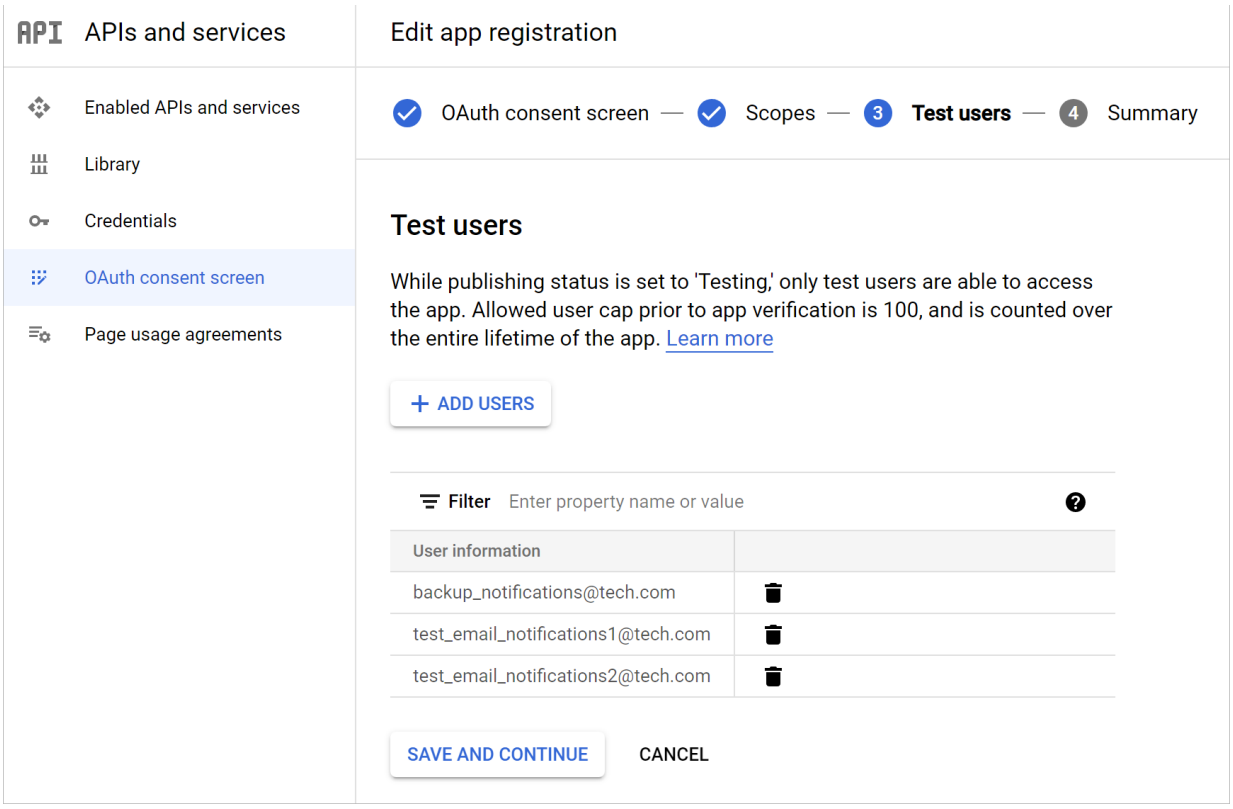


After you finish the registration, specify custom application registration settings when configuring the mail server for Google Gmail OAuth 2.0 authentication. For more information, see Configuring Mail Server.

> **NOTE**
>
> You can leave your application in the **Testing** status and do not publish it. In that case, you will receive a warning message *Google hasn't verified this app* when connecting to your application. If you want to verify it, see this Google article.

## Registering Application in Microsoft Azure Portal

If you want to use your own web application for email notifications, you need to configure it in the Microsoft Azure portal. To do this, perform the following steps:

1. Log in to the Microsoft Azure portal under Exchange Online credentials that has permissions to register Azure AD applications.
2. Register the application. To do this, open **Azure Active Directory** > **App registrations** and click **New registration:**
   a. In the **Name** field, specify the name of your application.
   b. In the **Supported account types** section, select the **Accounts in this organizational directory only** option.

c. Click **Register**.

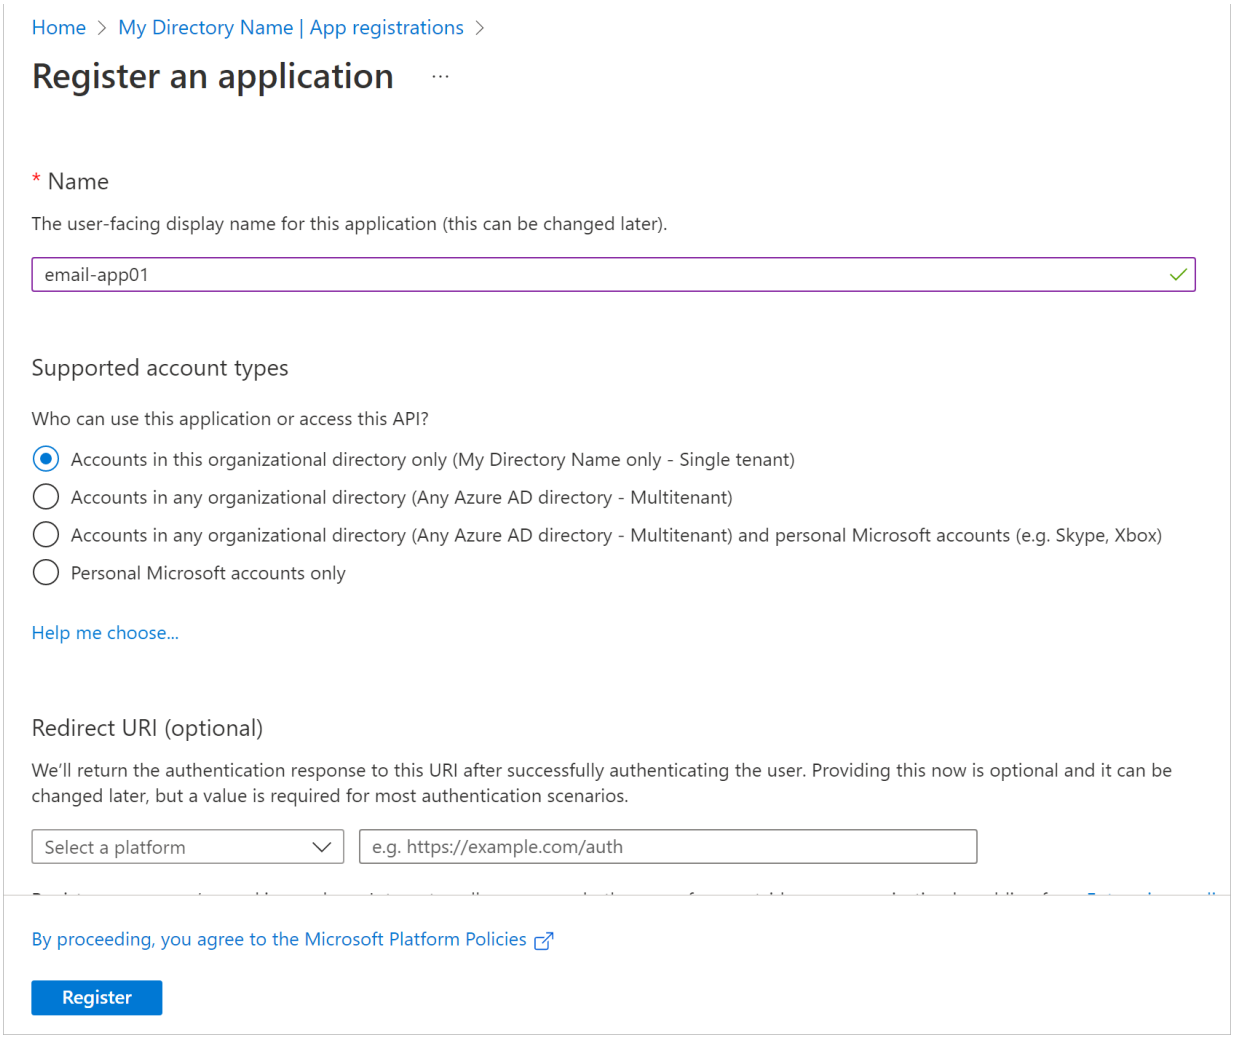


After registration, you can copy application (client) ID and directory (tenant) ID. You can also find these credentials later in the **Overview** section of you application properties.

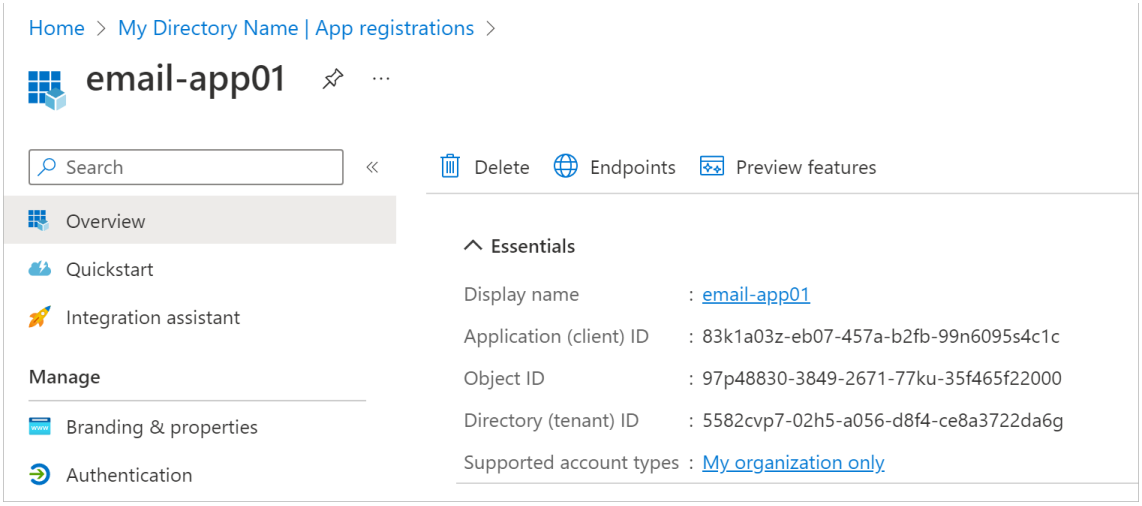


3. Add a platform configuration for your application. To do this, open **Authentication** > **Platform configurations** and click **Add a platform**:

   a. Select **Mobile and desktop applications**.

   b. Select the MSAL redirect URI generated in the following format: `msal<applicationid>://auth`.

c. Click **Configure**.

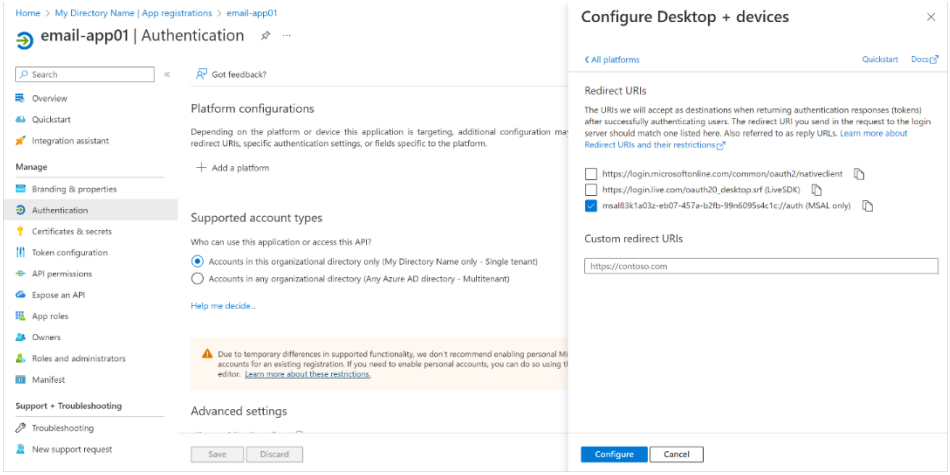


4. Grant the application the *Mail.Send* permission of Microsoft Graph. This will allow Veeam Backup & Replication to call the Microsoft Graph API for sending email notifications. To do this, open **API permissions** and click **Add a permission**:

   a. Select **Microsoft Graph** > **Application Permissions**.

   b. Select the *Mail.Send* permission from the list and click **Add permissions**.

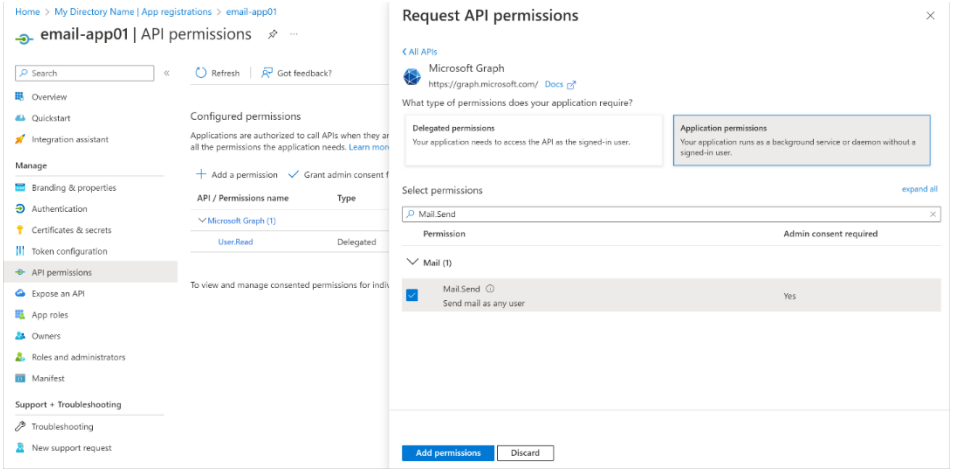


5. Click **Grant admin consent for <Your Directory Name>**. In the displayed window, click **Yes** to confirm the operation.

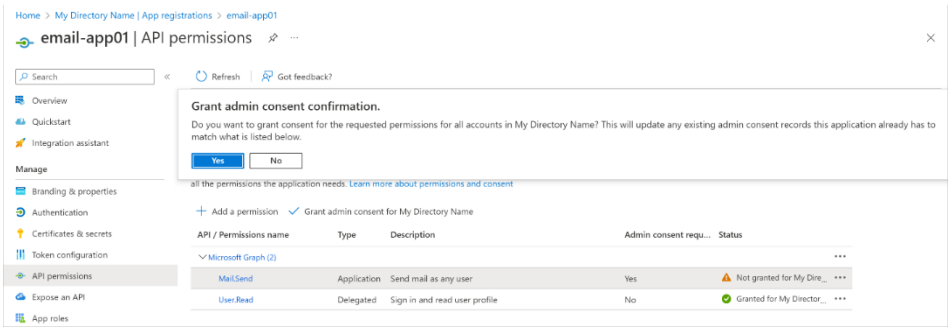


After you finish the registration, specify custom application registration settings when configuring the mail server for Microsoft 365 OAuth 2.0 authentication. For more information, see Configuring Mail Server .

# Configuring Job Notification Settings

To configure job notification settings:

1. Open advanced settings of the job.
2. On the **Notifications** tab, select the **Send email notifications to the following recipients** check box.
3. In the field under the **Send email notifications to the following recipients** check box, enter an email address to which a notification must be sent. You can enter several email addresses separated with a semicolon.

   > **IMPORTANT**
   > If you specify the same email recipient in both job notification and global notification settings, Veeam Backup & Replication will send the job notification only.

4. You can choose to use global notification settings for the job or specify custom notification settings.
   - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see Configuring Global Email Notification Settings.
   - To configure a custom notification for the job, select **Use custom notification settings** and specify notification settings as required.

5. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

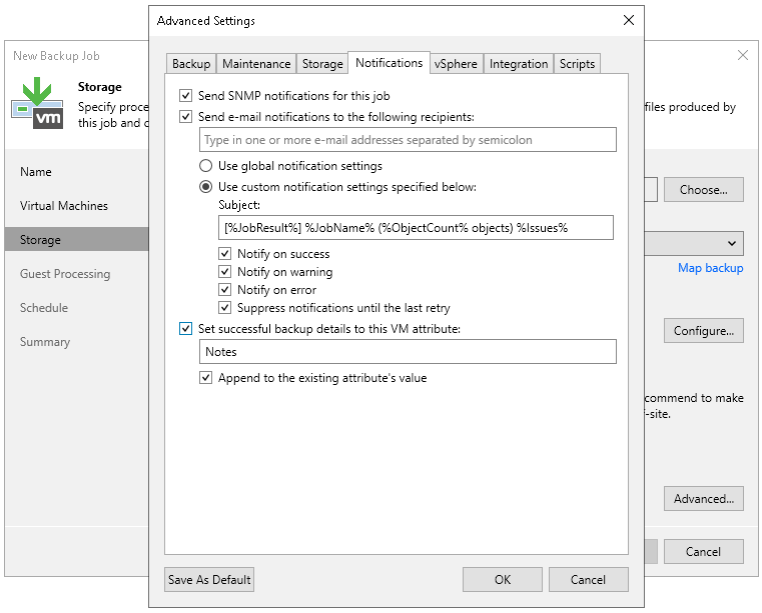

# Specifying SNMP Settings

You can receive SNMP traps with results on jobs performed on the backup server. You can use SNMP traps to feed data to other monitoring systems such as CA Unicenter, BMC Patrol, IBM Tivoli or HPE OneView. SNMP traps can be sent to 5 different destinations.

Veeam Backup & Replication supports SNMP versions 1 and 2 (including including v2c and v2p).

> **TIP**
>
> You can find the list of available SNMP traps in the `VeeamBackup.mib` file. The backup server stores this file in the `<vbr_installation_folder>\Backup\` folder. The installation folder is specified at the Program features step of the installation wizard. To interpret the traps incoming from the backup server, import the `VeeamBackup.mib` file to your recipient systems.

To receive SNMP traps, you must perform the following tasks:

- Configure global SNMP settings.
- Configure SNMP service properties.
- Configure SNMP settings for jobs.

# Configuring Global SNMP Settings

To configure global SNMP settings:

1. From the main menu, select **Options** > **Event Forwarding**.
2. In the **SNMP servers** window, Click **Add**.
3. In the **Receiver** field, specify an IPv4 or IPv6 address of the SNMP recipient. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
4. In the field on the right, enter the port number to be used.
5. In the **Community string** field, enter the community identifier.

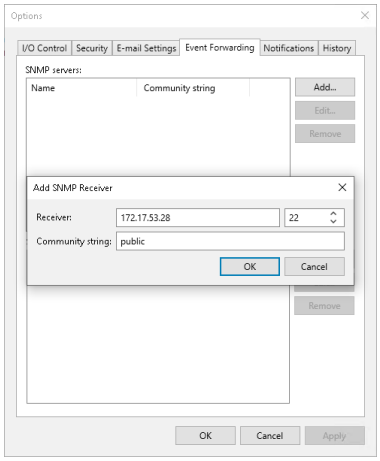

# Configuring SNMP Service Properties

To configure SNMP service properties on recipient systems:

1. Install a standard Microsoft SNMP agent from the Microsoft Windows distribution on the computer.
2. From the **Start** menu, select **Control Panel** > **Administrative Tools** > **Services**.
3. Double-click **SNMP Service** to open the **SNMP Service Properties** window.
4. Click the **Traps** tab.
5. Add the public string to the **Community name** list and name of the necessary host to the **Trap destinations** list.

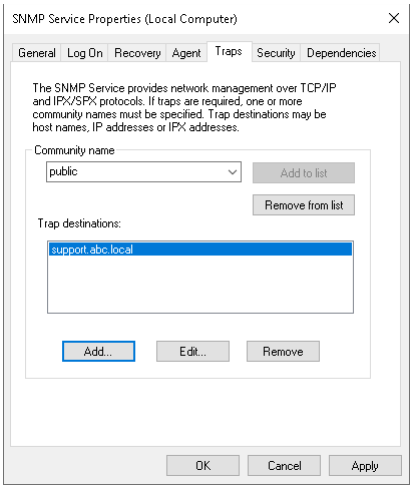


6. Click the **Security** tab.
7. Make sure the **Send authentication trap** check box is selected.
8. Add the public string to the **Accepted community names** list.
9. Select the **Accept SNMP packets from any host** check box.
10. Click **OK** to save changes.

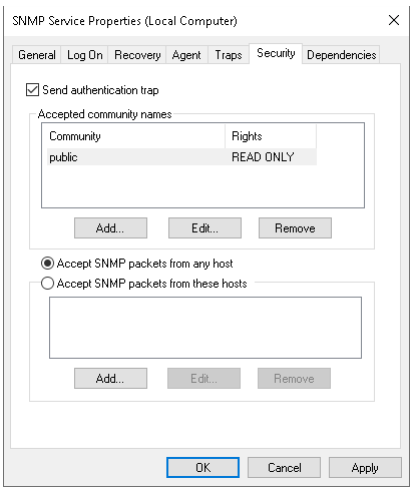

# Configuring SNMP Settings for Jobs

To receive SNMP traps with results of a specific job:

1. Open advanced settings of the job.
2. On the **Notifications** tab, select the **Send SNMP notifications for this job** check box.

3. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

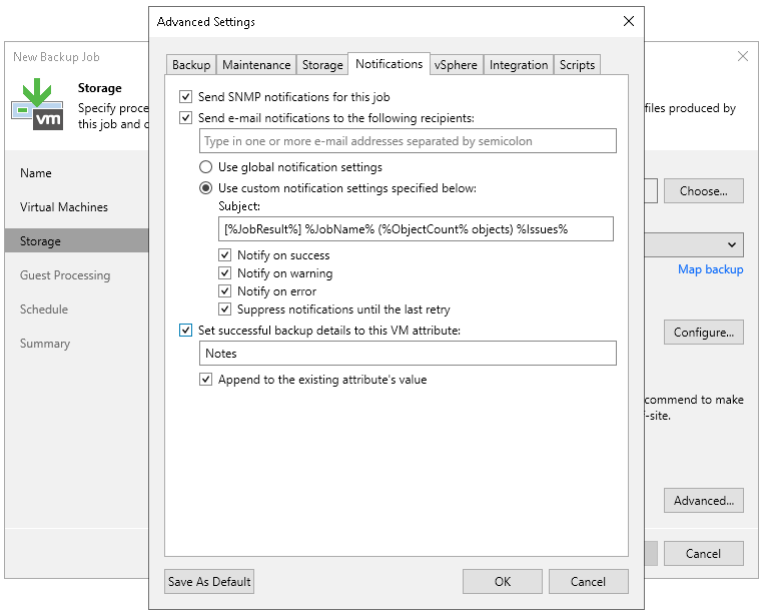

# Specifying Syslog Servers

You can use an external syslog server to manage events written by Veeam Backup & Replication.

Consider the following:

- To use this functionality, you must have a paid license.
- You can add only one syslog server.

## How integration with Syslog Server Works

When you add the syslog server in the Veeam Backup & Replication console, Veeam Backup & Replication sends a test event to check if it can communicate with the syslog server. Further, all events that Veeam Backup & Replication writes to Microsoft Windows Event Log will also be sent to the syslog server.

Each event contains a syslog message. The format of the message is defined by RFC 5424. For example:

```
2023-11-03T13:30:25.182677+01:00 <14> VBRSRV01 Veeam_MP  [categoryId=0 instance
Id=110 JobSessionID="58df29d6-a21b-43b2-a397-4c44ed1e05c1" JobID="cd13e656-8be9
-445a-bf9e-513b24293d35" JobType="0" Platform="0" Flags="0" Version="1" Descrip
tion="Backup job 'Backup Job 2' has been started."]
```

| Field | Description | Example |
|---|---|---|
| TIMESTAMP | Date and time. For more information about the format, see RFC 3339. | `2023-10-23T15:23:23.259882+02:00` |
| PRI | Message priority. | `<14>` |
| HOSTNAME | Host name of the backup server. | `VBRSRV01` |
| APP-NAME | Name of the application that generates events. | `Veeam_MP` |
| STRUCTURED-DATA | Event metadata. May include message details in the `Description` parameter. | `[categoryId=0 instanceId=110 JobSessionID="58df29d6-a21b-43b2-a397-4c44ed1e05c1" JobID="cd13e656-8be9-445a-bf9e-513b24293d35" JobType="0" Platform="0" Flags="0" Version="1" Description="Backup job 'Backup Job 2' has been started."]` |

| | | |
|---|---|---|
| 2023-11-03T13:30:25.182677+01:00 <14> VBRSRV01 Veeam_MP [categoryId=0 instanceId=110 JobSessionID="58df29d6-a21b-43b2-a397-4c44ed1e05c1" JobID="cd13e656-8be9-445a-bf9e-513b24293d35" JobType="0" Platform="0" Flags="0" Version="1" Description="Backup job 'Backup Job 2' has been started."] | | |
| MSG | Message details. May not be sent if message details are included in the STRUCTURED-DATA field. | Backup job 'Backup Job 2' has been started. |

**NOTE**

The structure and the content of the syslog message may vary for different syslog servers. For the full list of fields that can be sent in a syslog message, see the Syslog Message Format section in RFC 5424.

# Adding Syslog Server

To add a syslog server, do the following:

1. From the main menu, select **Options** > **Event Forwarding**.
2. In the **Syslog servers** section, click **Add**.
3. In the **Server** field, specify the FQDN or IPv4 address of the server. You cannot specify the IPv6 address in this field.
4. In the **Transport** field, specify the transport protocol: TCP, UDP or TLS. Default port numbers are 514 (for TCP and UDP) and 6514 (for TLS).
5. Click **OK**.

   **NOTE**

   If the syslog server is unavailable, the error message will be displayed.

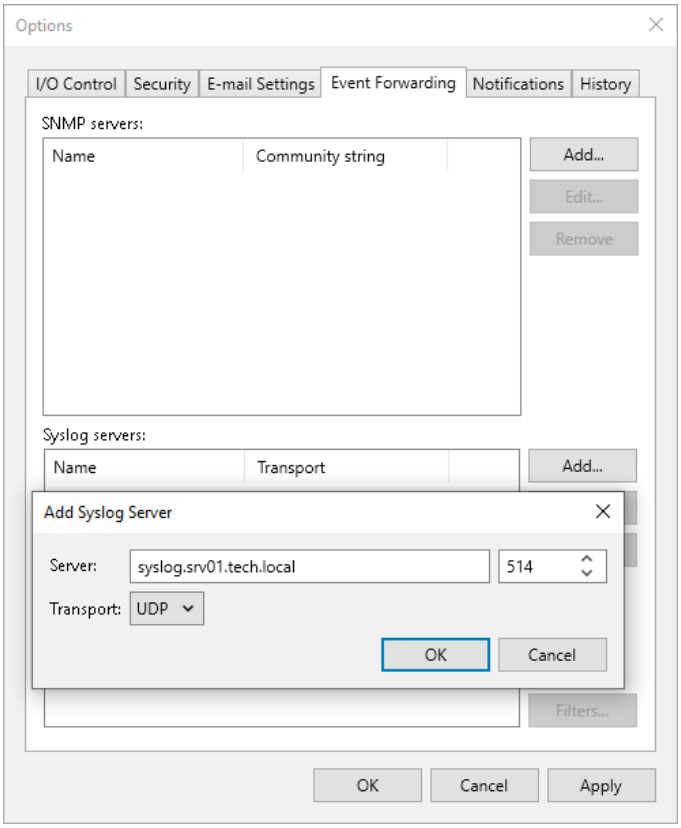


**TIP**

If required, you can also configure the following specific parameters on the backup server:

- Add the Unicode byte order mask (BOM) before the MSG field
- [For TCP or TLS connections] Use the octet count prefix as a syslog message delimiter instead of \n character
- [For TCP connections] Specify custom connection timeout

To configure these parameters, see this Veeam KB article.

# Filtering Events

You can exclude specific events from sending them to the syslog server. To do this, perform the following steps:

1. From the main menu, select **Options** > **Event Forwarding**.
2. In the **Syslog servers** section, click **Filters**.
3. In the **Syslog Event Filtering** window, click **Add**.
4. Enter the event ID and select the severity level. For more information about event IDs, see Event Reference.

5. Click **OK**.

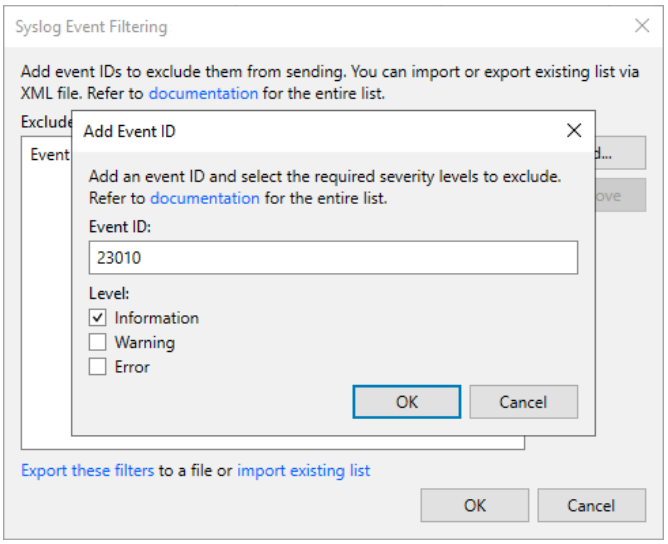


You can also export and import the list of event IDs in the XML format.

# Exporting Event IDs

To export event IDs, do the following:

1. From the main menu, select **Options** > **Event Forwarding**.
2. In the **Syslog servers** window, click **Filters**.
3. Click **Export these filters**.
4. Click **Browse** and select the folder to save the file.
5. Specify the name of the file and click **Save**.
6. Click **OK**.

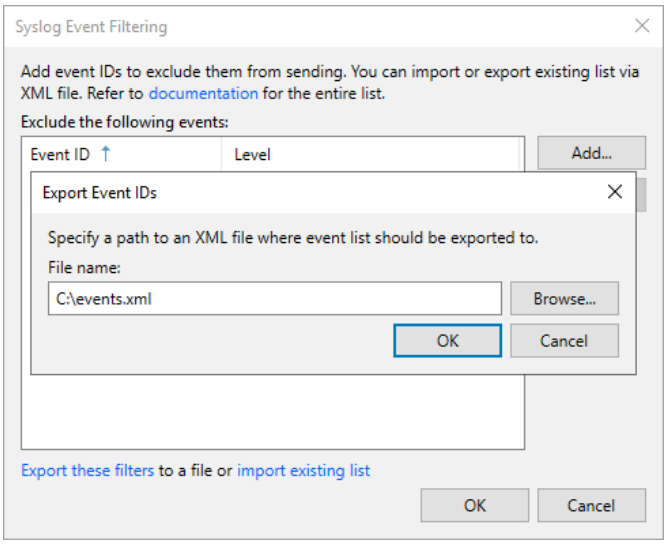

# Importing Event IDs

To import event IDs, do the following:

1. From the main menu, select **Options** > **Event Forwarding**.
2. In the **Syslog servers** window, click **Filters**.
3. Click **Import existing list**.

   > **IMPORTANT**
   >
   > The import operation will override event IDs you added earlier. It is recommended to export your current list before you import a new one.

4. Click **Browse** and select the folder where the file is located.
5. Select the file and click **Open**.
6. Click **OK**.

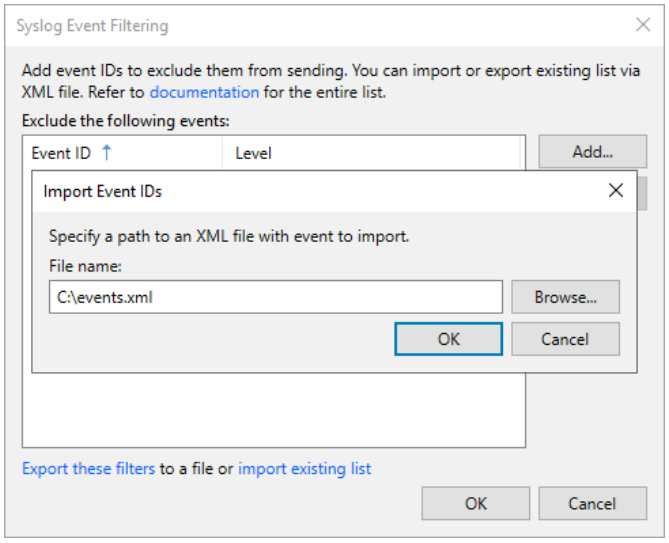


A valid XML Schema for the list must contain the following elements:

| Element | Description |
|---|---|
| `SyslogFilteredEvents` | Root element. |

| Element | Description |
| --- | --- |
| `FilteredEvent` | Child element of the `SyslogFilteredEvents` element. Each `FilteredEvent` contains an event that must be excluded from sending.<br><br>The element has the following attributes:<br>• `EventId` — Event ID.<br>• `FilterInfo` — Defines whether to exclude this event if it has the *Information* severity level. Possible values: `True`, `False`.<br>• `FilterWarning` — Defines whether to exclude this event if it has the *Warning* severity level. Possible values: `True`, `False`.<br>• `FilterError` — Defines whether to exclude this event if it has the *Error* severity level. Possible values: `True`, `False`. |

Example:

```
<SyslogFilteredEvents>
 <FilteredEvent EventId="23010" FilterInfo="True" FilterWarning="False" FilterE
rror="False" />
</SyslogFilteredEvents>
```

# Specifying Other Notification Settings

You can configure Veeam Backup & Replication to automatically notify you about the following events:

- Low disk space
- Support contract expiration
- New product versions, available updates

## Low Disk Space Notification

When you run a job, Veeam Backup & Replication checks disk space in the target backup repository and production storage. If the disk space is less than a specific value, Veeam Backup & Replication will display a warning message in the job session details.

To specify the disk space threshold:

1. From the main menu, select **Options**.
2. Click the **Notifications** tab.
3. In the **Backup storage** and **Production datastores** sections, select the **Warn me when free disk space is below <N> %** options and specify a desired disk space threshold.
4. In the **Production datastores** section, select the **Skip VMs when free disk is below <N> %** option and specify a desired disk space threshold. When the threshold is reached, Veeam Backup & Replication will terminate backup and replication jobs working with production datastores before VM snapshots are taken. Such behaviour helps ensure that production datastores do not run out of space.

   Veeam Backup & Replication also terminates jobs if the amount of free space on the datastore is less than 2 GB. You can change this threshold limit with registry values. For more information, see this Veeam KB article.

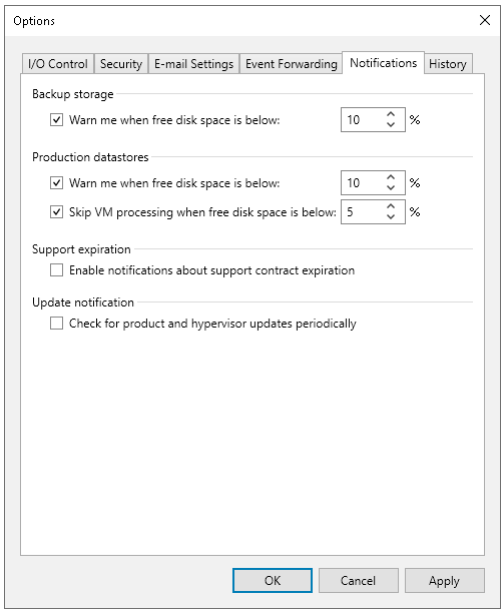

# Support Contract Expiration Notification

By default, Veeam Backup & Replication informs email recipients specified in global notification settings about the support expiration date in every email notification. Veeam Backup & Replication starts sending such notifications 14 days before the expiration date. Expiration information is also shown on the splash screen and on the **License Information** window (to display the **License Information** window, select **Help** > **License** from the main menu).

To stop receiving notifications about support contract expiration:

1. From the main menu, select **Options**.
2. Click the **Notifications** tab.
3. Clear the **Enable notifications about support contract expiration** check box.

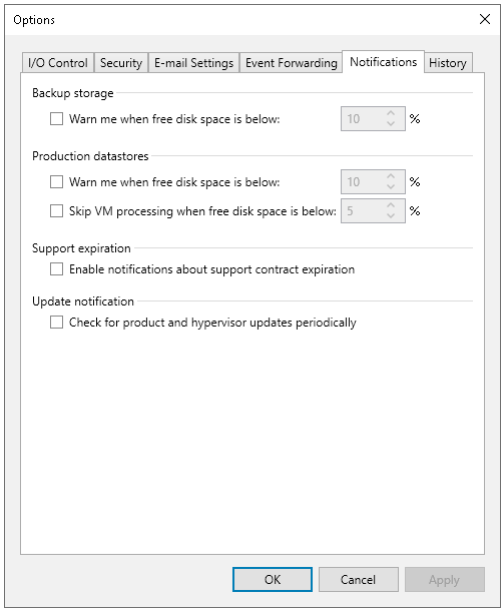


# New Product Versions

Veeam Backup & Replication automatically checks and notifies you about new product versions and updates available on the Veeam website. To receive the notifications, make sure that the backup server is connected to the internet. Otherwise, you will not receive the notifications about updates. For more information on notifications, see Update Notification.

To disable the notifications:

1. From the main menu, select **Options**.
2. Click the **Notifications** tab.
3. Clear the **Check for product and hypervisor updates periodically** check box.

> **IMPORTANT**
>
> We recommend you to leave the update notifications enabled. This will help you not to miss critical updates.

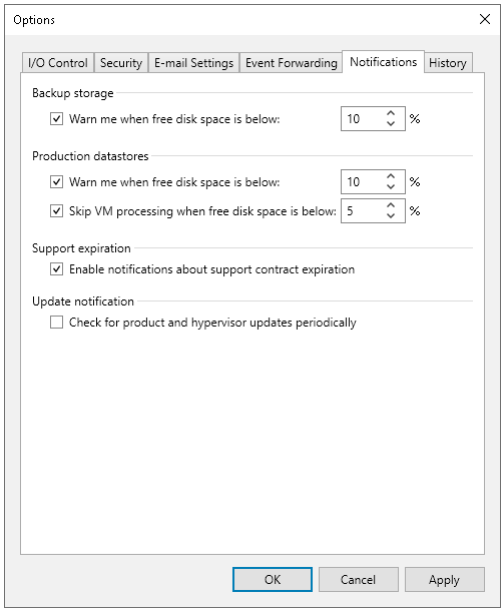
Options
I/O Control
Security
E-mail Settings
Event Forwarding
Notifications
History
Backup storage
Warn me when free disk space is below:
10
%
Production datastores
Warn me when free disk space is below:
10
%
Skip VM processing when free disk space is below:
5
%
Support expiration
Enable notifications about support contract expiration
Update notification
Check for product and hypervisor updates periodically
OK
Cancel
Apply

# Specifying Session History Settings

You can specify session history settings for jobs performed on the backup server.

1. From the main menu, select **Options**.
2. Click the **History** tab.
3. In the **Sessions** section, configure how many sessions you want to display in the **Sessions** list of the History view:
   - Select **Show all sessions** if you want to show all sessions.
   - Select **Show only last sessions** if you want to show a limited number of sessions. In the field, also specify the total number of sessions that you want to display.
4. In the **Session history retention** section, configure for how long you want to keep session information in the database:
   - Select **Keep all sessions** if you do not want to delete sessions.
   - Select **Keep only last weeks** if you want to keep sessions for a limited period of time. In the field, also specify this period in weeks.

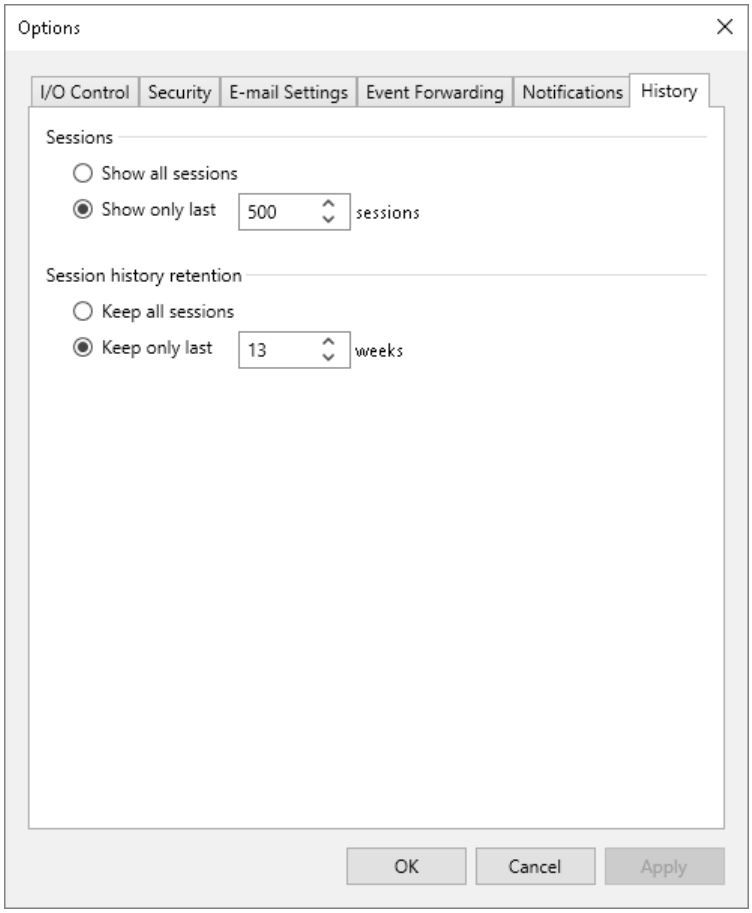

# Managing Users and Roles

This section describes how to configure Veeam Backup & Replication users and user groups and how to set up 4-eyes authorization for sensitive operations.

# Configuring Users

To perform Veeam Backup & Replication operations, you can add users and user groups with the following roles:

| Role | Operations |
|---|---|
| **Veeam Backup Administrator** | Can perform all administrative activities in the Veeam Backup & Replication console. Note that Veeam Backup Administrator has full access to all files on servers and hosts added to the backup infrastructure. |
| **Veeam Security Administrator** | Can perform the following operations:<br>• Add, edit and delete all types of credential records supported by Veeam Backup & Replication. For more details, see Managing Credentials.<br>• Manage Security & Compliance Analyzer: run security checks, configure scan scheduling, exclude parameters from the checklist. For more details, see Security & Compliance Analyzer.<br>• Approve four-eyes authorization requests. For more details, see Four-Eyes Authorization. |
| **Incident API Operator** | Can perform Veeam Backup & Replication REST API requests to manage malware detection events. For more details, see the Malware Detection section in the Veeam Backup & Replication REST API Reference.<br>Incident API Operators do not have access to the Veeam Backup & Replication console. Since they interact only with Veeam Backup & Replication REST API, make sure that multi-factor authentication is disabled for the user you add as Incident API Operator. For more details, see Disabling MFA for Service Accounts. |
| **Veeam Restore Operator** | Can perform restore operations using existing backups and replicas.<br>Consider the following:<br>• Veeam Restore Operator can restore data from any backups. This allows them to restore disks and files with specially crafted malicious content. To reduce the risk of privilege escalation attacks and the entire system takeovers, assign this role with caution.<br>• Veeam Restore Operator can overwrite existing VM, disk and file instances during VM, disk and file-level restore operations.<br>• Veeam Restore Operator cannot migrate a recovered VM to the production environment during Instant Recovery. |
| **Veeam Backup Operator** | Can start and stop existing jobs, export backups, copy backups and create VeeamZip backups. |

| Role | Operations |
| --- | --- |
| **Veeam Backup Viewer** | Has the read-only access to the Veeam Backup & Replication console. Can view a list of existing jobs and review the job session details. |
| **Veeam Tape Operator** | Can manage tapes and perform the following operations:<br>• Rescan tape libraries and tape servers.<br>• Run tape backup jobs, tape catalog jobs, tape inventory jobs, tape verification jobs.<br>• Eject tapes.<br>• Import and export tapes.<br>• Mark tapes as free.<br>• Move tapes to a media pool.<br>• Erase tapes.<br>• Set a tape password.<br>• Copy tapes. |

You can assign several roles to the same user. For example, if you want a user to start jobs and perform restore operations, you can assign both the *Veeam Backup Operator* and *Veeam Restore Operator* roles to this user.

# Requirements and Limitations

When you configure users, consider the following requirements and limitations:

- For security reasons, the account used to run Veeam services should be a LocalSystem account. If a Veeam service runs under a user account other than LocalSystem, this user will have full access to the Veeam Backup & Replication console even if they are not specified in the **Users and Roles** > **Security** settings.
- The user account under which the Veeam Backup Service runs must have the *Veeam Backup Administrator* role. By default, during installation the *Veeam Backup Administrator* role is assigned to all members of the *Administrators* group on the machine where Veeam Backup & Replication is installed.

  If you change the default settings, make sure that you assign the *Veeam Backup Administrator* role to the necessary user account. It is recommended to assign the *Veeam Backup Administrator* role to the user account explicitly rather than the group to which the user belongs.
- If multi-factor authentication (MFA) is disabled, consider the following:
  - Built-in administrator accounts (*Domain|Administrator* and *Machine|Administrator*) will have full access to the Veeam Backup & Replication console.
  - Local and domain members of the *Administrators* group will still have full access to Veeam Backup & Replication even if you delete this group in the **Users and Roles** > **Security** settings.

  To protect administrator accounts from being compromised, it is strongly recommended to enable multi-factor authentication. In that case, even users with administrator privileges must pass the additional verification. For more information, see Multi-Factor Authentication.

- If multi-factor authentication (MFA) is enabled, consider the following:
  - All users including built-in administrator accounts (*Domain|Administrator* and *Machine|Administrator*) must pass the additional verification.
  - If you do not add a local or domain member of the *Administrators* group in the **Users and Roles** > **Security** settings, this user will not have access to the Veeam Backup & Replication console.
  - If you add a local or domain member of the *Administrators* group in the **Users and Roles** > **Security** settings, this user will have full access to the Veeam Backup & Replication console regardless of the assigned role.
  - If a Veeam service runs under a user account other than LocalSystem, you must disable MFA for this account. For more information, see Disabling MFA for Service Accounts.

# Adding Users

To add a user or user group:

- From the main menu, select **Users and Roles** > **Security**.
- Click **Add**.
- In the **User or group** field, enter the name of a user or user group in the *DOMAIN|USERNAME* format.
- From the **Role** list, select the necessary role to be assigned.
- Click **OK**.

To reduce the number of user sessions opened for a long time, you can set the idle timeout to automatically log off users. To do this, select the **Enable auto logoff after <number> min of inactivity** check box and set the number of minutes.

To use additional user verification, you can enable multi-factor authentication. For more information, see Multi-Factor Authentication.

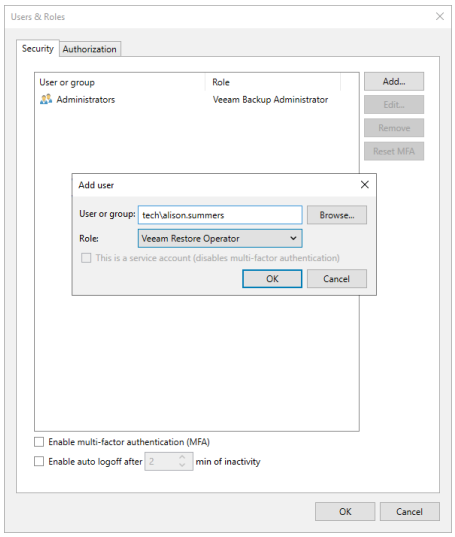


# Editing Users

To edit a user or user group:

1. From the main menu, select **Users and Roles** > **Security**.

2. Select a user or user group.
3. Click **Edit**.
4. In the **Edit** window, you can do the following:
   - Select a new role to be assigned to the account.
   - Disable multi-factor authentication for the account. For more information, see Disabling MFA for Service Accounts.
- Click **OK**.

## Removing Users

To remove a user or user group:

1. From the main menu, select **Users and Roles** > **Security**.
2. Select a user or user group.
3. Click **Remove**.

# Four-Eyes Authorization

You can enable four-eyes authorization to reduce the risk of accidental actions affecting sensitive data. This functionality requires additional approval for certain operations in Veeam Backup & Replication given by another user. To approve the request, the user must have the *Veeam Backup Administrator* or *Veeam Security Administrator* role.

> **IMPORTANT**
>
> Before you enable the feature, make sure that you have at least two users (added to a user group or separate ones) with the *Veeam Backup Administrator* or *Veeam Security Administrator* role assigned.

When enabled, four-eyes authorization is required to perform the following operations:

- Delete backup files or snapshots from the disk or configuration database.
- Delete information about unavailable backups from the configuration database.
- Remove backup repositories and storage from the backup infrastructure.
- Add, update and delete users or user groups.
- Enable and disable multi-factor authentication (MFA) for all users and user groups.
- Reset MFA for a specific user.
- Enable, update and disable automatic logoff for all users and user groups.
- Perform operations in the Veeam Cloud Connect infrastructure:
    - [For service providers] Remove cloud repositories and delete imported tenant backup files. Tenant backup files stored in Veeam Cloud Connect repositories cannot be deleted by service providers.
    - [For tenants] Remove service providers and delete backup files.

Consider that four-eyes authorization cannot protect the backup infrastructure if the Veeam Backup & Replication server is compromised. To build a more secure environment, follow security guidelines. For more details, see General Security Considerations and Securing Backup Infrastructure.

## How Four-Eyes Authorization Works

Veeam Backup & Replication supports the following scenario for four-eyes authorization:

1. A backup administrator tries to delete backup data or remove a repository from the backup infrastructure.
2. The request for additional approval is displayed in the **Home** view, under the **Pending approvals** node. Recipients specified in the global email notification settings also get email notifications. For more information, see Configuring Global Email Notification Settings.
3. A backup or security administrator approves or rejects the request. If there are multiple requests, the administrator can approve or reject them simultaneously. Recipients specified in the global email notification settings also get email notifications.

    > **NOTE**
    >
    > The backup administrator that created the request can only reject their own requests.

If no administrators process the request till the end of the specific time period (7 days by default), it will be automatically rejected.

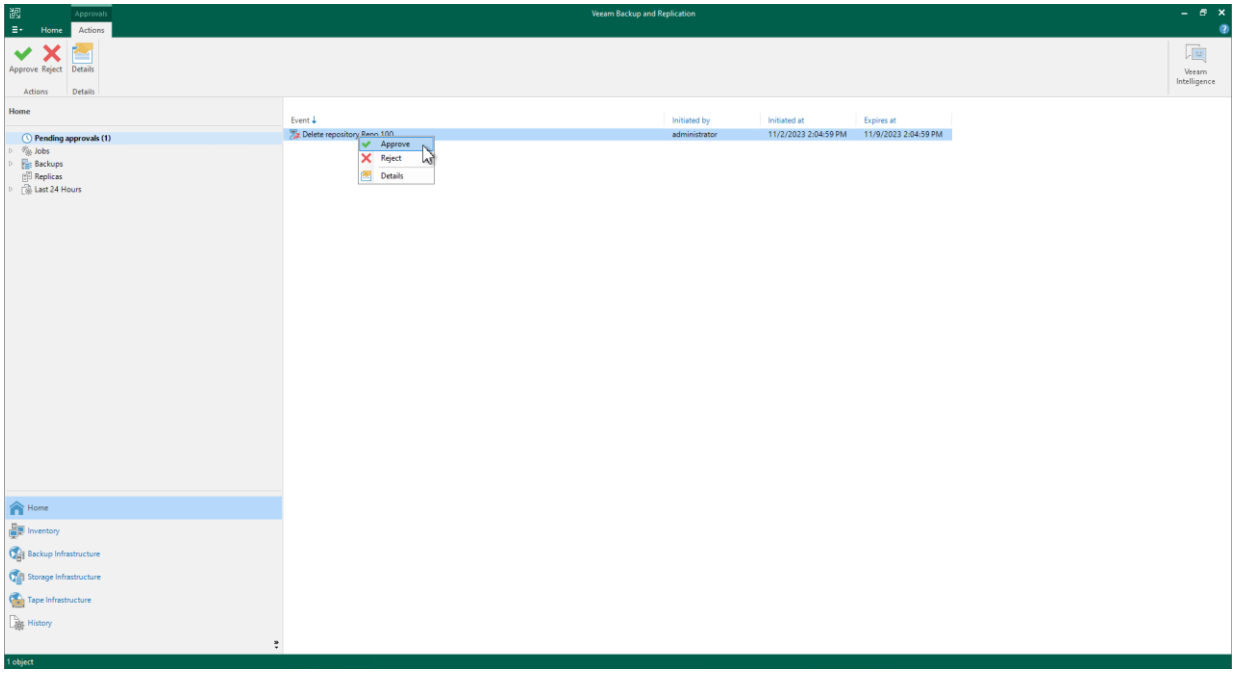


# Requirements and Limitations

Four-eyes authorization has the following requirements and limitations:

- The functionality is included only in the Veeam Universal License or the Enterprise Plus edition. If the license expires, you will still be able to process already created requests but not to create new ones.
- If four-eyes authorization is enabled, you cannot perform the following operations:
  - Delete operations using PowerShell cmdlets, REST API, and Veeam Backup Enterprise Manager.
  - Specific operations in the **Files** view:
    - Edit, rename and delete files
    - Overwrite files
    - Rename and delete folders
- If you try to approve or reject the request and the object that you want to delete is blocked by another operation, for example, by the job session, the operation will not be performed. In this case, you need to process the request later, when the object will not be blocked.
- Immutable backup files cannot be deleted even with the four-eyes authorization enabled.

# Enabling Four-Eyes Authorization

To enable four-eyes authorization, perform the following steps:

1. Make sure that you have at least two users (added to a user group or separate ones) with the *Veeam Backup Administrator* or *Veeam Security Administrator* role assigned.
2. From the main menu, select **Users and Roles** > **Authorization**.
3. Select the **Require additional approval for sensitive operations** check box.

- Specify the time period during which the requested operation must be approved or rejected (minimum 1 day, maximum — 30).

> **NOTE**
>
> To disable four-eyes authorization, you will also need an additional approval from another backup or security administrator.

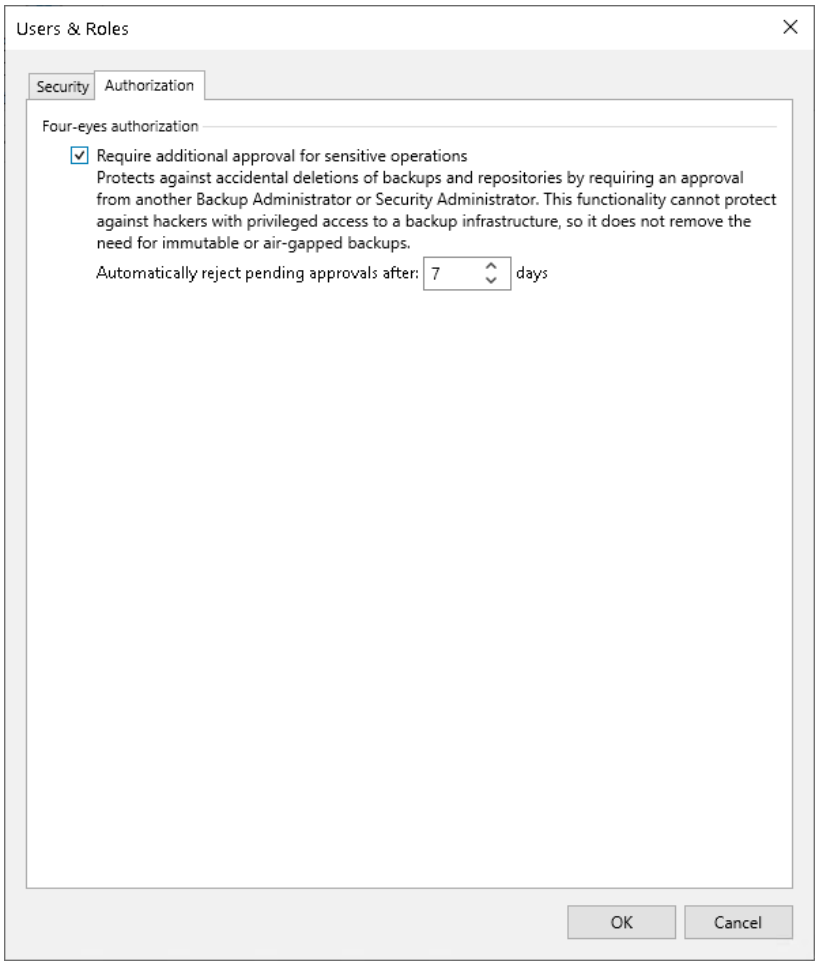


# Viewing Authorization Events

To view events related to four-eyes authorization, open the **History** view and select the **Authorization Events** node. These events include information about:

- Approved and rejected requests
- Updated four-eyes authorization settings
- Updated settings for users and user groups
- Assigned roles

- Added or deleted users and user groups

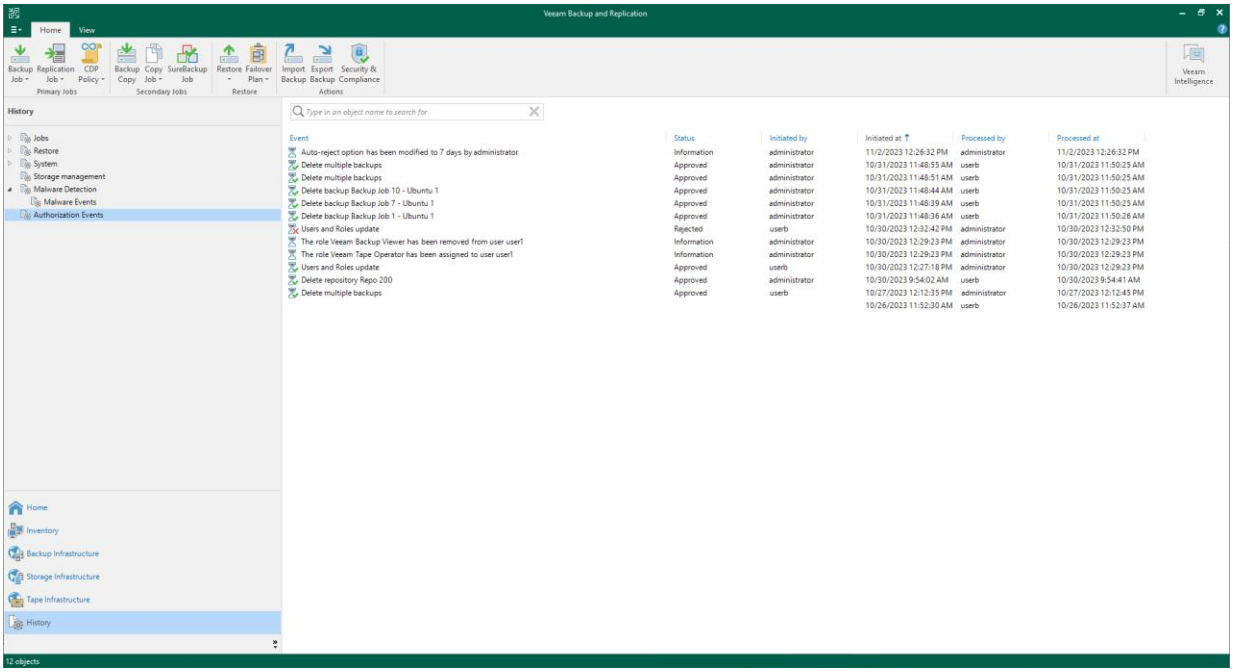

# Managing Credentials

This section describes different types of credential records that Veeam Backup & Replication supports.

# Credentials Manager

You can use the Credentials Manager to create and maintain a list of credentials records that you plan to use to connect to components in the backup infrastructure.

The Credentials Manager lets you create the following types of credentials records:

- Standard account
- Group Managed Service Accounts
- SSH credentials
- SSH private keys

Veeam Backup & Replication encrypts credentials stored in the configuration database using Data Protection API (DPAPI) mechanisms. For more information, see Microsoft Docs.

## Standard Accounts

By default, the Credentials Manager includes the following system standard credentials records:

- A credentials record for the provider-side network extension appliance
- A credentials record for the tenant-side network extension appliance
- A credentials record for the Microsoft Azure helper appliance (former Azure proxy)

You can create a credentials record for an account that you plan to use to connect to a Microsoft Windows server, vCenter server, ESXi host, a VM running Microsoft Windows OS, a storage system and others.

To create a new standard credentials record:

1. From the main menu, select **Credentials and Passwords** > **Datacenter Credentials**.
2. Click **Add** > **Standard account**.
3. In the **Username** field, enter a user name for the account that you want to add. You can also click **Browse** to select an existing user account.
4. In the **Password** field, enter a password for the account that you want to add. To view the entered password, click and hold the eye icon on the right of the field.

5. In the **Description** field, enter a description for the created credentials record. As there can be a number of similar account names, for example, *Administrator*, it is recommended that you provide a meaningful unique description for the credentials record so that you can distinguish it in the list. The description is shown in brackets, following the user name.

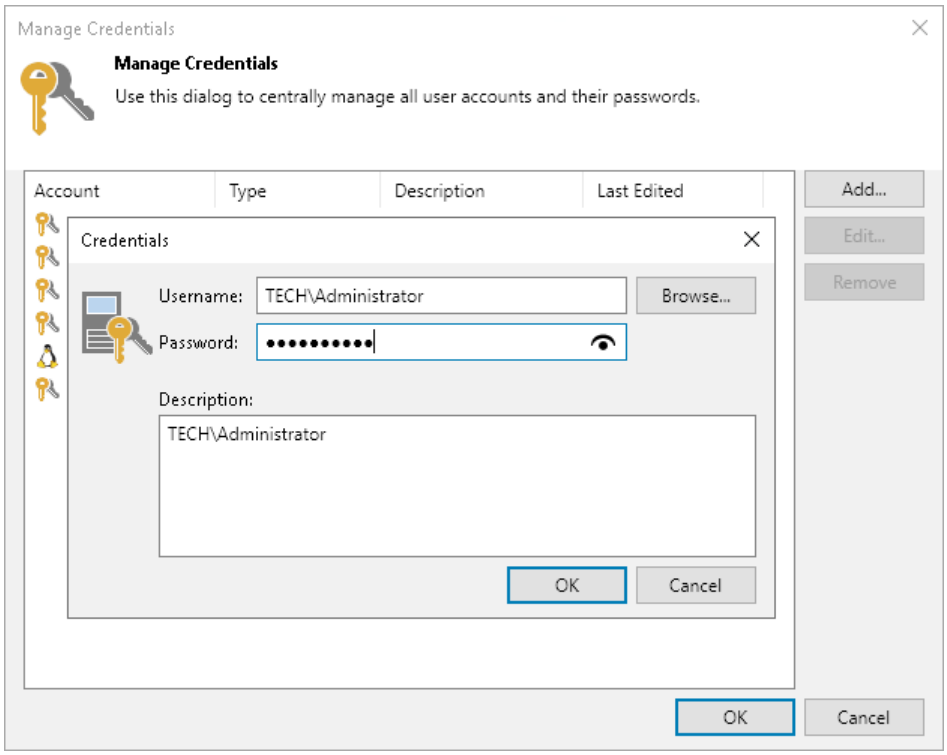


# Group Managed Service Accounts

You can add a credentials record for a Group Managed Service Account (gMSA) that you plan to use to run guest processing tasks. Note that before you add a record you should check requirements and limitations and perform actions described in section Using Group Managed Service Accounts.

To add a new credentials record with the gMSA:

1. From the main menu, select **Credentials and Passwords** > **Datacenter Credentials**.
2. Click **Add** > **Managed service account**.
3. In the **Username** field, enter a gMSA that you want to add. Note that the name of the account must be specified in the following format: DOMAIN\User or user@domain.xxx.

4. In the **Description** field, enter a description for the created credentials record. As there can be a number of similar account names, for example, *gmsa01*, it is recommended that you provide a meaningful unique description for the credentials record so that you can distinguish it in the list. The description is shown in brackets, following the user name.

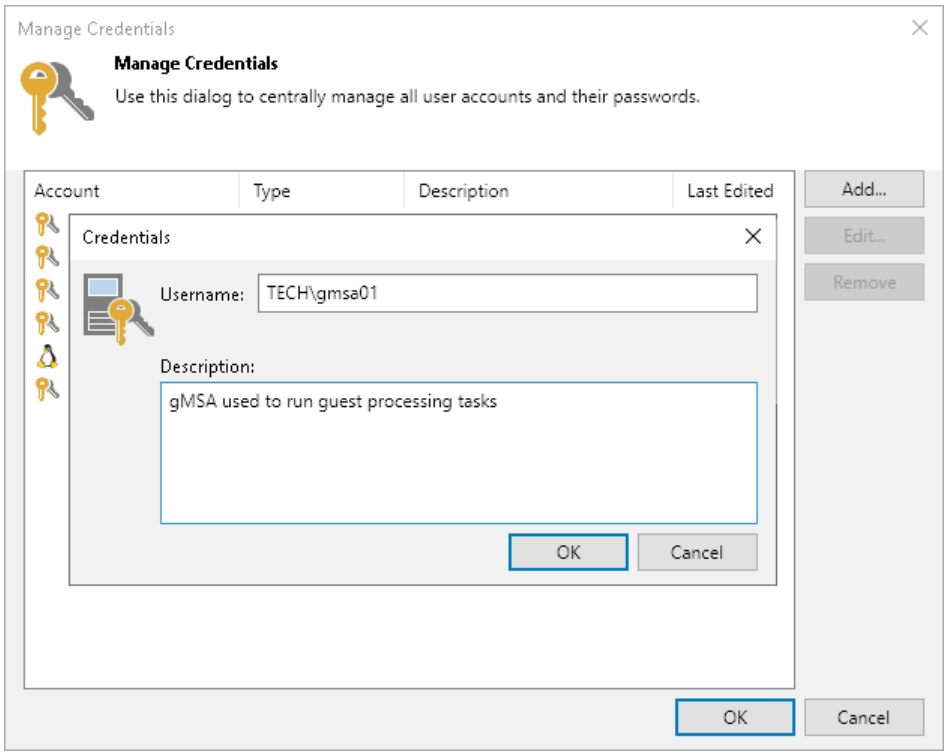


# SSH Credentials

By default, the Credentials Manager includes a system SSH credentials record for the helper appliance. You can create a credentials record for the account that you plan to use to connect to a Linux server or VM running Linux OS.

To create a new credentials record with a user name and password for a Linux server:

1. From the main menu, select **Credentials and Passwords** > **Datacenter Credentials**.
2. Click **Add** > **SSH credentials** (**Linux account** - before Veeam Backup & Replication 12.1 (build 12.1.0.2131)).
3. In the **Username** field, enter a user name for the account that you plan to add.
4. In the **Password** field, enter a password for the account that you want to add. To view the entered password, click and hold the eye icon on the right of the field.
5. In the **SSH port** field, specify the SSH port over which you want to connect to a Linux server. By default, port 22 is used.
6. If you specify data for a non-root account that does not have root permissions on a Linux server, you can use the **Non-root account** section to grant sudo rights to this account.
   a. To provide a non-root user with root account privileges, select the **Elevate account privileges automatically** check box.
   b. To add the user account to `sudoers` file, select the **Add account to the sudoers file** check box. In the **Root password** field, enter the password for the root account.

      If you do not enable this option, you will have to manually add the user account to the `sudoers` file.

c. If you plan to use the account to connect to Linux servers where the `sudo` command is not available or fails, you have an option to use the `su` command instead. To enable the su command, select the **Use "su" if "sudo" fails** check box and in the **Root password** field, enter the password for the root account.

   Veeam Backup & Replication will first try to use the `sudo` command. If the attempt fails, Veeam Backup & Replication will use the `su` command.

7. In the **Description** field, enter a description for the created credentials record. As there can be a number of similar account names, for example, *Root*, it is recommended that you provide a meaningful unique description for the credentials record so that you can distinguish it in the list. The description is shown in brackets, following the user name.

**IMPORTANT**

Consider the following:

- You can create a separate user account intended for work with Veeam Backup & Replication on a Linux-based VM, grant root privileges to this account and specify settings of this account in the Credentials Manager. It is recommended that you avoid additional commands output for this user (like messages echoed from within ~/.bashrc or command traces before execution) because they may affect Linux VM processing.
- Cases when root password is required to elevate account rights to root using sudo are no longer supported.

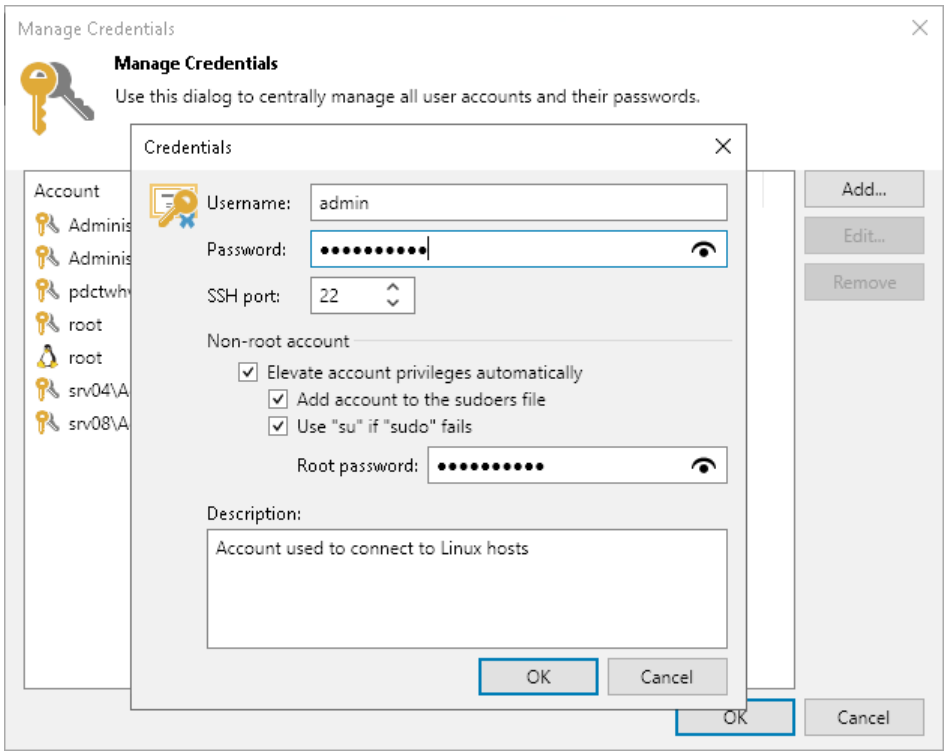


# SSH Private Keys

You can log on to a Linux server or VM running Linux OS using the Identity/Pubkey authentication method. The Identity/Pubkey authentication method helps protect against malicious applications like keyloggers, strengthens the security level and simplifies launch of automated tasks.

To use the Identity/Pubkey authentication method, you must generate a pair of keys — a public key and private key:

- Public key is stored on Linux servers to which you plan to connect from the backup server. The key is kept in a special `authorized_keys` file containing a list of public keys.
- Private key is stored on the client machine — backup server. The private key is protected with a passphrase. Even if the private key is intercepted, the eavesdropper will have to provide the passphrase to unlock the key and use it.

For authentication on a Linux server, the client must prove that it has the private key matching the public key stored on the Linux server. To do this, the client generates a cryptogram using the private key and passes this cryptogram to the Linux server. If the client uses the "correct" private key for the cryptogram, the Linux server can decrypt the cryptogram with a matching public key.

Veeam Backup & Replication has the following limitations for the Identity/Pubkey authentication method:

- Veeam Backup & Replication does not support keys that are stored as binary data, for example, in a file of DER format.
- Veeam Backup & Replication supports only keys whose passphrase is encrypted with algorithms supported by PuTTY:
  - AES (Rijndael): 128-bit, 192-bit and 256-bit CBC or CTR (SSH-2 only)
  - Blowfish: 128-bit CBC
  - Triple-DES: 168-bit CBC

  > **TIP**
  >
  > Veeam Backup & Replication 12 supports PPK file versions 2 and 3.

- Passphrases generated in PuTTY must only contain ASCII characters. Unicode characters can create decoding issues in Veeam Backup & Replication.

Veeam Backup & Replication supports the following key algorithms: RSA, DSA, ECDSA, EdDSA (ED25519). For these algorithms you can use the following key formats:

| Key Formats | Key Algorithms | | | |
|---|---|---|---|---|
| | RSA | DSA | ECDSA | EdDSA (ED25519) |
| PEM | ● | ● | ● | ○ |
| private | ● | ● | ● | ● |
| private-openssh | ● | ● | ● | ○ |
| sshcom | ● | ● | ○ | ○ |

| Key Formats | Key Algorithms | | | |
|---|---|---|---|---|
| | RSA | DSA | ECDSA | EdDSA (ED25519) |
| PKCS8 | ● | ○ | ● | ○ |
| RDC4716 (private-openssh-new) | ● | ● | ● | ● |

**IMPORTANT**

If you use VMware VIX/vSphere Web Services, Veeam Backup & Replication does not support usage of public keys for guest processing on Linux guest servers.

To add a credentials record using the Identity/Pubkey authentication method, do the following:

1. Generate a pair of keys using a key generation utility, for example, ssh-keygen.

   Note that keys generated as ED448 as are not supported.

2. Place the public key on a Linux server. To do this, add the public key to the `authorized_keys` file in the `.ssh/` directory in the home directory on the Linux server.
3. Place the private key in some folder on the backup server or in a network shared folder.
4. In Veeam Backup & Replication, from the main menu select **Credentials and Passwords** > **Datacenter Credentials**.
5. Click **Add** > **SSH private key**.
6. In the **Username** field, specify a user name for the created credentials record.
7. In the **Password** field, specify the password for the user account. The password is required in all cases except when you use root or a user with enabled *NOPASSWD:ALL* setting in */etc/sudoers*.
8. In the **Private key** field, enter a path to the private key or click **Browse** to select a private key.
9. In the **Passphrase** field, specify a passphrase for the private key on the backup server. To view the entered passphrase, click and hold the eye icon on the right of the field.
10. In the **SSH port** field, specify a number of the SSH port that you plan to use to connect to a Linux server. By default, port 22 is used.
11. If you specify data for a non-root account that does not have root permissions on a Linux server, you can use the **Non-root account** section to grant sudo rights to this account.

    a. To provide a non-root user with root account privileges, select the **Elevate account privileges automatically** check box.

    b. To add the user account to sudoers file, select the **Add account to the sudoers file** check box. In the **Root password** field, enter the password for the root account.

    If you do not enable this option, you will have to manually add the user account to the sudoers file.

c. When registering a Linux server, you have an option to failover to using the su command for distros where the sudo command is not available.

   To enable the failover, select the **Use "su" if "sudo" fails** check box and in the **Root password** field, enter the password for the root account.

12. In the **Description** field, enter a description for the created credentials record. As there can be a number of similar account names, for example, *Root*, it is recommended that you supply a meaningful unique description for the credentials record so that you can distinguish it in the list. The description is shown in brackets, following the user name.

**IMPORTANT**

Cases when root password is required to elevate account rights to root using sudo are no longer supported.

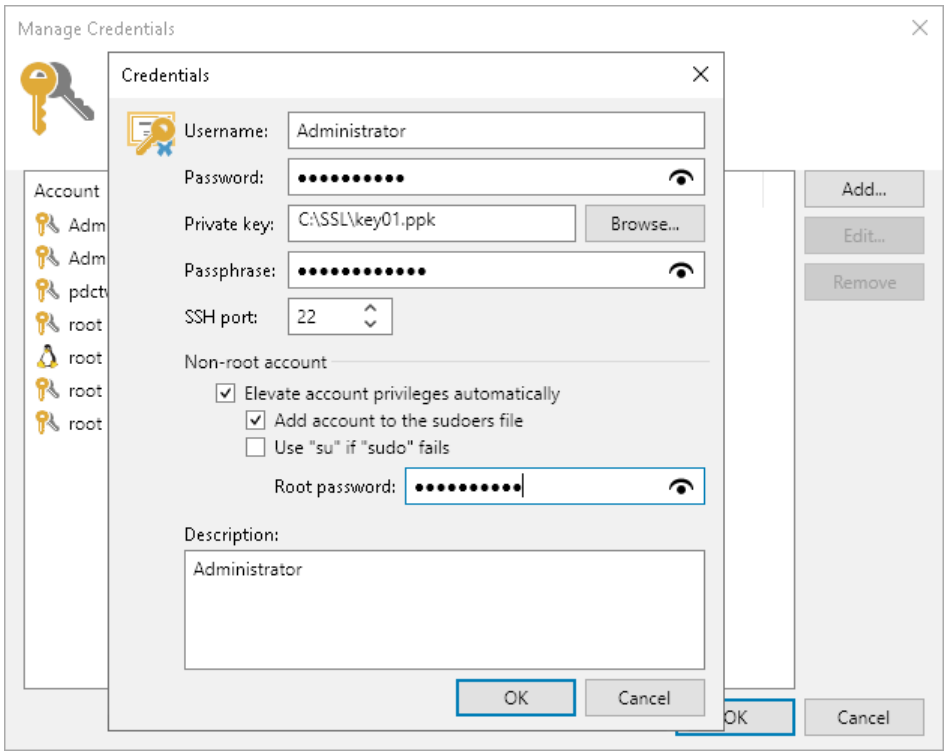


# Editing and Deleting Credentials Records

You can edit or delete credentials records that you have created. For the system credentials records, you can only change a password and record description. These credentials records cannot be deleted.

To edit a credentials record:

1. From the main menu, select **Credentials and Passwords** > **Datacenter Credentials**.
2. Select the credentials record in the list and click **Edit**.
3. If the credentials record is already used for any component in the backup infrastructure, Veeam Backup & Replication will display a warning. Click **Yes** to confirm your intention.
4. Edit settings of the credentials record as required.

To delete a credentials record:

1. From the main menu, select **Credentials and Passwords > Datacenter Credentials**.

2. Select the credentials record in the list and click **Remove**.

> **NOTE**
>
> You cannot delete a record that is already used for any component in the backup infrastructure. If you still need to do it, use a temporary record with dummy credentials for the required component and perform the rescan operation for this component. After that, you will be able to delete the record.

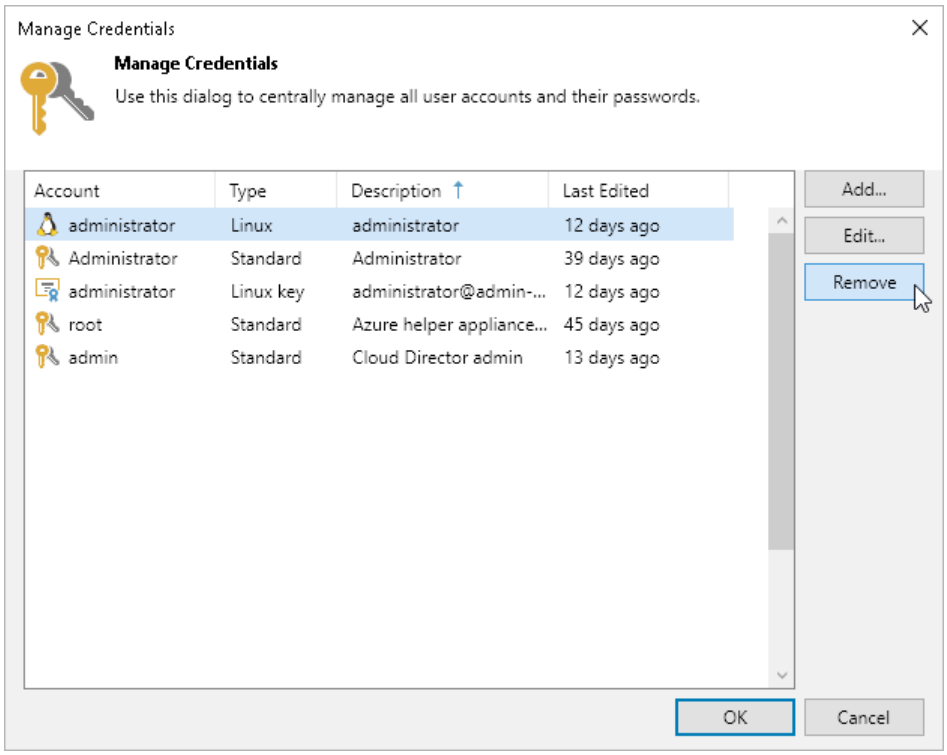

# Cloud Credentials Manager

You can use the Cloud Credentials Manager to create and maintain a list of credentials records that you plan to use to connect to cloud services.

The Cloud Credentials Manager lets you create the following types of credentials records:

- Veeam Cloud Connect Accounts
- Access Keys for AWS Users
- Microsoft Azure Storage Accounts (Shared Key)
- Microsoft Azure Compute Accounts
- Microsoft Azure Storage Accounts (Entra ID)
- Microsoft Azure Stack Hub Compute Accounts
- Google Cloud Accounts
- Google Cloud Service Accounts

Veeam Backup & Replication encrypts credentials stored in the configuration database using Data Protection API (DPAPI) mechanisms. For more information, see Microsoft Docs.

## Veeam Cloud Connect Accounts

You can add a credentials record for a tenant account — an account that you plan to use to connect to a service provider (SP).

Before you add a credentials record, the SP must register a tenant account on the SP Veeam backup server. Tenants without accounts cannot connect to the SP and use Veeam Cloud Connect resources. For more information, see the Registering Tenant Accounts section in the Veeam Cloud Connect Guide.

To create a credentials record for a tenant account:

1. From the main menu, select **Credentials and Passwords** > **Cloud Credentials**.
2. Click **Add** > **Veeam Cloud Connect service provider account**.
3. In the **Username** field, enter a user name for the account that the SP has provided to you.

   > **NOTE**
   >
   > If the SP used VMware Cloud Director to allocate replication resources to you, you must enter a user name for the VMware Cloud Director tenant account in the following format: *Organization*|*Username*. For example: *TechCompanyOrg*|*Administrator*.

4. In the **Password** field, enter a password for the account that the SP has provided to you. To view the entered password, click and hold the eye icon on the right of the field.

5. In the **Description** field, enter a description for the created credentials record.

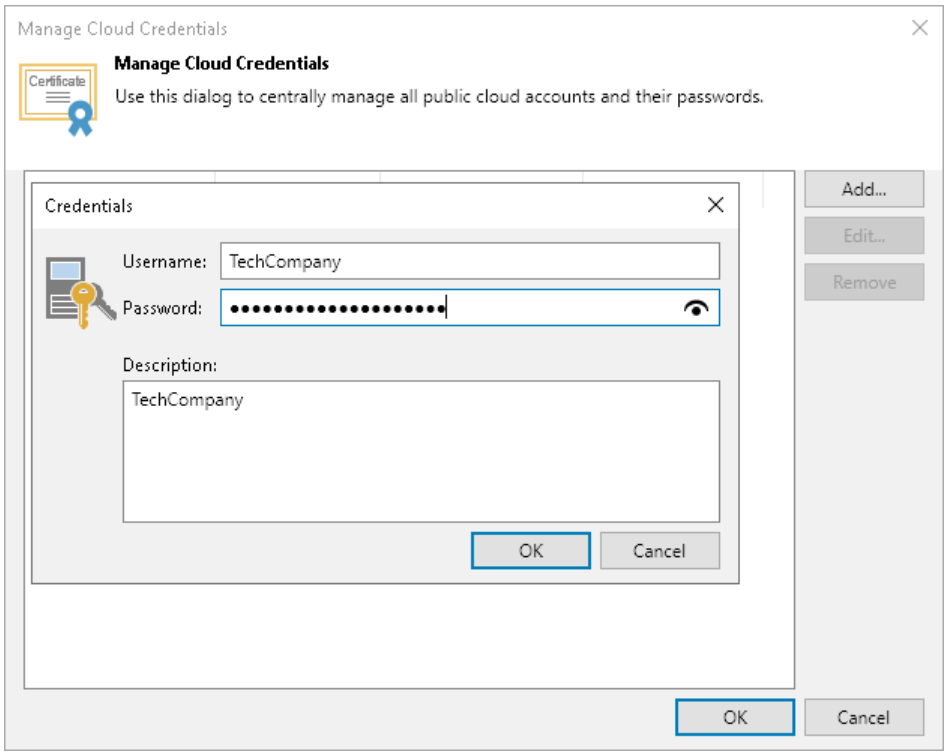


# Access Keys for AWS Users

You can create a record for credentials that you plan to use to connect to AWS.

To access AWS resources, you can use *Identity and Access Management (IAM) user* credentials or *AWS account root user* credentials. However, AWS recommends that you use the IAM user credentials. For details, see the AWS Account Root User Credentials vs. IAM User Credentials section in the AWS General Reference. The permissions that you must provide to the user are listed in section Permissions.

Since Veeam Backup & Replication uses AWS CLI commands to perform operations in AWS, instead of a user name and password you must specify an AWS access key. AWS access keys are long-term user credentials that consists of two parts: an access key ID and a secret access key. For details, see the Managing Access Keys for IAM Users section in the AWS IAM User Guide.

To create a credentials record:

1. From the main menu, select **Credentials and Passwords** > **Cloud Credentials**.
2. Click **Add** > **AWS access key**.
3. In the **Access key** field, enter an access key ID.
4. In the **Secret key** field, enter a secret access key. To view the entered secret key, click and hold the eye icon on the right of the field.
5. In the **Description** field, enter a description for the created credentials record.

> **IMPORTANT**
>
> It is recommended that the user whose credentials you plan to use to connect to AWS S3 has administrative permissions — access to all AWS S3 actions and resources.

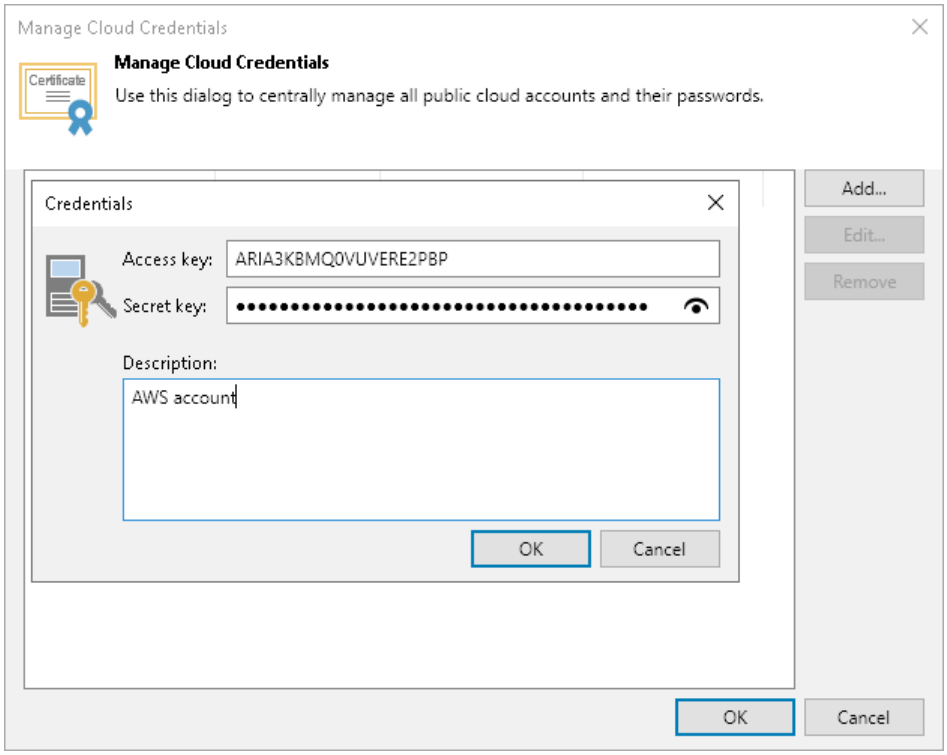


# Microsoft Azure Storage Accounts (Shared Key)

You can create a credentials record for a Microsoft Azure storage account to connect to the following types of accounts:

- Azure Blob storage added as an object storage repository, a performance extent or capacity extent of a scale-out backup repository. Use this option to store data on Azure Blob storage.
- Azure Archive storage added as an archive extent of a scale-out backup repository. Use this option to store data on Azure Blob storage.
- Azure Blob storage added as an external repository. Use this option to copy, import and restore backups created by Veeam Backup for Microsoft Azure from external to on-premises repositories.
- Microsoft Azure Blob storage added as a source of unstructured data. Use this option to backups data located on Azure Blob storage repository and restore backed-up data.
- Veeam Data Cloud Vault added as an object storage repository, a performance extent or capacity extent of a scale-out backup repository. Use this option to store data on Veeam Data Cloud Vault.

# Storage Accounts Supported Types

The following types of storage accounts are supported.

| Storage account type | Supported services | Supported performance tiers | Supported access tiers |
|---|---|---|---|
| General-purpose V2 | Blob (block blobs only) | Standard | |

| Storage account type | Supported services | Supported performance tiers | Supported access tiers |
|---|---|---|---|
| BlobStorage | | | The following Azure Storage access tiers are supported:<br>• **Hot**: to store data that you access frequently. Supported in Azure Blob Storage added as an object storage repository.<br>• **Cool**: to store data that you access infrequently. Supported in Azure Blob Storage added as an object storage repository.<br>• **Cold**: to store data that you access infrequently. Supported in Azure Archive Storage added as an object storage repository.<br>• **Archive**: to store data that you access rarely. Supported in Azure Archive Storage added as an object storage repository.<br>Consider the following:<br>• For Azure Blob storage<br>Veeam Backup & Replication will use the access tier that you specified at the Container step of the New Object Storage wizard.<br>• For Azure Archive storage<br>Veeam Backup & Replication will use the access tier that you specified at the Access Tier step of the New Object Storage wizard. |
| BlockBlobStorage | Blob (block blobs only) | Premium | N/A |

For more information about the types of storage accounts in Azure, see Microsoft Docs.

> **IMPORTANT**
>
> Microsoft Azure Blob storage accounts with the hierarchical namespace are not supported.

# Adding Microsoft Azure Storage Account

To create a record for a Microsoft Azure storage account:

1. From the main menu, select **Credentials and Passwords** > **Cloud Credentials**.
2. Click **Add** > **Microsoft Azure storage account**.
3. In the **Account** field, enter the storage account name.

4. In the **Shared key** field, enter the storage account shared key. To view the entered key, click and hold the eye icon on the right of the field.

> **NOTE**
>
> The **Allow storage account key access** option for Shared Key authorization must be enabled in the storage account. For more information on how to find this option, see Microsoft Docs.

5. In the **Description** field, enter an optional description for the credentials record.

> **TIP**
>
> If you do not have a Microsoft Azure storage account, you can create it in the Azure portal, as described in the Azure Storage Documentation.

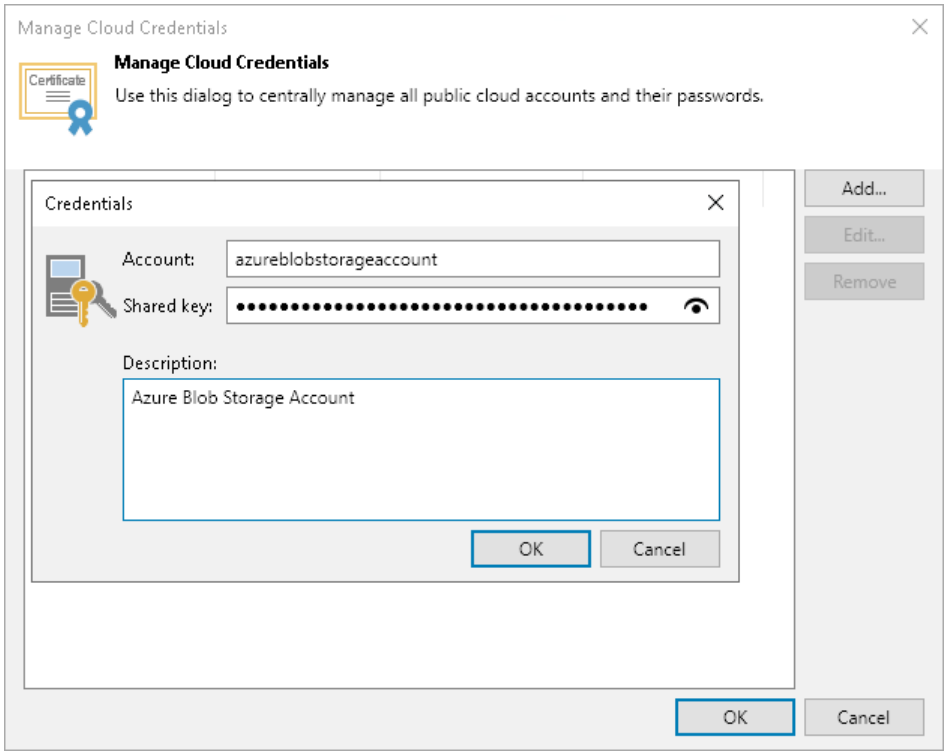


# Microsoft Azure Storage Accounts (Entra ID)

You can create a credentials record for a Microsoft Azure Storage account with Microsoft Entra authorization to connect to the following types of accounts:

- Azure Blob Storage added as an object storage repository, a performance extent or capacity extent of a scale-out backup repository. Use this option to store data on Azure Blob Storage.
- Azure Archive Storage added as an archive extent of a scale-out backup repository. Use this option to store data on Azure Blob Storage.
- Microsoft Azure Blob Storage added as a source of unstructured data. Use this option to backups data located on Azure Blob Storage repository and restore backed-up data.
- Veeam Data Cloud Vault added as an object storage repository, a performance extent or capacity extent of a scale-out backup repository. Use this option to store data on Veeam Data Cloud Vault.

To add a Microsoft Azure storage account with Microsoft Entra authorization, use the **Microsoft Azure Storage Account (Entra ID)** wizard:

## Before you Begin

Before you add a Microsoft Azure storage account with Microsoft Entra authorization to Veeam Backup & Replication, check the following prerequisites:

- If you plan to use a new Microsoft Entra application to access storage account (at the Account Type step), consider the following:
  - Make sure that you already have a user account in Microsoft Entra ID. This account must have privileges listed in the Permissions section.
  - If you have multiple tenants associated with the Microsoft Entra user account that you plan to use to create a new application, Veeam Backup & Replication will create the application in the home tenant of the account. As a result, the application will have access only to the subscriptions of the home tenant, as well as the Azure Compute account will. If you want to use another tenant and its subscriptions, follow the instructions in this Veeam KB article.
  - The created Microsoft Entra application is assigned with the following roles for a storage account which name was entered at the Name step of the wizard:
    - Storage Account Contributor
    - Storage Blob Data Contributor
    - Storage Blob Data Owner

    For more information on roles, see Microsoft Azure Docs.

    You can limit the subscriptions to which Veeam Backup & Replication assigns the privileges as described in this Veeam KB article.
- If you plan to use an existing Microsoft Entra application to access storage account (at the Account Type step), consider the following:
  - Create a Microsoft Entra application beforehand. For more information on how to create it, see Microsoft Docs. Note that you do not need to configure redirect URI.
  - The Microsoft Entra application must have several role privileges assigned on a storage account. For more information, see Permissions.
  - Only storage accounts from subscriptions that belong to the applications tenant can be used.
- On the backup server, you must set the correct time according to the timezone where the backup server is located. Otherwise, you may not be able to add a Microsoft Entra user account to Veeam Backup & Replication.
- When the internet access is possible only through HTTP/HTTPS proxy, you must configure the proxy settings for the Local System account or account under which the Veeam Backup Service is running. For more information, see this Microsoft article.

# Step 1. Launch Microsoft Azure Storage Account (Entra ID) Wizard

To launch the **Microsoft Azure Storage Account (Entra ID)** wizard, do the following:

1. In the main menu, click **Credentials and Passwords** > **Cloud Credentials**.
2. In the **Manage Cloud Credentials** window, click **Add** and select **Microsoft Azure storage account (Entra ID)**.

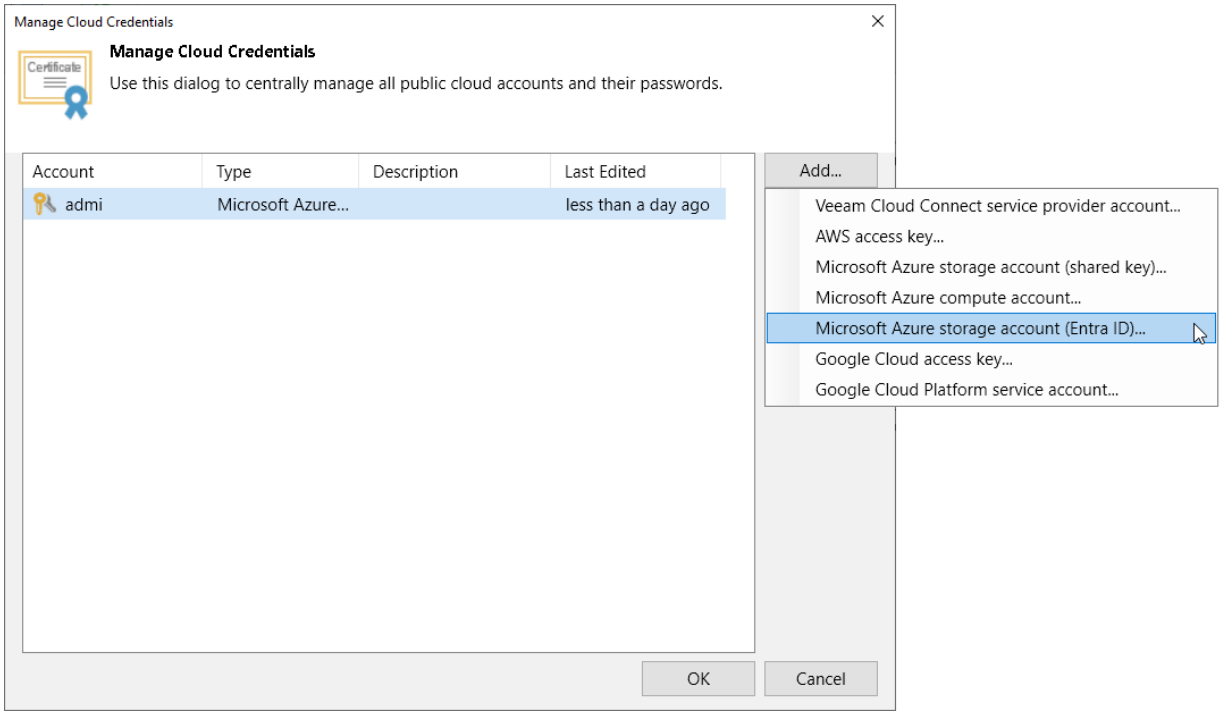

## Step 2. Specify Account Name

At the **Name** step of the wizard, specify an Azure storage account name. From the **Region** drop-down list, select a Microsoft Azure region where the storage account is located.

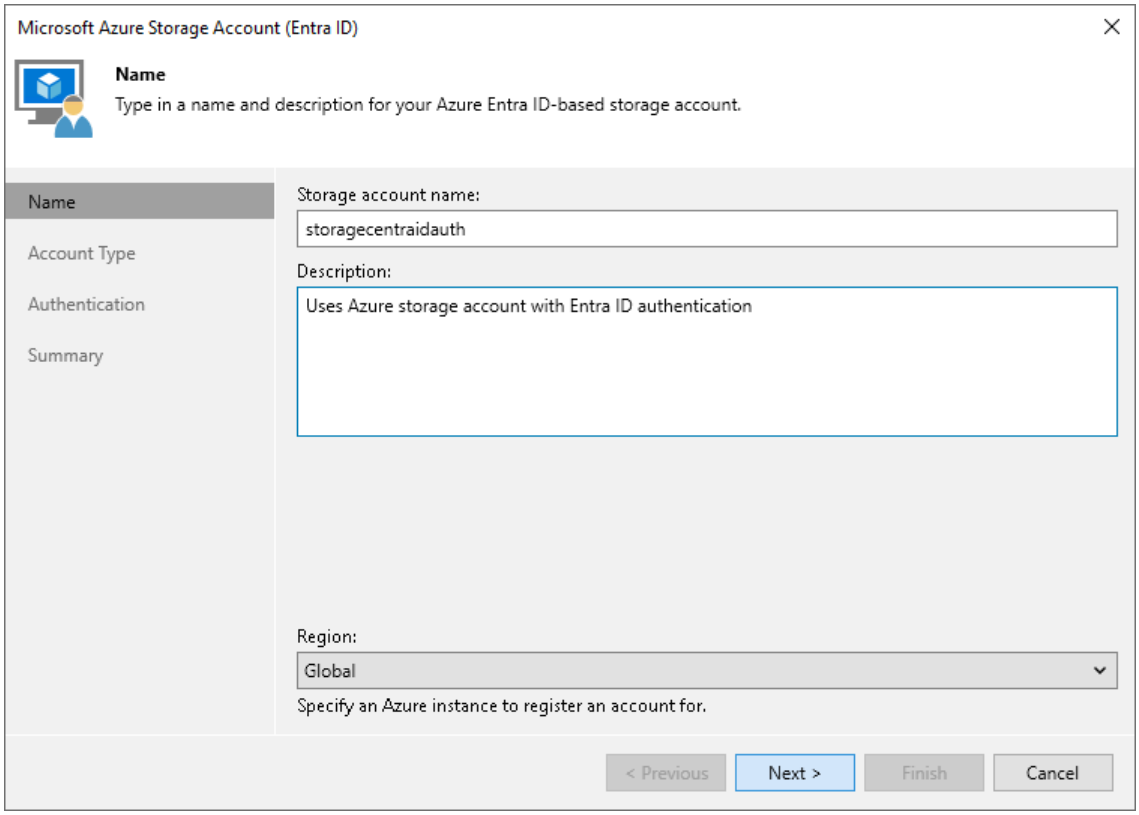

## Step 3. Select Account Type

At the **Account Type** step of the wizard, choose whether you want to connect to Microsoft Azure using an existing or a newly created Microsoft Entra application. In the latter case, Veeam Backup & Replication will create a new Microsoft Entra application automatically.

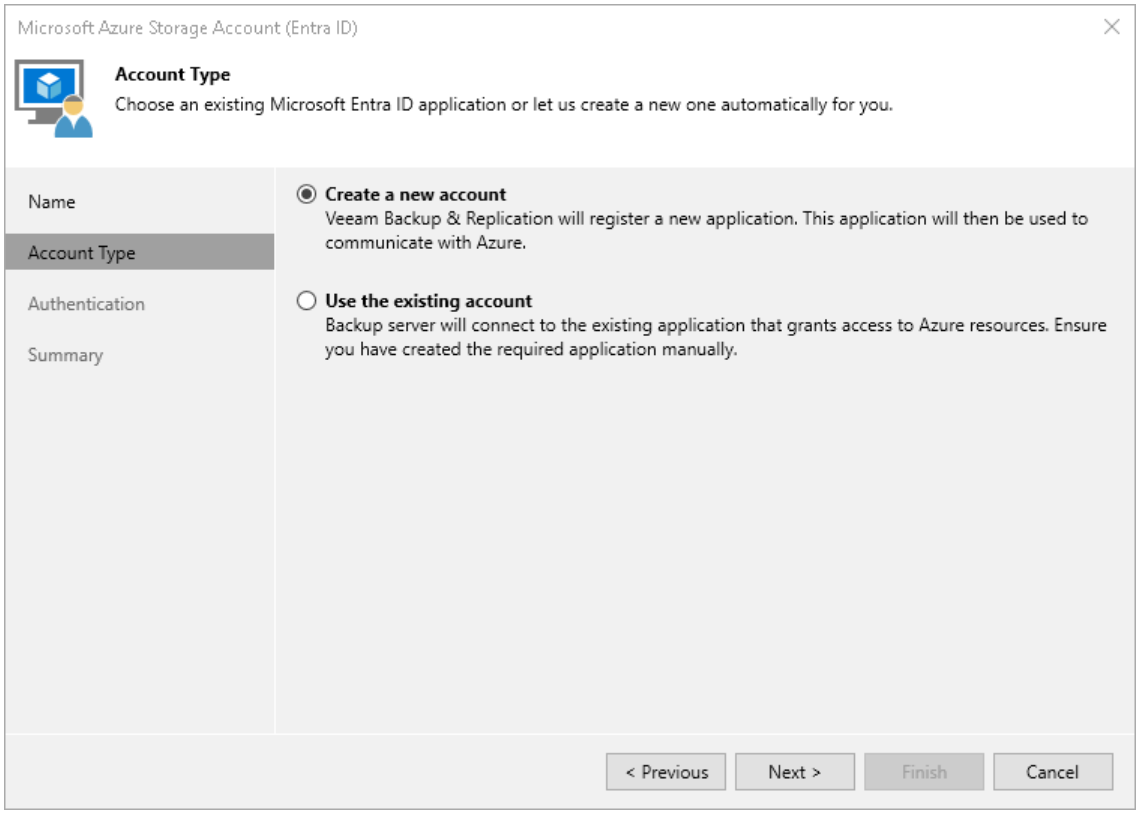

## Step 4. Specify Authentication Settings

At the authentication step of the wizard, create a new Microsoft Entra application or specify settings for an existing Microsoft Entra application.

### Creating New Microsoft Entra Application

This step applies only if you have selected the **Create a new account** option at the Account Type step of the wizard.

# Configuring Microsoft Entra Application

When you choose to create a new account, Veeam Backup & Replication registers a new Microsoft Entra application with a Microsoft Entra tenant. Veeam Backup & Replication will use this application to authenticate to Microsoft Azure. For more information on Microsoft Entra applications, see Microsoft Docs. To create the Microsoft Entra application, you must use a single-use verification code that Veeam Backup & Replication provides you.

At the **Authentication** step of the wizard, do the following:

1. Click **Copy to clipboard** to copy the verification code.
2. Click the https://microsoft.com/devicelogin link.
3. On the Microsoft Azure device authentication page, do the following:
   a. Paste the code that you have copied and click **Next**. Note that the code will expire in 15 minutes.
   b. Specify a Microsoft Entra account that will be used to create an application. Note that the user name must be specified in the user principal name format (username@domain). The account must have permissions described in section Permissions.

      Veeam Backup & Replication will create Microsoft Entra application in the tenant of the account.
4. Go back to the **Add Azure Account** wizard and check whether any errors occurred during the authentication process.

> **NOTE**
>
> Consider the following:
>
> - If you have multiple tenants associated with the Microsoft Entra user account that you plan to use to create a new application, Veeam Backup & Replication will create the single-tenant application in the home tenant of the account. As a result, the application can have access only to resources in the subscriptions of the home tenant. If you want to use storage account in another tenant subscriptions, follow the instructions this Veeam KB article.
> - The created Microsoft Entra application is assigned with the following roles for a storage account which name was entered at the Name step of the wizard:
>   - Storage Account Contributor
>   - Storage Blob Data Contributor
>   - Storage Blob Data Owner
>
>   For more information on roles, see Microsoft Azure Docs. You can limit the subscriptions to which Veeam Backup & Replication assigns the privileges as described in this Veeam KB article.

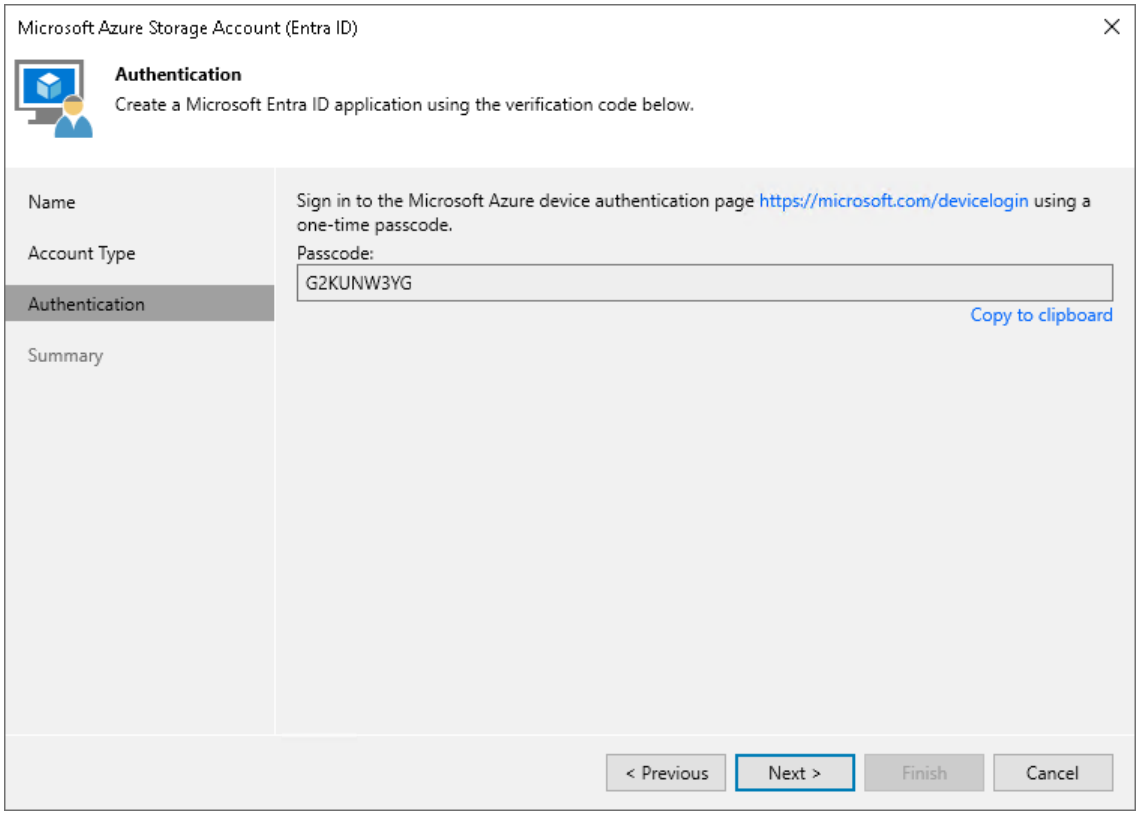


### Specifying Existing Microsoft Entra Application

This step applies only if you have selected the **Use the existing account** option at the Account Type step of the wizard.

# Configuring Microsoft Entra Application

To use an existing Microsoft Entra application:

1. In the **Tenant ID** specify an ID of a tenant (directory) where the Microsoft Entra application resides.
2. In the **Application ID** field, specify the ID of the necessary application. The Microsoft Entra application must have privileges listed in section Permissions.
3. In the **Select authentication type** area, choose whether you want to use password-based authentication (application secret) or certificate-based authentication. Then provide the necessary information.

For more information on how to get tenant and application IDs, secret and certificate, see Microsoft Docs.

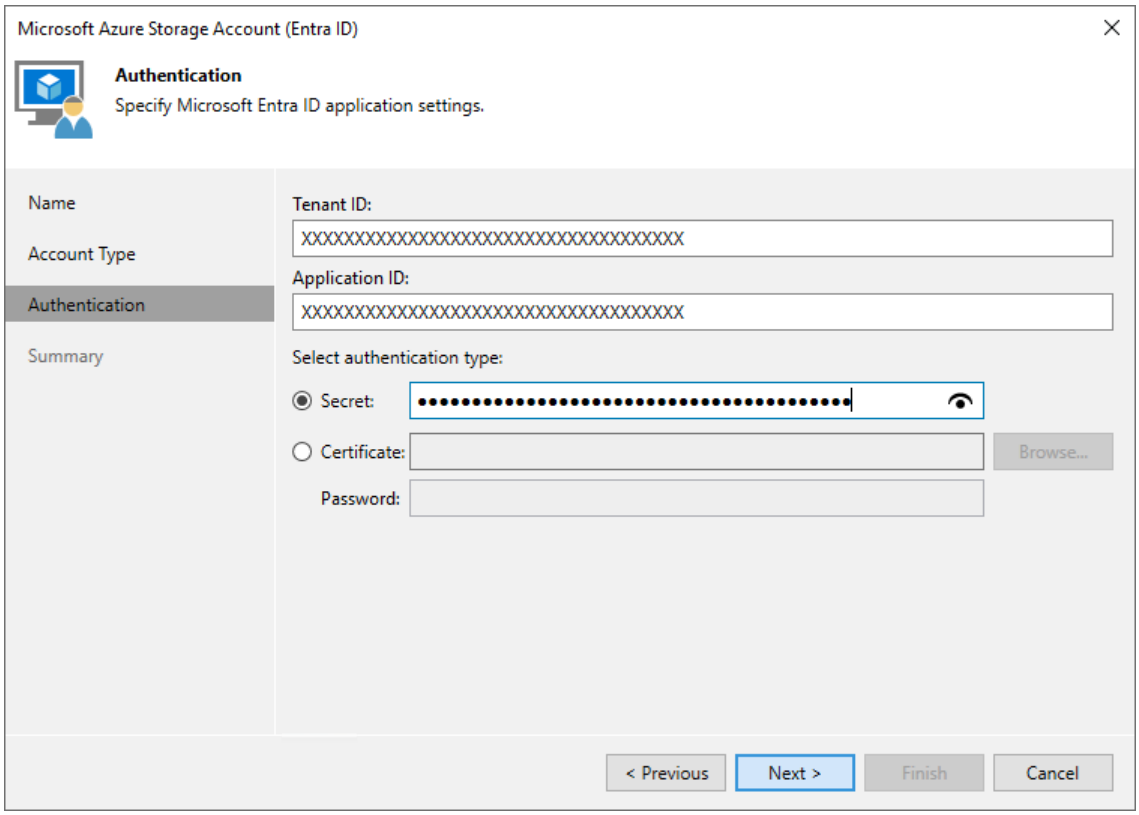

## Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of configured settings and click **Finish** to close the wizard.

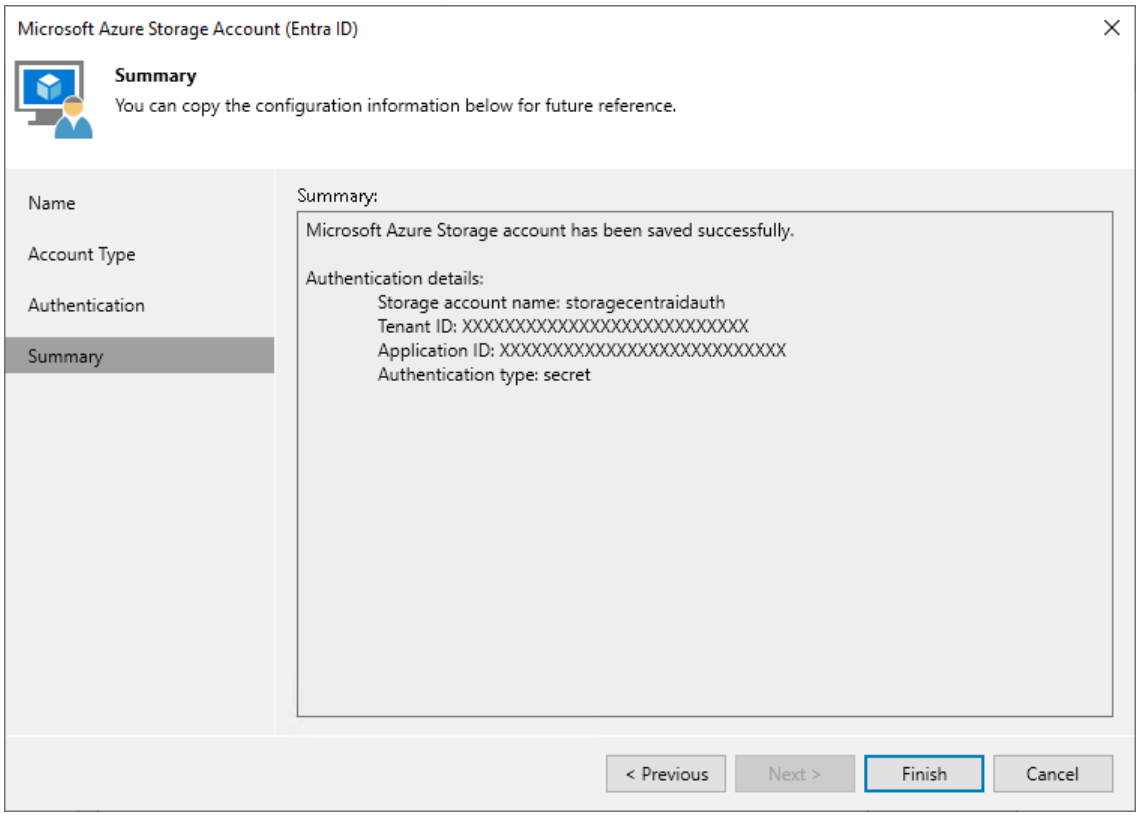


# Microsoft Azure Compute Accounts

You can use a Microsoft Azure Compute account for the following operations in Veeam Backup & Replication:

- Restore workloads to Microsoft Azure.
- Add Azure archive storage.
- Add a Microsoft Azure Plug-in for Veeam Backup & Replication appliance.

When you add a Microsoft Azure Compute account, Veeam Backup & Replication imports information about subscriptions and resources associated with this account. During the restore process, Veeam Backup & Replication accesses these resources and uses them to register new VMs in Microsoft Azure.

To add a Microsoft Azure Compute account, use the **Microsoft Azure Compute Account** wizard.

## Before You Begin

Before you add a Microsoft Azure Compute account to Veeam Backup & Replication, check the following prerequisites:

- If you plan to use a new Microsoft Entra ID (formerly Azure Active Directory) application to access Microsoft Azure (at the Account Type step), consider the following:
  - Make sure that you already have a user account in Microsoft Entra ID. This account must have privileges listed in section Permissions.

- If you have multiple tenants associated with the Microsoft Entra user account that you plan to use to create a new application, Veeam Backup & Replication will create the application in the home tenant of the account. As a result, the application will have access only to the subscriptions of the home tenant, as well as the Azure Compute account will. If you want to use another tenant and its subscriptions, follow the instructions in this Veeam KB article.
- The created Microsoft Entra application is assigned the *Contributor, Key Vault Crypto User* and *Storage Queue Data Contributor* role privileges for the subscriptions for which the following conditions are met:
  - The subscriptions are linked to the home tenant of the Microsoft Entra user.
  - The Microsoft Entra user has access to these subscriptions and can assign roles on the subscription level for the registered application.

  You can limit the subscriptions to which Veeam Backup & Replication assigns the privileges as described in this Veeam KB article. For more information on roles, see Microsoft Azure Docs.

- If you plan to use an existing Microsoft Entra ID (formerly Azure Active Directory) application to access Microsoft Azure (at the Account Type step), consider the following:
  - Make sure that you already have a Microsoft Entra application. For more information on how to create it, see Microsoft Docs. Note that you do not need to configure redirect URI.
  - The Microsoft Entra application must have privileges listed in section Permissions.
  - Only subscriptions that belong to the selected tenant will be added.
- [If you plan to restore Linux workloads using helper appliances] Veeam Backup & Replication uses its built-in credentials record to work with all helper appliances. For security reasons, we recommended that you change a password for this account before you set up the helper appliances. Changing credentials is required only once. For more information, see Changing Credentials for Helper Appliances.
- On the backup server, you must set the correct time according to the timezone where the backup server is located. Otherwise, you may not be able to add a Microsoft Azure user account to Veeam Backup & Replication.
- When the internet access is possible only through HTTP/HTTPS proxy, you must configure the proxy settings for the Local System account or account under which the Veeam Backup Service is running. For more information, see this Microsoft article.

# Step 1. Launch Microsoft Azure Compute Account Wizard

To launch the **Microsoft Azure Compute Account** wizard, do the following:

1. In the main menu, click **Credentials and Passwords** > **Cloud Credentials**.
2. In the **Manage Cloud Credentials** window, click **Add** and select **Microsoft Azure compute account**.

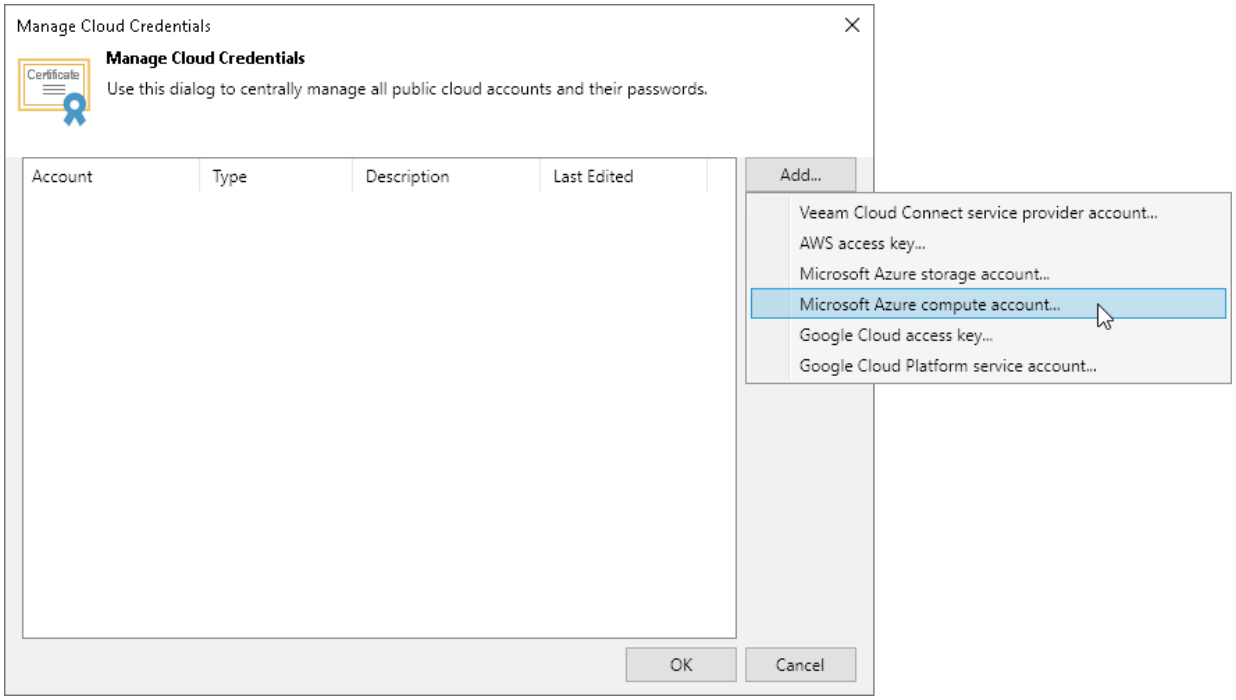

## Step 2. Specify Account Name

At the **Name** step of the wizard, specify a name under which this credentials record will be shown in the Cloud Credential Manager.

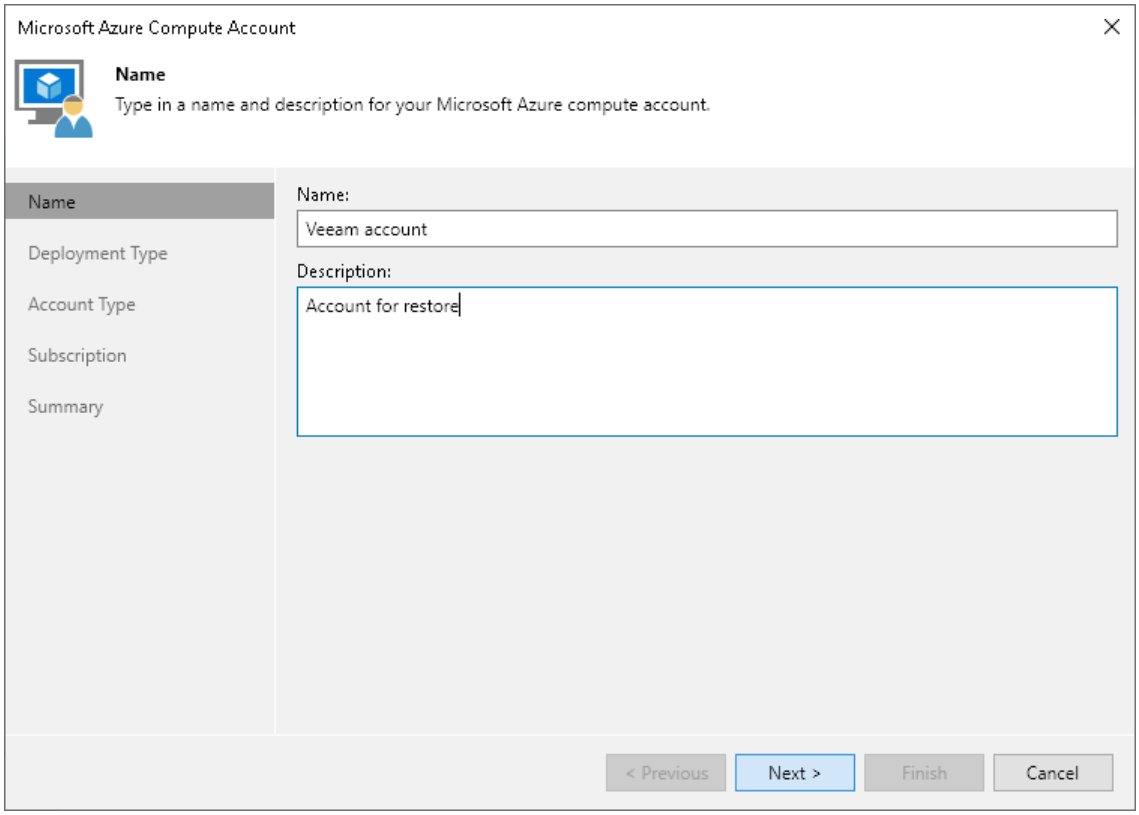

## Step 3. Select Deployment Type and Region

To select a deployment type and a region where your Microsoft Azure Compute account is located:

1. At the **Deployment Type** step of the wizard, select **Microsoft Azure**.
2. From the **Region** drop-down list, select a Microsoft Azure region.

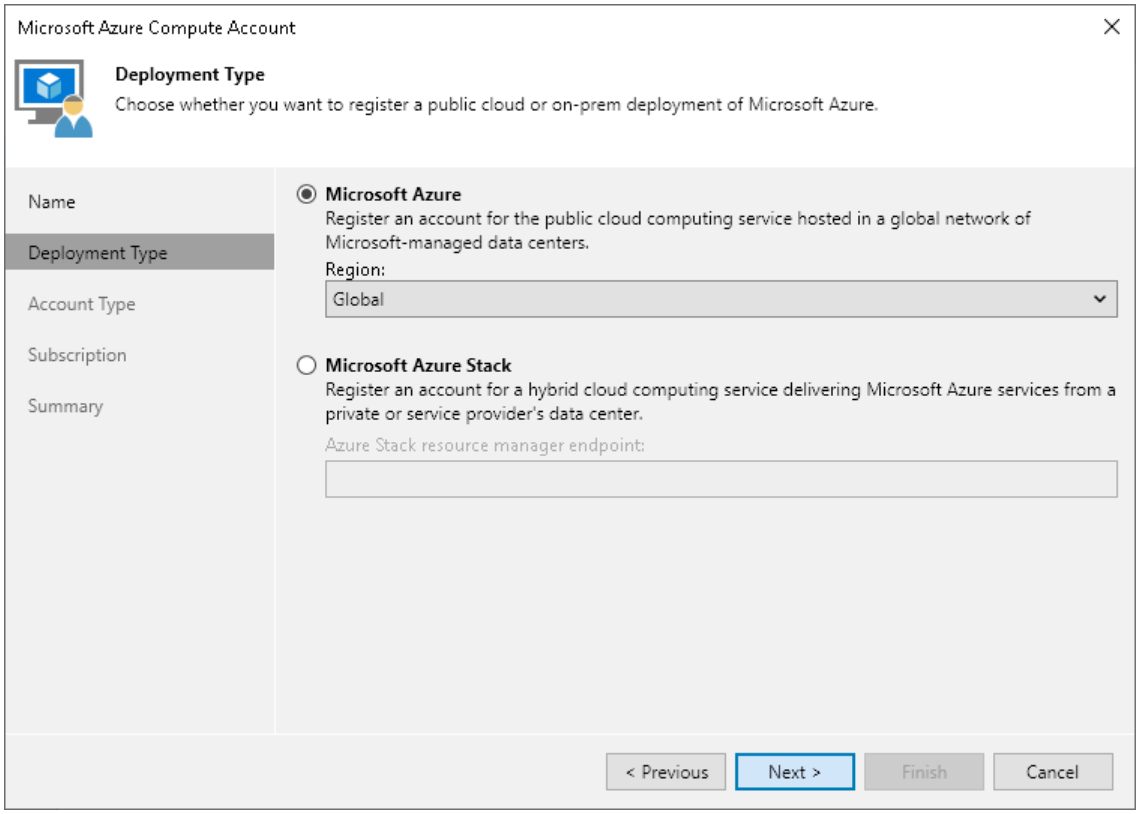

## Step 4. Select Access Type

At the **Account Type** step of the wizard, choose whether you want to connect to Microsoft Azure using an existing or newly created Microsoft Entra ID (formerly Azure Active Directory) application. In the latter case, Veeam Backup & Replication will create the new Microsoft Entra application automatically.

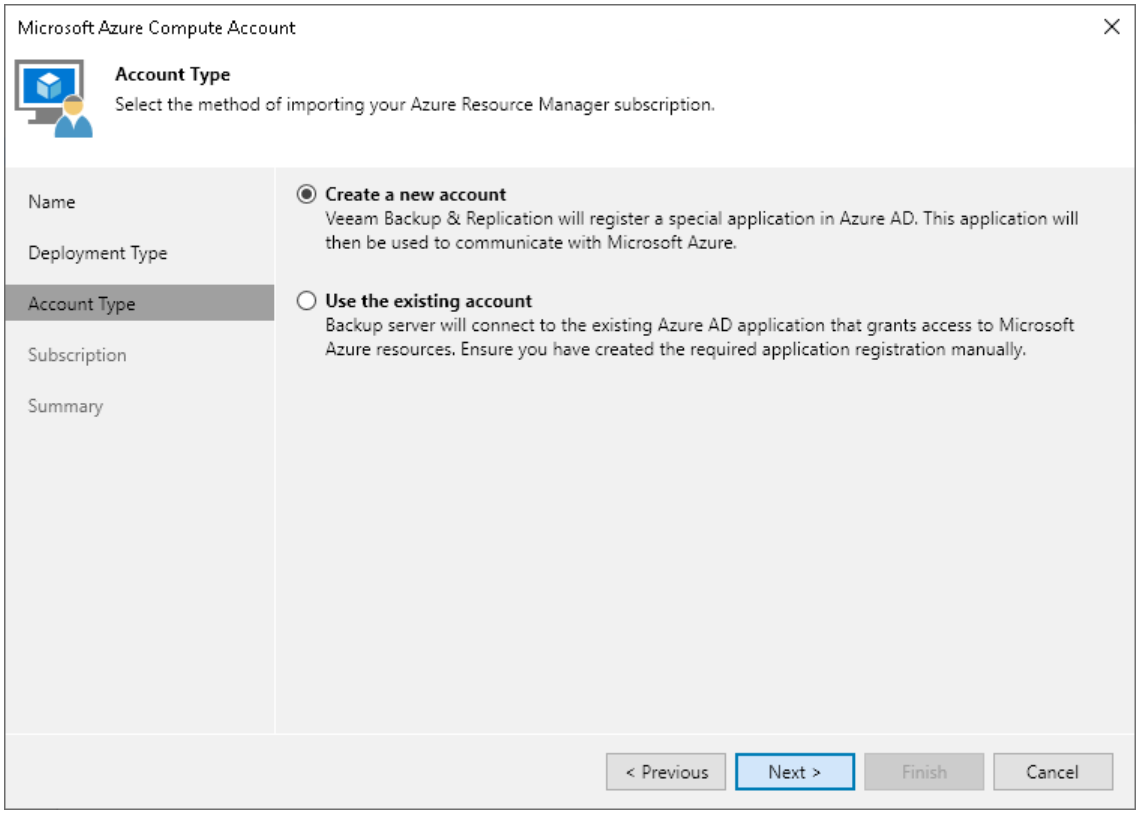


### Creating New Entra ID Application

This step applies only if you have selected the **Create a new account** option at the Account Type step of the wizard.

# Configuring Microsoft Entra Application

When you choose to create a new account, Veeam Backup & Replication registers a new Microsoft Entra ID (formerly Azure Active Directory) application in Microsoft Azure. Veeam Backup & Replication will use this application to authenticate to Azure. For more information on Microsoft Entra applications, see Microsoft Azure Docs. To create the Microsoft Entra application, you must use a single-use verification code that Veeam Backup & Replication provides you.

At the **Account Type** step of the wizard, do the following:

1. Click **Copy to clipboard** to copy the verification code.
2. Click the https://microsoft.com/devicelogin link.
3. On the Microsoft Azure device authentication page, do the following:
   a. Paste the code that you have copied and click **Next**. Note that the code will expire in 15 minutes.

b. Specify a Microsoft Entra user account that will be used to create an application. Note that the user name must be specified in the user principal name format (username@domain). The account must have permissions described in section Permissions.

   Veeam Backup & Replication will retrieve information about subscriptions to which the Microsoft Entra user has access and will create Microsoft Entra application in the tenant of the account.

4. Back to the **Add Azure Account** wizard, check whether any errors occurred during the authentication process.

> **NOTE**
>
> Consider the following:
>
> - If you have multiple tenants associated with the Microsoft Entra user account that you plan to use to create a new application, Veeam Backup & Replication will create the application in the home tenant of the account. As a result, the application will have access only to the subscriptions of the home tenant, as well as the Azure Compute account will. If you want to use another tenant and its subscriptions, follow the instructions in this Veeam KB article.
> - The created Microsoft Entra application is assigned the *Contributor, Key Vault Crypto User* and *Storage Queue Data Contributor* role privileges for the subscriptions for which the following conditions are met: the subscriptions are linked to the home tenant of the Microsoft Entra user; the Microsoft Entra user has access to these subscriptions and can assign roles on the subscription level for the registered application.
>
>   You can limit the subscriptions to which Veeam Backup & Replication assigns the privileges as described in this Veeam KB article. For more information on roles, see Microsoft Azure Docs.

# Enabling Direct Restore of Linux Workloads

If you plan to restore Linux-based workloads, select the **Enable direct restore of Linux-based computers**. When selected, this check box enables the **Helper Appliance** step where you will configure helper appliances. Note that the helper appliance is required if you restore Linux-based workloads that are not Azure-native.

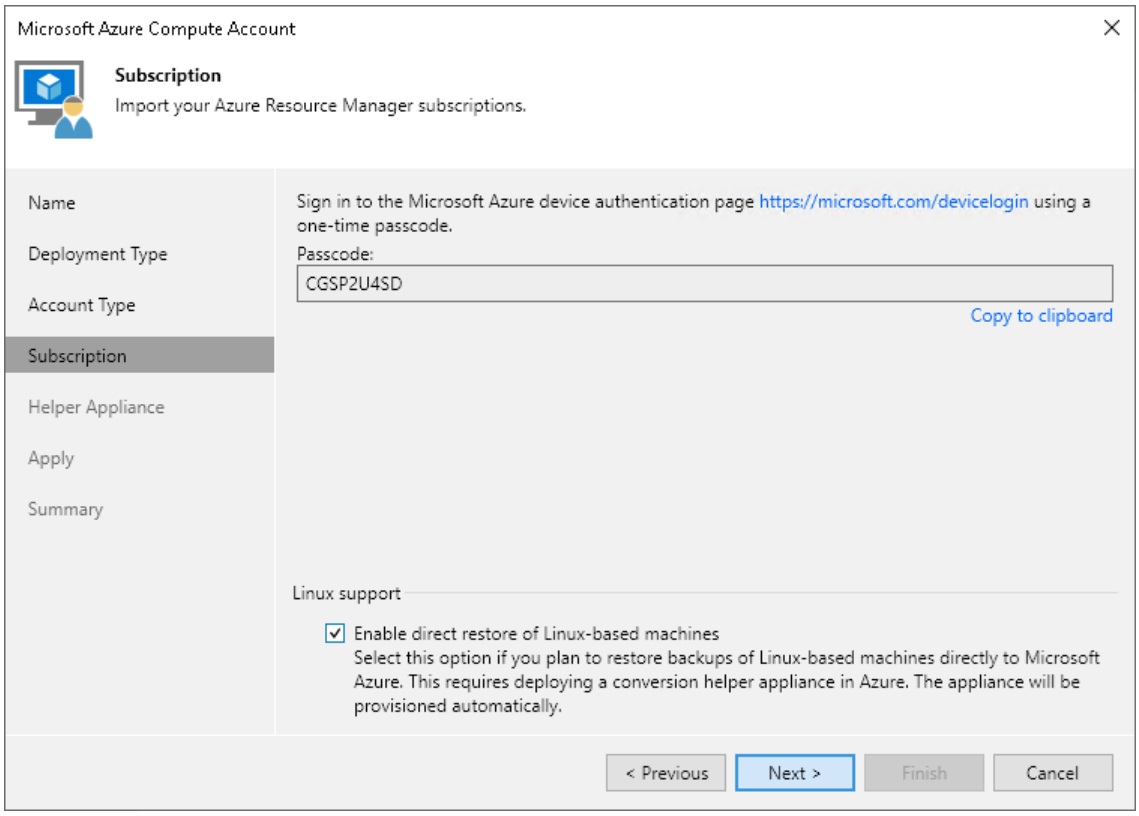


## Specifying Existing Entra ID Application

This step applies only if you have selected the **Use the existing account** option at the Account Type step of the wizard.

# Configuring Microsoft Entra Application

To use an existing Microsoft Entra ID (formerly Azure Active Directory) application:

1. In the **Tenant ID** specify an ID of a tenant (directory) where the Microsoft Entra application resides.
2. In the **Application ID** field, specify the ID of the necessary application. The Microsoft Entra application must have privileges listed in section Permissions.
3. In the **Select authentication type** area, choose whether you want to use password-based authentication (application secret) or certificate-based authentication. Then provide the necessary information.

For more information on how to get tenant and application IDs, secret and certificate, see Microsoft Docs.

# Enabling Direct Restore of Linux Workloads

If you plan to restore Linux-based workloads, select the **Enable direct restore of Linux-based computers**. When selected, this check box enables the **Helper Appliance** step where you will configure helper appliances. Note that the helper appliance is required if you restore Linux-based workloads that are not Azure-native.

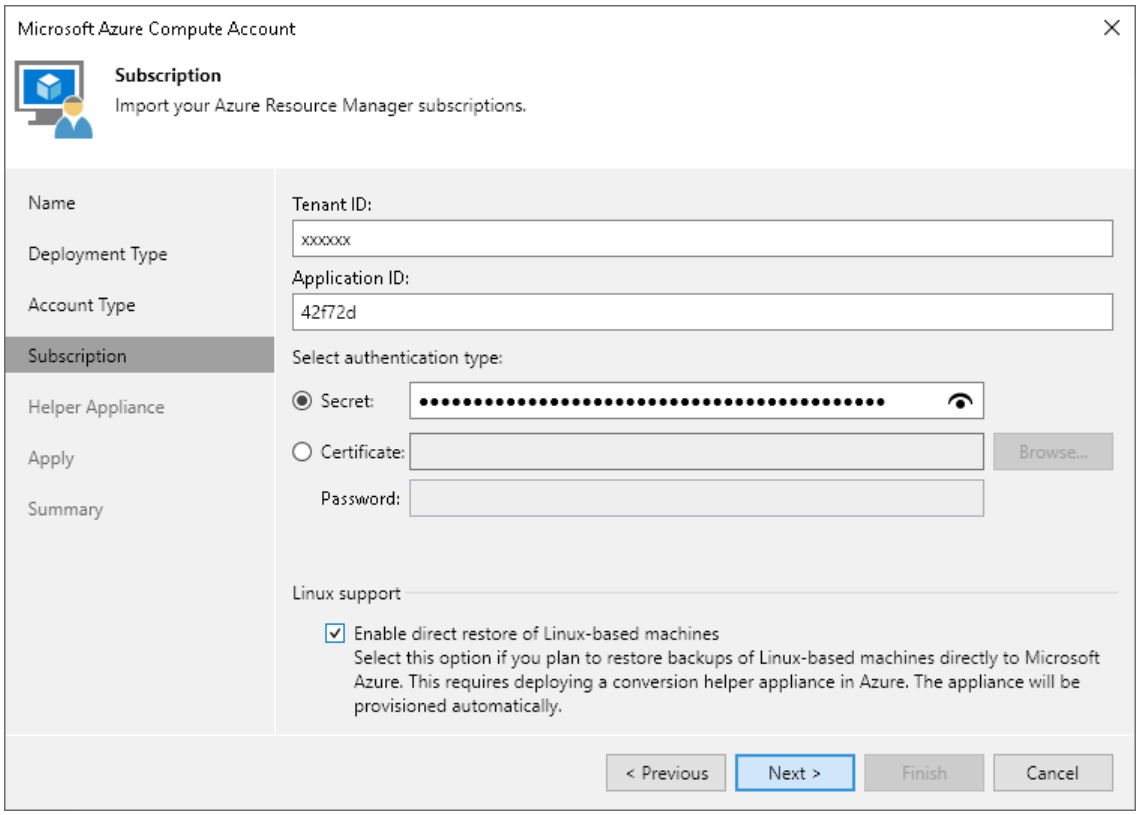

## Step 5. Configuring Helper Appliance

This step is available if you have selected **Enable direct restore of Linux-based computers** at the **Subscription** step of the wizard.

If you plan to restore Linux workloads to multiple locations, you must configure a helper appliance in each location.

To configure a helper appliance, do the following:

1. On the right of the **Helper appliances** list, click **Add**.
2. From the **Subscription** list, select a subscription whose resources you want to use to configure the helper appliance. The subscription list contains all subscriptions that are associated with the Azure Compute account.

   To be displayed in the **Subscription** list, a subscription must be created in advance and associated to the Azure account as described in Microsoft Docs.
3. From the **Location** list, select a location where you want to configure the helper appliance. Make sure that you select a geographic region with which at least one storage account of the subscription is associated.
4. From the **Storage account** list, select a storage account whose resources you want to use to store disks of the helper appliance.

   To be displayed in the **Storage account** drop-down list, a storage account must be created in advance as described in Microsoft Docs.

   > **NOTE**
   >
   > You cannot use a storage account with the ZRS or GZRS replication option. For details, see Microsoft Docs.
5. Click **Choose** if you do not want Veeam Backup & Replication to create a new resource group and select the required group.

   To be displayed in the **Resource group** drop-down list, a resource group must be created in advance as described in Microsoft Docs.
6. From the **Virtual network** list, select a network to which the helper appliance must be connected.

   To be displayed in the **Virtual network** drop-down list, a virtual network must be created in advance as described in Microsoft Docs.
7. From the **Subnet** list, select a subnet for the helper appliance.

   To be displayed in the **Subnet** drop-down list, a subnet must be created in advance as described in Microsoft Docs.
8. At the **SSH port** field, specify a port over which Veeam Backup & Replication will communicate with the helper appliance.
9. Click **OK**.

After you have configured all the helper appliances, click **Apply** and wait while Veeam Backup & Replication deploys the configured helper appliances.

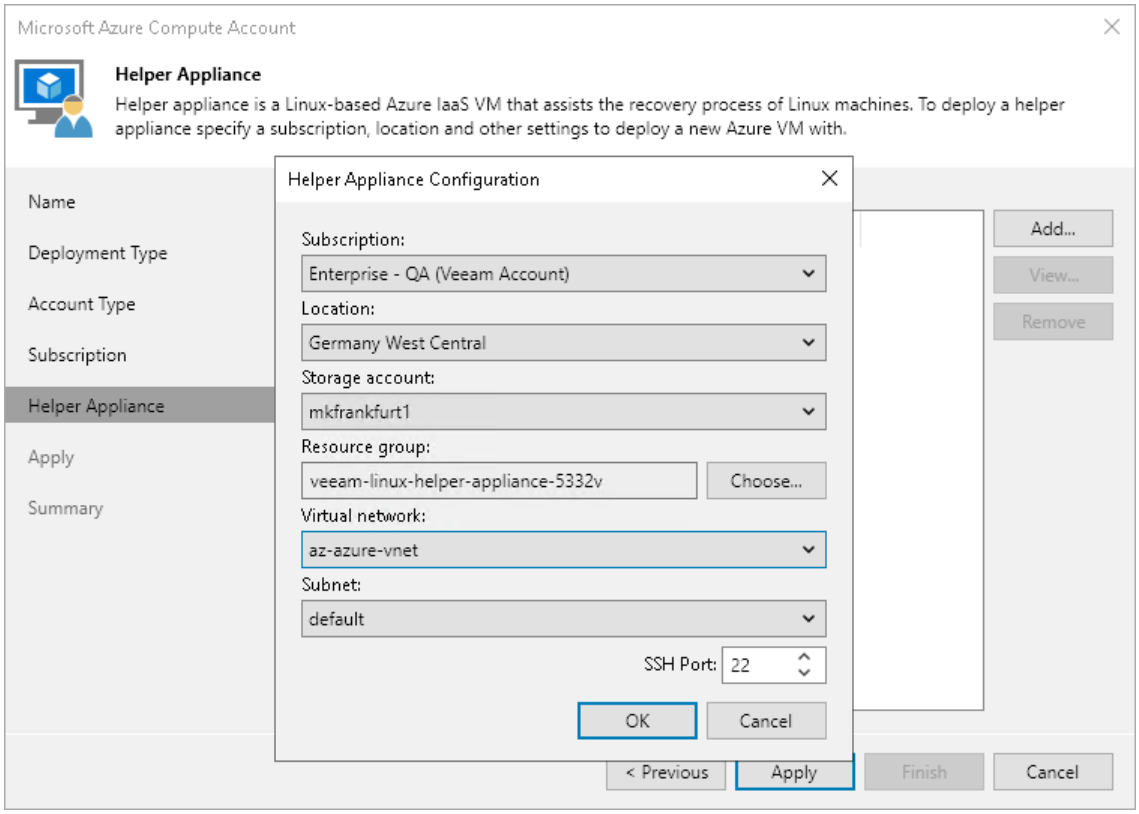

## Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of configured settings and click **Finish** to close the wizard.

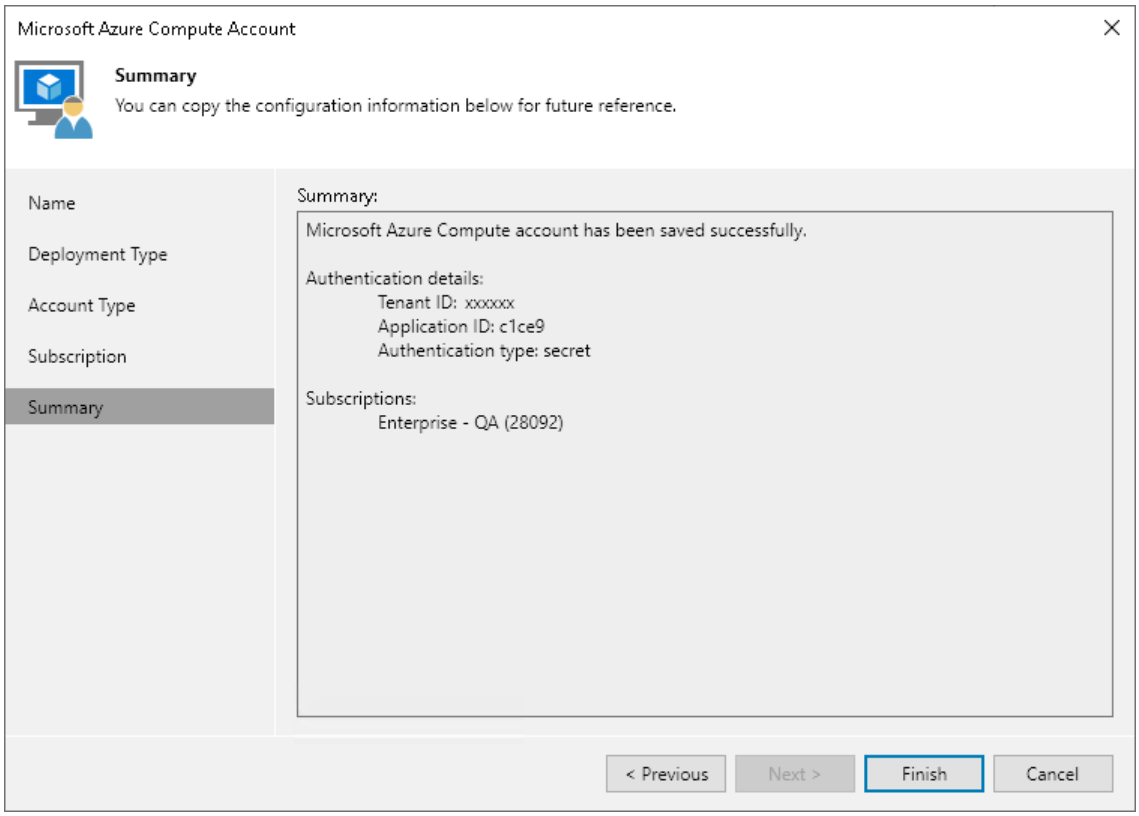


# Microsoft Azure Stack Hub Compute Accounts

To restore workloads to Microsoft Azure Stack Hub, you must add an Azure Stack Hub Compute account to Veeam Backup & Replication. When you add an Azure Stack Hub Compute account, Veeam Backup & Replication imports information about subscriptions and resources associated with this account. During the restore process, Veeam Backup & Replication accesses these resources and uses them to register new VMs in Azure Stack Hub.

If necessary, you can add different Azure Stack Hub Compute accounts to Veeam Backup & Replication. In this case, Veeam Backup & Replication will import information about all subscriptions and resources associated with provided accounts, and you will be able to use these resources for restore.

Information about subscriptions and resources is saved to the Veeam Backup & Replication configuration database. You can re-import this information at any time.

# Prerequisites

Before restoring workloads to Microsoft Azure Stack Hub, consider the following:

- [For restore to Microsoft Azure Stack Hub version 1808 and later] You must configure the backup server as described in this Veeam KB article.
- Check the Permissions section.
- [If you plan to restore Linux workloads using helper appliances] Veeam Backup & Replication uses its built-in credentials record to work with all helper appliances. For security reasons, we recommended that you change a password for this account before you set up the helper appliances. Changing credentials is required only once. For more information, see Changing Credentials for Helper Appliances.

- If you plan to use an existing Microsoft Entra ID (formerly Azure Active Directory) application to access Microsoft Azure (at the Account Type step), only subscriptions that belong to the selected tenant will be added.

# Adding Microsoft Azure Stack Hub Account

To add a Microsoft Azure Stack Hub Compute account, do the following.

1. From the main menu, select **Credentials and Passwords** > **Cloud Credentials**.
2. In the **Manage Cloud Credentials** window, click **Add** and select **Microsoft Azure compute account**.

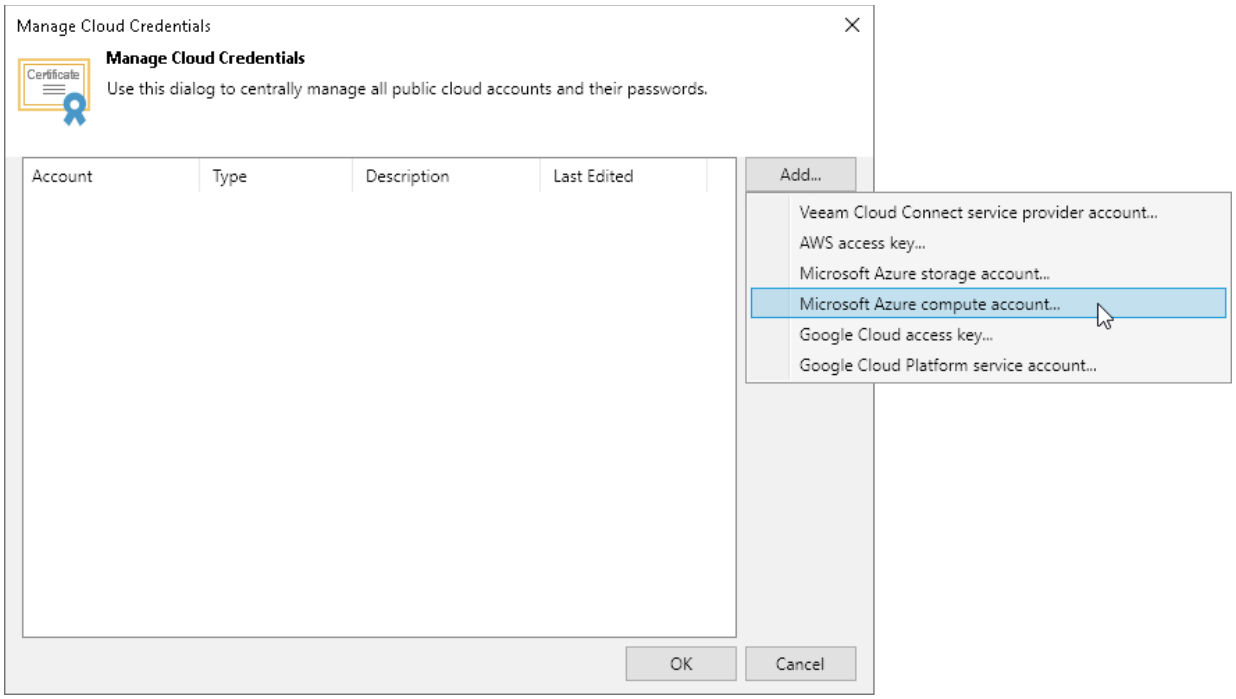


3. At the **Name** step of the wizard, specify a name under which this credentials record will be shown in the Cloud Credential Manager.

4. At the **Deployment Type** step of the wizard, select **Microsoft Azure Stack.** In the **Azure Stack resource manager endpoint** field, specify the virtual IPv4 address of Azure Resource Manager in the following format: `management.<region>.<FQDN>`. Note that IPv6 addresses are not supported.

   To learn about Azure Stack Hub virtual IP addresses, see Microsoft Docs.

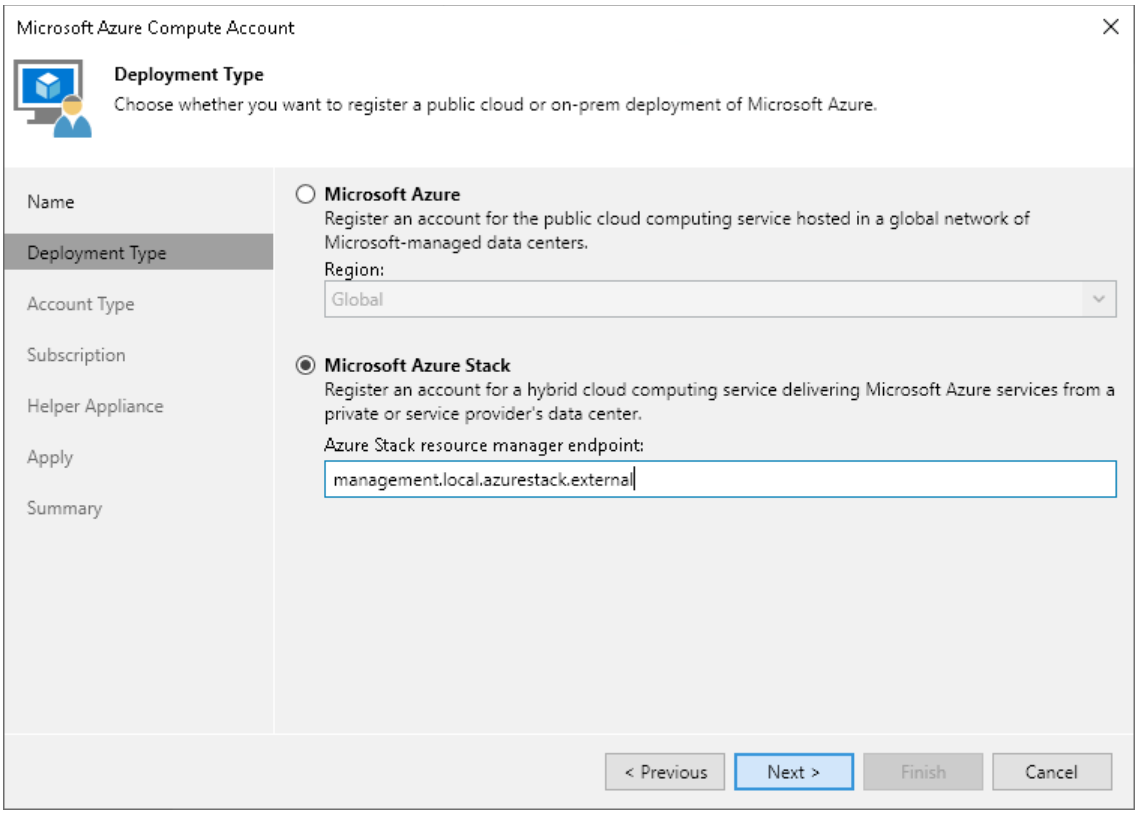


5. The rest of the procedure for adding the Azure Stack account does not differ from the procedure for adding the Microsoft Azure account. Follow the steps described in Microsoft Azure Compute Account.

# Google Cloud Accounts

You can add to the backup infrastructure a credentials record for the Google Cloud account that will be used to connect to Google Cloud Storage. Veeam Backup & Replication uses these credentials to add a Google Cloud storage repository to the backup infrastructure as an external repository.

To create a record for a Google Cloud account:

1. From the main menu, select **Credentials and Passwords** > **Cloud Credentials**.
2. Click **Add** > **Google Cloud access key**.
3. In the **Access Key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) key associated with the Google Cloud account. Veeam Backup & Replication will use the HMAC key to authenticate requests to the Google Cloud Storage. To view the entered key, click and hold the eye icon on the right of the field. For more information on Google Cloud accounts, see the Google Cloud documentation.

   If you have not created the HMAC key beforehand, you can do one of the following:

   - Create the HMAC key in the Google Cloud console, as described in the Google Cloud documentation.
   - Create the HMAC key in the Veeam Backup for Google Cloud Web UI, as described in the Veeam Backup for Google Cloud User Guide.

4. In the **Description** field, enter an optional description for the credentials record.

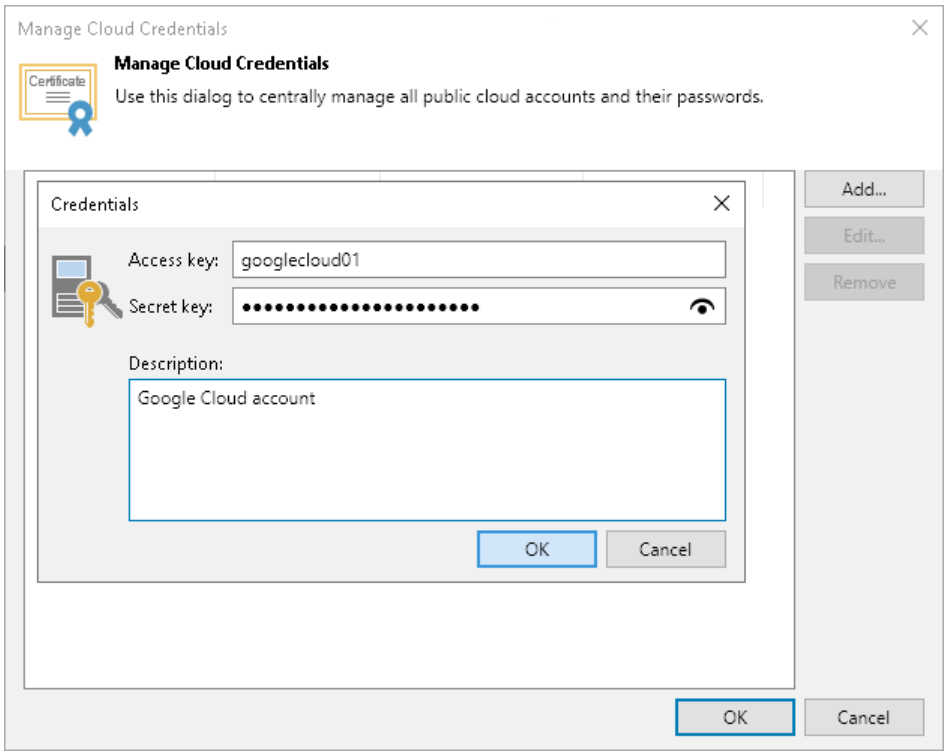


# Google Cloud Service Accounts

You can create a record for credentials that you plan to use to connect to Google Compute Engine within Google Cloud. This Google Cloud service account is used by Veeam Backup & Replication to perform direct restore to Google Compute Engine and backup and restore operations available with Google Cloud Plug-in for Veeam Backup & Replication. For more information on the latter, see the Veeam Backup for Google Cloud User Guide.

To create a credentials record for a Google Cloud service account:

1. From the main menu, select **Credentials and Passwords** > **Cloud Credentials**.
2. Click **Add** > **Google Cloud Platform service account**.
3. At the **Type** step of the wizard, select if you want to create a new service account automatically or use an existing service account.

> **NOTE**
>
> If you select **Create a new service account**, the created service account will be granted the Owner IAM role with a wide scope of permissions and capabilities. If you want to limit the list of permissions granted to the service account, create a user-managed service account, as described in the Google Cloud documentation, with the limited set of permissions:
>
> - For the information on permissions required to restore to Google Compute Engine, see Google Compute Engine IAM User Permissions.
> - For the information on permissions required to deploy Google Cloud Plug-in for Veeam Backup & Replication, see the Permissions section in the Veeam Backup for Google Cloud User Guide.

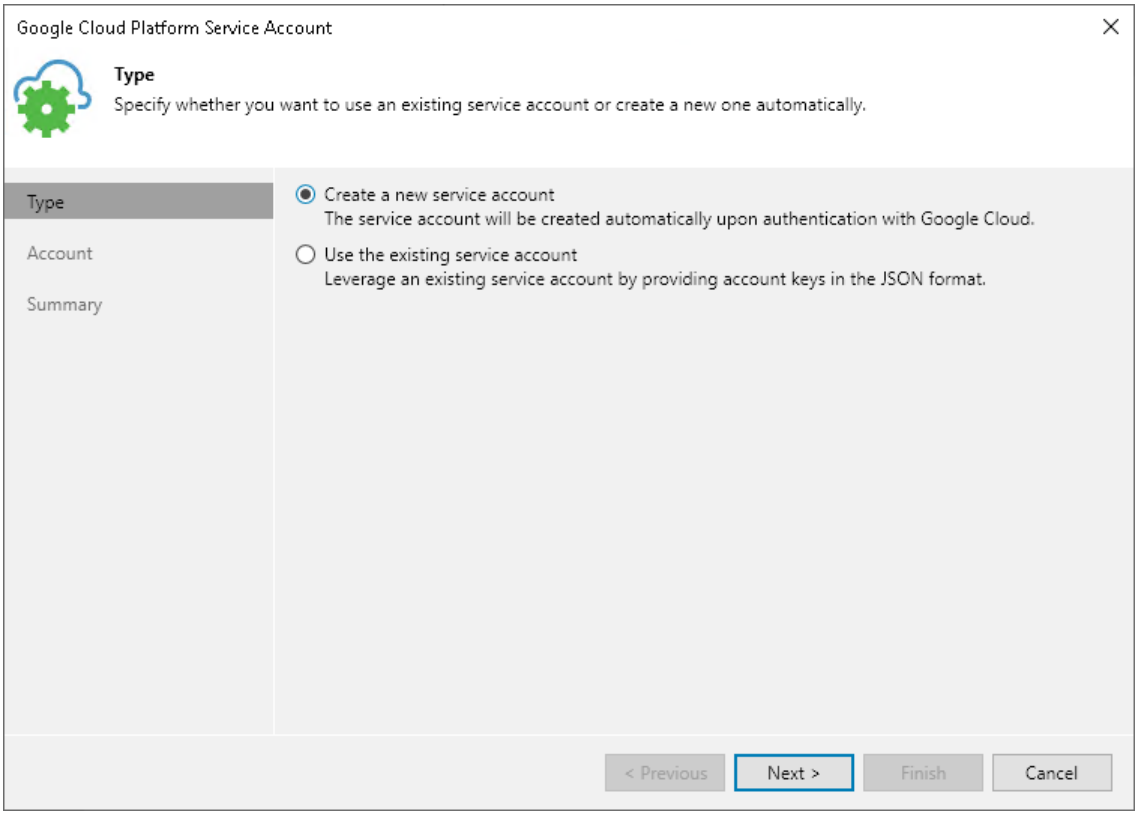


4. At the **Account** step of the wizard, specify credentials required for accessing the service account:
   - If you have selected **Create a new service account**, do the following:
     i. Log into your Google Cloud account. Read and accept the Google Terms of Service and the Google Privacy Policy.
     ii. Allow Veeam Backup & Replication to access your Google account. After that, Veeam Backup & Replication can manage your Identity and Access Management (IAM) policies, and see, edit, configure and delete your Google Cloud data.

iii. Return to the wizard and select the project with which you want the created service account to work.

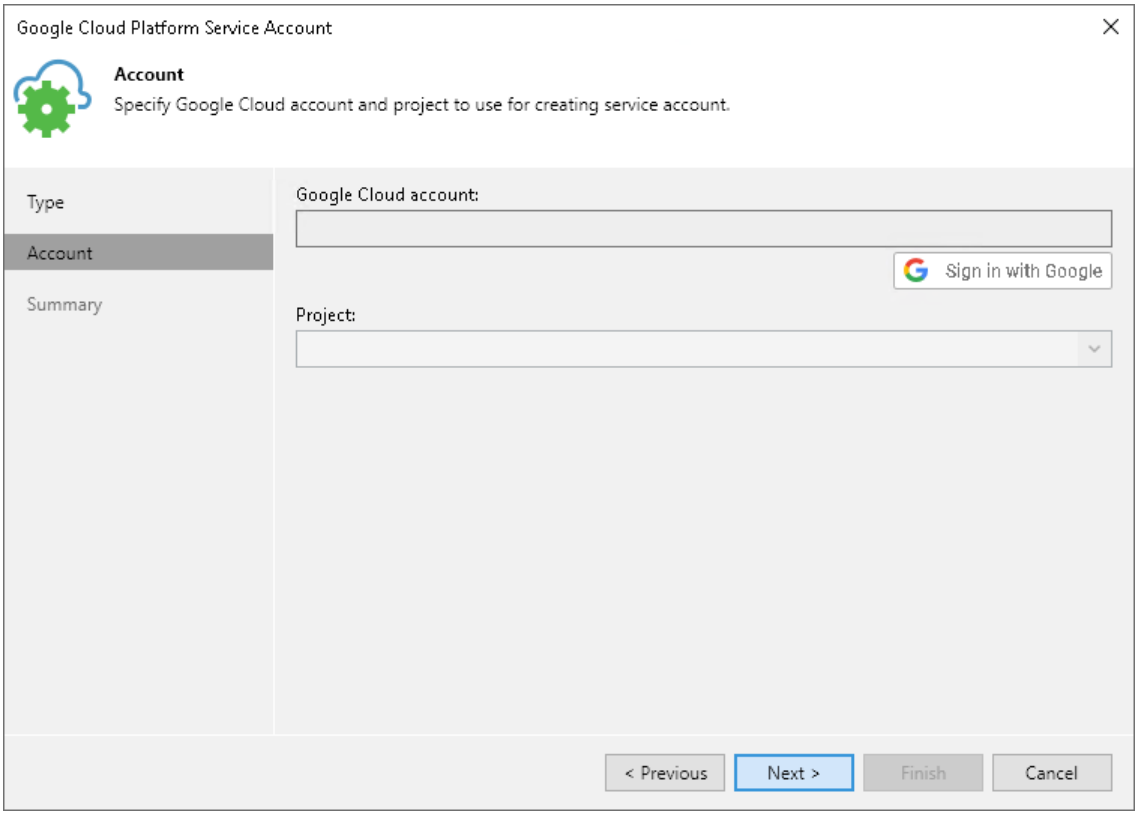


- If you have selected **Use the existing service account**, do the following:

  i. Download the service account key in the JSON format, created as described in Google Cloud documentation.

  > **IMPORTANT**
  >
  > Depending on the scenarios that the service account will be used for, make sure that the service account meets all requirements and limitations.
  >
  > For restoring virtual workloads from backups to Google Cloud, consider the requirements and limitations listed in section Restore to Google Compute Engine.

ii. At the **Account** step of the wizard, select the downloaded service account key.

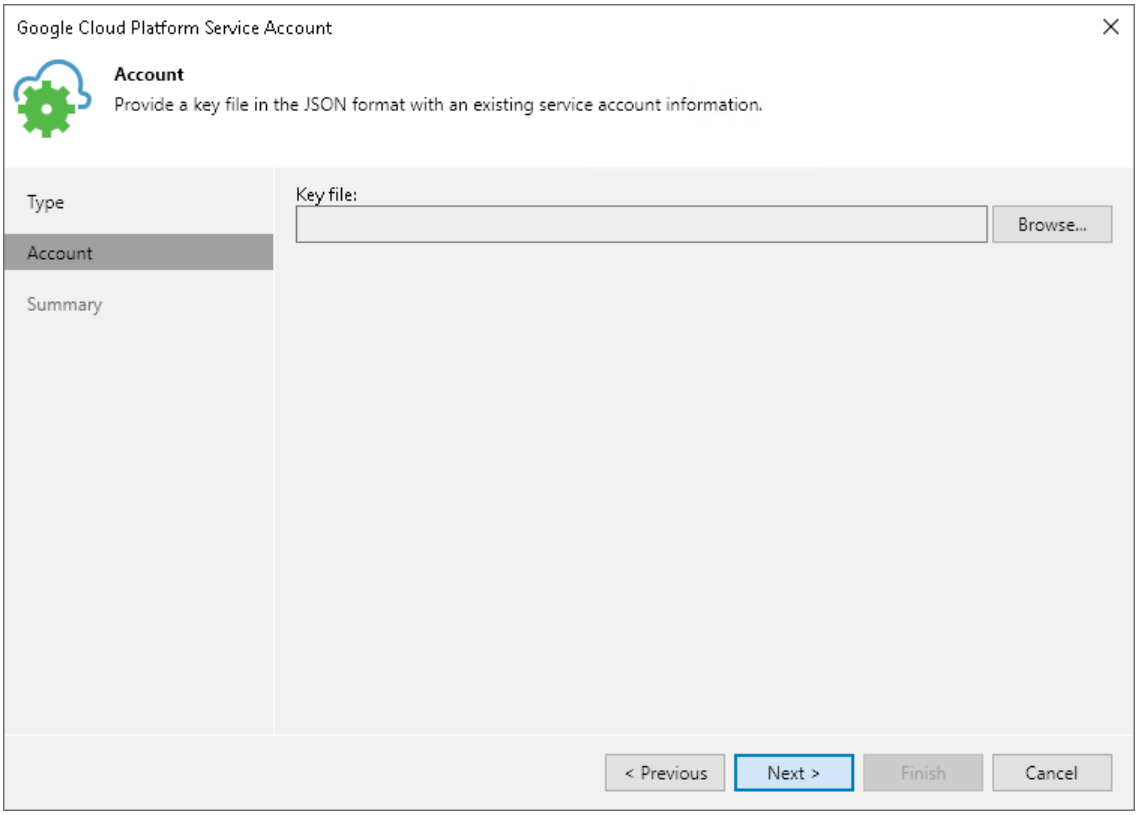


5. At the **Summary** step of the wizard, review details of the configured account and click Finish to close the wizard.

# Editing and Deleting Credentials Records

You can edit or delete existing cloud credentials records.

# Editing Credentials

To edit a credentials record:

1. From the main menu, select **Credentials and Passwords** > **Cloud Credentials**.
2. Select the credentials record in the list and click **Edit**.
3. Edit settings of the credentials record as required.

> **TIP**
>
> You can use the Cloud Credentials Manager to change the password for a tenant account provided by the SP. For more information, see the Changing Password for Tenant Account section in the Veeam Cloud Connect Guide.

# Creating New Google Service Account

If you created a Google Cloud service account using the Create a new service account option, you may need to create a new service account instead of the already created one. For example, the service account was deleted.

To cretae a new service account, do the following:

1. Edit the necessary credential as described in the Editing Credentials section.
2. At the **Account** step, click the **Log In** link.
3. In the opened window, log into your Google Cloud account.
4. Back to the **Google Cloud Platform Service Account** wizard, click **Finish**.

Veeam Backup & Replication will create a new service account and will use it to communicate with Google Cloud.

# Deleting Credentials

To delete a credentials record:

1. From the main menu, select **Credentials and Passwords > Cloud Credentials**.
2. Select the credentials record in the list and click **Remove**. You cannot delete a record that is already used for any component in the backup infrastructure.

# Password Manager

You can use the Password Manager to create and maintain a list of passwords that you plan to use for data encryption. Password management can be helpful in the following situations:

- You want to create new passwords. You can use one password per job or share the same password between several jobs on the backup server.
- You want to edit an existing password, for example, change its hint, or delete a password.

> **TIP**
>
> Periodical change of passwords is a security best practice. You can create new passwords as often as you need based on your company security needs and regulatory requirements.

Veeam Backup & Replication encrypts credentials stored in the configuration database using Data Protection API (DPAPI) mechanisms. For more information, see Microsoft Docs.

## Creating Passwords

You can use the Password Manager to create one or more passwords.

To create a new password:

1. From the main menu, select **Credentials and Passwords** > **Encryption Passwords**. Alternatively, you can use job properties to create a new password:
   a. Open the **Home** view.
   b. In the inventory pane, select **Jobs**.
   c. In the working area, right-click the backup or backup copy job and select **Edit**.
   d. At the **Storage** step of the wizard (for backup job) or **Target** step of the wizard (for backup copy job), click **Advanced**.
   e. Click the **Storage** tab.
   f. In the **Encryption** section of the **Advanced Settings** window, select the **Enable backup file encryption** check box and click the **Manage passwords** link.

   Veeam Backup & Replication will open the Password Manager.
2. In the Password Manager, click **Add**.
3. In the **Hint** field, specify a hint for the created password. It is recommended that you provide a meaningful hint that will help you recall the password. The password hint is displayed when you import an encrypted file on the backup server and access this file.
4. In the **Password** field, enter a password. The recommended minimum for the length of a password is 12 characters.

   To view the entered password, click and hold the eye icon on the right of the field.

> **IMPORTANT**
>
> Always save a copy of the password you create in a secure place. If you lose the password, you will not be able to restore it.

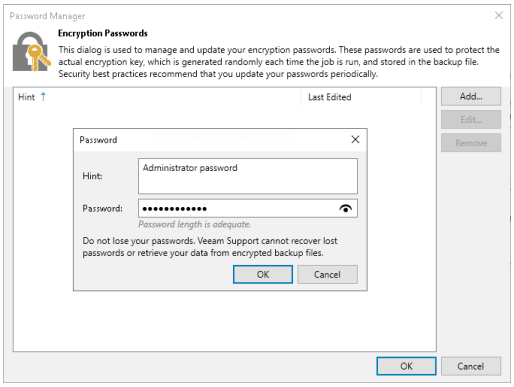


# Editing Passwords

You can edit passwords you have created using the Password Manager.

To edit a password:

1. From the main menu, select **Credentials and Passwords** > **Encryption Passwords**. Alternatively, you can use job properties to edit the password:
   a. Open the **Home** view.
   b. In the inventory pane, select **Jobs**.
   c. In the working area, right-click the backup or backup copy job and select **Edit**.
   d. At the **Storage** step of the wizard (for backup job) or **Target** step of the wizard (for backup copy job), click **Advanced**.
   e. Click the **Storage** tab.
   f. In the **Encryption** section of the **Advanced Settings** window, select the **Enable backup file encryption** check box and click the **Manage passwords** link.

   Veeam Backup & Replication will open the Password Manager.
2. In the Password Manager, select the password and click **Edit**.
3. Edit the password data: hint and password, as required. The recommended minimum for the length of a password is 12 characters.

After you edit the password, you do not need to perform any other actions. Veeam Backup & Replication will start using the changed password after a job runs for the next time. Before the job run, the old password is still used. For more information on which password to use when you restore data, see Restoring Data from Encrypted Backups.

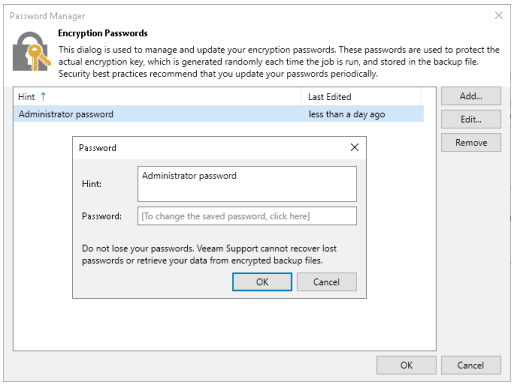

# Deleting Passwords

You can delete passwords using the Password Manager.

You cannot remove a password that is currently used by any job on the backup server. To remove such password, you first need to delete a reference to this password in the job settings.

To delete a password:

1. From the main menu, select **Credentials and Passwords** > **Encryption Passwords**. Alternatively, you can use job properties to delete passwords:
   a. Open the **Home** view.
   b. In the inventory pane, select **Jobs**.
   c. In the working area, right-click the backup or backup copy job and select **Edit**.
   d. At the **Storage** step of the wizard (for backup job) or **Target** step of the wizard (for backup copy job), click **Advanced**.
   e. Click the **Storage** tab.
   f. In the **Encryption** section of the **Advanced Settings** window, select the **Enable backup file encryption** check box and click the **Manage passwords** link.

   Veeam Backup & Replication will open the Password Manager.
2. In the Password Manager, select the password and click **Remove**.

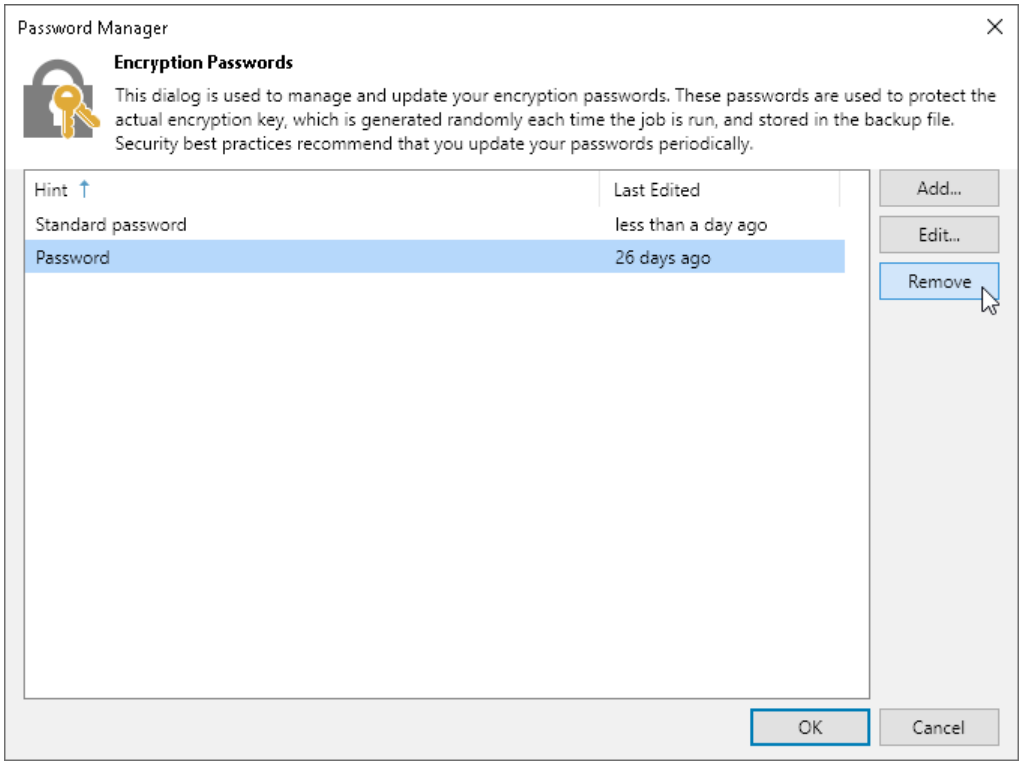

# Key Management System Keys

You can use Key Management System (KMS) keys for data encryption instead of secret keys based on a password. KMS keys are based on an asymmetric key encryption algorithm. They are managed and rotated by an external KMS server and provide a higher level of security.

You can use KMS keys to encrypt backup files on the following encryption levels:

- Job-level encryption:
  - Backup and backup copy jobs
  - Veeam Agent backup jobs managed by Veeam Backup & Replication
  - File backup jobs and object storage backup jobs
  - Transaction log backup and backup copy jobs
  - VeeamZIP jobs

  For more information about job-level encryption, see Job Encryption.

  > **NOTE**
  >
  > If you use Veeam Cloud Connect repositories as a target backup storage, you can also use KMS keys for the following jobs:
  >
  > - Backup and backup copy jobs
  > - Veeam Agent backup jobs managed by Veeam Backup & Replication
  > - Transaction log backup copy jobs

- Storage-level encryption:
  - Backup repositories that store backup files created by:
    - Veeam Backup for Nutanix AHV
    - Veeam Backup for OLVM and RHV
    - Veeam Kasten for Kubernetes

    For more information about storage-level encryption for Veeam Backup & Replication additional solutions, see Encrypting Standalone Application Backups in Backup Repositories.
  - Capacity tier repositories. For more information about storage-level encryption for capacity tier repositories, see Encryption for Capacity Tier.
  - Media pools and GFS media pools. For more information about storage-level encryption for tape devices, see Tape Encryption.
  - External repositories (decryption only).

> **IMPORTANT**
>
> The following jobs and repositories do not support data encryption with KMS keys:
>
> - Configuration backup jobs
> - Veeam Agent backup jobs managed by Veeam Agents
> - Backup repositories that store backup files created by Veeam Agents operating in the standalone mode

# How KMS Works

When you add the KMS server in the Veeam Backup & Replication console and start using KMS keys for data encryption, Veeam Backup & Replication asks the KMS server to generate an asymmetric KMS key for the required job or repository. Veeam Backup & Replication stores a public key in the configuration database and uses it for data encryption. The KMS server stores a private key and uses it for data decryption.

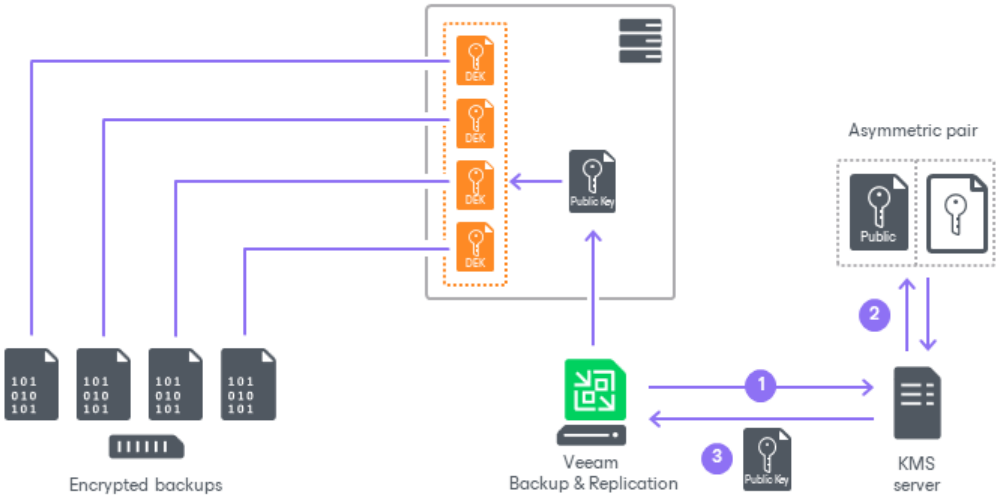


Backup files encrypted with the KMS solution are decrypted automatically when you import them. You need to decrypt backup files manually only in the following situations:

- You import backup files to a new Veeam Backup & Replication installation that is not connected to the KMS server yet.
- You use VBK files to import backups. Also, a part of the backup chain is encrypted with the KMS solution, the other part is encrypted with the password-based keys.
- You use tapes and have specific considerations. For more information, see Decrypting Tapes with KMS Keys in the Tape Device Support Guide.

> **NOTE**
>
> In case of a KMS server failure, backup jobs that use the KMS keys for data encryption will fail. To decrypt backup files, you can use Veeam Backup Enterprise Manager if the encryption password loss protection is enabled. For more information, see Password Loss Protection.

The KMS server rotates KMS keys at a time interval specified in the KMS policies. To get updates from the KMS server, Veeam Backup & Replication runs a system job. During the job session, Veeam Backup & Replication performs the following steps:

1. Sends a request to the KMS server and gets information about recently rotated KMS keys if there are any.
2. Updates public keys in the Veeam Backup & Replication configuration database.

By default, the Veeam Backup & Replication system job runs every 24 hours. If you want to change the default time period, contact Veeam Customer Support.

# Requirements and Limitations

The KMS feature has the following requirements and limitations:

- The feature is included in the Veeam Data Platform Advanced or Premium License. For more details about all license types, see Veeam Data Platform Feature Comparison. Data decryption is available for all licenses.
- Veeam Backup & Replication supports KMS servers that meet the following requirements:
  - Key Management Interoperability Protocol (KMIP) Profile v1.4. Other versions of KMIP Profiles are not supported by Veeam Backup & Replication.
  - Requirements for a baseline server. For more information, see the Baseline Server section in the KMIP Profile standard.
  - Requirements for an asymmetric key lifecycle server. For more information, see the Asymmetric Key Lifecycle Server section in the KMIP Profile standard.

  > **TIP**
  >
  > The list of tested KMS solutions includes the following vendor product lines:
  >
  > - Thales CipherTrust Manager k170v 2.10.0+7973 and later
  > - Fortanix Data Security Manager KMS 4.20.2274 and later (Public Cloud solution)
  > - IBM Security Guardium Key Lifecycle Manager (GKLM) 4.1.1.0 and later
- To decrypt data, the KMS server must support:
  - Requirements for a basic cryptographic server. For more information, see the Basic Cryptographic Server section in the KMIP Profile standard.
  - Optimal Asymmetric Encryption Padding (OAEP) with SHA-1.

  In other cases, Veeam Backup & Replication will retrieve private keys from the KMS server to decrypt backup files. These keys are not stored in the configuration database and deleted immediately after decryption.
- [For Cloud Connect] To use the KMS feature in the Veeam Cloud Connect environment, both a service provider and a tenant must run Veeam Backup & Replication 12.1 (build 12.1.0.2131) or later.
- [For Cloud Connect] If a tenant uses the same KMS server as a service provider, backup files stored in the tenant quota cannot be decrypted on the service provider side.

# KMS Certificates

The KMS server certificate must meet the following requirements:

- The **Subject** extension must be equal to the fully qualified domain name (FQDN) of the KMS server. For example: `kms.domain.local`.
- The server certificate must have valid CRL distribution points specified in the **CRL Distribution Points** extension.
- If the Veeam Backup & Replication server does not trust the Certificate Authority (CA) of the server certificate, it should be added to the Trusted Root Certification Authority store.

The client certificate issued by the KMS administrator for Veeam Backup & Replication must be exportable.

# Adding KMS Server

To add a KMS server, do the following:

1. From the main menu, select **Credentials and Passwords** > **Key Management Servers**.
2. In the **Key Management Servers** window, click **Add**.
3. In the **Server** field, specify the FQDN, IPv4 or IPv6 address of the server. By default, the port number 5696 is used.
4. In the **Server certificate** field, click **Browse** and specify a KMS server certificate. You can select one of the following options:
   a. **Select an existing certificate from the certificate store**. You can specify a KMS server certificate if it is located in the Microsoft Windows certificate store.
   b. **Import certificate from a file**. You can import a KMS server certificate from a file in the PFX, CER, or PEM format.

   > **NOTE**
   >
   > If you use a server certificate in the PEM format, it must contain the `-----BEGIN CERTIFICATE-----` header at the beginning of the file and the `-----END CERTIFICATE-----` footer at the end of the file.

   For more information about requirements that a server certificate must meet, see KMS Certificates.

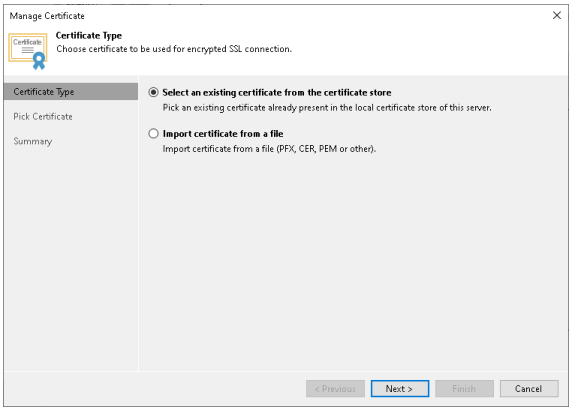


5. In the **Client certificate** field, click **Browse** and specify the client certificate issued by the KMS administrator for Veeam Backup & Replication. You can select one of the following options:
   a. **Select an existing certificate from the certificate store**. You can specify a client certificate issued by the KMS administrator for Veeam Backup & Replication if it is located in the Microsoft Windows certificate store.
   b. **Import certificate from a file**. You can import a client certificate from a file in the PFX format.

   > **NOTE**
   >
   > If you use a PEM-encoded file, select the **Import certificate from a PEM file** option.

   c. **Import certificate from a PEM file**. You can import a client certificate from a PEM-encoded file. Consider the following:
      - You must have two separate PEM files for the certificate and private key.

- The certificate must contain the `-----BEGIN CERTIFICATE-----` header at the beginning of the file and the `-----END CERTIFICATE-----` footer at the end of the file.
- The private key must be in the PKCS#1 format. Also, it must contain the `-----BEGIN RSA PRIVATE KEY-----` header at the beginning of the file and the `-----END RSA PRIVATE KEY-----` footer at the end of the file.

For more information about requirements that a client certificate must meet, see KMS Certificates.

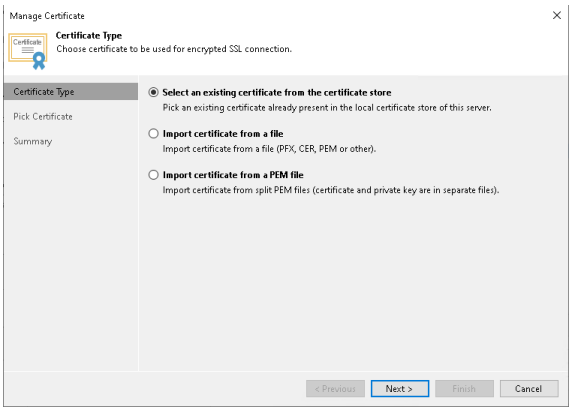


6. [Optional] In the **Description** field, provide the description for future reference.
7. Click **OK**.

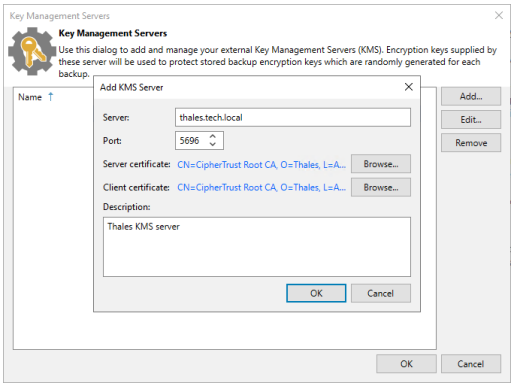


# Using KMS Keys

To use KMS keys for data encryption, go to the backup job or storage settings and select the KMS server in the **Password** field. For more information about data encryption levels, see Job Encryption and Storage Encryption.

Note that If the KMS server is unavailable, the job session will fail.

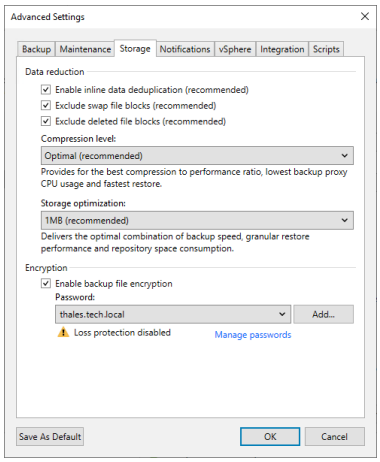


# Decrypting Backups

If you encrypt data encryption keys with KMS keys, import an encrypted backup file to the Veeam Backup & Replication console. The file will be automatically decrypted and appear under the **Backups > Disk (Imported)** node in the inventory pane.

If you need to decrypt backup files manually, do the following:

1. In the inventory pane, select the **Backups** > **Disk (Encrypted)** node.
2. In the working area, select an encrypted backup file and click **Specify Password** on the ribbon. Alternatively, you can right-click the backup and select **Decrypt backup**.
3. Check the encryption key ID and selected KMS server, and click **OK**.

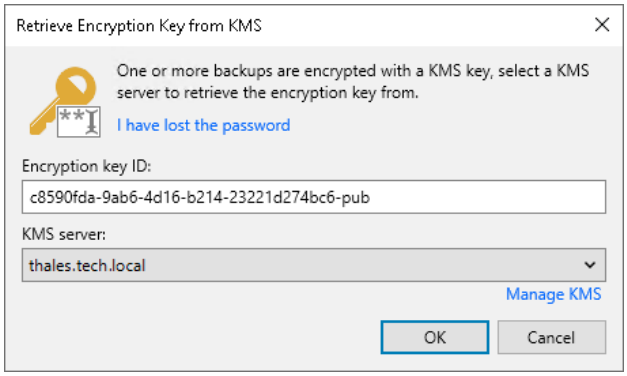


Consider the following:

- If the KMS key has changed several times or you have changed the encryption method to the password, you need to specify them in the following manner:
  - If you select a metadata file (VBM) for import, you must specify the latest KMS key or password that was used to encrypt data encryption keys.

- If you select a full backup file (VBK) for import, you must specify the whole set of KMS keys and passwords that were used to encrypt data encryption keys.

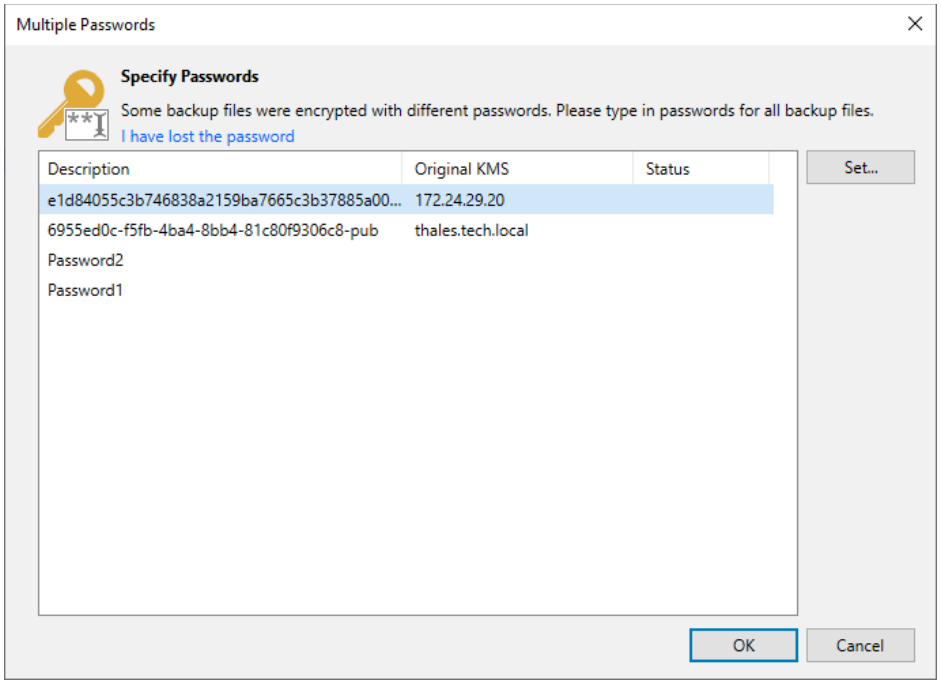


- If you cannot use KMS keys due to a KMS server failure, you can issue a request to Veeam Backup Enterprise Manager and restore data from an encrypted file using Enterprise Manager keys. This option works only if the encryption password loss protection is enabled. For more information, see Password Loss Protection.

# Managing Locations

To control data migration in the virtual infrastructure, Veeam Backup & Replication introduces a notion of location. A location defines a geographic region, or country, in which an infrastructure object resides. You can create a list of locations, and assign to backup infrastructure objects information about locations to which they belong.

Veeam Backup & Replication allows you to assign information about locations to the following infrastructure objects:

- Virtual infrastructure objects: vCenter Servers, datacenters, clusters and hosts.
- Backup infrastructure objects: backup repositories, scale-out backup repositories, tape libraries and tape vaults. Locations of external repositories are determined and assigned automatically based on the datacenter region.
- Agent management objects: protection groups.
- Veeam Cloud Connect for service providers: cloud repositories and hardware plans.

Information about infrastructure objects location is stored in the Veeam Backup & Replication configuration database. When VM data in the virtual infrastructure is migrated from one location to another, Veeam Backup & Replication displays a warning and stores a record about data migration to job or task session details. In addition to it, Veeam Backup & Replication logs this information to Microsoft Windows event logs. For example, if you back up VMs from a host that resides in Germany to a backup repository that resides in Australia, Veeam Backup & Replication will display a warning that VM data changes its location in the backup job wizard, display information about data migration in the backup job session details and log it to Microsoft Windows event logs.

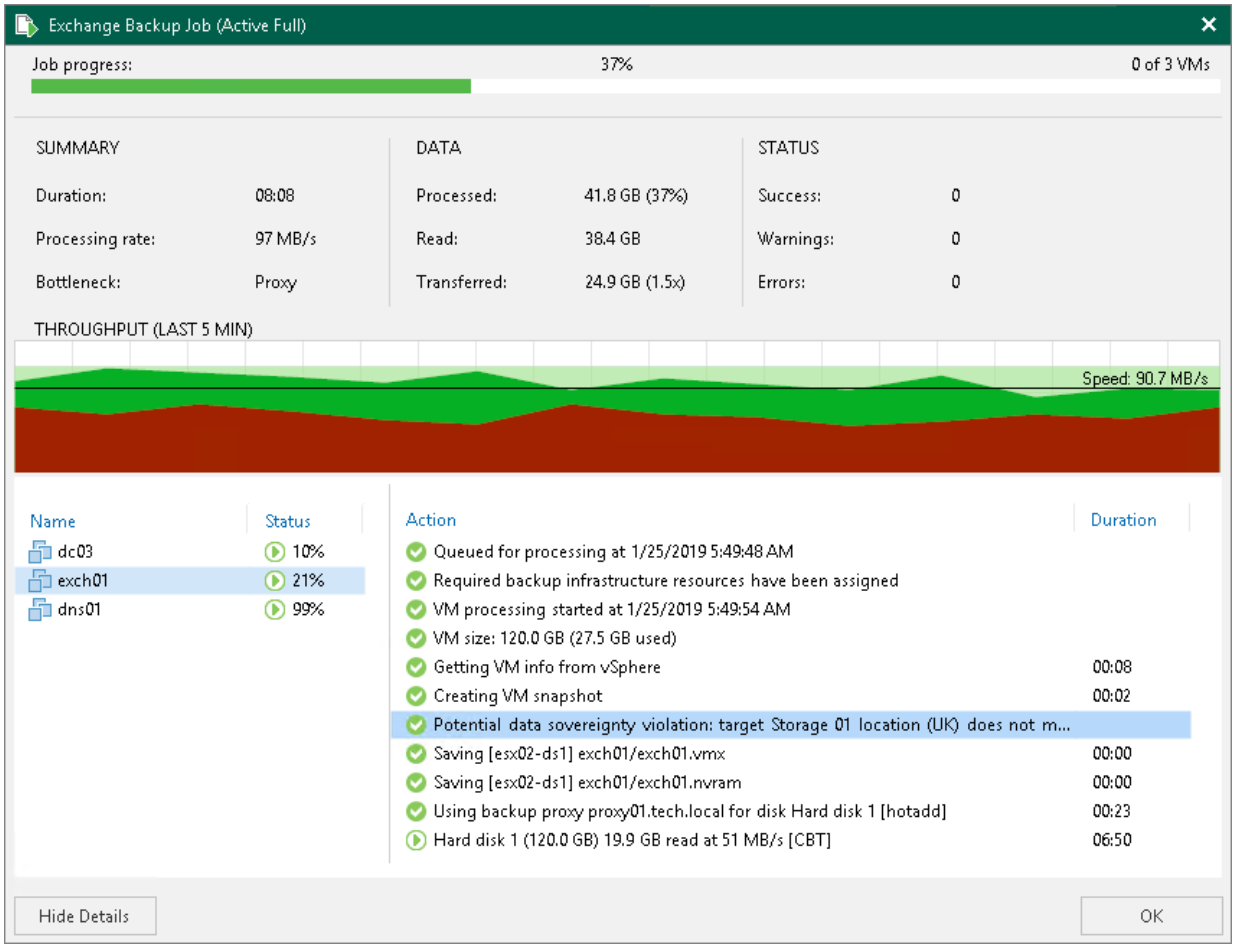

Veeam Backup & Replication displays information about VM and file share data migration in statistics for the following types of jobs:

- Backup jobs — Veeam Backup & Replication compares the location of the source host on which VMs are registered with the location of the target backup repository or cloud repository.
- Backup copy jobs — Veeam Backup & Replication compares the location of the source host with the location of the target host.
- File backup jobs — Veeam Backup & Replication compares the location of the source file share with the location of the target backup repository or cloud repository. If the secondary repository is specified, Veeam Backup & Replication compares the location of the source file share with the location of the secondary target host.
- VeeamZIP tasks (except the cases when you select to store the VeeamZIP file in a local or shared folder) — Veeam Backup & Replication compares the location of the source host on which VMs are registered with the location of the target backup repository.
- Replication jobs — Veeam Backup & Replication compares the location of the source host on which VMs are registered with the location of the target host.
- Replica failback tasks — Veeam Backup & Replication compares the location of the source host with the location of the host to which the VM is restored.
- VM copy jobs — Veeam Backup & Replication compares the location of the source host on which VMs are registered with the location of the target backup repository or target host.
- Quick migration tasks — Veeam Backup & Replication compares the location of the source host on which VMs are registered with the location of the target host.
- File share backup copy tasks — Veeam Backup & Replication compares the location of the source file share with the location of the target backup repository or target host.
- Entire VM restore tasks — Veeam Backup & Replication compares the location of the source host with the location of the host to which VMs are restored.
- File share data recovery tasks — Veeam Backup & Replication compares the location of the source file share with the location of the file share to which files are restored.
- External repository tasks:
  - Backup copy jobs: Veeam Backup & Replication compares the location of the source external repository with the location of the target backup repository.
  - Restore to Amazon EC2: Veeam Backup & Replication compares the geographic region of the backed-up Amazon EC2 instance with the geographic region of the target EC2 instance.
  - Restore to Microsoft Azure: Veeam Backup & Replication always displays a warning about VM data migration when restore to Microsoft Azure is performed from external repositories.
- SureBackup jobs — Veeam Backup & Replication compares the source location with the target location. The target location is always a host on which the virtual lab is registered. The source location may be one of the following:
  - If a VM is added to the application group, Veeam Backup & Replication compares the host on which the VM is registered (or was registered at the moment of backup) with the target location.
  - If a VM is added to the SureBackup job from the linked job, Veeam Backup & Replication compares the backup repository on which the backup file resides with the target location.

- Tape tasks:
  - Backup to tape jobs: In backup to tape jobs, Veeam Backup & Replication compares the location of the source job or repository with the location of the tape library in the target media pool. If the media pool spans multiple tape libraries, Veeam Backup & Replication analyzes locations of all libraries in the media pool.
  - Vaults: If a tape job exports offline backups to a vault, Veeam Backup & Replication compares the location of the source job or repository with the location of the vault. If a GFS tape job exports tapes to multiple vaults, Veeam Backup & Replication analyzes all vaults configured for target media pools of the GFS tape job.
  - Media pools: Veeam Backup & Replication compares locations of all tape libraries added to the media pool. If the media pool exports tapes to a vault, Veeam Backup & Replication analyzes all vaults configured for the media pool.

# Limitations for Locations

Consider the following:

- You cannot assign or edit locations of external repositories. Veeam Backup & Replication automatically determines and sets locations for external repositories based on the datacenter region. You can check the datacenter region at the Bucket step of the **Edit External Repository** wizard.
- For SureReplica jobs, Veeam Backup & Replication does not compare information about source and target hosts location.
- Veeam Backup & Replication does not display a warning about VM data migration for file copy jobs.

# Creating and Assigning Locations to Infrastructure Objects

You can create a list of locations in Veeam Backup & Replication and assign locations to infrastructure objects. If you assign a location to a vCenter Server, it will be applied to all child hosts (clusters and ESXi hosts). You can also assign the location to a child host.

To create a location:

1. In the **Inventory** or **Backup Infrastructure** view, right-click the infrastructure object and select **Location** > **Manage locations**.
2. In the **Manage Locations** window, click **Add**.
3. In the **Name** field, enter a name of the location.

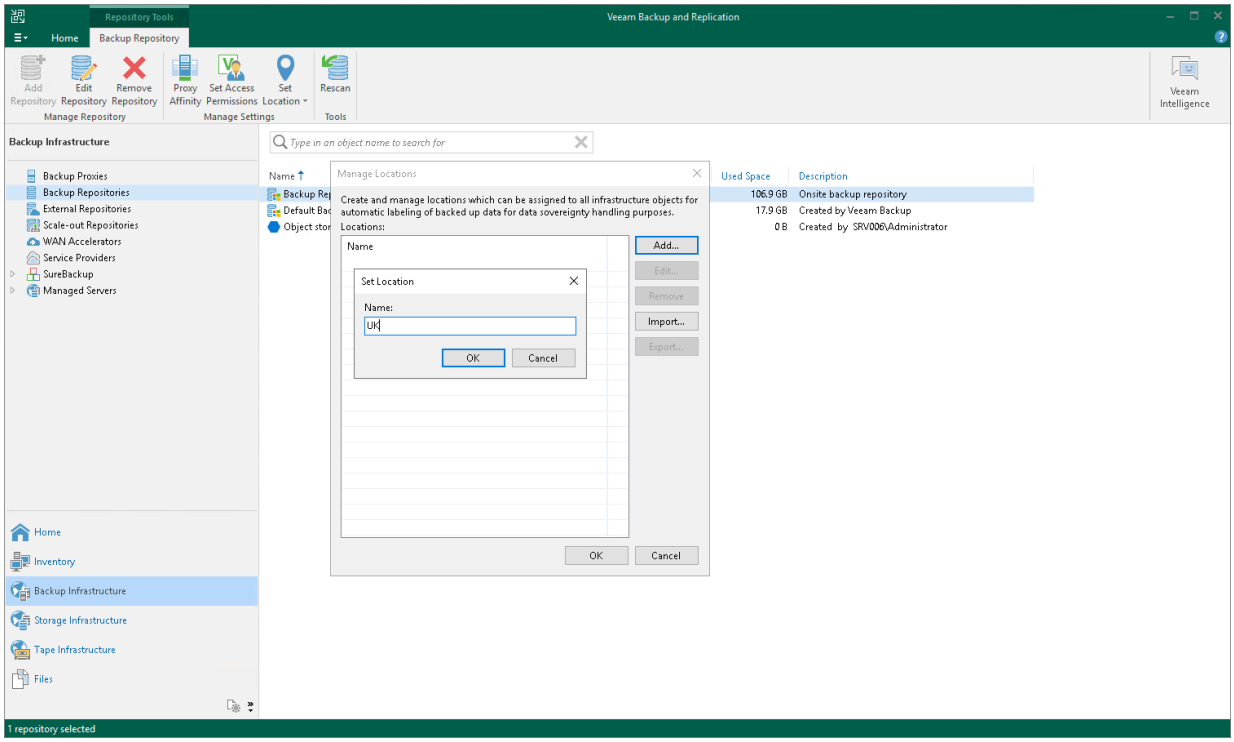


To assign a location to an infrastructure object, in the **Inventory** or **Backup Infrastructure** view, right-click the infrastructure object and select Location > *<Location name>*. If the location is not in the list, select **Location** > **Manage Locations** and add the location to the list.

> **NOTE**
>
> When assigning a location to a scale-out backup repository, the location will be global for all extents. If you add an extent whose location differs from the global location, it will be changed in favor of the location of the scale-out repository.

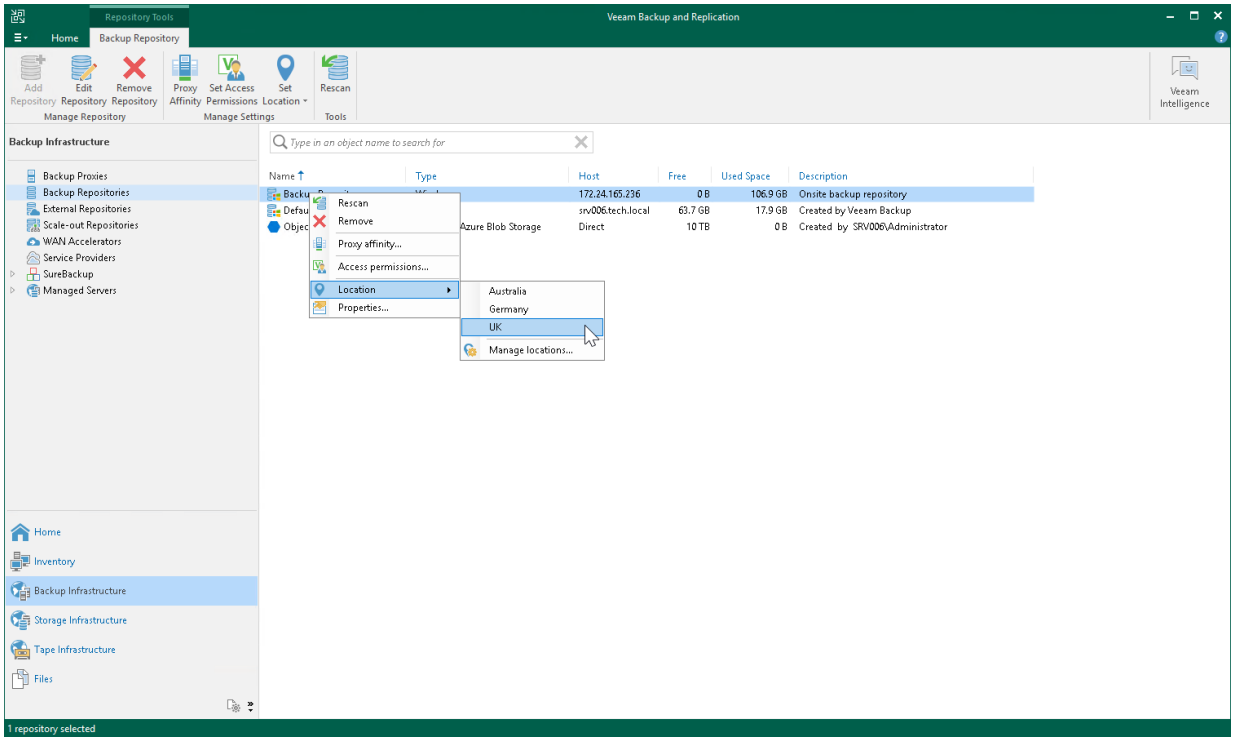
Repository Tools
Veeam Backup and Replication
Home
Backup Repository
Add Repository
Edit Repository
Remove Repository
Manage Repository
Proxy Affinity
Set Access Permissions
Set Location
Manage Settings
Rescan
Tools
Veeam Intelligence
Backup Infrastructure
Backup Proxies
Backup Repositories
External Repositories
Scale-out Repositories
WAN Accelerators
Service Providers
SureBackup
Managed Servers
Type in an object name to search for
Name
Type
Host
Free
Used Space
Description
172.24.165.236
0 B
106.9 GB
Onsite backup repository
srv006.tech.local
63.7 GB
17.9 GB
Created by Veeam Backup
Azure Blob Storage
Direct
10 TB
0 B
Created by SRV006\Administrator
Rescan
Remove
Proxy affinity...
Access permissions...
Location
Properties...
Australia
Germany
UK
Manage locations...
Home
Inventory
Backup Infrastructure
Storage Infrastructure
Tape Infrastructure
Files
1 repository selected

# Editing Locations

You can edit a location in the locations list, for example, if you want to change the location name.

To edit a location:

1. In the **Inventory** or **Backup Infrastructure** view, right-click the infrastructure object and select **Location** > **Manage locations**.
2. In the **Manage Locations** window, select the location and click **Edit**.
3. In the **Name** field, change the location name as required.

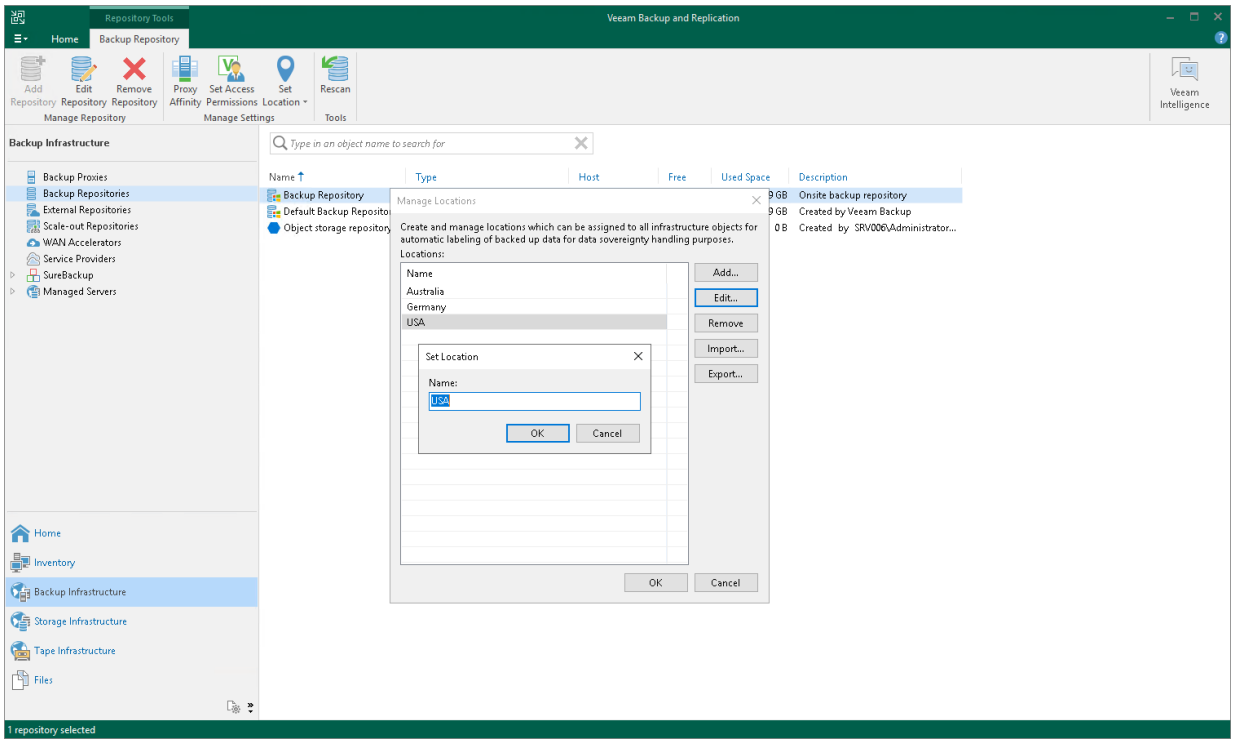

# Deleting Locations

You can delete a location from the locations list, for example, if you no longer host infrastructure objects in this location.

To delete a location:

1. In the **Inventory** or **Backup Infrastructure** view, right-click the infrastructure object and select **Location** > **Manage locations**.
2. In the **Manage Locations** window, select the location and click **Remove**. If the location is currently assigned to some infrastructure objects, Veeam Backup & Replication will display a warning with the list of objects that belong to this location. Click **Yes** to confirm the location deletion.

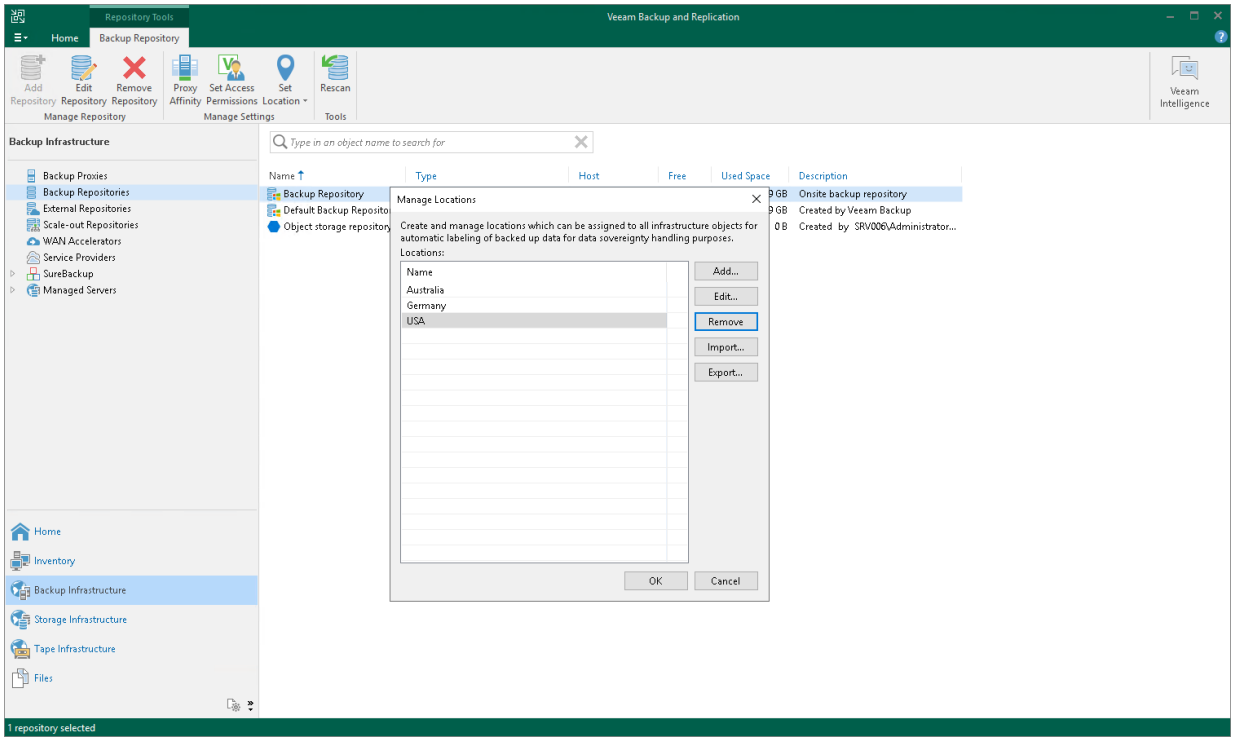

# Exporting and Importing Locations List

You can export and import the list of locations to/from a file of XML format.

The import and export functionality facilitates the process of locations creation and maintenance. For example, if you need to set up the same list of locations throughout the whole backup infrastructure, you can create a list of locations on one backup server manually, export this list to an XML file, and then import the list on other backup servers and machines running the Veeam Backup & Replication console.

> **TIP**
>
> If you delete and recreate a location, Veeam Backup & Replication will create an object with a new ID in the database and consider it as a new location. Thus, to preserve the uniqueness of the location, use the location export/import operations.

To export the locations list:

1. In the **Inventory** or **Backup Infrastructure** view, right-click an infrastructure object and select **Location** > **Manage locations**.
2. In the **Manage Locations** window, click **Export** and specify a name of the XML file to which the locations list must be exported.

To import the locations list:

1. In the **Inventory** or **Backup Infrastructure** view, right-click an infrastructure object and select **Location** > **Manage locations**.
2. In the **Manage Locations** window, click **Import** and browse to the XML file from which the locations list must be imported.

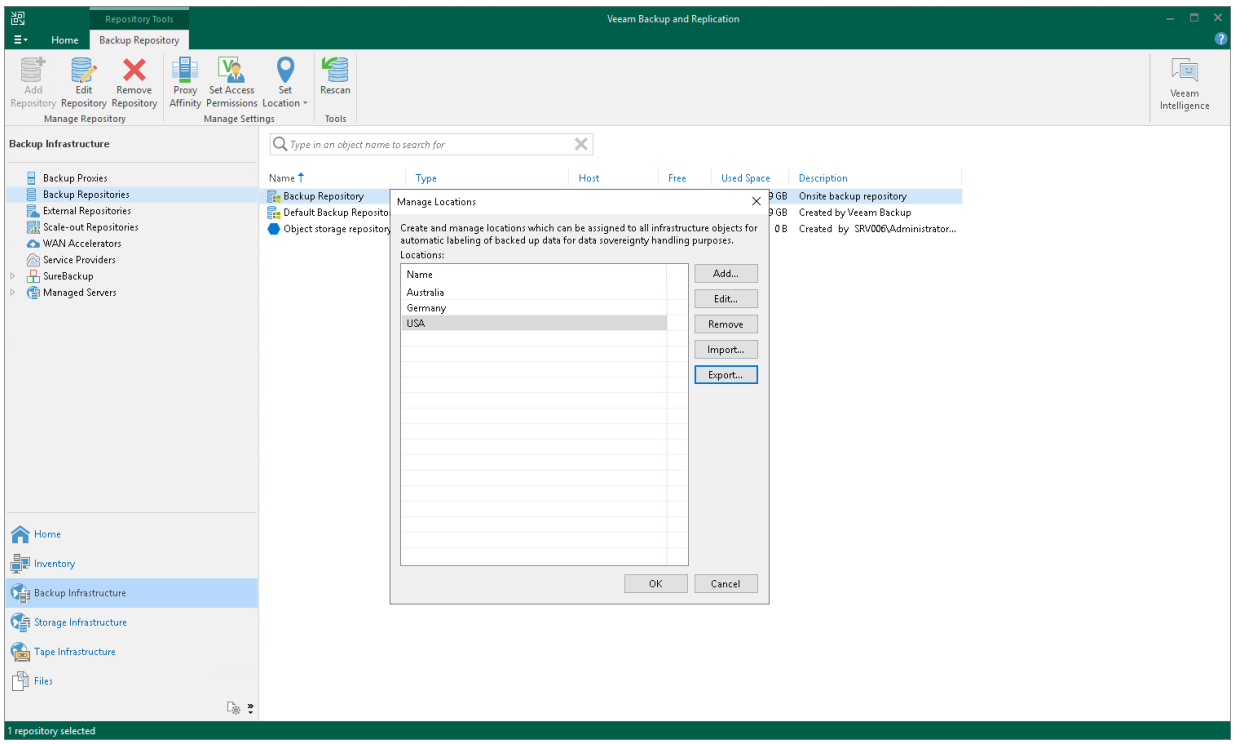

# Managing Network Traffic

Data protection requires large amount of data to be transferred through networks. This may cause heavy network loads. Veeam Backup & Replication offers the following capabilities to manage network traffic and reduce the load:

- Configuring Network Traffic Rules
- Managing Upload Streams
- Specifying Preferred Networks
- IPv6 Support

# Configuring Network Traffic Rules

Network traffic rules control traffic transferred between backup infrastructure components. These rules allow you to do the following:

- Throttle network traffic
- Encrypt transferred data

The rules apply only to traffic sent between the backup infrastructure components, so you do not have to change your network infrastructure.

## How Network Rules Work

Each network rule contains IP address ranges for source and target components. When a job starts, Veeam Backup & Replication checks the rules against the components involved in the job. If the IP addresses of the components fall into the IP address ranges of a rule, the rule applies.

For example, 192.168.0.1-192.168.0.255 is the source range, and 172.16.0.1-172.16.0.255 is the target range. 192.168.0.12 is the IP address of one component, and 172.16.0.31 is the IP address of another component. Both IP addresses fall into the ranges, so the rule will apply.

Note that the rules are reversible. The rule from the example will also apply to the specified components if you swap the ranges: make 192.168.0.1-192.168.0.255 the target range and 172.16.0.1-172.16.0.255 the source range.

**NOTE**

If the configured network rule covers at least one of the connection interfaces to a Veeam agent, Veeam Backup & Replication applies network traffic throttling even if the traffic is transferred to/from this Veeam agent through another interface.

**TIP**

You can define a rule for specific components. For this, specify a single IP address in the source range and in the target range.

## Creating Network Rules

You must create network rules at the backup server level. For details, see Enabling Traffic Throttling and Enabling Traffic Encryption. Veeam Backup & Replication also has a predefined rule for traffic transferred between public networks. For more information, see Adjusting Internet Rule.

After you create or edit network rules, they are applied almost immediately.

**TIP**

If you created a rule for a VMware backup proxy, you can check whether it applies. For this, open the Traffic Rules step of the backup proxy wizard. The rule must be in the list of rules.

# Adjusting Internet Rule

The internet rule is a predefined network rule. This rule manages traffic transferred through public networks — all IPv4 networks whose IP ranges differ from 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16 and all IPv6 networks whose IP ranges differ from fc00::/7. The internet rule also encrypts traffic for such networks. You cannot delete the internet rule — you can only adjust or turn it off.

> **NOTE**
>
> The internet rule does not apply to restores to cloud platforms (restore to Amazon EC2, restore to Microsoft Azure, and so on).

# Turning Off Rule

To turn off the internet rule, clear the **Throttle network traffic to** and **Encrypt network traffic** check boxes.

# Adjusting Rule

To adjust the rule:

1. From the main menu, select **Network Traffic Rules**.
2. In the **Global Network Traffic Rules** window, select **Internet** from the list and click **Edit**.
3. In the **Edit Network Traffic Rule** window:
   - To disable encryption, clear the **Encrypt network traffic** check box.
   - To enable network traffic throttling, select the **Throttle network traffic to** check box. For details, see Enabling Traffic Throttling.

> **TIP**
>
> You can create custom network traffic rules targeted to public networks as described in section Enabling Traffic Throttling.

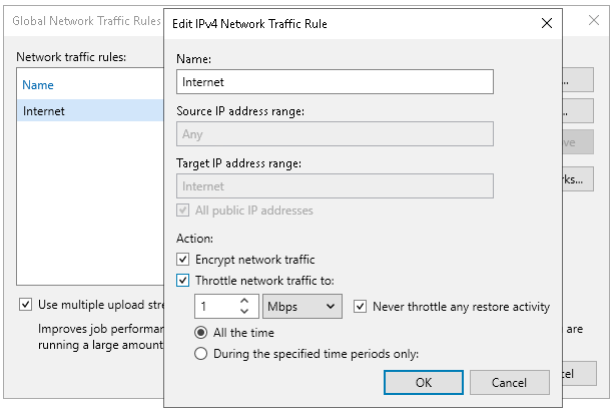

# Enabling Traffic Throttling

Traffic throttling setting in a network rule allows you to limit the impact of Veeam Backup & Replication tasks on network performance. Traffic throttling prevents jobs from utilizing the entire bandwidth available in your environment and makes sure that other network operations get enough traffic.

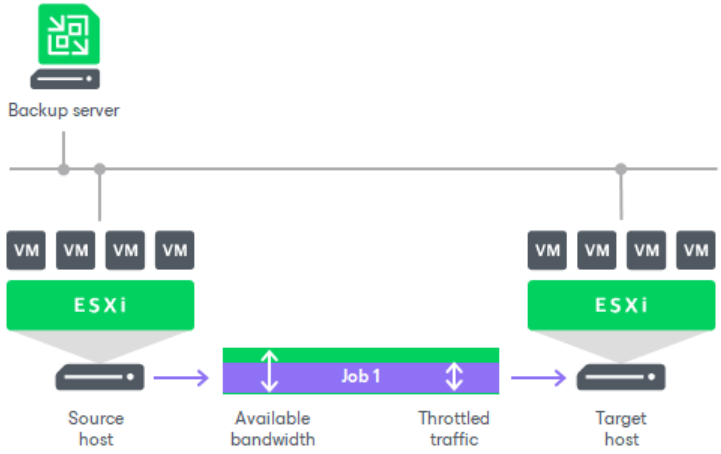


When several jobs that fall into the same network rule run simultaneously, Veeam Backup & Replication equally splits the throttled traffic between them. For example, if two jobs run at a time, each job gets half of the throttled traffic. Once one job finishes, the other gets the entire bandwidth allowed by the rule.

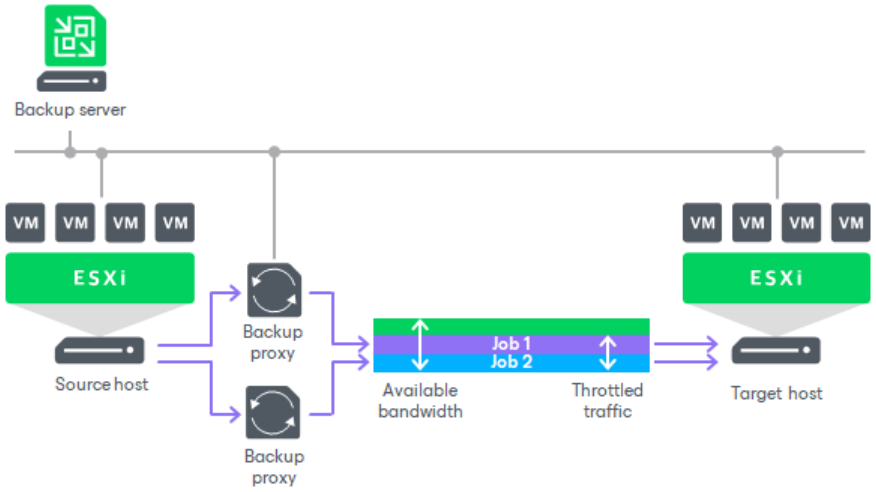


**NOTE**

Consider the following:

- Traffic can be throttled only if the infrastructure components are located on different servers and have different IP addresses. For example, if you use the same server as a backup proxy and a backup repository, traffic throttling rules do not apply to them.
- It is recommended that you throttle network traffic if you perform off-site backup or replicate VMs to a DR site over slow WAN links.
- Traffic throttling rules apply to the transfer of virtual disks. At the beginning of the job, you may see bandwidth usage peaks caused by transferring other VM files.

# Infrastructure Components

Traffic can be throttled between the backup infrastructure components on which Veeam Data Movers are deployed, also capacity tier of the scale-out backup repository and object storage repository. The components differ depending on a scenario. The following table shows this dependency.

| Scenario | Components |
|---|---|
| Backup to a Microsoft Windows or Linux backup repository* | VMware backup proxy and backup repository. |
| Backup to an SMB share, Dell Data Domain and HPE StoreOnce* | VMware backup proxy and gateway server. |
| Backup to an object storage repository (the direct connection mode)** | Backup proxy and object storage repository. |
| Backup to an object storage repository (using a gateway server)** | Backup proxy and gateway server. |
| Backup to a Veeam Cloud Connect repository | Backup proxy and cloud gateway server. |
| CDP | Source and target VMware CDP proxies. |
| VM copy | Backup proxy and backup repository. |
| Backup copy* | Source and target backup repositories, gateway servers or WAN accelerators (if they are involved in the backup copy process). |
| Replication | Source and target VMware backup proxies or WAN accelerators (if they are involved in the replication process). |
| File backup from a managed file server | General-purpose backup proxy and backup repository. |
| Object storage backup | General-purpose backup proxy and backup repository. |
| Backup to tape* | Backup repository and tape server. For more information about backup to tape jobs, see the Machines Backup to Tape section in the Veeam Backup & Replication User Guide. |
| SOBR data offload | Gateway server and object storage repository. |
| Virtual Disk Restore | Backup proxy and backup repository. |

| Scenario | Components |
| --- | --- |
| Entire VM Restore | Backup proxy and backup repository. |

* Veeam Backup & Replication throttles traffic between the listed components also if backups are created with Veeam Agents (Windows, Linux and so on) operating in the standalone or managed mode.

** Veeam Backup & Replication throttles traffic for all object storage repositories. For on-premises object storage repositories Veeam Backup & Replication uses network rules.

# Configuring Traffic Throttling

To configure traffic throttling settings in a rule:

1. From the main menu, select **Network Traffic Rules**.
2. In the **Global Network Traffic Rules** window, click **Add** and select an IPv4 or IPv6 rule. Note that you can add the IPv6 rule only if IPv6 communication is enabled as descried in IPv6 Support.
3. In the **Name** field, specify a name for the rule.
4. In the **Source IP address range** section, specify a range of IP addresses for the backup infrastructure components on the source site.
5. Specify ranges of IP addresses on the target site:
   - To specify IP addresses of the backup infrastructure components, use the **Target IP address range** section.
   - To throttle traffic to public networks, select the **All public IP addresses** check box. Public networks are all IPv4 networks whose IP ranges differ from 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16 and all IPv6 networks whose IP ranges differ from fc00::/7.
6. Select the **Throttle network traffic to** check box.
7. In the **Throttle network traffic to** field, specify the maximum speed that must be used to transfer data from source to target.
8. If you want to throttle traffic during the restore activities, clear the **Never throttle any restore activity** check box.

   > **NOTE**
   >
   > The **Never throttle any restore activity** check box applies only to the restore methods listed in the Infrastructure Components table.

9. In the section under the **Throttle network traffic to** field, specify time periods when traffic is throttled:
   - Select the **All the time** option if traffic throttling rules must be applied continuously.

- Select **During the specified time periods only** and use the diagram under it to mark the periods when the traffic must be throttled or not. Use the **Unthrottle to** field to increase or decrease the throttling limit for specific hours and days.

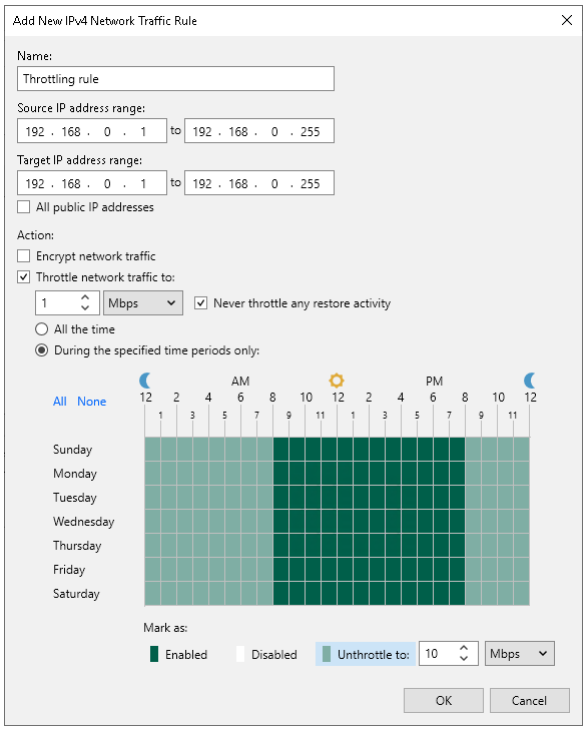


# Several Rules with Traffic Throttling

If you create several rules with the same ranges of IP addresses, make sure that time intervals of the rules do not overlap. For example, to manage network traffic during business and non-business hours, you can create the rules as in the following example. These rules have the same ranges of IP addresses.

- **Rule 1**. Speed limit: 1 Mbps; time interval: Monday through Friday from 7 AM to 7 PM.
- **Rule 2**. Speed limit: 10 Mbps; time interval: Saturday through Sunday from 7 AM to 7 PM.

With such rules, Veeam Backup & Replication will limit the speed up to 1 Mbps during business hours and up to 10 Mbps during non-business hours.

If several rules have the same target/source IP address range but different speed limits, the lowest limit is used. For example, if you configure the following rules:

- **Rule 1**. Source IP range: 192.168.0.1-192.168.0.30; target IP range: 192.168.0.1-192.168.0.255; speed limit: 4 Mbps.
- **Rule 2**. Source IP range: 192.168.0.1-192.168.0.255; target IP range: 192.168.0.1-192.168.0.255; speed limit: 1 Mbps.

In this case, Veeam Backup & Replication will use the lowest speed limit — 1 Mbps.

The principle of several rules and the lowest speed limit also applies if some rules are created on the Veeam Backup & Replication side and others on the Veeam Agent for Microsoft Windows side. For more information on how to throttle traffic by Veeam Agent for Microsoft Windows, see the Veeam Agent Management Guide and Veeam Agent for Microsoft Windows User Guide.

# Enabling Traffic Encryption

By default, Veeam Backup & Replication encrypts network traffic transferred between public networks. For details, see Adjusting Internet Rule.

Network rules also allow you to encrypt backup data transfer connections between Veeam Data Movers in private networks. Network traffic encryption is provided by TLS connection and configured as the part of global network traffic rules that are set for backup infrastructure components.

For more information about supported TLS versions and cipher suites, see Supported Cipher Suites and Protocols.

To create a network rule with traffic encryption:

1. From the main menu, select **Network Traffic Rules**.
2. In the **Name** field, specify a name for the rule.
3. In the **Global Network Traffic Rules** window, click **Add** and select an IPv4 or IPv6 rule. Note that you can add the IPv6 rule only if the **Enable IPv6 communication** check box is selected. For more information, see IPv6 Support.
4. In the **Source IP address range** section, specify a range of IP addresses for backup infrastructure components on the source site.
5. In the **Target IP address range** section, specify a range of IP addresses for backup infrastructure components on the target site.
6. Select the **Encrypt network traffic** check box.

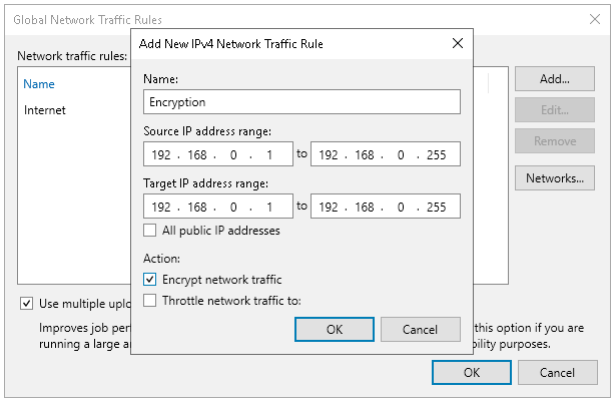


> **NOTE**
>
> If encryption is enabled on a backup job level, backup data will be encrypted before sending. For more information, see Job Encryption.

# Related Topics

Communications Encryption

# Managing Upload Streams

By default, Veeam Backup & Replication uses multithreaded data transfer for every job session. VM data going from source to target is transferred over 5 TCP/IP connections. However, if you schedule several jobs to run at the same time, load on the network may be heavy. If the network capacity is not sufficient to support multiple data transfer connections, you can disable multithreaded data transfer or change the number of TCP/IP connections.

To change the number of connections:

1. From the main menu, select **Network Traffic Rules**.
2. In the **Global Network Traffic Rules** window, specify new data transfer settings:
   - To disable multithreaded data transfer, clear the **Use multiple upload streams per job** check box. Veeam Backup & Replication will use only one TCP/IP transfer connection for every job session.
   - To change the number of TCP/IP connections, leave the **Use multiple upload streams per job** check box selected and specify the necessary number of connections in the field on the right.

**NOTE**

Consider the following:

- The minimum number of TCP/IP connections supported by Veeam Backup & Replication is 1, the maximum number is 100.
- If the **Use multiple upload streams per job** check box is enabled, Veeam Backup & Replication performs a CRC check for the TCP traffic going between the source and the target. The check works when you perform data protection and recovery operations, except for restores that involve mount activities such as guest OS file restore and Instant Recovery. Veeam Backup & Replication calculates checksums for data blocks going from the source. On the target, it recalculates checksums for received data blocks and compares them to the checksums created on the source. If the CRC check fails, Veeam Backup & Replication automatically re-sends data blocks without any impact on the job. If you disable the **Use multiple upload streams per job** check box, the CRC check also becomes disabled.
- [For Veeam Plug-ins for Enterprise Applications] The multithreaded data transfer setting does not affect backup jobs created by Veeam Plug-in for Oracle RMAN/SAP HANA/SAP on Oracle. To configure multiple channels for backup and restore operations for these plug-ins, see the Veeam Plug-ins for Enterprise Applications Guide.

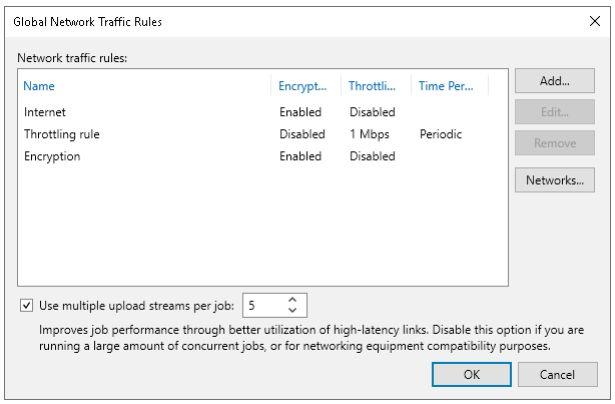

# Specifying Preferred Networks

You can choose networks over which Veeam Backup & Replication must transport data when you perform data protection and disaster recovery tasks. This option can be helpful if you have a non-production network and want to route data traffic over this network instead of the production one.

Preferred network rule applies only to traffic between the following backup infrastructure components:

- Backup server
- WAN accelerator[1]
- Gateway server[2]
- VMware CDP proxy
- VMware backup proxy
- Backup repository
- Log shipping server
- Tape server
- Storage system (backup from storage snapshots scenario)[3]
- Veeam Agent
- Veeam Plug-ins for Enterprise Applications

[1] The rule applies only to traffic between the source and target WAN accelerators.
[2] The rule does not apply to traffic between the gateway server and backup repository. For the list of backup repositories with which the gateway server communicates, see Gateway Servers.
[3] The rule does not apply when backing up from Cisco HyperFlex systems or when the transport mode is Direct storage access.

To define networks for data transfer, you must create a list of preferred networks. When Veeam Backup & Replication needs to transfer data, it uses networks from this list. If a connection over preferred networks cannot be established for some reason, Veeam Backup & Replication will automatically fail over to the production network.

To specify a preferred network:

1. From the main menu, select **Network Traffic Rules**.
2. In the **Global Network Traffic Rules** window, click **Networks**.
3. In the **Preferred Networks** window, select the **Prefer the following networks for backup and replication traffic** check box.
4. Click **Add** and select IPv4 or IPv6 network. Note that you can add the IPv6 network only if the **Enable IPv6 communication** check box is selected. For more information, see IPv6 Support.

5. [For the IPv4 network] Specify a network address and a subnet mask using a CIDR notation and click **OK**.

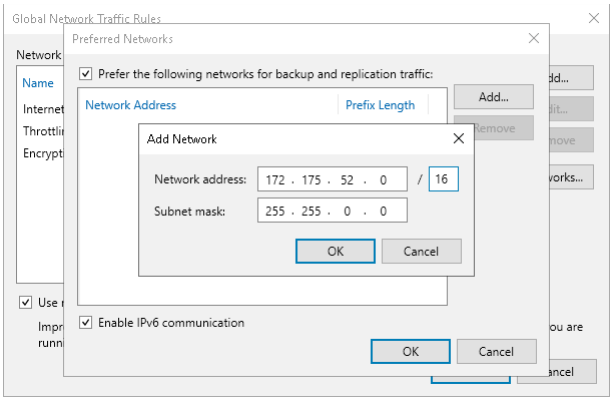


6. [For the IPv6 network] Specify a network address and prefix length using a CIDR notation and click **OK**.

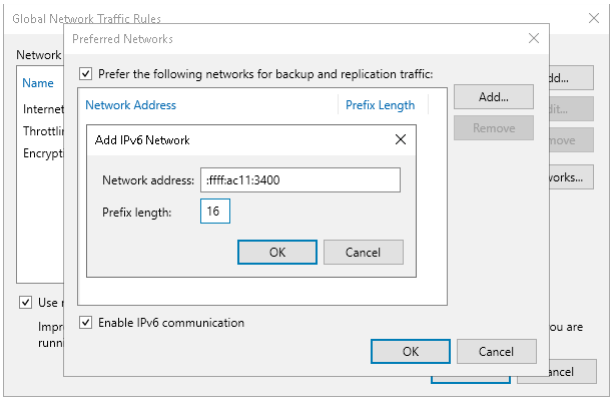


**NOTE**

[For multiple preferred network] The order of the networks in the preferred networks list does not specify the order of how Veeam Backup & Replication connects to these networks.

# IPv6 Support

Veeam Backup & Replication supports IPv6 communication for all backup infrastructure components. The following options are available:

- **IPv4/IPv6 dual stack**. This mode is enabled by default for new Veeam Backup & Replication installations.
- **IPv4 only**. For compatibility, this mode is used by default if you upgrade to Veeam Backup & Replication 12.

> **NOTE**
>
> Temporary IPv6 addresses are not supported for backup infrastructure components and backed-up machines. For more information about using temporary addresses, see this RFC section.

To manage IPv6 communication, perform the following steps:

1. From the main menu, select **Network Traffic Rules**.
2. Click **Networks**.
3. To enable the IPv4/IPv6 dual stack mode, select the **Enable IPv6 communication** check box. To enable the IPv4 only mode, clear the check box.

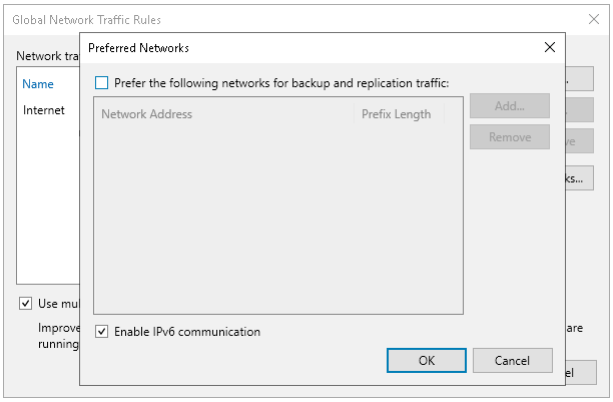


After you enable IPv6 communication, you can add backup infrastructure components with IPv6 addresses to the Veeam Backup & Replication console and configure IPv6 networks and traffic rules. You can use any text representation format of the IPv6 address, although RFC 5952 describes the canonical format as recommended. For more information, see this RFC section.

Note that if you want to switch back to the IPv4 only mode, you must remove backup infrastructure components added only with IPv6 addresses from the Veeam Backup & Replication console.

> **NOTE**
>
> If you use the IPv4/IPv6 dual stack mode and there are backup infrastructure components added to the Veeam Backup & Replication console using FQDN, Microsoft Windows itself determines which source address and the network protocol (IPv4 or IPv6) will be used for connection between Veeam Backup & Replication and these components. For more information, see this Microsoft article.
>
> This case is also applied to your backup infrastructure if you use Kerberos authentication. For more information, see Kerberos Authentication.

# Managing Logs

Veeam Backup & Replication provides detailed logging of performed activities, data protection and disaster recovery tasks.

Logs are stored on the backup server and all servers added to the backup infrastructure:

- On the backup server logs are stored in the following folders:
  - Installation and upgrade logs - the `%ProgramData%\Veeam\Setup\Temp` folder.
  - Logs for the Veeam Backup & Replication console and other user specific logs - the `%UserProfile%\AppData\Local\Veeam\Backup` folder.
  - Logs for jobs, components and services - the `%ProgramData%\Veeam\Backup` folder. Veeam Backup & Replication creates separate folders for jobs and also separate log files or folders for components and services.
  - Logs for PowerShell cmdlets - the `%UserProfile%\AppData\Local\Veeam\Backup\VeeamPowerShell.log` file.
- On Linux servers and ESXi hosts, logs are stored in the `/var/log/VeeamBackup` or `/tmp/VeeamBackup` directory.
- On Microsoft Windows servers, logs are stored in the `%ProgramData%\Veeam\Backup` folder.

> **TIP**
>
> You can change default log files settings, for example, logs location and size. For more information, see this Veeam KB article.

To collect log files from the backup server and servers managed by Veeam Backup & Replication, use the **Export Logs** wizard as described in section Exporting Logs.

# Exporting Logs

You can use log files to submit a support ticket. It is recommended that you send all log files when submitting a support ticket to ensure that overall and comprehensive information is provided to Veeam Support Team.

To aggregate all log files in the same location, use the **Export Logs** wizard.

> **NOTE**
>
> If you do not have access to the Veeam Backup & Replication console and you cannot use the built-in log export, export logs manually as described in this Veeam KB article.

# Step 1. Launch Export Logs Wizard

To launch the **Export Logs** wizard, in the main menu of Veeam Backup & Replication select **Help** > **Support Information**.

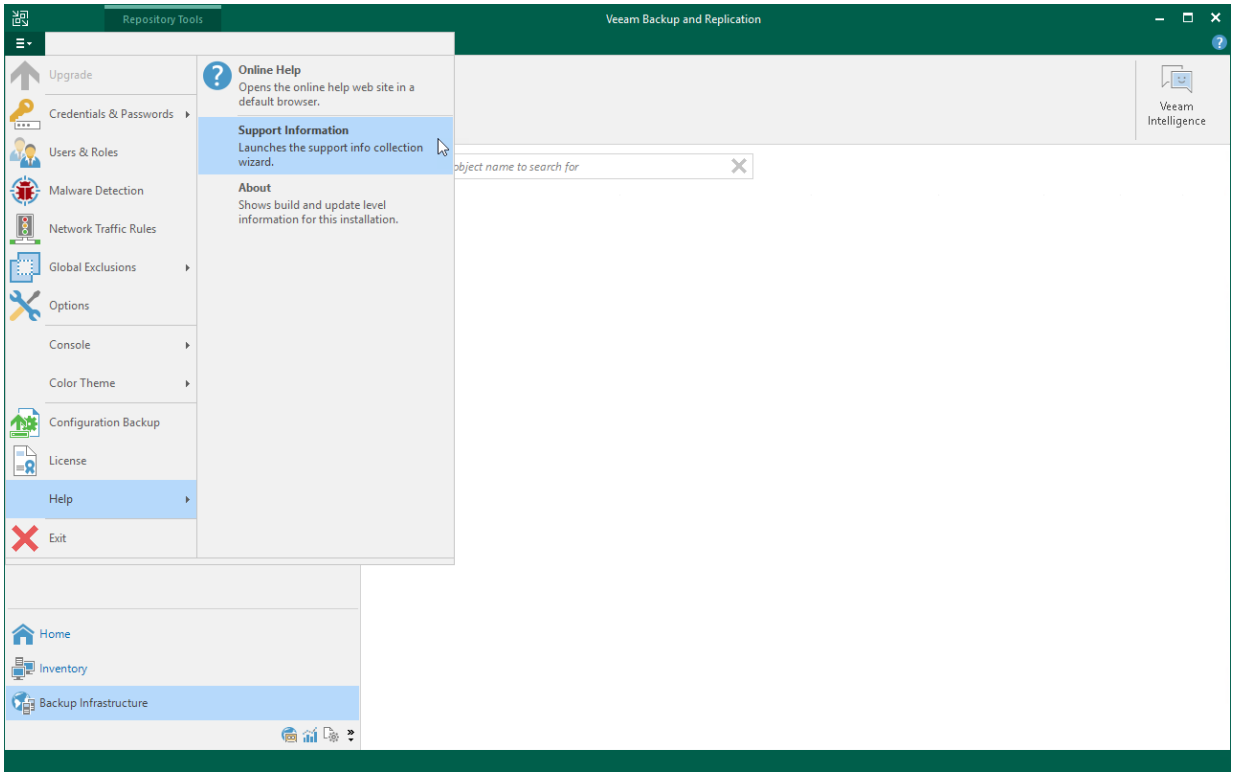

# Step 2. Select Virtual Infrastructure Scope

At the **Scope** step of the wizard, define the scope for logs export. You can export logs for the following objects:

- Specific jobs on the backup server.
- Specific objects (VMs, physical machines, CDP clusters, backups, scale-out backup repositories).
- Specific components in the backup infrastructure.

Select **Collect local PostgreSQL instance logs** to include PostgreSQL instance logs into the log archive that you may provide to Veeam Customer Support. The logs will contain information about all databases contained on the selected instance.

> **NOTE**
>
> If you export logs from the Veeam Backup & Replication console, the exported logs will be copied to the machine where the console is installed. The log archive will also contain logs from the console machine.

> **TIP**
>
> To select multiple jobs or objects in one bundle, do one of the following:
>
> - Hold [Ctrl] and click items to add to your selection.
> - Hold [Shift] and select a range of items between the currently selected item and the one you click.

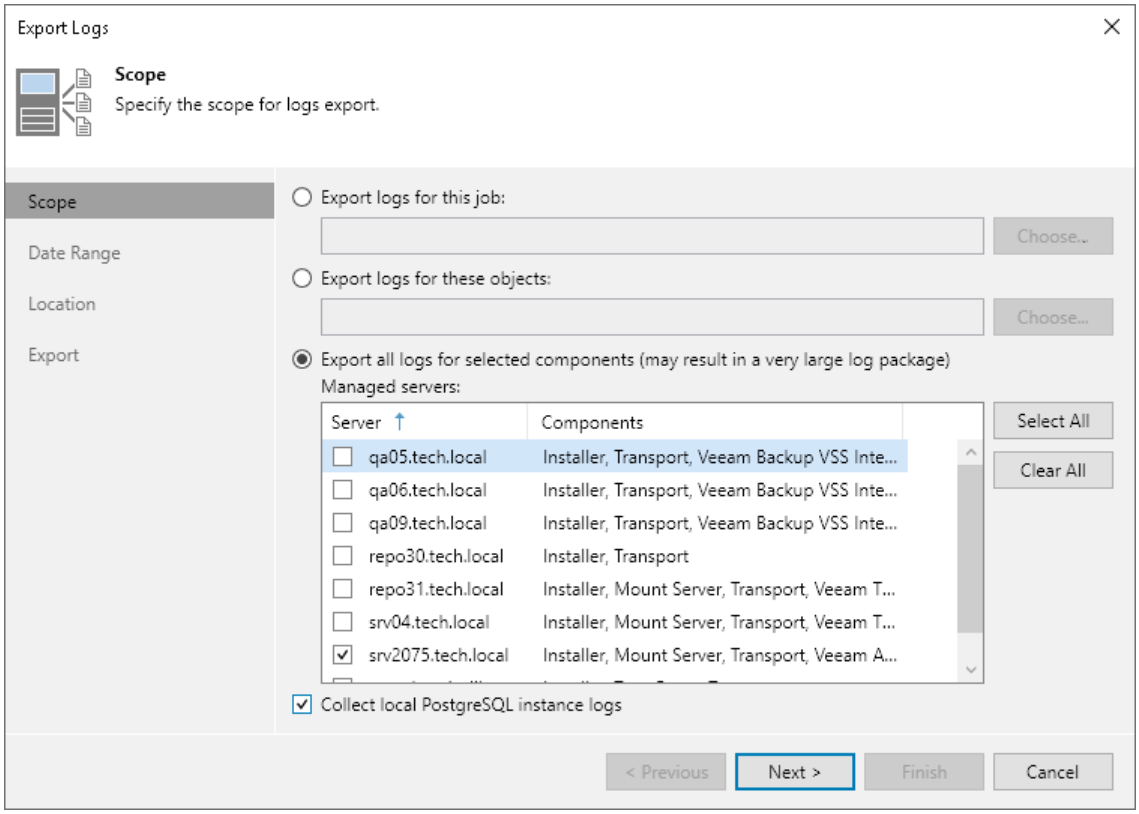

# Step 3. Specify Time Interval

At the **Date Range** step of the wizard, define the time interval for which logs must be collected. You can select one of the following options:

- Collect logs for the last N days
- Collect all available logs

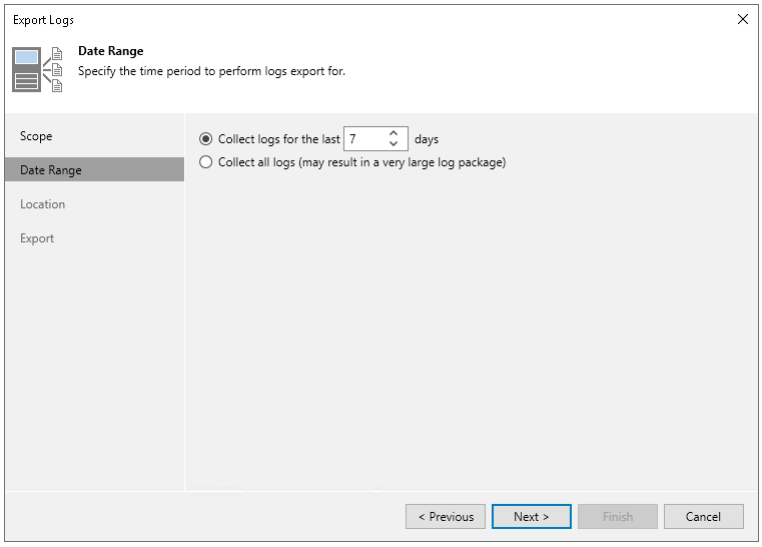

# Step 4. Specify Destination Folder

At the **Location** step of the wizard, specify the destination folder to which the logs will be exported.

In the **Path to folder** field, specify a path to an archive with log files that will be created. By default, the archive is placed to the `C:\temp\logs` folder on the backup server.

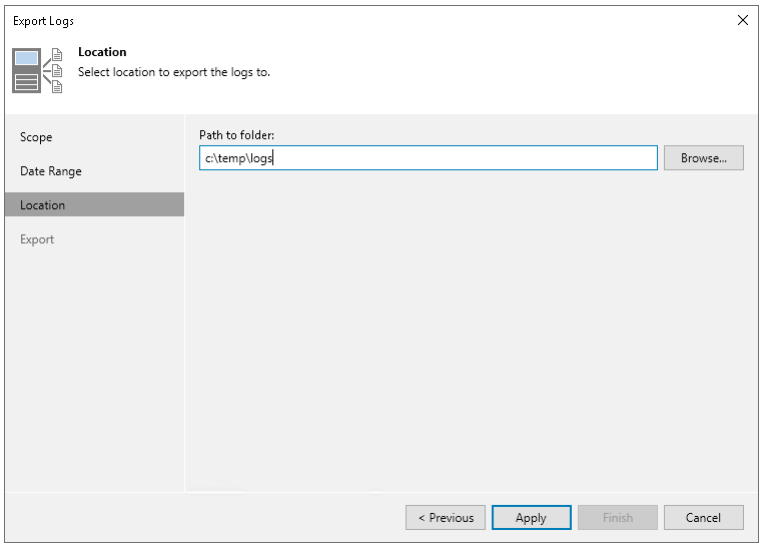

# Step 5. Review Results

At the **Export** step of the wizard, Veeam Backup & Replication will collect specified logs and create a log archive. Wait for the export process to complete, review the results and click the **Open folder** link to browse to exported log files and log package.

> **TIP**
>
> During the log export process, Veeam Backup & Replication locks the console, so you cannot close the **Export Logs** wizard. If you do not want to wait until the log export process completes, you can run another session of the Veeam Backup & Replication concurrently and continue work with it.

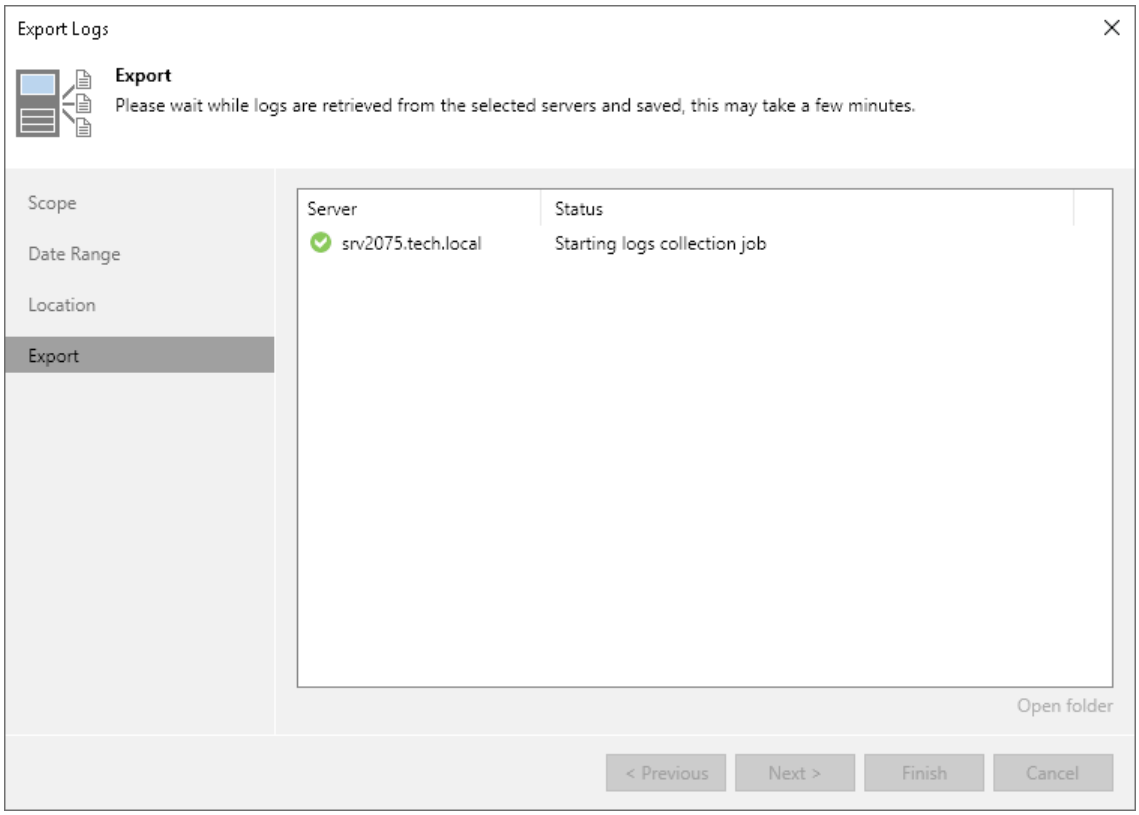

# Configuring Global VM Exclusions

Global VM exclusion allows you to stop processing VMs even if they are included in jobs. When excluding VMs globally, you do not need to change job settings. Global exclusion applies to all types of jobs that process VMs except backup copy jobs and SureBackup jobs. Note that if Veeam Backup & Replication has started to process a VM at the moment when you exclude the VM from processing, Veeam Backup & Replication finishes processing and only then excludes the VM.

At any moment, you can enable processing of VMs so that jobs will continue processing the VMs.

> **IMPORTANT**
>
> Note that you can exclude only VMs from processing. It is not possible for VM containers (folder, resource pool, VirtualApp, datastore, tag and so on).

## Excluding VMs from Processing

To exclude VMs globally, do the following:

1. From the main menu, select **VM Exclusions**.
2. In the **VM Exclusions** window, click **Add** and select the required platform.
3. In the **Add Objects** window, select VMs that you want to exclude from processing. Click **OK**.
4. If you want to add a note why a VM is excluded from processing, select a VM from the list and click **Note**. In the **Edit Note** window, enter the note text. Click **OK**.

Alternatively, you can open the Inventory view and switch to the required view. In the working area, select the VMs that you want to exclude, right-click one of them and select Disable processing. In the Inventory view, VMs excluded from processing will be displayed with the icon.

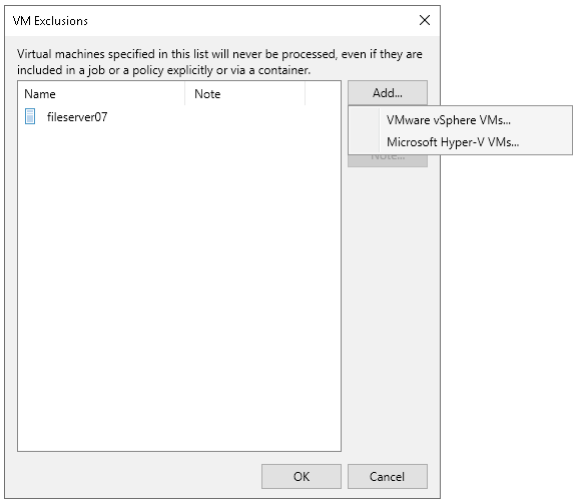


## Resuming VM Processing

To resume processing for VMs, do the following:

1. From the main menu, select **VM Exclusions**.
2. In the **VM Exclusions** window, select VMs for which you want to resume processing.

3. Click **Remove**.

Alternatively, you can open the **Inventory** view and switch to the required view. In the working area, select the VMs that were excluded from processing, right-click one of them and select **Disable processing**.

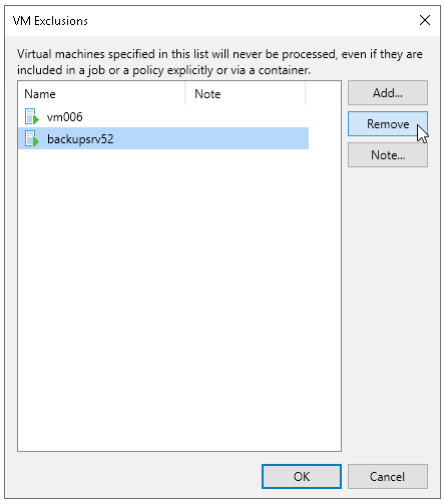

# Configuring Analytics View

> **NOTE**
>
> Consider the following:
>
> - The **Analytics** view is available if you have the Microsoft Edge WebView2 Runtime component installed. The component is not installed for Microsoft Windows Server 2012 and 2012 R2 due to the version incompatibility, so the **Analytics** view is not available for the backup server running these Microsoft Windows versions.
> - If Veeam ONE lacks the license required for the integration with Veeam Backup & Replication, you will see a placeholder with the license information in the **Analytics** view. For more information, see the Licensing Veeam ONE section in the Veeam ONE Deployment Guide.
> - You can use the **Analytics** view only to view statistical data. You cannot make any changes to them.

The **Analytics** view is designed to provide easy access to dashboards configured in Veeam ONE for monitoring the backup infrastructure and data protection operations in the virtual environment. The **Analytics** view displays the following Veeam ONE dashboards:

- Veeam Threat Center
- Veeam Backup & Replication Overview
- Backup Heatmap
- Jobs Calendar

## Enabling Analytics Feature

To enable the Analytics feature, do the following:

1. Add your Veeam Backup & Replication server to the Veeam ONE infrastructure, as described in the Veeam ONE Deployment Guide. Do not forget to select the **Provide access to embedded dashboards** check box.
2. Wait till Veeam ONE collects the required analytical data. After that, the Veeam ONE dashboards can be opened in Veeam Backup & Replication.

3. When opening any dashboard in the **Analytics** view for the first time, Veeam Backup & Replication will prompt you to install the certificate issued by the server where Veeam ONE is installed.

   In the **Certificate Security Alert** dialog box, click **Continue** to install it. Alternatively, you can click View to view certificate details; after that, in the **Certificate** window, click **Install Certificate**, then **OK**.

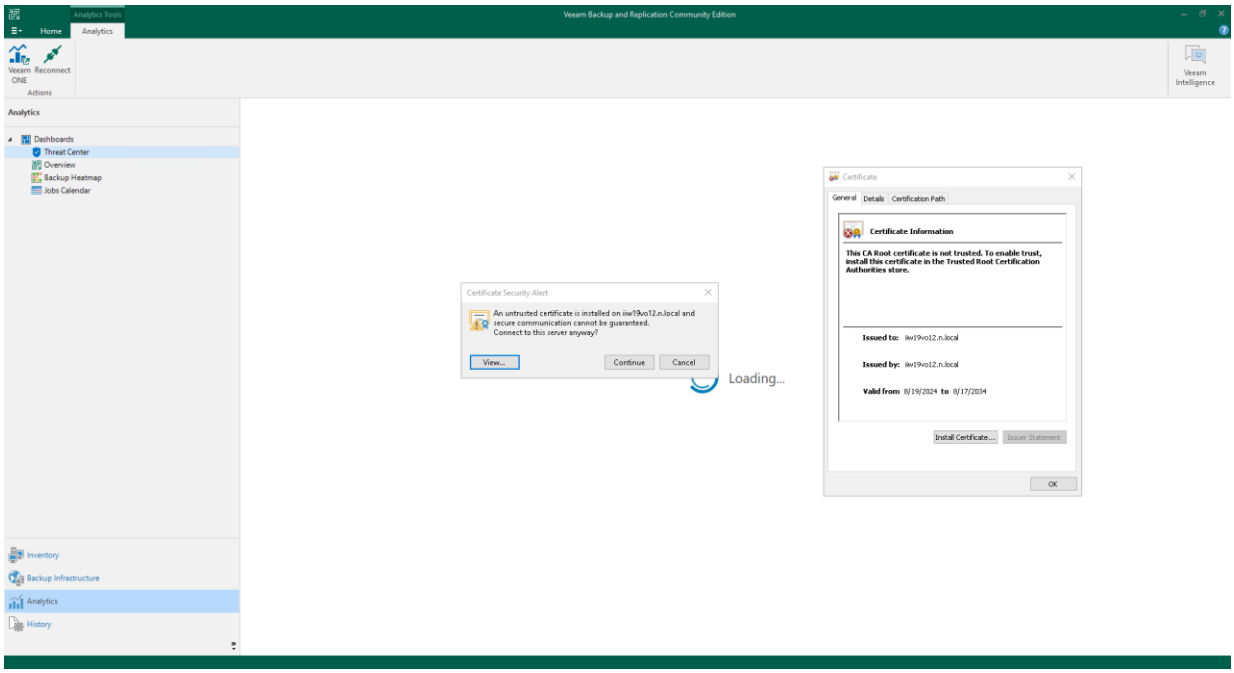


Now, you can use the **Analytics** view to work view Veeam ONE dashboards for the current backup server.

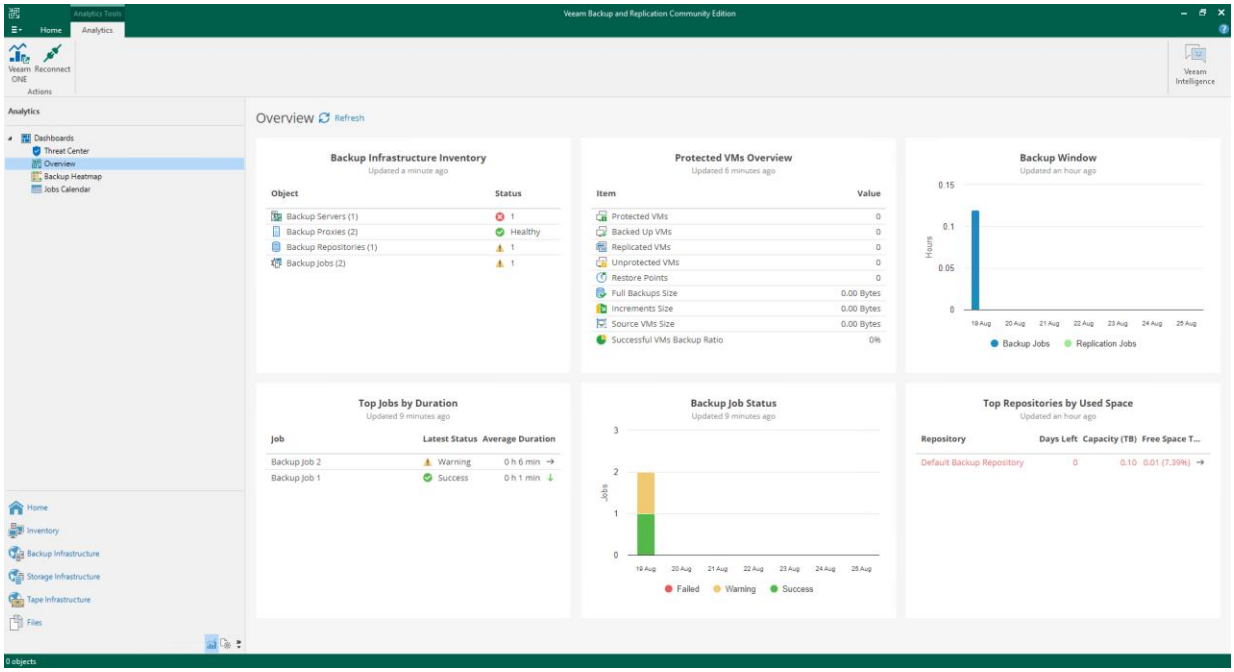


If you move a Veeam Backup & Replication server to a different Veeam ONE instance and delete it from the current instance, the Veeam Backup & Replication servers become unregistered for all integrations. To solve this, clear the **Provide access to embedded dashboards for added backup servers** check box in **Connection Settings** and select it in the new Veeam ONE instance. For details, see the Changing Server Connection Settings section in the Veeam ONE Deployment Guide.

# Switching to Veeam ONE Web Client

You can quickly switch to the Veeam Threat Center dashboard of the Veeam ONE web client. To do that, click Veeam ONE in the ribbon. Alternatively, you can open the **Dashboards** node in the **Analytics** view and click **Open Veeam ONE** in the working area.

# Reconnecting to Veeam ONE Instance

In case of any problems with the connection to the Veeam ONE instance, you can reconnect to it. To do that, click **Reconnect** in the ribbon. Alternatively, you can click another dashboard under the **Dashboards** node on the inventory pane.

# Backup Infrastructure Components

Veeam Backup & Replication is a modular solution that lets you build a scalable backup infrastructure for environments of different sizes and configuration. The installation package of Veeam Backup & Replication includes a set of components that you can use to configure the backup infrastructure. Some components are mandatory and provide core functionality; some components are optional and can be installed to provide additional functionality for your business and deployment needs.

You can co-install Veeam Backup & Replication components on the same machine, physical or virtual, or you can set them up separately for a more scalable approach. Some components can be deployed with the help of the setup file. Other components can be deployed using the Veeam Backup & Replication console.

# Backup Server

The backup server is a Windows-based physical or virtual machine on which Veeam Backup & Replication is installed. It is the core component in the backup infrastructure that fills the role of the "configuration and control center". The backup server performs all types of administrative activities:

- Coordinates backup, replication, recovery verification and restore tasks.
- Controls job scheduling and resource allocation.
- Is used to set up and manage backup infrastructure components as well as specify global settings for the backup infrastructure.

In addition to its primary functions, a newly deployed backup server also performs the roles of the default VMware backup proxy and the backup repository (it manages data handling and data storing tasks).

# Backup & Replication Console

The Veeam Backup & Replication console is a client-side component that provides access to the backup server. The console allows you to log in to Veeam Backup & Replication and perform all kinds of data protection and disaster recovery operations on the backup server.

To connect to the backup server, the Veeam Backup & Replication console uses Simple and Protected GSS-API Negotiation Mechanism (SPNEGO), the Microsoft Windows built-in authentication mechanism. SPNEGO supports both NTLM and Kerberos authentication.

> **NOTE**
>
> Using NTLM increases the attack surface of your backup infrastructure. To build a more secure environment, disable NTLM and leave Kerberos the only domain authentication protocol. For more information, see Kerberos Authentication.

The console does not have a direct access to the backup infrastructure components and configuration database. Such data as user credentials, passwords, roles and permissions are stored on the backup server side. User credentials and passwords are stored in the configuration database encrypted with Data Protection API (DPAPI) mechanisms. To access the data, the console connects to the backup server and queries this information periodically during the work session.

To make users work as uninterrupted as possible, the remote console maintains the session for 5 minutes if the connection is lost. If the connection is re-established within this period, you can continue working without re-logging to the console.

## Backup & Replication Console Deployment

The console is installed locally on the backup server by default. You can also use it in a standalone mode — install the console on any machine and access Veeam Backup & Replication remotely over the network.

You can install as many remote consoles as you need so that multiple users can access Veeam Backup & Replication simultaneously. Veeam Backup & Replication prevents concurrent modifications on the backup server. If several users are working with Veeam Backup & Replication at the same time, the user who saves the changes first has the priority. Other users will be prompted to reload the wizard or window to get the most recent information about the changes in the configuration database.

If you have multiple backup servers in the infrastructure, you can connect to any of them from the same console. For convenience, you can save several shortcuts for these connections.

> **IMPORTANT**
>
> You cannot use the same console to connect to backup servers with different versions of Veeam Backup & Replication. Note this if you have more than one backup server in your backup environment, and these backup servers run different versions of Veeam Backup & Replication. For example, if one of your backup servers run version 11, and another backup server runs version 12.1.2, you will need to use 2 separate consoles for connecting to these servers.

The console supports automatic update. Every time you connect to the backup server locally or remotely, the console checks for updates. If the backup server has updates installed, the console will be updated automatically.

Consider the following:

- Upgrade to another Veeam Backup & Replication major product version is supported starting from Veeam Backup & Replication version 11a (build 11.0.1.1261). For example, if your console has this version, you can upgrade it automatically upon connecting to the Veeam backup server version 12 (build 12.0.0.1420). Automatic upgrade is not supported for Preview, Beta or RTM versions of Veeam Backup & Replication.
- Downgrade of the console is not supported. If the console is of a higher version than the backup server (for example, you have upgraded the console manually), the connection to the server will fail.

If other Veeam Backup & Replication components, such as Veeam Cloud Connect Portal or Veeam Backup Enterprise Manager, are installed on the machine where the console runs, these components will also be upgraded.

# Backup & Replication Console Components

When you install a remote console on a machine, Veeam Backup & Replication installs the following components:

- Veeam Backup PowerShell Module
- Veeam Explorer for Microsoft Active Directory
- Veeam Explorer for Microsoft Exchange
- Veeam Explorer for Microsoft OneDrive for Business
- Veeam Explorer for Microsoft SharePoint
- Veeam Explorer for Microsoft SQL Server
- Veeam Explorer for Microsoft Teams
- Veeam Explorer for Oracle
- Veeam Explorer for PostgreSQL
- Veeam Explorer for SAP HANA
- Veeam Data Mover Service
- Veeam Explorers Recovery Service
- Veeam Installer Service
- Veeam Mount Service

# Backup & Replication Console User Access Rights

To log in to Veeam Backup & Replication using the console, the user must be added to the Local Users group on the backup server or a group of domain users who have access to the backup server. The user can perform the scope of operations permitted by his or her role in Veeam Backup & Replication. For more information, see Users and Roles.

# Requirements for Backup & Replication Console

A machine on which you install the Veeam Backup & Replication console must meet the following requirements:

- The machine must meet the system requirements. For more information, see System Requirements.
- The remote console can be installed on a Microsoft Windows machine (physical or virtual).
- If you install the console remotely, you can deploy it behind NAT. However, the backup server must be outside NAT. The opposite type of deployment is not supported: if the backup server is deployed behind NAT and the remote console is deployed outside NAT, you will not be able to connect to the backup server.

# Limitations for Backup & Replication Console

The Veeam Backup & Replication console has the following limitations:

- You cannot perform restore from the configuration backup using the remote console.
- The machines on which the remote console is installed are not added to the list of managed servers automatically. For this reason, you cannot perform some operations, for example, import backup files that reside on the remote console machine or assign roles of backup infrastructure components to this machine. To perform these operations, you must add the remote console machine as a managed server to Veeam Backup & Replication. For more information, see Managing Servers.

# Veeam Backup & Replication Configuration Database

Veeam Backup & Replication Configuration Database stores data about the backup infrastructure, jobs, sessions and other configuration data. The database instance can be located on a Microsoft SQL Server or PostgreSQL installed either locally (on the same machine where the backup server is running) or remotely.

Veeam Backup & Replication maintains the configuration database. Veeam Backup & Replication runs the DatabaseMaintenance system job once a week and when the Veeam Backup Service is restarted. The job updates the database internal statistics, defragments indexes and clears unused data. For details, see the `Job.DatabaseMaintenance.log` file in the `%ProgramData%\Veeam\Backup` folder.

# Managing Configuration Database

You can back up and restore the configuration database that Veeam Backup & Replication uses. If the backup server fails for some reason, you can re-install the backup server and quickly restore its configuration from the configuration backup. You can also use configuration backups to apply the configuration of one backup server to another backup server in the backup infrastructure. During configuration backup, Veeam Backup & Replication exports data from the configuration database and saves it to a backup file in the backup repository.

> **NOTE**
>
> If you use cloud plug-ins to protect VMs in Google Cloud, AWS and other environments, you can also back up configurations of cloud backup appliances. To back up the configurations, you must enable encryption as described in section Creating Encrypted Configuration Backups.
>
> For more information on appliance configuration backup and restore, see the following guides:
>
> - Veeam Backup for Google Cloud, the Performing Configuration Backup and Restore section.
> - Veeam Backup for AWS, the Performing Configuration Backup and Restore section.
> - Veeam Backup for Microsoft Azure, the Performing Configuration Backup and Restore section.

It is recommended that you regularly perform configuration backup for every backup server in the backup infrastructure. Periodic configuration backups reduce the risk of data loss and minimize the administrative overhead if any problem with backup servers occurs.

> **NOTE**
>
> If you have a hardened repository in your backup infrastructure, you must enable data encryption for configuration backup. Otherwise, you cannot perform configuration backups. For more information, see Creating Encrypted Configuration Backups.

Do not back up the backup server configuration using backup or replication jobs in Veeam Backup & Replication. For backup and replication, Veeam Backup & Replication uses VM snapshots. During snapshot creation and commit, the VM freezes for some time, which can potentially lead to the following consequences:

- Disconnection from the configuration database. For more information, see this Veeam KB article.
- Disconnection from remote Veeam Backup & Replication agents.
- Disconnection from network storage (for example, storage presented through iSCSI) and so on.

For this reason, you must always use the configuration backup functionality to back up and restore configuration of the backup server.

## Creating Configuration Backups

By default, Veeam Backup & Replication creates a configuration backup daily. You can change the schedule or create a configuration backup manually. You can choose the backup repository in which the configuration backup must be stored, specify the necessary retention settings. Veeam Backup & Replication allows you to create immutable configuration backups.

For more information on managing configuration backup, see the following topics:

- Scheduling Configuration Backups
- Configuring Notification Settings for Configuration Backups
- Running Configuration Backups Manually

- Creating Encrypted Configuration Backups
- Creating Immutable Configuration Backups

# Configuration Backup Files

When you perform configuration backup, Veeam Backup & Replication retrieves data for the backup server from the configuration database, writes this data into a set of XML files and archives these XML files to a backup file of the BCO format.

Veeam Backup & Replication exports information about the following objects:

- **Backup infrastructure components and objects**: hosts, servers, backup proxies, repositories, WAN accelerators, virtual lab configurations, global settings configured on the backup server and so on.
- **Backups**: backups and backup copies, replicas, CDP policies created on the backup server.
- **Sessions**: job sessions performed on the backup server.
- **Tapes**: tape libraries connected to the backup server.

**NOTE**

Consider the following:

- If a configuration backup job has a remote repository as a target, a temporary BCO file is created in the `%TEMP%` folder during the job (by default, this is `C:\Windows\Temp`, but it can be modified in the environment variables). After the BCO file is copied to the remote repository, it is removed from `%TEMP%`. If the repository is local, the BCO file is saved directly to it without using `%TEMP%`.

  Therefore, there must be enough free space on the disk for this temporary BCO file if the remote repository is used for storing the configuration backup. The size of the BCO file depends on the size of the database and its content (for example, sessions are compressed better than tapes in BCO). The compression ratio can be checked in the configuration backup session details.
- If you use custom configuration registry values, note that configuration backup will not apply to them. You may want to back them up manually.
- The configuration backup job creates a snapshot of the configuration database and retrieves data required for successful restore from it. If the database size is large, the job may produce significant load on the Microsoft SQL Server. Make sure that you schedule the configuration backup job for a period of low operation intensity on the backup server.

# Backup Repository Target

The resulting configuration backup file is stored in the `\VeeamConfigBackup\%BackupServer%` folder on the default backup repository. However, for security sake, it is recommended that you do not store configuration backups on the default backup repository or in any other folder on the backup server. In this case, if the backup server fails, its configuration data will remain, and you will be able to recover the failed backup server. You can store configuration backups on repositories of different types, including Veeam Cloud Connect repositories.

**IMPORTANT**

You cannot store configuration backups in external repositories and scale-out backup repositories.

When you configure a new backup repository, Veeam Backup & Replication offers you to change the configuration backup file location from the default backup repository to the new backup repository. Click **Yes**, and Veeam Backup & Replication will automatically change the backup target in the configuration backup job settings and will use this target in future.

Configuration backups that were created before the target change will remain in the default backup repository. You can manually copy them to the new backup repository to have all restore points of the configuration backup in one place.

## Scheduling Configuration Backups

You can instruct Veeam Backup & Replication to perform configuration backup automatically by schedule.

> **IMPORTANT**
>
> If you plan to migrate configuration data to the database used by another backup server, stop all running jobs and disable scheduled jobs before creating the configuration backup. In the opposite case, job sessions may be failing after configuration restore. For more information, see Migrating Veeam Backup & Replication to Another Backup Server.

To schedule a configuration backup:

1. From the main menu, select **Configuration Backup**.
2. Make sure that the **Enable configuration backup to the following repository** check box is selected.
3. From the **Backup repository** list, choose a backup repository on which the configuration backup must be stored.
4. In the **Restore points to keep** field, specify the number of restore points that you want to maintain in the backup repository.
5. Click **Schedule** next to the **Perform backup on** field and specify the time schedule according to which the configuration backup must be created.

6. To create an encrypted backup, select the **Enable backup file encryption** check box. From the **Password** field, select a password you want to use for encryption. If you have not created a password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Creating Encrypted Configuration Backups.

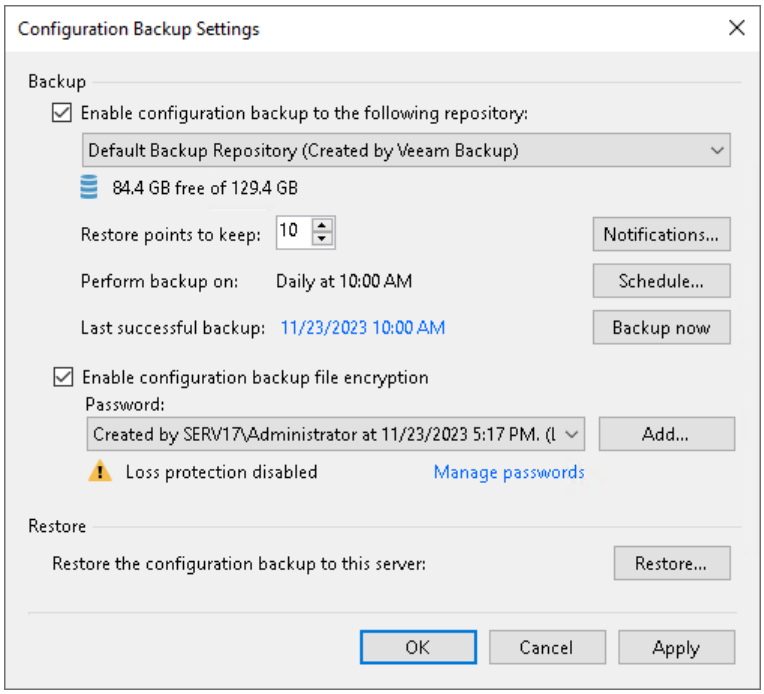


## Configuring Notification Settings for Configuration Backups

You can configure notifications for the configuration backup:

1. From the main menu, select **Configuration Backup**.
2. Click **Notifications**.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully.

   SNMP traps will be sent if you specify global SNMP settings in Veeam Backup & Replication and configure software on recipient's machine to receive SNMP traps. For more information, see Specifying SNMP Settings.
4. Select the **Send email notifications to the following recipients** check box if you want to receive notifications about the job completion status by email. In the field under the check box, specify the recipient email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication. For more information, see Configuring Global Email Notification Settings.
5. You can choose to use global notification settings or specify custom notification settings.
   - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see Configuring Global Email Notification Settings.
   - To configure a custom notification for the job, select **Use custom notification settings specified below** check box. You can specify the following notification settings:
     i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%JobResult%*, *%JobName%*, *%Time%* (completion time).

ii. Select the **Notify on success**, **Notify on warning** and **Notify on error** check boxes to receive email notification if the job completes successfully, fails or completes with a warning.

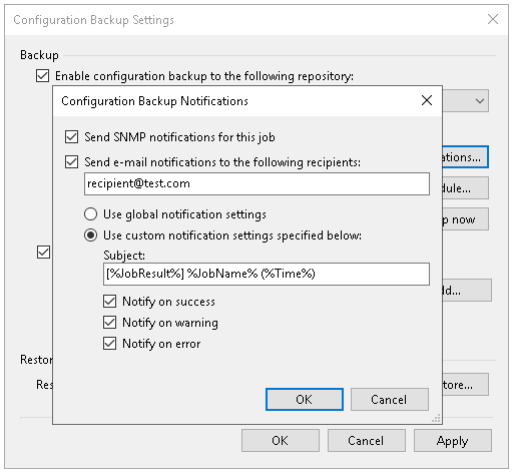


## Running Configuration Backups Manually

You can create a configuration backup manually when you need it, for example, if you want to capture a state of the configuration database at a specific point in time.

> **IMPORTANT**
>
> If you plan to migrate configuration data to the database used by another backup server, stop all running jobs and disable scheduled jobs before creating the configuration backup. In the opposite case, job sessions may be failing after configuration restore. For more information, see Migrating Veeam Backup & Replication to Another Backup Server.

To create a configuration backup manually:

1. From the main menu, select **Configuration Backup**.
2. Make sure that the **Enable configuration backup to the following repository** check box is selected.
3. From the **Backup repository** list, choose a backup repository on which the configuration backup must be stored.
4. In the **Restore points to keep** field, specify the number of restore points that you want to maintain in the backup repository.
5. To create an encrypted backup, select the **Enable backup file encryption** check box. From the **Password** field, select a password you want to use for encryption. If you have not created a password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Creating Encrypted Configuration Backups.
6. Click **Backup now**.

Veeam Backup & Replication will back up the configuration database and store a new restore point to the selected backup repository. The resulting configuration backup file (.BCO) is stored in the `\VeeamConfigBackup\%BackupServer%` folder on the repository you have selected. The file name contains the date when the configuration backup was performed.

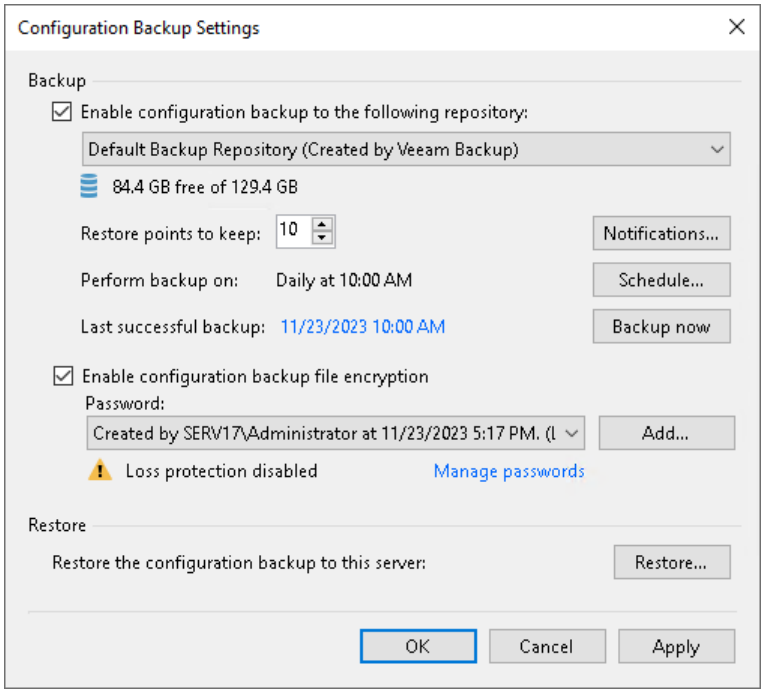


## Creating Encrypted Configuration Backups

Veeam Backup & Replication requires that you encrypt the configuration backup if you have created at least one password in the Password Manager on the backup server. We strongly recommend that you encrypt the configuration backup, even if you have not created any passwords.

> **IMPORTANT**
>
> Configuration backup jobs do not support data encryption with KMS keys.

Veeam Backup & Replication encrypts credentials, encryption keys, and certificates stored in the configuration database using Data Protection API (DPAPI) mechanisms. For more information, see Microsoft Docs. Data Protection API mechanisms in the configuration backup are replaced with password encryption.

If the configuration backup is unencrypted, data from it can be restored to any backup server. Credentials, encryption keys, and certificates will not be stored in the configuration backup. The content of encrypted files and access to the backup infrastructure will become unavailable after the configuration restore. Additionally, sensitive information such as machine or server IP addresses may be exposed to unintended audiences.

If the Password Manager contains at least one password, and you do not enable encryption for the configuration backup, Veeam Backup & Replication disables configuration backup. To enable the configuration backup, you must enable encryption in the configuration backup job settings.

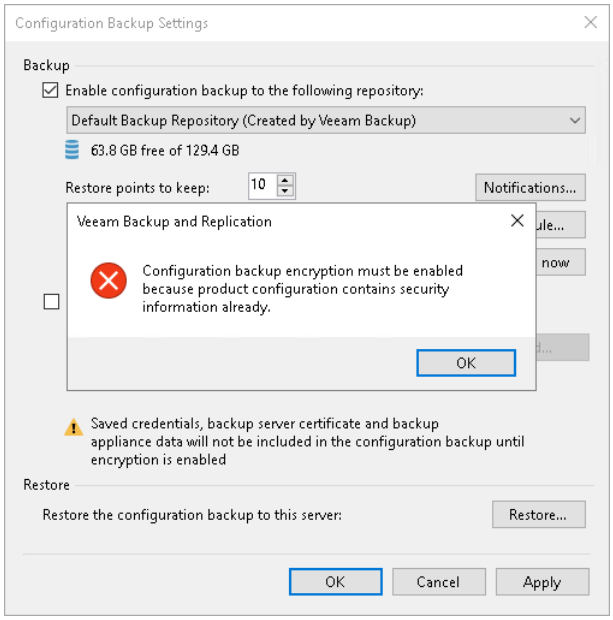


After you enable the encryption option, Veeam Backup & Replication will create encrypted configuration backups. Beside encryption keys, the created backups capture credential records specified in the Credentials Manager. When you restore data from such backup, you will not have to enter passwords for credentials records again (unless the passwords for credentials records have changed by the time of restore).

To encrypt the configuration backup:

1. From the main menu, select **Configuration Backup**.
2. Select the **Enable backup file encryption** check box.
3. From the **Password** field, select a password you want to use for encryption. If you have not created the password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Password Manager.

   For more information about data encryption, see How Data Encryption Works.

   If you do not use password loss protection, Veeam Backup & Replication will display a warning about it. For more information, see Password Loss Protection.

   > **NOTE**
   >
   > If you enable encryption and use cloud plug-ins to protect VMs in Google Cloud, AWS and other environments, Veeam Backup & Replication will also create backups for cloud backup appliances. For more information on appliance configuration backup and restore, see the following guides:
   >
   > - Veeam Backup for Google Cloud, the Performing Configuration Backup and Restore section.
   > - Veeam Backup for AWS, the Performing Configuration Backup and Restore section.
   > - Veeam Backup for Microsoft Azure, the Performing Configuration Backup and Restore section.

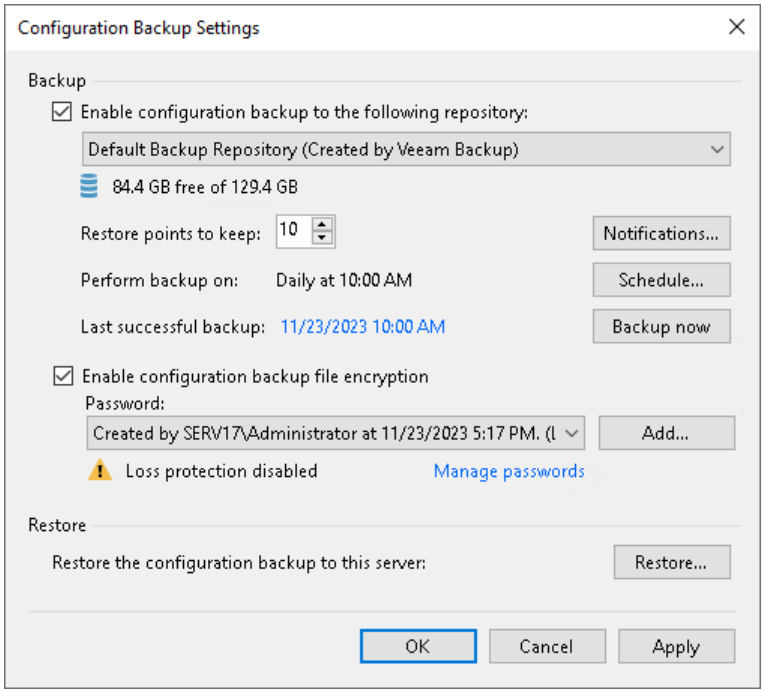


## Creating Immutable Configuration Backups

Veeam Backup & Replication allows you to keep configuration backups on object storage repositories that support immutability. Immutability makes data temporarily immutable and prohibits deletion of configuration backups from object storage repositories. Therefore, it protects your data against loss as a result of attacks, malware activity or any other injurious actions.

> **IMPORTANT**
>
> The version of Veeam Backup & Replication installed on a backup server for a Service Provider and a tenant must be 12.1 (build 12.1.0.2131). Otherwise, the configuration backup stored on an object storage repository will not be immutable.

# Retention Policy for Immutable Configuration Backups

By default, object storage repositories keeps the number of restore points that you define in the configuration backup settings. In case, a total number of restore point exceeds the retention policy settings, Veeam Backup & Replication deletes the oldest restore point. If you keep configuration backups on immutable object storage repositories, Veeam Backup & Replication will not delete the older restore points. Instead, it will mark them as immutable and will keep them on object storage repositories until immutability period passes.

# Creating Immutable Configuration Backups

To make configurations backups immutable, perform the following steps:

1. Configure an object storage repository where you will keep configuration backups. For more information, see Immutability for Object Storage Repositories.
2. From the main menu, select **Configuration Backup**.
3. Make sure that the **Enable configuration backup to the following repository** check box is selected.

4. From the **Backup repository** list, choose an object storage repository on which the configuration backup must be stored.
5. In the **Restore points to keep** field, specify the number of restore points that you want to maintain in the object storage repository.
6. To create an encrypted backup, select the **Enable backup file encryption** check box. From the **Password** field, select a password you want to use for encryption. If you have not created a password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Creating Encrypted Configuration Backups.
7. Click **Backup now**.

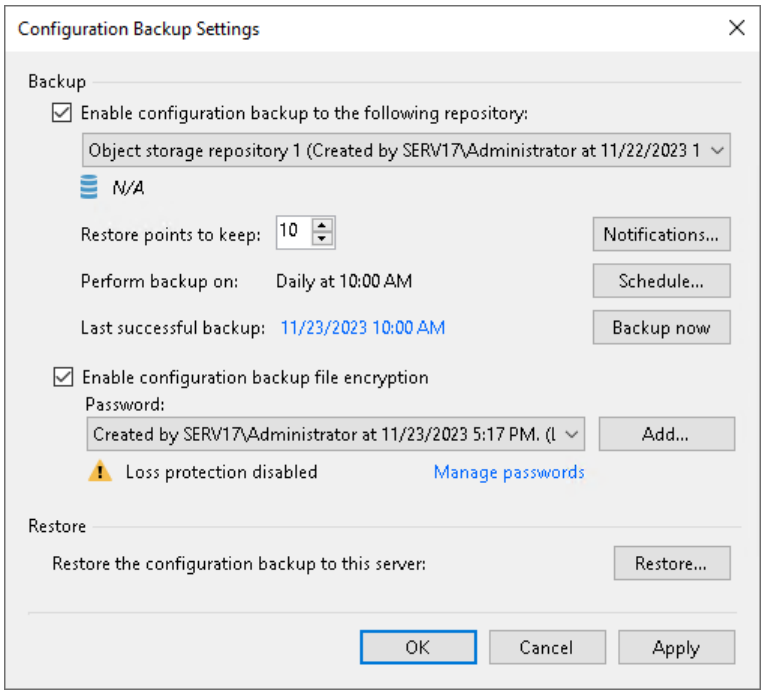


# Restoring Configuration Database

Restore of the configuration database is helpful in the following situations:

- The configuration database got corrupted and you want to recover data from the configuration backup.
- You want to deploy the configuration database on a new Microsoft SQL Server or PostgreSQL and restore data from the configuration backup to it.
- You want to roll back the configuration database to a specific point in time.
- You want to restore data to a new configuration database on the same database instance, for example, for testing purposes.

You can restore a configuration backup on the same backup server where the backup was created or on another backup server.

> **NOTE**
>
> If you use cloud plug-ins to protect VMs in Google Cloud, AWS and other environments and want to restore cloud backup server configurations, see the following guides:
>
> - Google Cloud, the Configuration Restore section.
> - Veeam Backup for AWS, the Configuration Restore section.
> - Veeam Backup for Microsoft Azure, the Configuration Restore section.

Before you start the restore process, check prerequisites. Then use the **Veeam Backup & Replication Configuration Restore** wizard to restore the configuration database.

## Before You Begin

Before you start the restore process, check the following prerequisites:

- Make sure the user, who initiates the backup configuration restore process, has the **Debug programs** policy applied. Otherwise, Veeam Backup & Replication returns the *Access is Denied* error. To learn how to apply the policy, see this Veeam KB article.
- Multi-factor authentication (MFA) for the user account must be disabled before launching configuration database restore wizard. Alternatively, you can manually stop the backup service. To learn how to disable MFA, see Multi-Factor Authentication.
- Stop all jobs that are currently running. During restore of configuration, Veeam Backup & Replication temporary stops the Veeam Backup Service and jobs.
- Save registry values that you changed or created on the backup server. After restore, you will need to recreate or change the keys manually because the configuration database does not store them.
- Check the version of the backup server. On the backup server running Veeam Backup & Replication 12, you can restore configuration backups created with Veeam Backup & Replication 11a or later.
- Make sure that the certificate chain restored from a configuration backup will successfully pass validation on the target backup server. This precaution is required if the following conditions are met:
  a. You want to restore configuration database of a backup server used in the Veeam Agent management scenario.
  b. The backup server whose configuration database you want to restore uses a custom certificate issued by a Certificate Authority instead of the default self-signed certificate to ensure a secure connection in the Veeam Agent management infrastructure.
- If you plan to restore configuration data to the database on another Microsoft SQL Server or PostgreSQL Server, make sure the account for using Veeam Backup & Replication has sufficient permissions. For more information, see Permissions.
- Veeam Backup & Replication supports configuration database migration between different database engines only within the same Veeam Backup & Replication version.
- After you run configuration restore for capacity tier, Veeam Backup & Replication will put the capacity tier extents into the Sealed mode. Rescan the backup infrastructure and then remove the extents from the Sealed mode.
- If you use a Veeam Backup & Replication server as a Veeam repository, after restore to a new host this repository will be located in the same file path as it was before the migration.
- If you use Veeam plug-ins (Veeam Backup for OLVM and RHV, Veeam Backup for Nutanix AHV and so on) and want to restore the configuration database to a new machine, make sure you have the plug-ins installed before the restoration process.

**IMPORTANT**

You can start configuration restore only from the Veeam Backup & Replication console installed locally on the backup server. You cannot start configuration restore from the console installed on a remote machine.

## Step 1. Launch Configuration Database Restore Wizard

To launch the **Veeam Backup and Replication Configuration Restore** wizard, do either one of the following:

- From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**. In the **Restore** section, click **Restore**.
- In the **Start** menu of the backup server, click **Configuration Restore**.
- Use the `Veeam.Backup.Configuration.Restore.exe` file located in the installation folder on the backup server. By default, the path to the folder is the following: `%PROGRAMFILES%\Veeam\Backup and Replication\Backup`.
- [If the configuration backup is stored on the backup server] In Microsoft Windows Explorer, open the folder where configuration backups are stored (by default, `Backup\VeeamConfigBackup\<BackupServerName>` on the volume with most disk space on the backup server) and double-click the necessary configuration backup file.

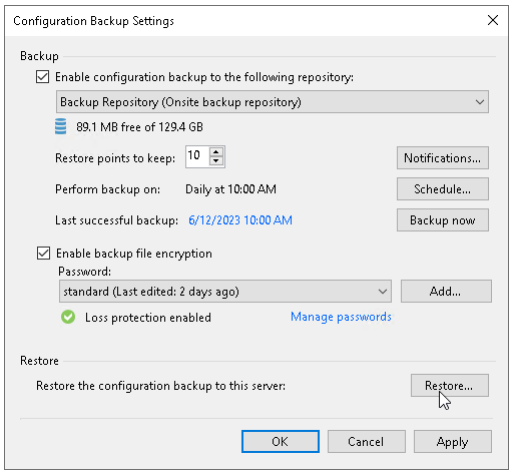

# Step 2. Select Restore Mode

At the **Restore Mode** step of the wizard, choose a restore mode that you want to use.

- Select **Restore** if you want to restore data from the configuration backup to the database used by the initial backup server.

  In the Restore mode, Veeam Backup & Replication retrieves configuration data from the backup and stores it to the target database. After that, Veeam Backup & Replication performs additional rescan of VM replicas and backup repositories connected to the backup server, and runs tape library cataloging process. Rescan helps synchronize potential changes between the backup infrastructure and restored database that took place from the moment when the configuration backup was created till the present time. As a result, the target configuration database will contain information about restore points that were created after the configuration backup was taken, and this information is displayed in the Veeam Backup & Replication console.

- Select **Migrate** if you want to restore data from the configuration backup to the database used by another backup server.

  In the Migrate mode, Veeam Backup & Replication retrieves configuration data from the backup and stores it to the target database. No rescan operation is performed.

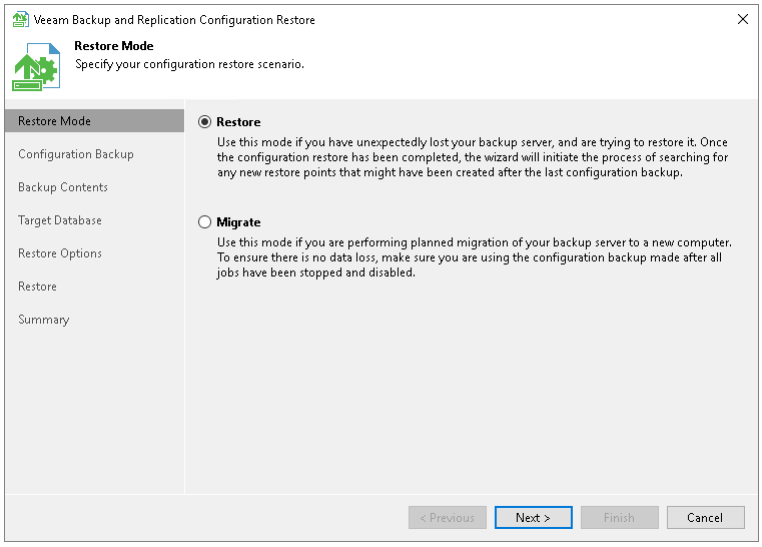

# Step 3. Select Configuration Backup

At the **Configuration Backup** step of the wizard, select a configuration backup from which you want to restore data:

1. From the **Backup repository** list, select a server or backup repository on which the configuration backup file is located.

   > **NOTE**
   >
   > The list of backup repositories is stored in the current configuration database. To get a full list of repositories, consider the following prerequisites:
   >
   > - The user, who initiates the backup configuration restore process, must have access to the current configuration database.
   > - Multi-factor authentication (MFA) must be disabled to see the list of object storage repositories. To disable MFA, see Multi-Factor Authentication.
   > - The Veeam Backup Service must be running.

2. Click **Browse** next to the **Backup file** field and select the backup file.

If you select to restore configuration data from a backup in a remote backup repository, during restore Veeam Backup & Replication will first copy the backup file to a temporary folder on the backup server. After you finish the restore process and close the wizard, Veeam Backup & Replication will automatically delete the configuration file from the temporary folder.

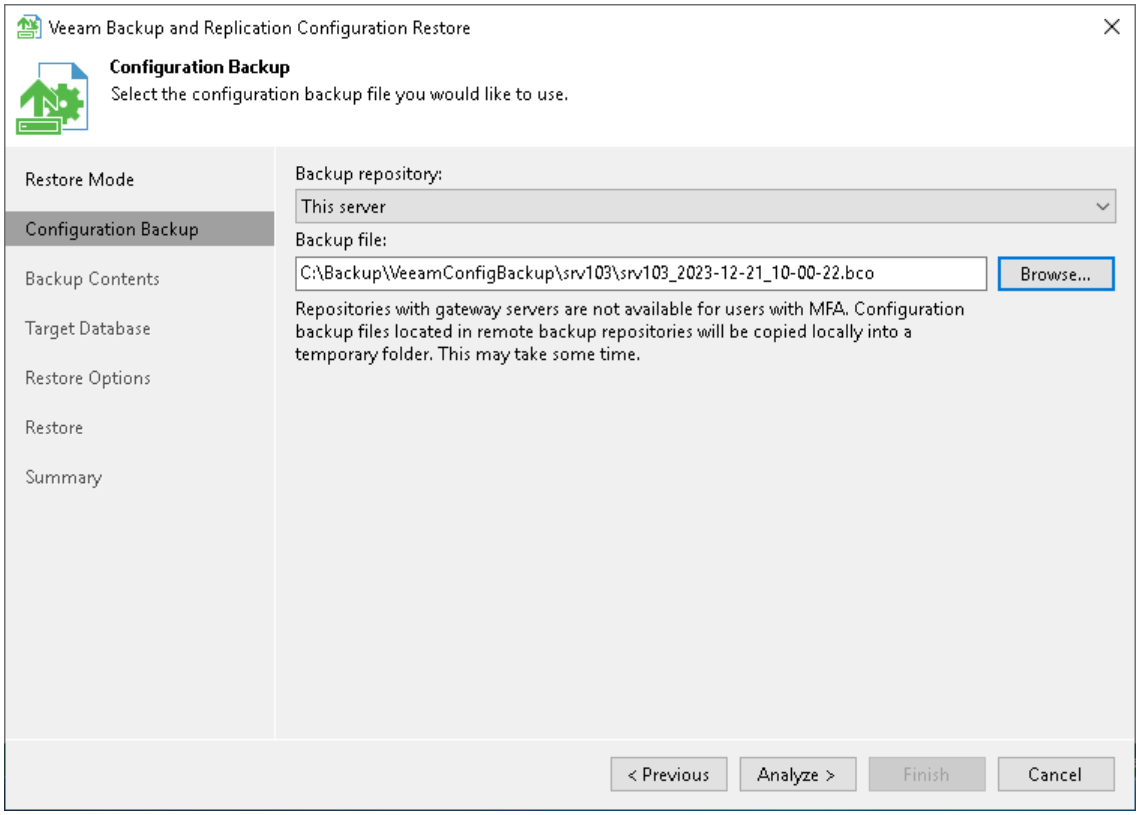

# Step 4. Review Configuration Backup Parameters

At the **Backup Contents** step of the wizard, Veeam Backup & Replication will analyze the content of the selected backup file and display the following information:

- **Backup file**: information about configuration backup file itself.
- **Product**: version of Veeam Backup & Replication installed on the initial backup server and configuration database version.
- **Database**: information about original database server, database version and compatibility.
- **Catalogs**: catalogs storing backup configuration data.

Review the displayed settings and click **Next**.

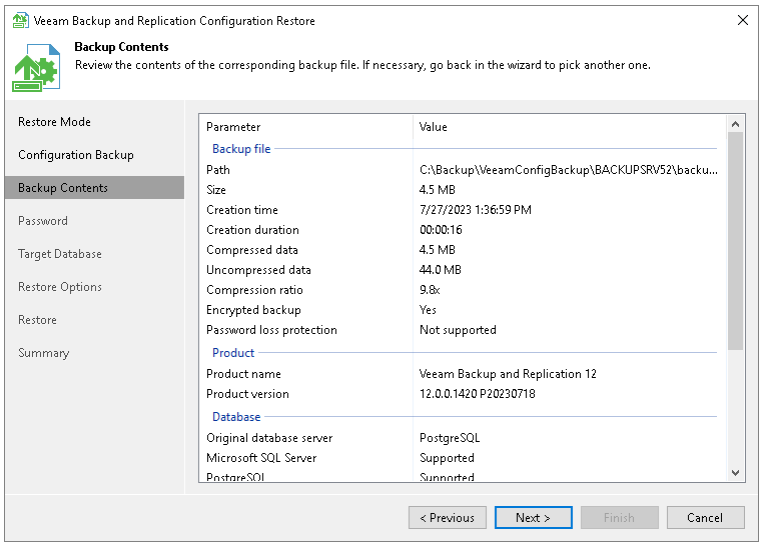

# Step 5. Specify Password

The **Password** step of the wizard is available if you have enabled the encryption option in the configuration backup properties.

Enter the password to decrypt configuration backup data:

1. Check the password hint to recall the password.
2. In the **Password** field, enter the password to decrypt the configuration backup file.

If you forgot or lost the password, click the **I forgot the password** link. For more information, see Decrypting Data Without Password.

> **NOTE**
>
> Restoring configuration data without a password is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.
>
> Also, your backup server must be connected to Veeam Backup Enterprise Manager. Otherwise, you will not see the **I forgot the password** link.

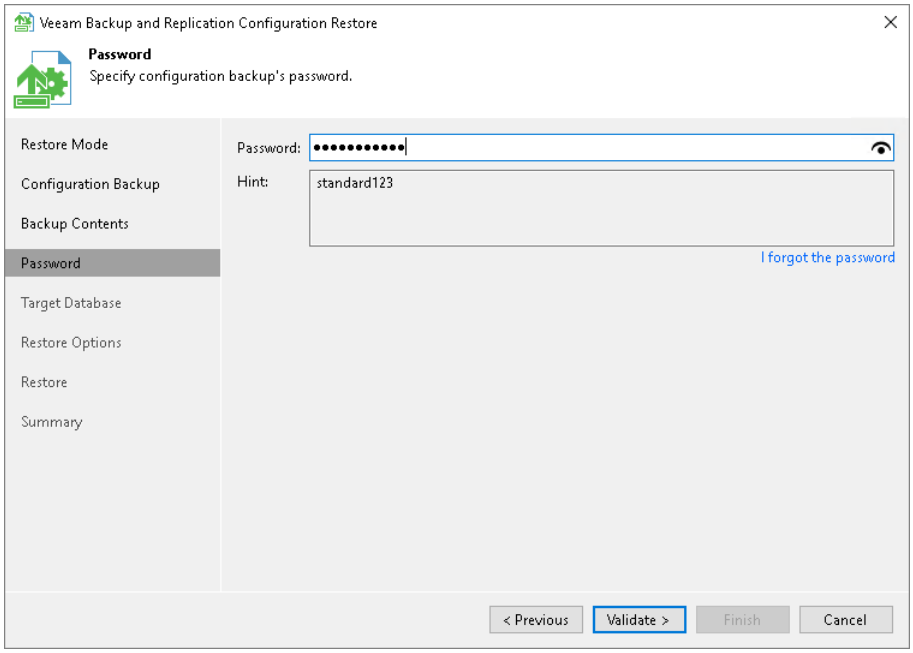

# Step 6. Specify Target Database

At the **Target Database** step of the wizard, specify the target Microsoft SQL Server or PostgreSQL instance to which configuration data must be restored.

1. In the **Database** field, select the database engine, PostgreSQL or Microsoft SQL Server.
2. In the **Instance** field, do one of the following:
   - For Microsoft SQL Server, select an instance on which the database is deployed or must be deployed. In the list of Microsoft SQL Server instances, Veeam Backup & Replication displays all servers from the network where the backup server resides. To update the list of servers, click **Refresh** on the right. To specify an instance manually, use the `SERVER_NAME\INSTANCE_NAME` format.
   - For PostgreSQL, specify an instance on which the database is deployed or must be deployed. Use the `hostname:port` format to specify an instance.
3. In the **Database name** field, specify a name of the database to which configuration data must be restored.

   By default, Veeam Backup & Replication uses the default name or port for the target database. If you specify a name of an existing target database, Veeam Backup & Replication will overwrite this database. If you specify a name of the database that does not exist, Veeam Backup & Replication will create it on the specified Microsoft SQL Server or PostgreSQL instance.

   > **NOTE**
   >
   > If a backup repository is located in the backup server, after configuration restore, this repository will point to the same path as it was before the migration but in a new host. For example, if you keep backed-up data on the D disk, after migration Veeam Backup & Replication will keep new backups on the D disk of the new host.

4. In the **Authentication** section, select the authentication mode to connect to the target database instance:
   - For Microsoft SQL Server, select **Windows authentication** or **SQL authentication** mode. If you select the **SQL authentication** mode, specify credentials that will be used to connect to the target Microsoft SQL Server instance.
   - For PostgreSQL, select **Windows authentication** or **Native authentication** mode. If you select the **Native authentication** mode, specify credentials that will be used to connect to the target PostgreSQL instance.

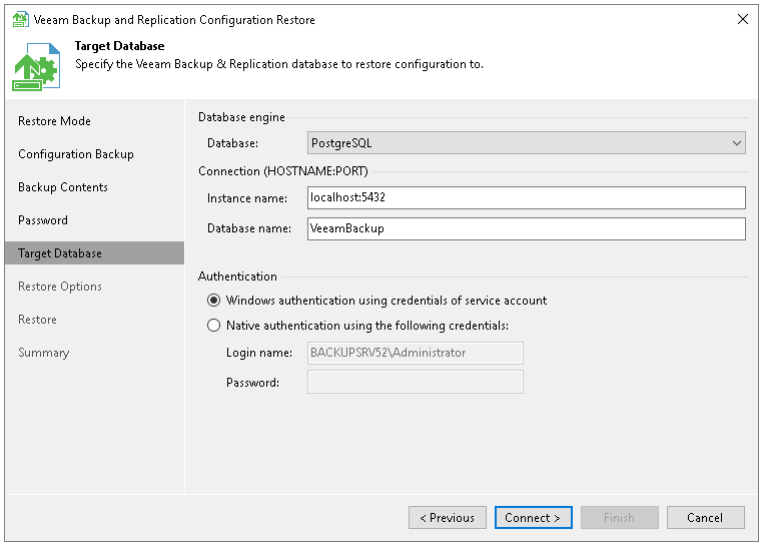

When you restore configuration to an existing database, the configuration restore process will delete the current state of the database contents and replace it with the restored data. Veeam Backup & Replication will display a warning. If you want to replace the contents, click **Yes** to confirm.

If you do not want to lose the current data, restore the configuration to a new database. To do this, click **No** to the warning and specify a non-existing database name in the **Database name** field.

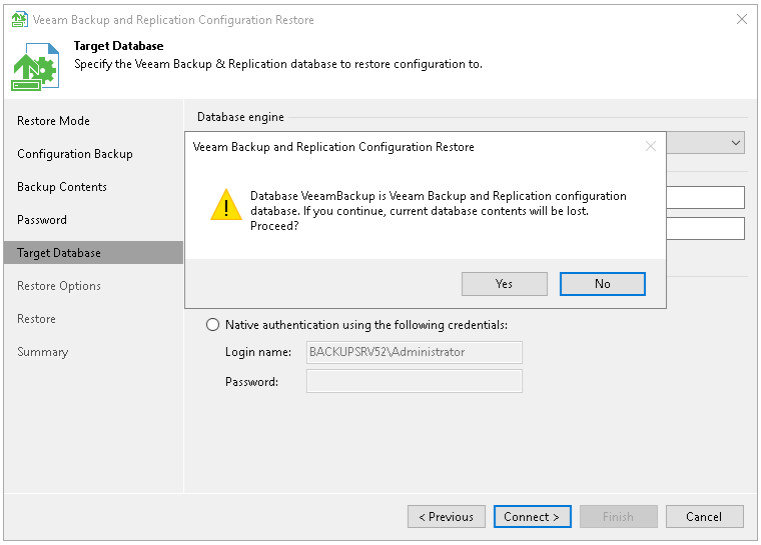

# Step 7. Specify Restore Options

At the **Restore Options** step of the wizard, specify additional restore options.

1. In the **Restore** section, select what data you want to restore from the configuration backup. Veeam Backup & Replication always restores configuration data for backup infrastructure components, jobs and global settings specified at the level of the backup server. You can additionally restore the following data:
   - **Backup and replica catalog**: data about all backups and replicas registered on the backup server and information about tapes to which backups were written and location of these tapes.
   - **Session history**: data about all sessions performed on the backup server.
2. If you plan to use PowerShell on the restored backup server, select the **Enable required PowerShell policy for SCVMM** check box. During restore, Veeam Backup & Replication will enable the PowerShell execution policy and you will not have to enable it manually afterwards. Enabling this option is identical to running the '*Set-ExecutionPolicy RemoteSigned*' command on the backup server.
3. If you are restoring configuration data to the same database, select the **Backup existing database before configuration restore** check box. This option will help you protect the current database from accidental errors during the restore process. During restore, Veeam Backup & Replication will first back up the current database using the native tools of Microsoft SQL Server or PostgreSQL. After that, Veeam Backup & Replication will purge the current database and import data from the configuration backup to it. In such scenario, if an error occurs during the restore process, you will be able to restore the current database from the Microsoft SQL backup using Microsoft SQL Management Studio or SQL scripts. For PostgreSQL, you will be able you will be able to restore the current database from the PostgreSQL backup using PGadmin or SQL scripts.

   The created Microsoft SQL database backup is named by the following pattern: *VeeamBackup<DatabaseName><date>.bak* and stored to the default Microsoft SQL backups location, for example: `%ProgramFiles%\Microsoft SQL Server\MSSQL11.MSSQLSERVER\MSSQL\Backup\`.

   The created PostgreSQL database backup is stored by the following path: `%ProgramData%\Veeam\Backup\Utils\PostgreSQLBackup`.
4. Click **Restore**. Veeam Backup & Replication will stop currently running jobs and Veeam Backup & Replication services and will restore the database to the specified location.

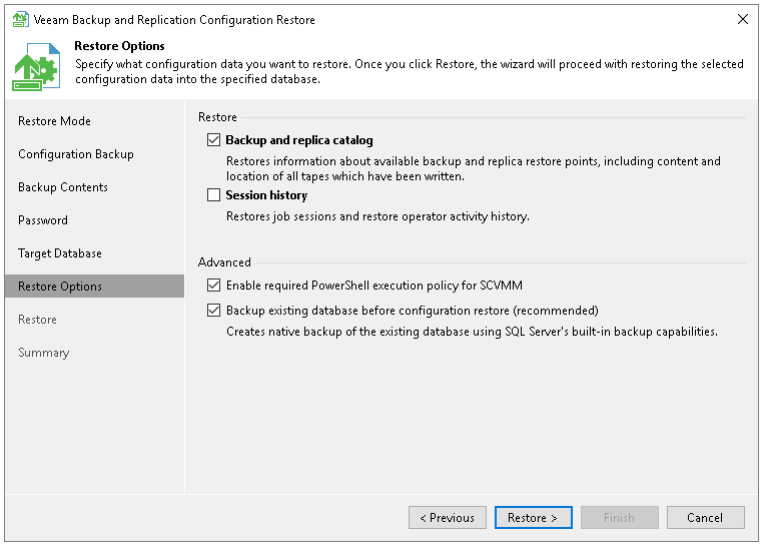

# Step 8. Review Restore Settings

At the **Restore** step of the wizard, Veeam Backup & Replication will display the progress on the restore process. Wait for the restore process to complete and click **Next**.

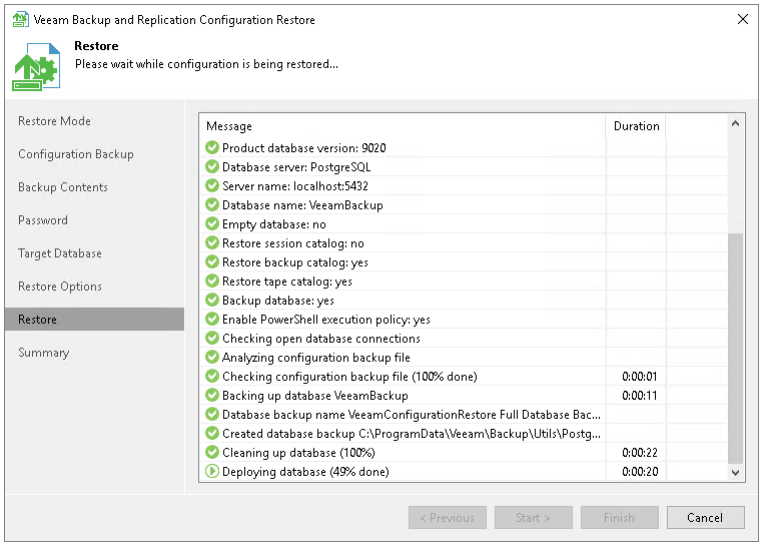


If you have chosen to restore data in the Migrate mode and the configuration backup file does not meet the Migrate mode requirements, Veeam Backup & Replication will display a warning and offer you to switch to the Restore mode. The Restore mode requires more time but guarantees that information about all new restore points will be available in the restored database.

- To switch to the **Restore** mode, in the warning window click **Yes**.
- To carry on data restore in the **Migrate** mode, in the warning window click **No**.
- To stop the restore process, in the warning window click **Cancel**.

For more information, see Migrating Veeam Backup & Replication to Another Backup Server.

# Step 9. Check Configuration

At the **Check Configuration** step of the wizard, Veeam Backup & Replication will analyze the content of the restored database and display issues which can potentially interfere with Veeam Backup & Replication functionality.

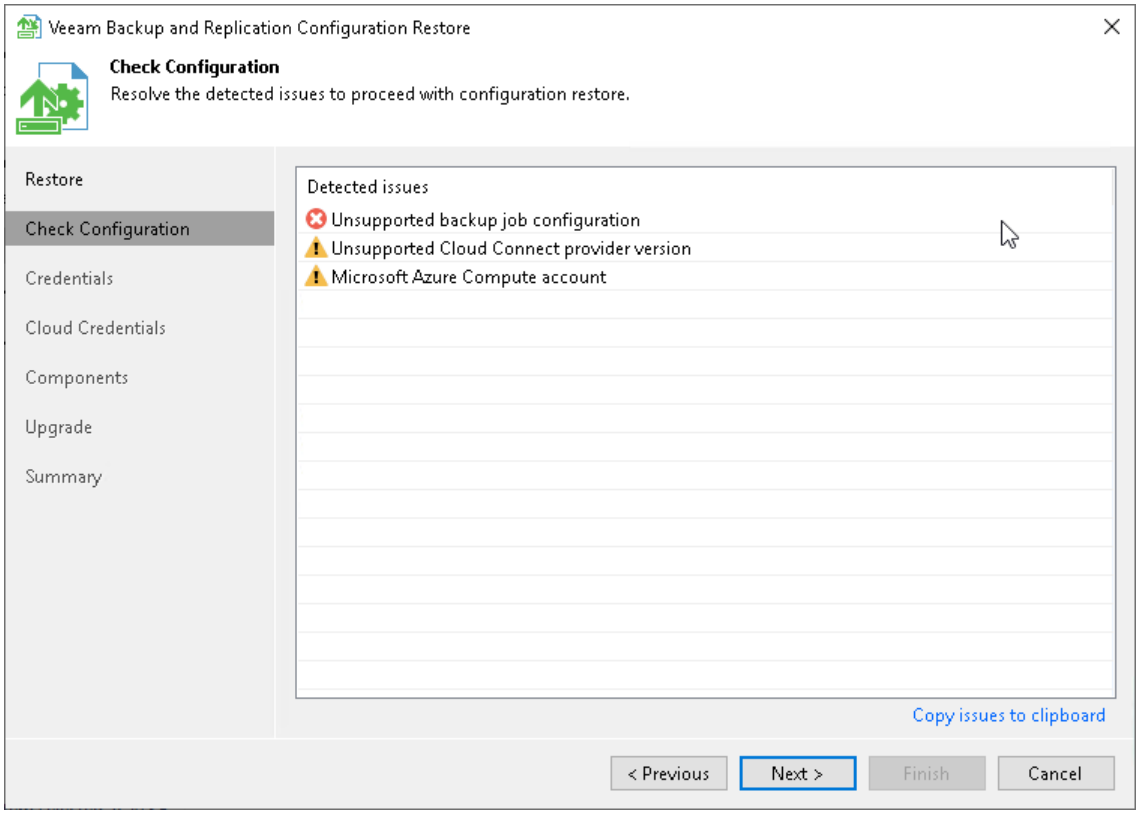

## Step 10. Finalize Restore Process

After the restore process has finished, you may need to perform the following actions to finalize the configuration database restore:

1. Specify credentials for backup infrastructure objects.
2. Specifying credentials for cloud services.
3. Perform components upgrade.

# Specifying Credentials

At the **Credentials** step of the wizard, Veeam Backup & Replication displays a list of credentials records that existed on the backup server at the time when the configuration backup was created. If by the time of restore passwords for credentials records have changed, you can specify new values for these records.

> **IMPORTANT**
>
> If you have not enabled encryption for configuration backups, Veeam Backup & Replication will not restore passwords for credentials records. You need to re-enter passwords for all credentials records to make sure that backup infrastructure components and jobs work in a proper way after you complete configuration restore.

To edit credentials records:

1. Select a record in the list and click **Edit**.
2. Edit settings of the record as required.
3. Repeat the procedure for all records in the list.

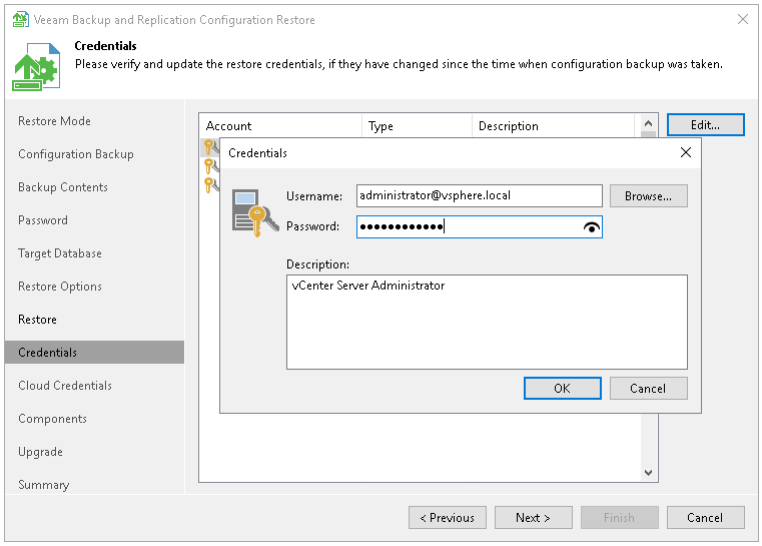


# Specifying Cloud Credentials

At the **Cloud Credentials** step of the wizard, Veeam Backup & Replication displays a list of cloud credentials records that existed on the backup server at the time when the configuration backup was created. If by the time of restore passwords for cloud credentials records have changed, you can specify new values for these records.

**IMPORTANT**

Consider the following:

- If you have not enabled encryption for configuration backups, Veeam Backup & Replication will not restore passwords for cloud credentials records. You need to re-enter passwords for all cloud credentials records to make sure that cloud services and jobs work in a proper way after you complete configuration restore.
- You cannot edit credentials of Microsoft Azure Compute, Google Cloud service account and Microsoft Azure storage (Entra ID) accounts in the configuration restore wizard. You can edit these credentials only after configuration restore in Cloud Credentials Manager. For details, see Editing and Deleting Credentials Records.
- Veeam Backup & Replication does not restore credentials records of Google Cloud service accounts that were created automatically using the Cloud Credentials Manager. For details, see Google Cloud Service Accounts.

To edit cloud credentials records:

1. Select a record in the list and click **Edit**.
2. Edit settings of the record as required.
3. Repeat the procedure for all records in the list.

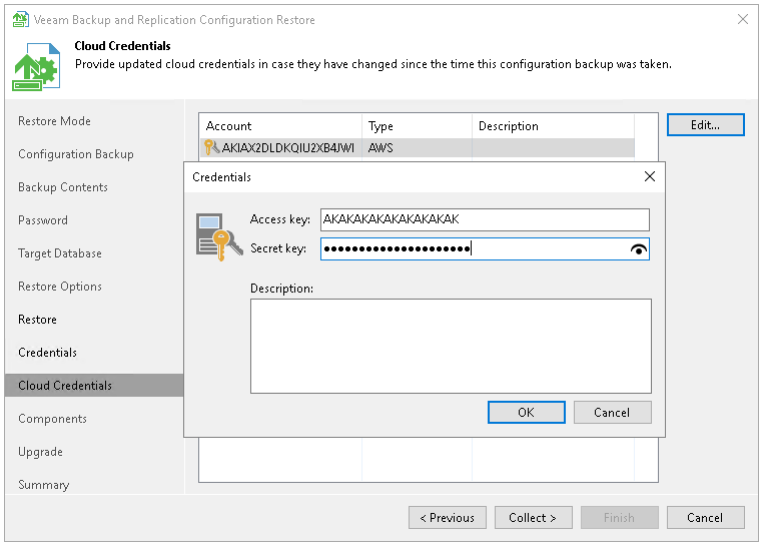


# Performing Components Upgrade

After the restore process is complete, Veeam Backup & Replication will check if services on backup infrastructure components must be upgraded and display a list of outdated components.

To upgrade backup infrastructure components, select check boxes next to the necessary components and click **Next**. If some component fails to upgrade, you can get back to a previous step of the wizard and repeat the procedure or close the wizard and upgrade the components manually. For more information, see Upgrading Infrastructure Components.

## Step 11. Synchronize Backups and Tape Libraries

After the configuration database is restored, Veeam Backup & Replication can perform a synchronization operation for backups and replicas created on the backup server and tape libraries connected to the backup server.

- The synchronization operation for backups and replicas is performed if you are restoring a database from a backup created by Veeam Backup & Replication in the Restore mode and you have selected to restore data from the backup and replica catalog.
- The synchronization operation for tape libraries is performed if you are restoring a database from a backup created by Veeam Backup & Replication in the Restore mode and you have selected to restore data from the backup and replica catalog.

Wait for the synchronization operation to complete.

## Step 12. Finish Working with Wizard

At the **Summary** step of the wizard, finalize the process of configuration data restore.

1. Review the restore process results.
2. If you want to start Veeam Backup & Replication after you finish working with the wizard, select the **Launch the Backup & Replication user interface** check box.
3. If you want to export the configuration file, click **Export answer file**.
4. Click **Finish** to exit the wizard.

> **NOTE**
>
> Consider the following:
>
> - If you created custom registry values or changed the existing ones on the backup server, you must recreate or change the registry values again.
> - If you restore data from the configuration backup in the Restore mode, all jobs on the backup server will be disabled after the restore process is complete. You need to enable them manually.

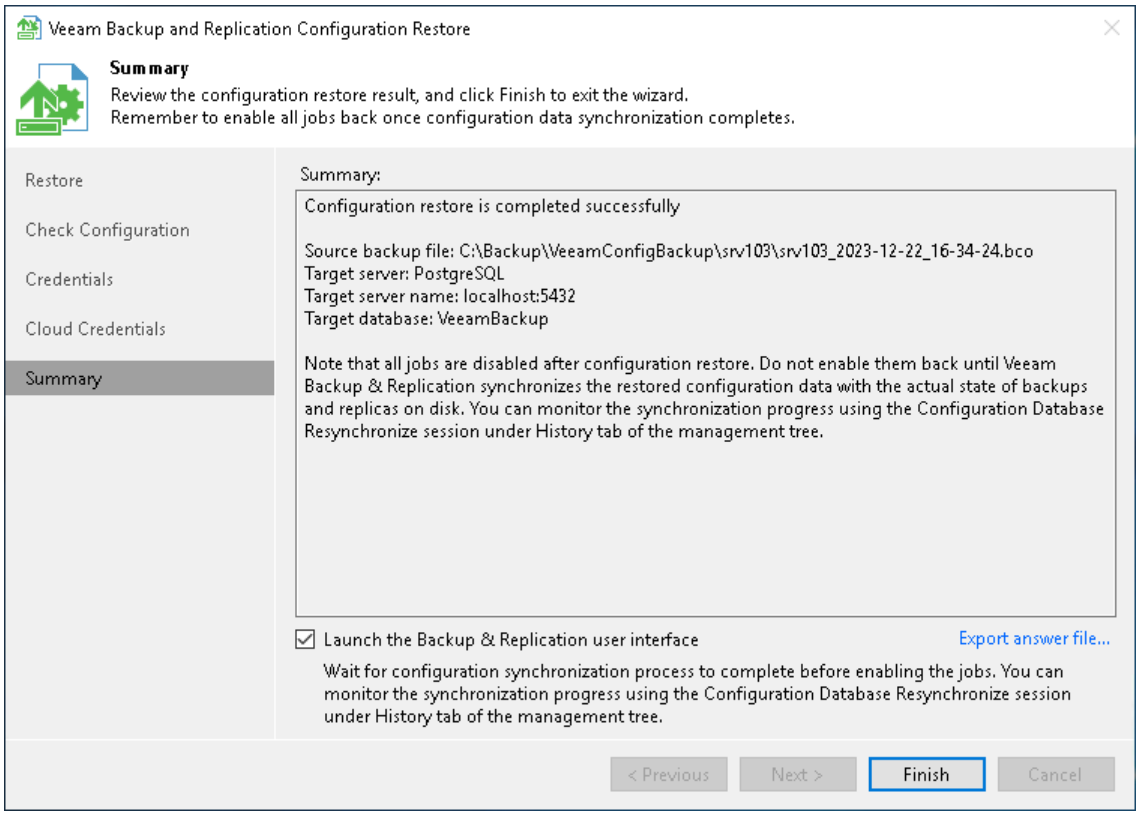


# Migrating Veeam Backup & Replication to Another Backup Server

If you need to migrate Veeam Backup & Replication to another backup server, you can back up its configuration database, install Veeam Backup & Replication on the target server and restore the configuration data from the backup. As a result, you will have a new Veeam Backup & Replication server with all settings, jobs and backup infrastructure elements from the old server.

# Limitations and Considerations

Before you migrate the configuration database of Veeam Backup & Replication to another backup server, consider the following limitations and considerations:

- If you want to specify the original database as the target database to restore configuration to, you must stop the Veeam Backup Service on the machine where the original Veeam Backup & Replication is installed before the restore process.
- If backup servers are not joined to an Active Directory domain, you must disable the Four-Eyes Authorization before the migration process.
- This section gives instructions on how to migrate Veeam Backup & Replication together with its configuration database to another server. If you need to migrate only the configuration database, see Migrating Configuration Database to Another SQL Server.
- Veeam Backup & Replication supports configuration database migration between different database engines only within the same Veeam Backup & Replication version.

> **IMPORTANT**
>
> We strongly recommend using only the configuration database restore process to migrate Veeam Backup & Replication configuration database to another backup server.

# Migrating Veeam Backup & Replication

If you want to migrate Veeam Backup & Replication to another backup server, perform the following steps:

1. Stop running jobs and disable scheduled jobs.
2. Save registry values that you changed or created.
3. Back up the configuration database of Veeam Backup & Replication.
4. Install Veeam Backup & Replication on the target machine.
5. Restore the configuration database from the backup.
6. Finish the configuration.

# Step 1. Stop and Disable Jobs

Before you start the migration process, stop all running jobs and disable all scheduled jobs on the source backup server before you create the configuration backup.

> **NOTE**
>
> Do not start or enable any jobs until the migration of Veeam Backup & Replication is finished. If you start a job before migration is completed, Veeam Backup & Replication will produce a new restore point in the chain and update the chain metadata. The created configuration backup will not contain information about this new restore point. When you migrate data from the configuration backup to the database and start the job again, Veeam Backup & Replication will fail to synchronize the chain metadata with data in the database. As a result, the job will fail.

## Step 2. [Optional] Save registry Values That You Changed or Created

If you have created new registry values or changed the existing keys on the backup server, you will need to recreate the keys or change the keys manually after the migration. The configuration database does not store registry values.

To save registry values, you can use the `reg export` command or you can use Registry Editor to save keys manually. For details, see Windows Documentation.

## Step 3. Back Up Configuration Database

Back up the configuration database of Veeam Backup & Replication. For instructions, see Running Configuration Backups Manually.

## Step 4. Install Veeam Backup & Replication on Target Machine

Install Veeam Backup & Replication on the machine on which you plan to move your source backup server. The machine must meet system requirements for a backup server.

During installation, you need to specify a new database name to store the configuration data. You can select a Microsoft SQL Server instance on the local or remote server. It does not matter whether you select an existing instance or create a new one. In the next steps of this guide, you will restore the previous configuration data to the database selected at this step.

For instructions on how to install Veeam Backup & Replication, see Installing Veeam Backup & Replication.

## Step 5. Restore Configuration Database from Backup

Use the native feature of Veeam Backup & Replication to restore the configuration database from the backup created in Step 3. Before you start the restore process, check the prerequisites.

1. On the target backup server, log in to the Veeam Backup & Replication console. For details, see Logging in to Veeam Backup & Replication.
2. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.

3. In the **Restore** section, click **Restore** to launch the **Configuration Database Restore** wizard.

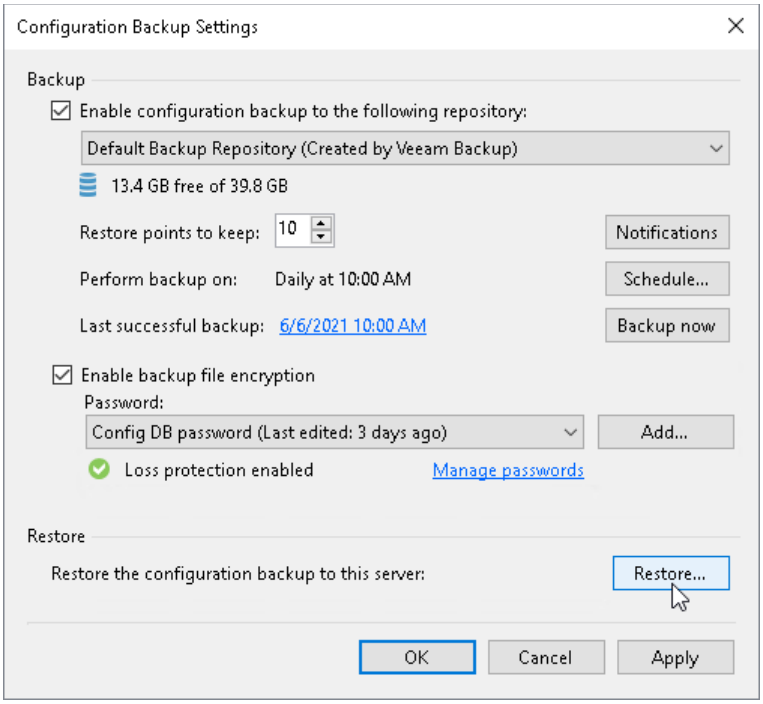


> **NOTE**
>
> **Loss protection disabled** warning is safe to ignore if you do not have Veeam Backup Enterprise Manager installed, your backup server is not registered with Veeam Backup Enterprise Manager server, or your system administrator chose not to enable loss protection functionality.

4. At the **Restore Mode** step of the wizard, select **Migrate**.

   Before migrating the configuration backup, Veeam Backup & Replication performs an additional check. If the configuration backup does not meet the requirements, Veeam Backup & Replication will offer you an option to switch to the **Restore** mode.

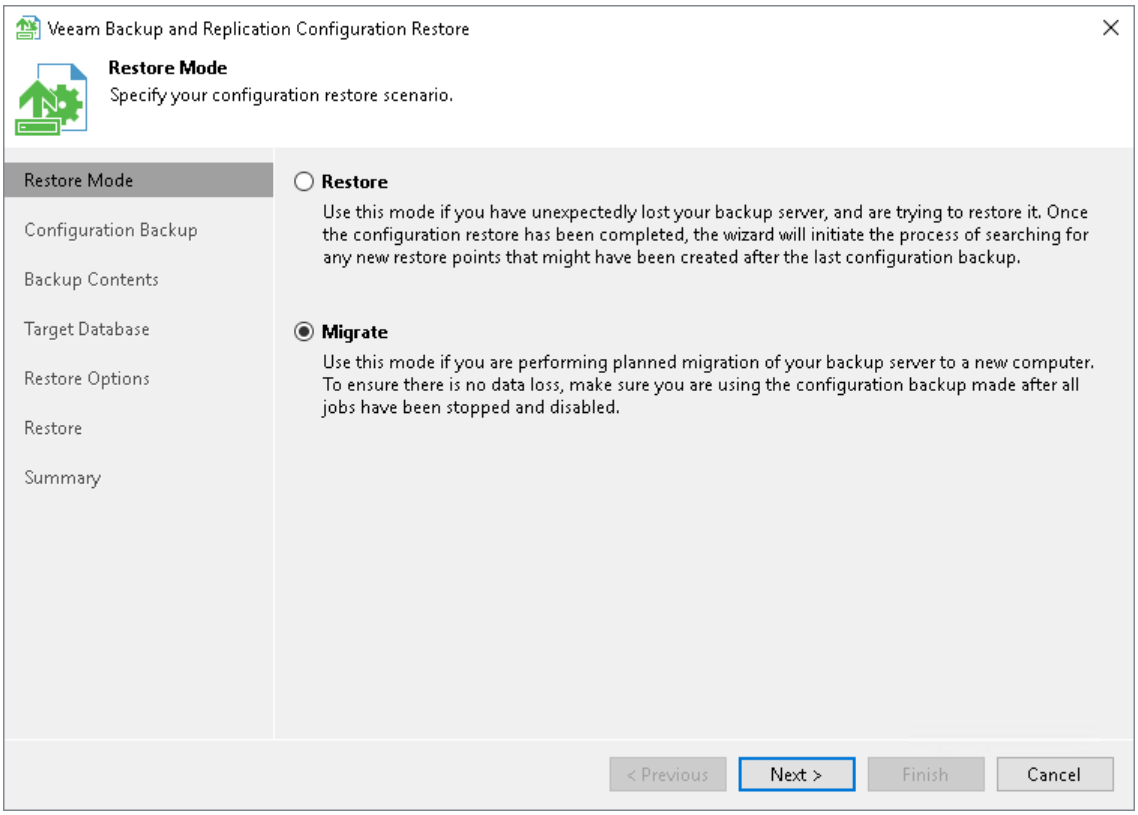


5. At the **Configuration Backup** step of the wizard, specify the backup file of the configuration database

   a. From the **Backup repository** drop-down list, select **This server**.

   b. Click **Browse** to specify the backup file location.

> **NOTE**
>
> This backup file has been created in Step 3. You must copy the backup file to the target Veeam Backup & Replication server beforehand.

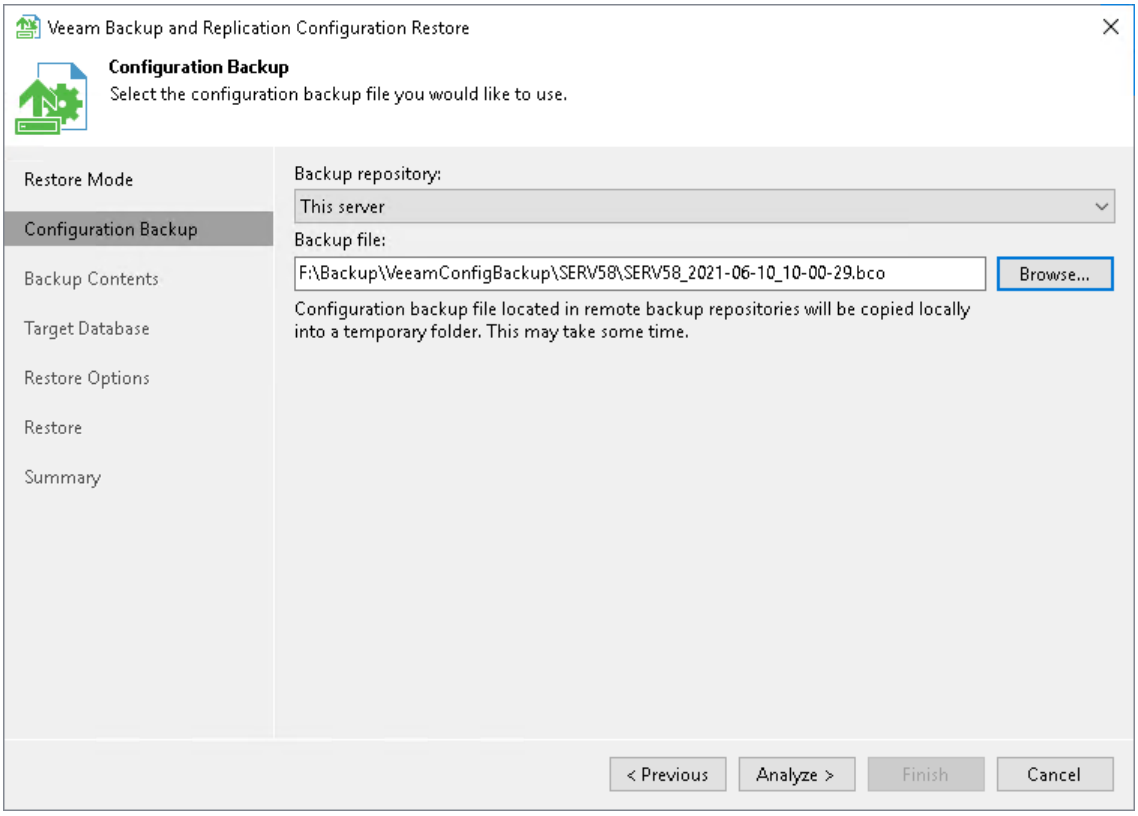


6. Complete the wizard as described in section Restoring Configuration Database.

   Veeam Backup & Replication will rescan VM replicas, backup repositories and tape libraries connected to the backup server. The database will be updated to include information about new restore points, and subsequent job sessions will work in a proper way.

# Step 6. Finish Configuration

After you restore the configuration backup, finalize the configuration:

- If you created custom registry values or changed the existing ones on the previous backup server, you must recreate or change the registry values again manually on the target backup server. You can import saved keys using the `reg import` command or Registry Editor.
- If you have local repositories, after migration to another machine they may be displayed as empty. In this case, add them again and remap the jobs.
- Enable your backup jobs and backup copy jobs. Take a closer look at your backup infrastructure to ensure that everything is working as expected.
- If you use Veeam Backup & Replication to back up storage systems, after migration they will not be added to the backup infrastructure. In this case, you must re-add them after migration completes. Storage System Snapshot Integration section in the Veeam Backup & Replication User Guide.
- If you use a hardened repository with immutability, after migration this server will not be available. In this case, you must specify single-use credentials for this repository again. For more information, see Editing Settings of Backup Repositories.
- If you use Linux hosts in your backup infrastructure, after migration these hosts and hosts that are associated with them will not be available. (For example, if you have a Linux host with the backup proxy role, the backup repositories to which this Linux backup proxy transfer during a backup job will not be available). In this case, you must open the Edit Linux Server wizard for the necessary Linux host, follow the steps of the wizard and click **Finish**.

- If you have Veeam Agent backup jobs managed by Veeam Agents, update the backup policies after the migration process.
- You can safely uninstall Veeam Backup & Replication from the old backup server after migration.
- If you use CDP, after the migration the I/O filter will be owned by the previous backup server. You must take ownership of the I/O filter on the new backup server. For more information, see Taking I/O Filter Ownership.

> **NOTE**
>
> If the Veeam Backup & Replication server was added to the Veeam Backup Enterprise Manager or Veeam ONE infrastructure, you must re-add the backup server after you migrate it to another server.

# Migrating Configuration Database to Another SQL Server

It is the best practice to keep the Veeam Backup & Replication application and its configuration database on the same server to maintain lowest latency and highest performance. However, in some scenarios a remote Microsoft SQL Server instance can be the better choice:

- High Availability. SQL Clustering and Always On Availability Group on external SQL Servers can be used for high availability of the configuration database. To learn about the configuration details, see this Veeam KB article.
- Licensing. Some enterprises have dedicated virtual clusters for SQL Servers due to licensing constraints. In such cases, you can place the Veeam configuration database on an existing instance to lower the total cost of ownership.

If you need to migrate the Veeam Backup & Replication configuration database to another SQL server, you can connect the configuration database to a Microsoft SQL Server instance deployed on another server and restore the configuration settings from the backup. As a result, you will be able to continue using the same Veeam Backup & Replication server but it will be connected to a configuration database on another server.

# Limitations and Considerations

Before you migrate the configuration database of Veeam Backup & Replication to another SQL server, consider the following limitations and considerations:

- This section gives instructions on how to migrate a configuration database to another SQL server. If you need to migrate the Veeam Backup & Replication application itself, see Migrating Veeam Backup & Replication to Another Backup Server.
- It is recommended that you use Veeam Backup & Replication tools to create configuration backups and migrate the configuration database. If you use native Microsoft SQL Server tools or others, after migration, some information, such as secure configuration data, may not be accessible.
- If you are migrating the configuration database to a remote SQL instance that uses Windows Authentication, all services that had access to the remote SQL instance before a migration must have the these permissions after migration. For more information on permissions, see Permissions.
- The account used to run Veeam Backup Service requires access to the database. For more information see, Permissions.
- If a backup server and a configuration database are located in different AD domains, the AD domain where the configuration database is located must have a trust relationship with the AD domain to which the backup server is added.

- When you migrate the configuration database to another SQL server, you must use the Microsoft SQL Server credentials that have CREATE ANY DATABASE permission on the target Microsoft SQL Server. For details, see Microsoft Docs. After database creation this account automatically gets the db_owner role and can perform all operations with the database. If the current account does not have this permission, a Database Administrator may create an empty database in advance and grant the db_owner role to the account that will be used for migration of the configuration database.
- Veeam Backup Enterprise Manager collects data from backup servers with configuration databases that run on the same database engine as the Enterprise Manager configuration database. The database engine used by Veeam Backup Enterprise Manager and all of the Veeam Backup & Replication Servers managed by this Veeam Backup Enterprise Manager must match. To migrate the Enterprise Manager configuration database, see Veeam Backup Enterprise Manager Guide.

# Migrating Configuration Database

If you want to migrate the configuration database of Veeam Backup & Replication to another SQL server, perform the following steps:

1. Stop and disable jobs.
2. Back up the configuration database.
3. Restore the configuration database from the backup.
4. [Optional] Reactivate the Enterprise Manager Keyset.
5. Finish the configuration.

# Step 1. Stop and Disable Jobs

Before you start the database migration, you must finish all jobs and restore sessions. If the job is scheduled, you must disable the job. For instructions on how to stop and disable jobs, see Managing Backup Jobs.

> **NOTE**
>
> Do not start or enable any jobs until the migration of Veeam Backup & Replication is finished. If you start a job before migration is completed, Veeam Backup & Replication will produce a new restore point in the chain and update the chain metadata. The created configuration backup will not contain information about this new restore point. When you migrate data from the configuration backup to the database and start the job again, Veeam Backup & Replication will fail to synchronize the metadata of the backup chain with data in the database. As a result, the job will fail.

# Step 2. Create Configuration Database Backup

To create a configuration database backup manually, perform the following steps:

1. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.
2. Make sure that the **Enable configuration backup to the following repository** check box is selected.
3. From the list of repositories, select a backup repository in which the configuration backup must be stored.

4. Click **Backup now**.

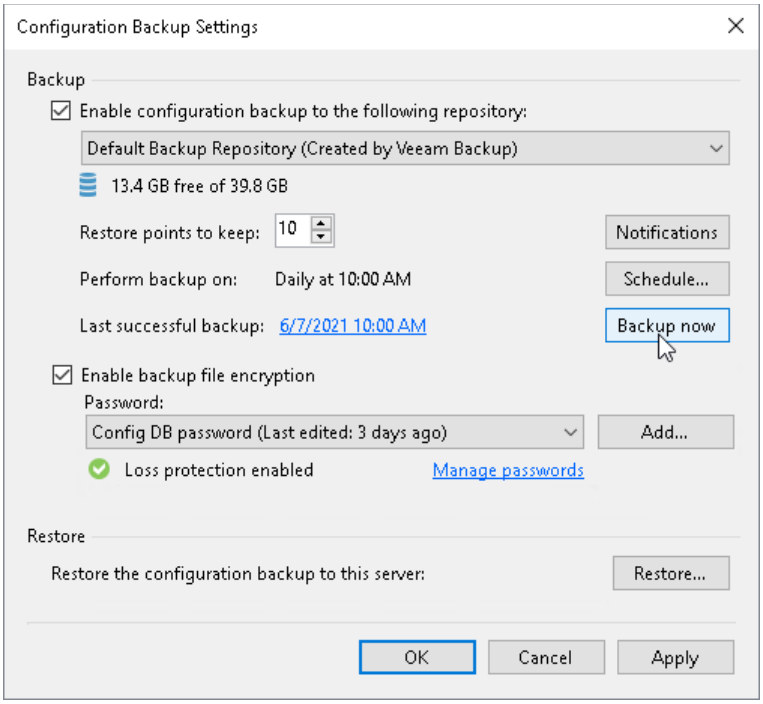


> **NOTE**
>
> **Loss protection disabled** warning is safe to ignore if you do not have Veeam Backup Enterprise Manager installed, your backup server is not registered with Veeam Backup Enterprise Manager server, or your system administrator chose not to enable loss protection functionality.

# Step 3. Restore Configuration Database from Backup

Before you start the restore process, check the prerequisites. To restore the configuration database, perform the following:

1. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.

2. In the **Restore** section, click **Restore**.

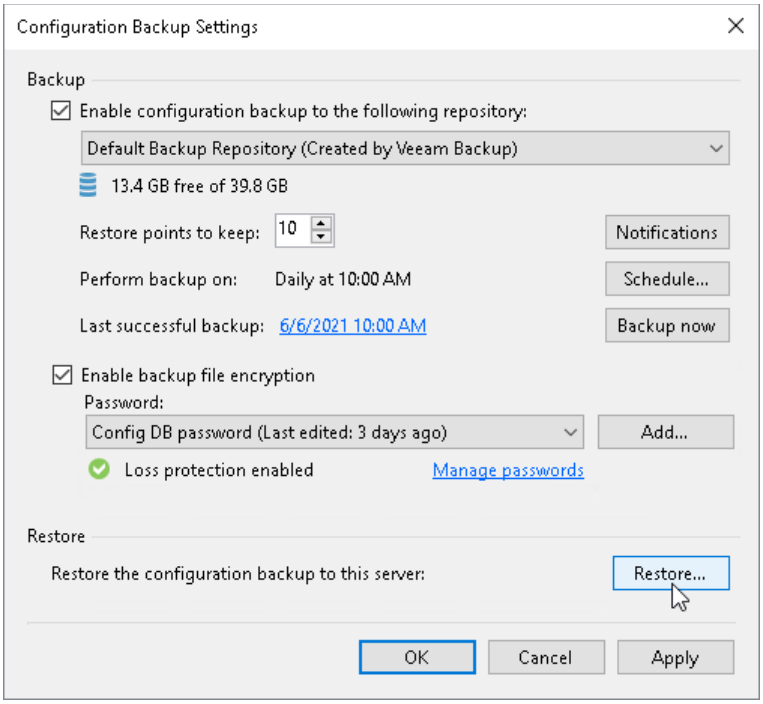


3. At the **Restore Mode** step of the **Veeam Backup & Replication Configuration Restore** wizard, select **Migrate**.

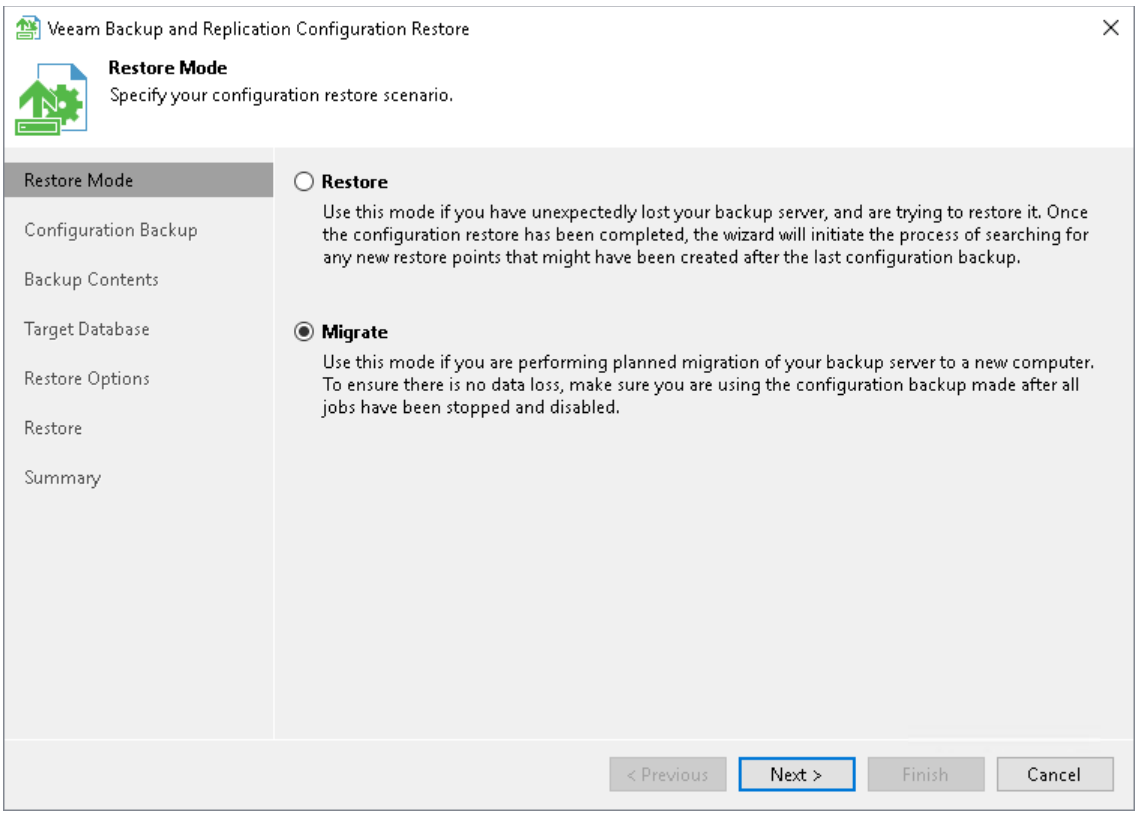


4. Complete the wizard as described in section Restoring Configuration Database.

## Step 4. [Optional] Reactivate Enterprise Manager Keyset

After you migrate the Veeam Backup & Replication configuration database to another server, Veeam Backup Enterprise Manager still sees the Veeam Backup & Replication server. However, you may need to reactivate encryption keys.

If you use the Data Encryption feature to encrypt backups and your Veeam Backup & Replication server is added to the Veeam Backup Enterprise Manager infrastructure, then you must reactivate the Enterprise Manager keyset.

To reactivate the Enterprise Manager key, perform the following steps:

1. In the Veeam Backup Enterprise Manager web console, open the **Settings** section of the **Configuration** view.
2. Open the **Key Management** tab.
3. In the **Managed keys** section, select the necessary keyset and click **Activate**.

For detailed instructions, see the Activating Enterprise Manager Keyset section in the Veeam Backup Enterprise Manager Guide.

## Step 5. Finish Configuration

After restoring the configuration database from the backup, finalize the configuration:

- Configure all necessary settings to ensure that you have a working configuration database backup. You can now perform a backup of your new configuration database in the **Configuration Backup Settings** window.

  Reschedule your configuration database backup. Also, check if you can see the *Loss protection enabled* label.

- If you have local repositories, after migration to another VM they may be displayed as empty. In this case, add them again and remap the jobs.
- Enable your backup jobs and backup copy jobs. Take a closer look at your backup infrastructure to ensure that everything is working as expected.

# Migrating Configuration Database to PostgreSQL Server

If you need to migrate the Veeam Backup & Replication configuration database from Microsoft SQL Server to PostgreSQL, you need to create a new configuration backup and restore it on the PostgreSQL instance. As a result, you will be able to continue using the same Veeam Backup & Replication server but it will be connected to a configuration database on PostgreSQL instead of Microsoft SQL Server.

## Limitations and Considerations

Before you migrate the configuration database of Veeam Backup & Replication to PostgreSQL, consider the following limitations and considerations:

- This section gives instructions on how to migrate a configuration database to PostgreSQL. If you need to migrate the Veeam Backup & Replication application itself, see Migrating Veeam Backup & Replication to Another Backup Server.
- If you want to migrate configuration database from PostgreSQL to another SQL server, see Migrating Configuration Database to Another SQL Server.

- It is recommended that you use Veeam Backup & Replication tools to create configuration backups and migrate the configuration database.
- If a backup server and a configuration database are located in different AD domains, the AD domain where the configuration database is located must have a trust relationship with the AD domain to which the backup server is added.
- Veeam Backup & Replication supports configuration database migration between different database engines only within the same Veeam Backup & Replication version.
- Veeam Backup Enterprise Manager collects data from backup servers with configuration databases that run on the same database engine as the Enterprise Manager configuration database. The database engine used by Veeam Backup Enterprise Manager and all of the Veeam Backup & Replication Servers managed by this Veeam Backup Enterprise Manager must match. To migrate the Enterprise Manager configuration database, see Veeam Backup Enterprise Manager Guide.
- PostgreSQL standard system database named `template1` must have UTF-8 encoding.

# Migrating Configuration Database

If you want to migrate the configuration database of Veeam Backup & Replication to another PostgreSQL, perform the following steps:

1. Stop and disable jobs.
2. Back up the configuration database.
3. Restore the configuration database from the backup.
4. [Optional] Reactivate the Enterprise Manager Keyset.
5. Finish the configuration.

# Step 1. Stop and Disable Jobs

Before you start the database migration, you must finish all jobs and restore sessions. If the job is scheduled, you must disable the job. For instructions on how to stop and disable jobs, see Managing Backup Jobs.

> **NOTE**
>
> Do not start or enable any jobs until the migration of Veeam Backup & Replication is finished. If you start a job before migration is completed, Veeam Backup & Replication will produce a new restore point in the chain and update the chain metadata. The created configuration backup will not contain information about this new restore point. When you migrate data from the configuration backup to the database and start the job again, Veeam Backup & Replication will fail to synchronize the metadata of the backup chain with data in the database. As a result, the job will fail.

# Step 2. Create Configuration Database Backup

To create a configuration database backup manually, perform the following steps:

1. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.
2. Make sure that the **Enable configuration backup to the following repository** check box is selected.
3. From the list of repositories, select a backup repository in which the configuration backup must be stored.

4. Click **Backup now**.

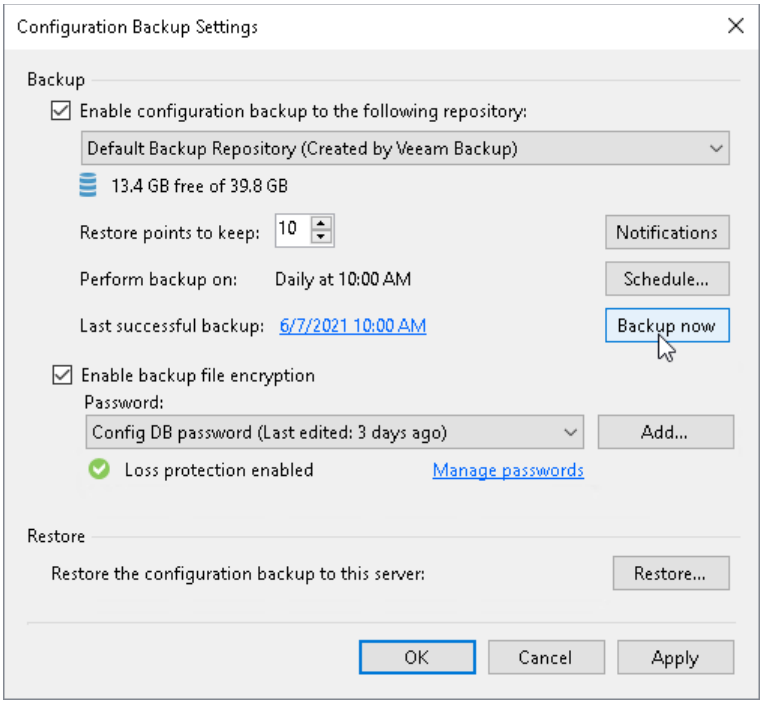


> **NOTE**
>
> **Loss protection disabled** warning is safe to ignore if you do not have Veeam Backup Enterprise Manager installed, your backup server is not registered with Veeam Backup Enterprise Manager server, or your system administrator chose not to enable loss protection functionality.

# Step 3. Restore Configuration Database from Backup

Before you start the restore process, check the prerequisites. To restore the configuration database, perform the following:

1. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.

2. In the **Restore** section, click **Restore**.

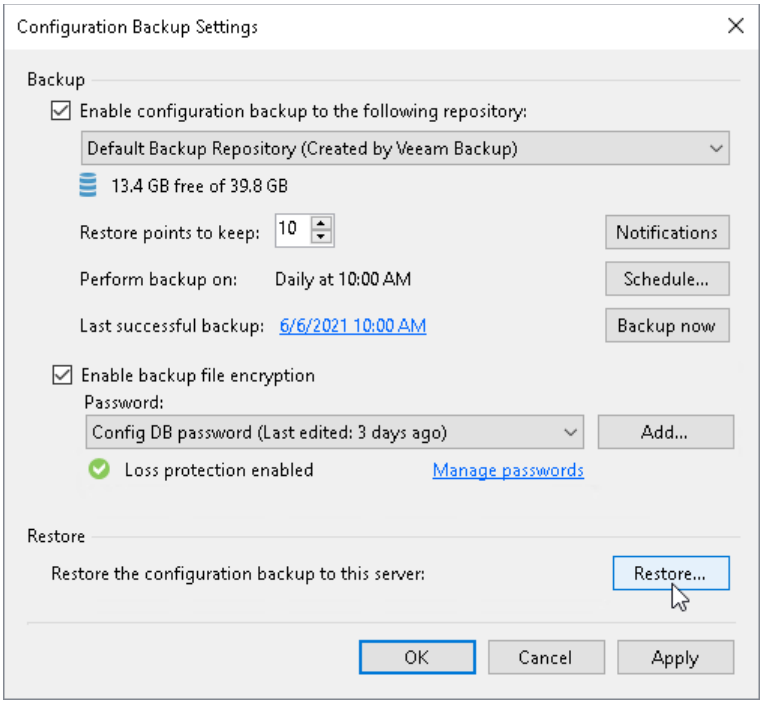


3. At the **Restore Mode** step of the **Veeam Backup & Replication Configuration Restore** wizard, select **Migrate**.

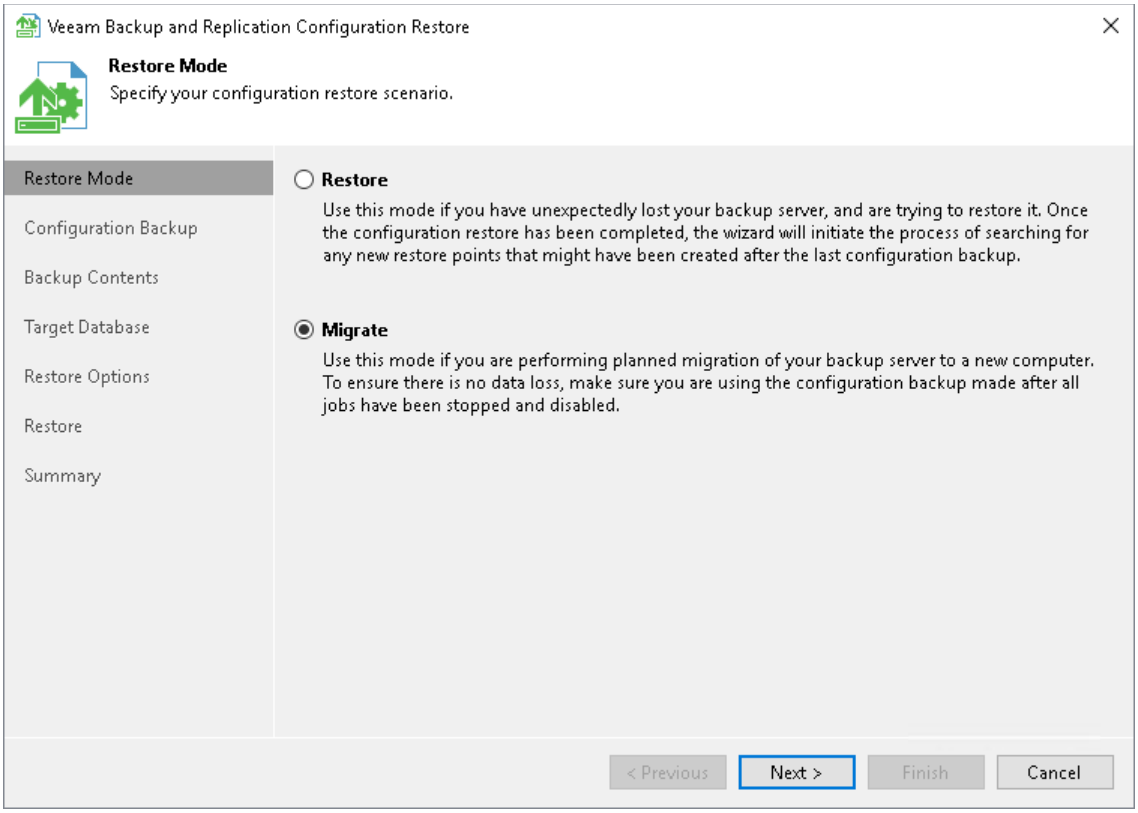


4. Complete the wizard as described in section Restoring Configuration Database. On the **Target Database** tab of the configuration restore wizard, select PostgreSQL as the database engine.

**IMPORTANT**

After the restore process, make sure that PostgreSQL target instance is configured according to the recommended hardware resources values. You can modify settings of PostgreSQL target instance using the `Set-VBRPSQLDatabaseServerLimits` cmdlet. For more information, see Veeam PowerShell Reference.

**NOTE**

After you complete the restore process, Veeam Backup & Replication will connect to the database you specified as a target.

# Step 4. [Optional] Reactivate Enterprise Manager Keyset

After you migrate the Veeam Backup & Replication configuration database to another server, Veeam Backup Enterprise Manager still sees the Veeam Backup & Replication server. However, you may need to reactivate encryption keys.

If you use the Data Encryption feature to encrypt backups and your Veeam Backup & Replication server is added to the Veeam Backup Enterprise Manager infrastructure, then you must reactivate the Enterprise Manager keyset.

To reactivate the Enterprise Manager key, perform the following steps:

1. In the Veeam Backup Enterprise Manager web console, open the **Settings** section of the **Configuration** view.
2. Open the **Key Management** tab.
3. In the **Managed keys** section, select the necessary keyset and click **Activate**.

For detailed instructions, see the Activating Enterprise Manager Keyset section in the Veeam Backup Enterprise Manager Guide.

# Step 5. Finish Configuration

After restoring the configuration database from the backup, finalize the configuration:

- Configure all necessary settings to ensure that you have a working configuration database backup. You can now perform a backup of your new configuration database in the **Configuration Backup Settings** window.

  Reschedule your configuration database backup. Also, check if you can see the *Loss protection enabled* label.
- If you have local repositories, after migration to another VM they may be displayed as empty. In this case, add them again and remap the jobs.
- Enable your backup jobs and backup copy jobs. Take a closer look at your backup infrastructure to ensure that everything is working as expected.

# Automating Configuration Database Restore

If you want to automate the process of the configuration database restore, you can do it using the `Veeam.Backup.Configuration.UnattandedRestore.exe` file.

# Limitations and Considerations

Before you start the restore process, check the prerequisites.

# Restoring Configuration Database

To restore Veeam Backup & Replication configuration database using the `Veeam.Backup.Configuration.UnattendedRestore.exe` file, perform the following steps:

1. Generate Configuration File.
2. Specify Restore Settings.
3. Restore Configuration Database.

## Step 1. Generate Configuration File

To generate the configuration file, do the following:

1. Open the command prompt.
2. Change the directory to the directory where the `Veeam.Backup.Configuration.UnattendedRestore.exe` file is stored. By default, `C:\Program Files\Veeam\Backup and Replication\Backup`.

   ```
   cd "C:\Program Files\Veeam\Backup and Replication\Backup"
   ```

3. Use the following command to generate the configuration file. In the `generate` parameter, specify the path where the configuration file will be saved.

   ```
   Veeam.Backup.Configuration.UnattendedRestore.exe /generate:C:\backup\unatt
   ended.xml
   ```

Alternatively, you can complete the configuration restore wizard and click **Export answer file** on the Summary step of the wizard. In this case, the configuration file will contain parameters you specified during the restore process.

## Step 2. Specify Restore Settings

After you execute the command, the configuration file will be generated with default preferences by the path you specified. You can manually edit the file to specify the necessary parameters.

## Step 3. Restore Configuration Database

To restore the configuration database, do the following:

1. Open the command prompt.
2. Change the directory to the directory where the `Veeam.Backup.Configuration.UnattendedRestore.exe` file is stored. By default, `C:\Program Files\Veeam\Backup and Replication\Backup`.

```
cd "C:\Program Files\Veeam\Backup and Replication\Backup"
```

3. Use the following command to restore the configuration file. In the `file` parameter, specify the path where the configuration file is stored.

```
Veeam.Backup.Configuration.UnattendedRestore.exe /file:C:\backup\unattended.xml
```

Veeam Backup & Replication will use the preferences specified in the `unattended.xml` file to perform the configuration database restore.

The path to the log file is `%ProgramData%\Veeam\Backup\Utils\Util.VeeamBackupConfiguration.UnattendedRestore.log`.

# Configuration File Parameters

| Parameter | Description | Required |
|---|---|---|
| **CONFIGURATION_FILE** | Specifies a full path to the configuration backup file.<br>Supported values:<br>• For local backup files: `C:\Backup\VeeamConfigBackup\VM\VM_2022-12-04_13-01-38.bco`.<br>• For network backup files: `\\VM\Backups\VM_2022-12-04_13-01-38.bco`.<br>• For repository backup files: `VeeamConfigBackup\VM\VM_2022-12-04_13-01-38.bco`. | Yes |
| **REPOSITORY_NAME** | Specifies Veeam Backup & Replication repository name where the configuration backup file is stored. If you do not specify this parameter, empty name will be used.<br>Supported values: String. | No |
| **BACKUP_PASSWORD** | Specifies the password to decrypt the configuration backup file.<br>Supported values: String. | No |
| **NETWORK_USER** | Specifies the user account for the network share.<br>Supported values: `domain\username`. | No |

| Parameter | Description | Required |
| --- | --- | --- |
| NETWORK_PASSWORD | Specifies the password for the network share.<br>Supported values: String. | No |
| SQLSERVER_ENGINE | Specifies the database engine.<br>Supported values: `mssql` or `postgresql`.<br>Default: `postgresql`. | No |
| DATABASE_SERVER | Specifies the database server and instance on which the configuration database will be deployed.<br>Supported values:<br>• Microsoft SQL Server: `MSSQLSERVER\DBINSTANCE:PORT`.<br>• PostgreSQL Server: `POSTGRESQLSERVER:PORT`. | Yes |
| SQLSERVER_DATABASE | Specifies the name for the configuration database.<br>Supported values: String.<br>Default: `VeeamBackup`. | No |
| SQLSERVER_AUTHENTICATION | Specifies the authentication mode to connect to the database server where the Veeam Backup & Replication configuration database will be deployed.<br>Supported values:<br>• Windows authentication: `0`.<br>• SQL native authentication: `1`.<br>Default: `0`. | No |
| VBR_SQLSERVER_USERNAME | Specifies a LoginID to connect to the SQL server in the native authentication mode<br>Supported values: String.<br>**Note**: The parameter is required if the `SQLSERVER_AUTHENTICATION` parameter value is `1`. | No |
| VBR_SQLSERVER_PASSWORD | Specifies a password to connect to the SQL server in the native authentication mode.<br>Supported values: String.<br>**Note**: The parameter is required if you specify the `VBR_SQLSERVER_USERNAME` parameter. | No |

| Parameter | Description | Required |
| --- | --- | --- |
| **PG_DUMP_PATH** | Specifies a path to the `pg_dump.exe` file.<br>Supported values: String. | No |
| **RESTORE_BACKUPS** | Defines that Veeam Backup & Replication will restore backup and replica restore points catalog.<br>Supported values:<br>• No: `0`.<br>• Yes: `1`.<br>Default: `1`. | No |
| **RESTORE_SESSIONS** | Defines that Veeam Backup & Replication will restore sessions history.<br>Supported values:<br>• No: `0`.<br>• Yes: `1`.<br>Default: `0`. | No |
| **ENABLE_POWERSHELL_POLICY** | Defines that PowerShell execution policy will be set to work for SCVMM.<br>Supported values:<br>• No: `0`.<br>• Yes: `1`.<br>Default: `1`. | No |
| **BACKUP_EXISTING_DATABASE** | Defines that the current database will be backed up.<br>Supported values:<br>• No: `0`.<br>• Yes: `1`.<br>Default: `0`. | No |

| Parameter | Description | Required |
| --- | --- | --- |
| SERVICES_AUTOSTART | Defines that Veeam Backup & Replication will start automatically after the migration.<br>Supported values:<br>• No: 0.<br>• Yes: 1.<br>Default: 1. | No |
| CREATE_NEW_DATABASE | Defines that the new database will be created if it does not exist.<br>Supported values:<br>• No: 0.<br>• Yes: 1.<br>Default: 1. | No |
| USE_EXISTING_DATABASE | Defines that Veeam Backup & Replication will use an existing database if it is not empty.<br>Supported values:<br>• No: 0.<br>• Yes: 1.<br>Default: 0. | No |
| USE_LOCKED_DATABASE | Defines that Veeam Backup & Replication will use an existing database owned by another backup server.<br>Supported values:<br>• No: 0.<br>• Yes: 1.<br>Default: 1. | No |
| OVERWRITE_EXISTING_DATABASE | Defines that the new database will overwrite the existing one.<br>Supported values:<br>• No: 0.<br>• Yes: 1.<br>Default: 0. | No |

| Parameter | Description | Required |
| --- | --- | --- |
| STOP_PROCESSES | Defines that Veeam Backup & Replication will terminate all running backup server processes.<br>Supported values:<br>• No: `0`.<br>• Yes: `1`.<br>Default: `0`. | No |
| SWITCH_TO_RESTORE_MODE | Defines that the configuration restore will switch to the restore mode if enabled backup jobs are found.<br>Supported values:<br>• No: `0`.<br>• Yes: `1`.<br>• Cancel: `2`.<br>Default: `2`. | No |
| RETRY_COUNT | Specifies the amount of retries that configuration restore should perform.<br>Supported values: Int32.<br>Default: `3`. | No |
| ACCEPT_FOUND_DATABASE_ISSUES | Defines that the configuration restore will proceed if database issues are found.<br>Supported values:<br>• No: `0`.<br>• Yes: `1`.<br>Default: `0`. | No |
| CREDENTIALS | Specify this parameter if you want to update stored passwords during the configuration restore. If you do not specify this parameter, passwords will not be updated<br>Supported values: `user=password;{hint}`. | No |
| PRIVATE_KEYS | Specify this parameter if you want to update stored private keys during the configuration restore. If you do not specify this parameter, private keys will not be updated.<br>Supported values: `user=privatekey;password;{hint}`. | No |

| Parameter | Description | Required |
| --- | --- | --- |
| HOSTS | Forces target server components upgrade if necessary. If you do not specify this parameter, all hosts will be upgraded.<br><br>**Note**: If you specify only several hosts out of all hosts in your backup infrastructure, only these hosts will be upgraded. Do not specify this parameter if you want to upgrade all hosts.<br><br>Supported values: `DNS name\IP address`. | No |

# Veeam Backup PowerShell Module

**Veeam Backup PowerShell** is an extension for Microsoft Windows PowerShell that adds a set of cmdlets to allow users to perform backup, replication and recovery tasks through the command-line interface of PowerShell or run custom scripts to fully automate operation of Veeam Backup & Replication. For more information on the Veeam Backup PowerShell module, see the Getting Started section of Veeam PowerShell Reference.

# Virtualization Servers and Hosts

You can add the following types of servers and hosts to the backup infrastructure:

- VMware vSphere Server
- VMware Cloud Director
- Microsoft Windows Server
- Linux Server

**NOTE**

We recommend that only one instance of a server or host is present in the backup infrastructure at a time. Do not add the same server or host multiple times, for example, by a DNS name and IP address, this can cause unexpected behavior.

You can add physical machines and VMs to the backup infrastructure and assign different roles to them. The following table describes which roles can be assigned to the different types of servers.

| Server Type | Source Host | Target Host | Backup Proxy | Backup Repository |
|---|---|---|---|---|
| **VMware vSphere Server**<br>(standalone ESXi host or vCenter Server) | ✓ | ✓ | × | × |
| **VMware Cloud Director** | ✓ | ✓<br>(for Cloud Director replication and continuous data protection) | × | × |
| **Microsoft Windows server** | × | × | ✓ | ✓ |
| **Linux server** | × | × | ✓ | ✓ |

NOTE

You can also add the following servers:

- Hyper-V servers. For more information on how to add servers, see the Adding Microsoft Hyper-V Servers section in the Veeam Backup & Replication User Guide for Microsoft Hyper-V.
- Proxmox VE servers (standalone hosts and clusters). For more information on how to add servers, see the Connecting Proxmox VE Server section in the Proxmox VE User Guide.
- Kasten instances. For more information on how to add Kasten instances, see the Deployment and Configuration section in the Veeam Kasten Integration Guide.
- Nutanix AHV servers (standalone clusters and Prism Centrals). For more information on how to add AHV servers, see the Connecting Nutanix AHV Server section in the Veeam Backup for Nutanix AHV User Guide.
- oVirt KVM managers. For more information on how to add oVirt KVM managers see the Connecting oVirt KVM Manager section in the Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization User Guide.
- Veeam Backup for AWS servers. For more information on how to add servers, see the Connecting to Existing Appliances section in the Veeam Backup for AWS User Guide.
- Veeam Backup for Microsoft Azure servers. For more information on how to add servers, see the Connecting to Existing Appliances section in the Veeam Backup for Microsoft Azure User Guide.
- Veeam Backup for Google Cloud servers. For more information on how to add servers, see the Connecting to Existing Appliances section in the Veeam Backup for Google Cloud Guide.

# Related Topics

- Veeam Data Mover Service
- Rescanning Servers
- Editing Server Settings
- Removing Servers

# Adding VMware vSphere Servers

You must add to the backup infrastructure VMware vSphere servers that you plan to use as source and target for backup, replication and other activities.

You can add vCenter Servers and ESXi hosts. If an ESXi host is managed by a vCenter Server, it is recommended that you add the vCenter Server, not a standalone ESXi host. If you move VMs between ESXi hosts managed by the vCenter Server, you will not have to re-configure jobs in Veeam Backup & Replication.
Veeam Backup & Replication will automatically locate migrated VMs and continue processing them as usual.

To add a VMware vSphere server, use the **New VMware Server** wizard.

# Step 1. Launch New VMware Server Wizard

To launch the **New VMware Server** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane, select the **Managed Servers** node and click **Add Server** on the ribbon or right-click the **Managed Servers** node and select **Add Server.** In the **Add Server** window, click **Virtualization Platforms** > **VMware vSphere** > **vSphere**.
- Open the **Inventory** view, in the inventory pane select the **VMware vSphere** node and click **Add Server** on the ribbon. You can also right-click the **VMware vSphere** node and select **Add Server**.

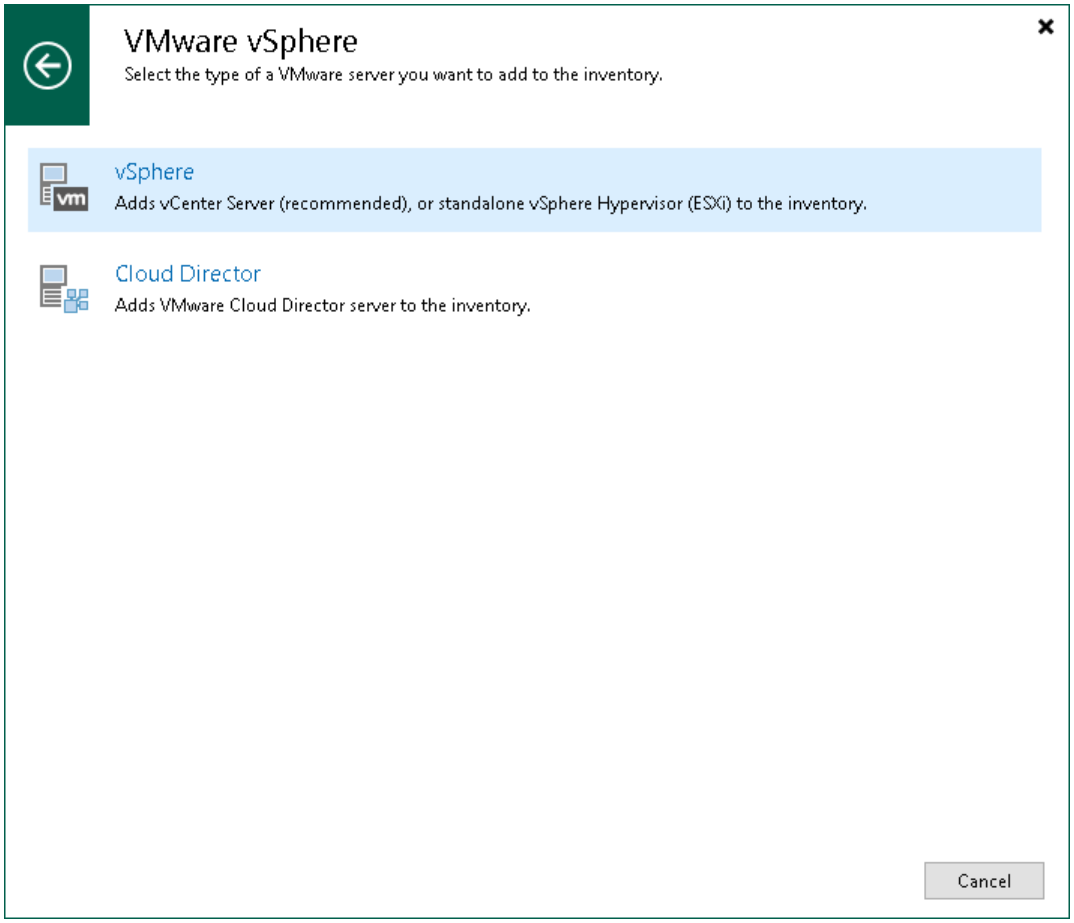

# Step 2. Specify Server Name or Address

At the **Name** step of the wizard, specify an address and description for the VMware vSphere server.

1. Enter a full DNS name, or IPv4 or IPv6 address of the vCenter Server or standalone ESXi host. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.

   If you add a VMware Cloud on AWS vCenter Server, use its Fully Qualified Domain Name (FQDN). Make sure the name you specify ends with <vmc.vmware.com>.

2. Provide a description for future reference. The default description contains information about the user who added the server, date and time when the server was added.

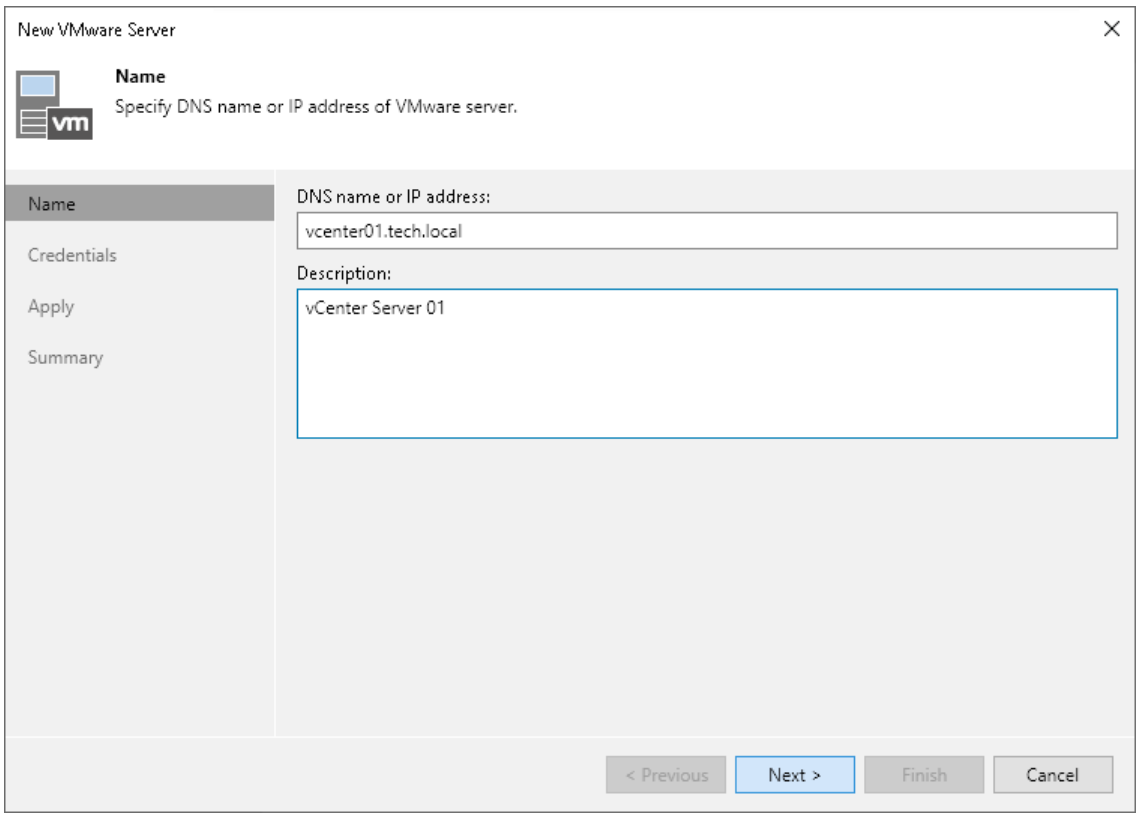

# Step 3. Specify Credentials

At the **Credentials** step of the wizard, specify credentials and port settings for the VMware vSphere server.

1. From the **Credentials** list, select credentials with the required permissions. For more information, see Permissions.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add the credentials. For more information, see Credentials Manager.

2. By default, Veeam Backup & Replication uses port 443 to communicate with vCenter Servers and ESXi hosts. If a connection with the vCenter Server or ESXi host over this port cannot be established, you can customize the port number in vCenter Server/ESXi host settings and specify the new port number in the **Port** field.

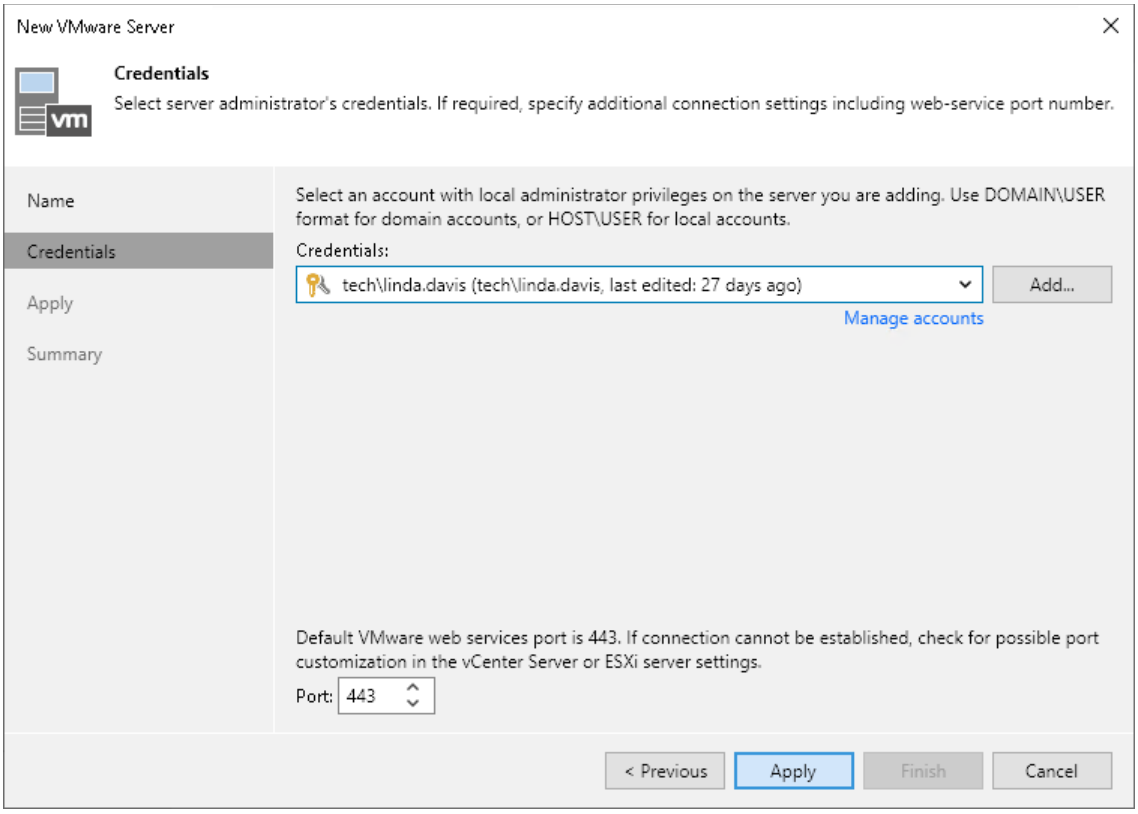


3. When you add a vCenter Server or ESXi host, Veeam Backup & Replication saves to the configuration database a thumbprint of the TLS certificate installed on the vCenter Server or ESXi host. During every subsequent connection to the server, Veeam Backup & Replication uses the saved thumbprint to verify the server identity and avoid the man-in-the-middle attack. For details on managing TLS Certificates, see Backup Server Certificate.

   If the certificate installed on the server is not trusted, Veeam Backup & Replication displays a warning.

   - To view detailed information about the certificate, click **View**.
   - If you trust the server, click **Continue**.
   - If you do not trust the server, click **Cancel**. Veeam Backup & Replication will display an error message, and you will not be able to connect to the server.

**NOTE**

When you update a certificate on a server, this server becomes unavailable in the Veeam Backup & Replication console. To make the server available again, acknowledge the new certificate at the **Credentials** step of the Edit Server wizard.

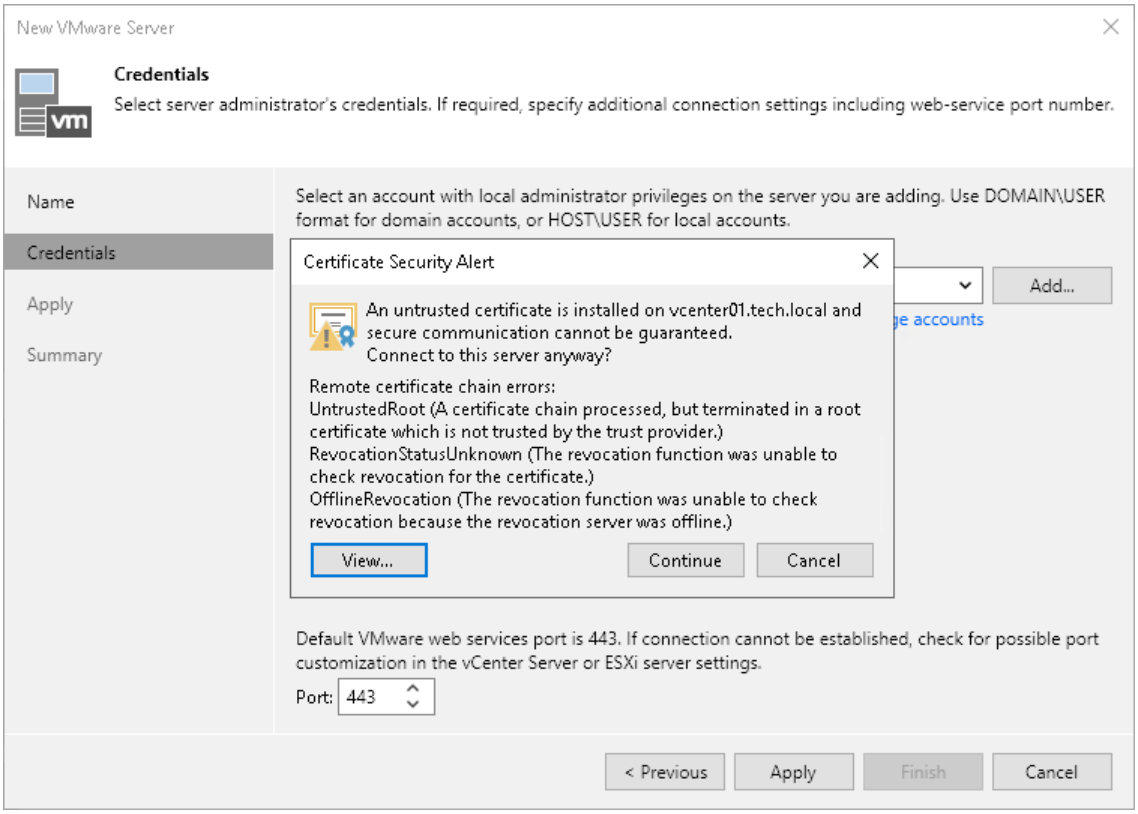

# Step 4. Apply Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all the required components. Click **Next** to complete the adding of the server.

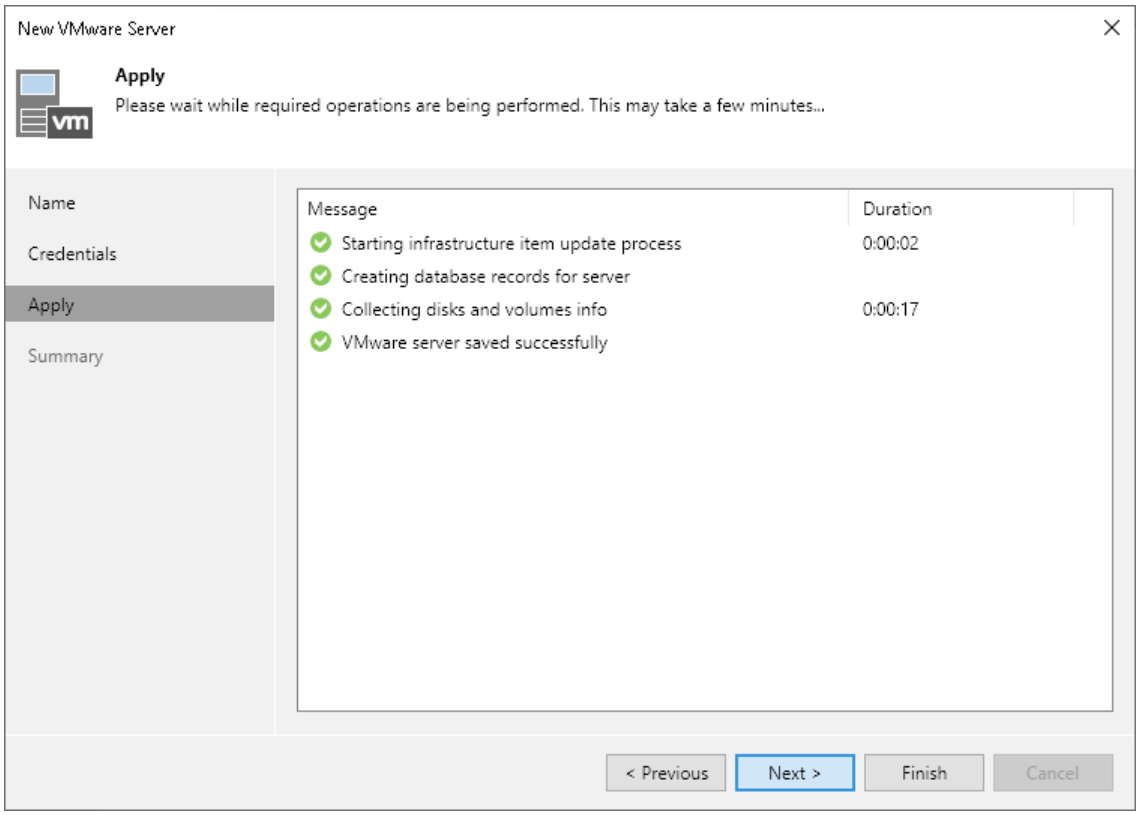

# Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the VMware vSphere server and click **Finish** to exit the wizard.

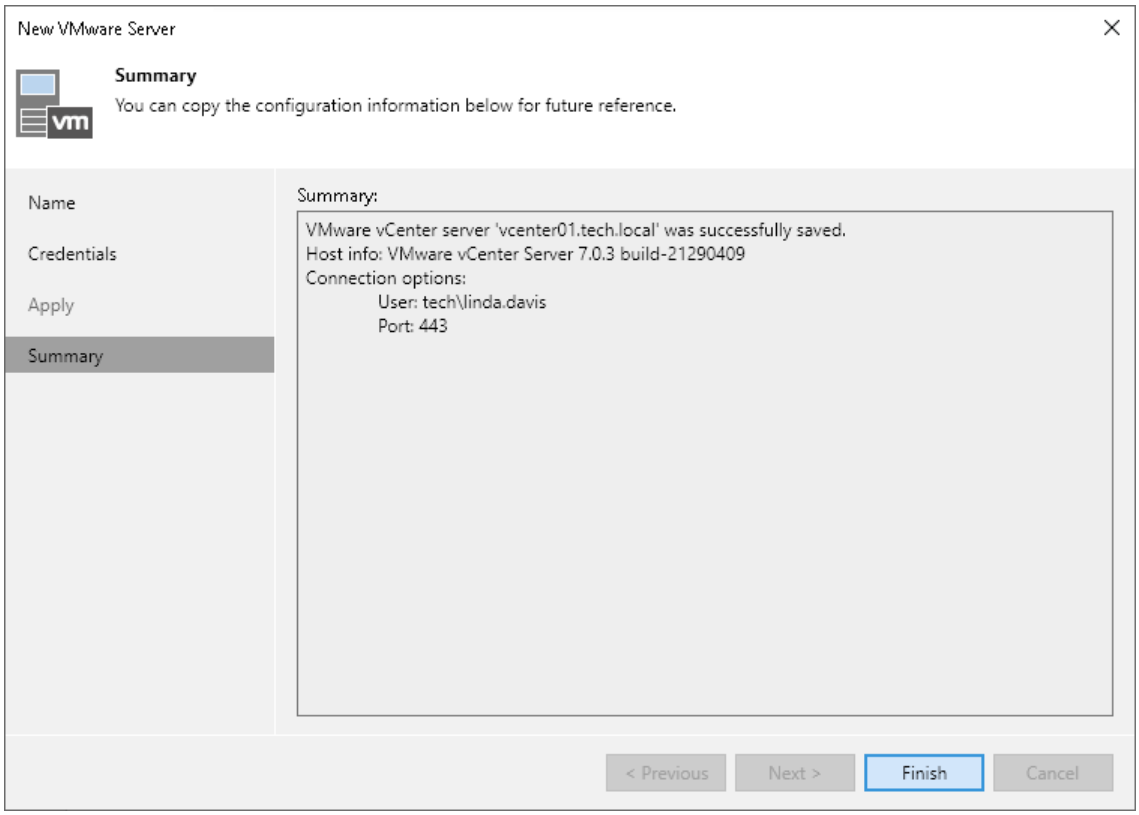

# Adding VMware Cloud Director Servers

To work with vApps and VMs managed by VMware Cloud Director, you must add VMware Cloud Director to the backup infrastructure.

When you add VMware Cloud Director to the backup infrastructure, the VMware Cloud Director hierarchy is displayed under the **Cloud Director** view. You can work with VMs managed by VMware Cloud Director directly in the Veeam Backup & Replication console. For more information on how to open the **Cloud Director** view, see Viewing VMware Cloud Director VMs.

To add the VMware Cloud Director server, use the **New VMware Cloud Director** wizard.

# Step 1. Launch New VMware Cloud Director Server Wizard

To launch the **New VMware Cloud Director Server** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Managed Servers** node and select **Add Server**. Alternatively, you can click **Add Server** on the ribbon.
3. In the **Add Server** window, click **Virtualization Platforms** > **VMware vSphere** > **Cloud Director**.

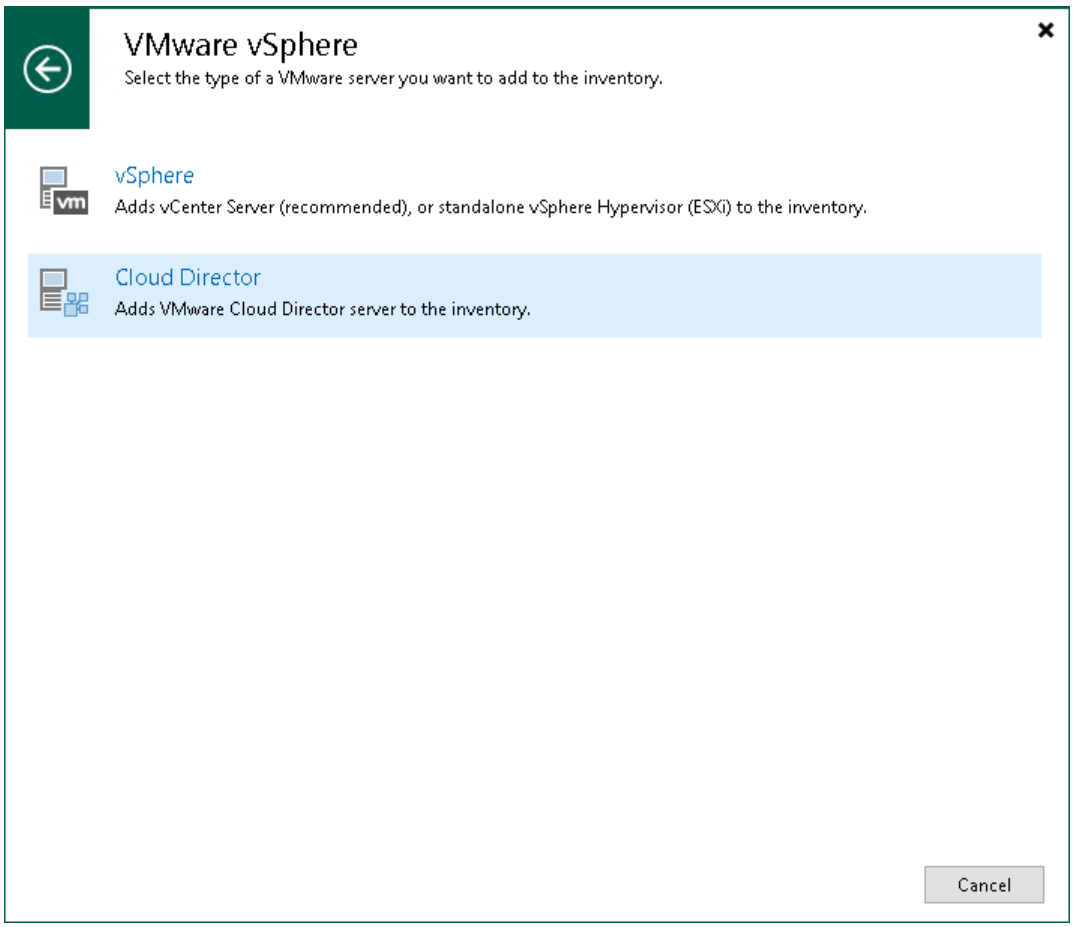

# Step 2. Specify Server Name or Address

At the **Name** step of the wizard, specify connection settings for VMware Cloud Director. If the VMware Cloud Director infrastructure comprises several cells, you can specify connection settings for any cell in the VMware Cloud Director hierarchy.

1. In the **DNS name or IP address** field, enter a full DNS name, or IPv4 or IPv6 address of the VMware Cloud Director server or any cell in the VMware Cloud Director infrastructure. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
2. In the **URL** field, enter a URL of the VMware Cloud Director server. By default, Veeam Backup & Replication uses the following URL: `https://<clouddirectorservername>:443`, where `<clouddirectorservername>` is the name or IP address of the VMware Cloud Director server that you have specified in the field above and `443` is the default port for communication with VMware Cloud Director.
3. In the **Description** field, provide a description for future reference. The default description contains information about the user who added the server, date and time when the server was added.

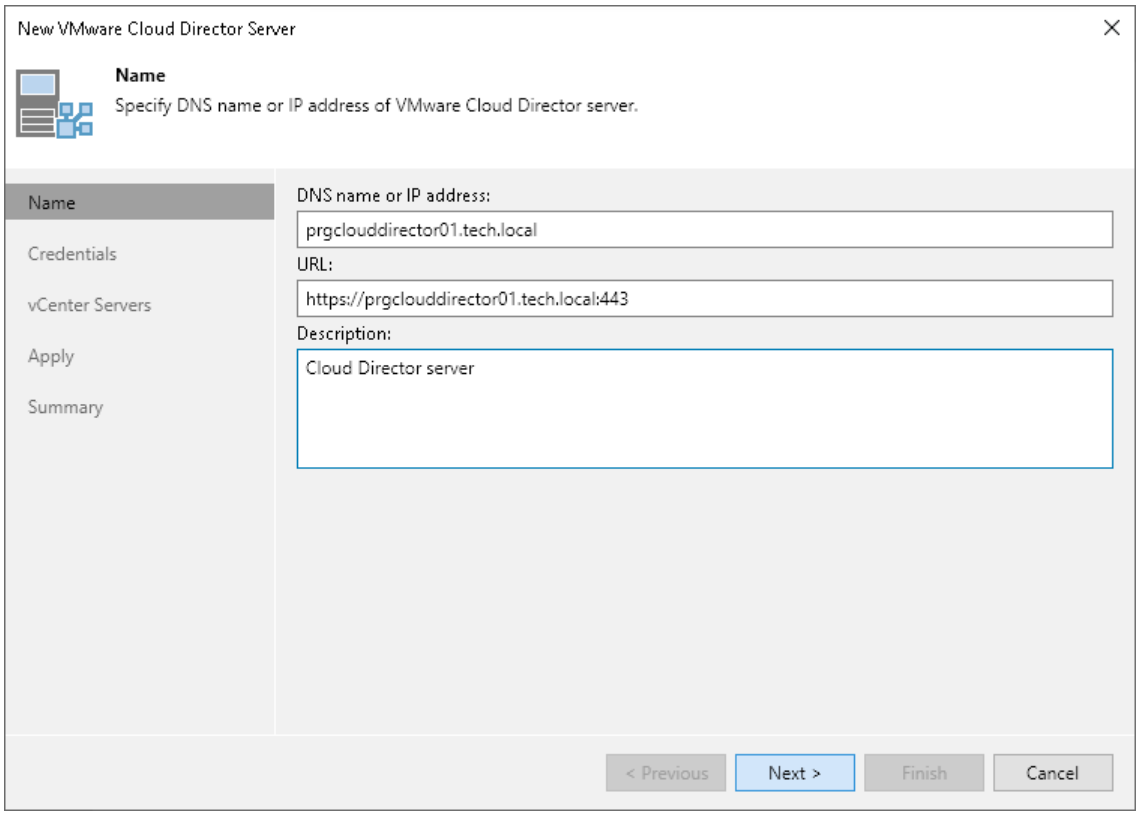

# Step 3. Specify VMware Cloud Director Credentials

At the **Credentials** step of the wizard, specify credentials to connect to VMware Cloud Director.

From the **Credentials** list, select credentials for the account that has permissions described in Permissions. If you have not set up credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Credentials Manager.

When you add a Cloud Director server, Veeam Backup & Replication saves a thumbprint of the TLS certificate installed on Cloud Director to the configuration database. During every subsequent connection to the server, Veeam Backup & Replication uses the saved thumbprint to verify the server identity and avoid the man-in-the-middle attack. For details on managing TLS Certificates, see Backup Server Certificate.

If the certificate installed on the server is not trusted, Veeam Backup & Replication displays a warning.

- To view detailed information about the certificate, click **View**.
- If you trust the server, click **Continue**.
- If you do not trust the server, click **Cancel**. Veeam Backup & Replication will display an error message, and you will not be able to connect to the server.

> **NOTE**
>
> When you update a certificate on a server, this server becomes unavailable in the Veeam Backup & Replication console. To make the server available again, acknowledge the new certificate at the **Credentials** step of the Edit Server wizard.

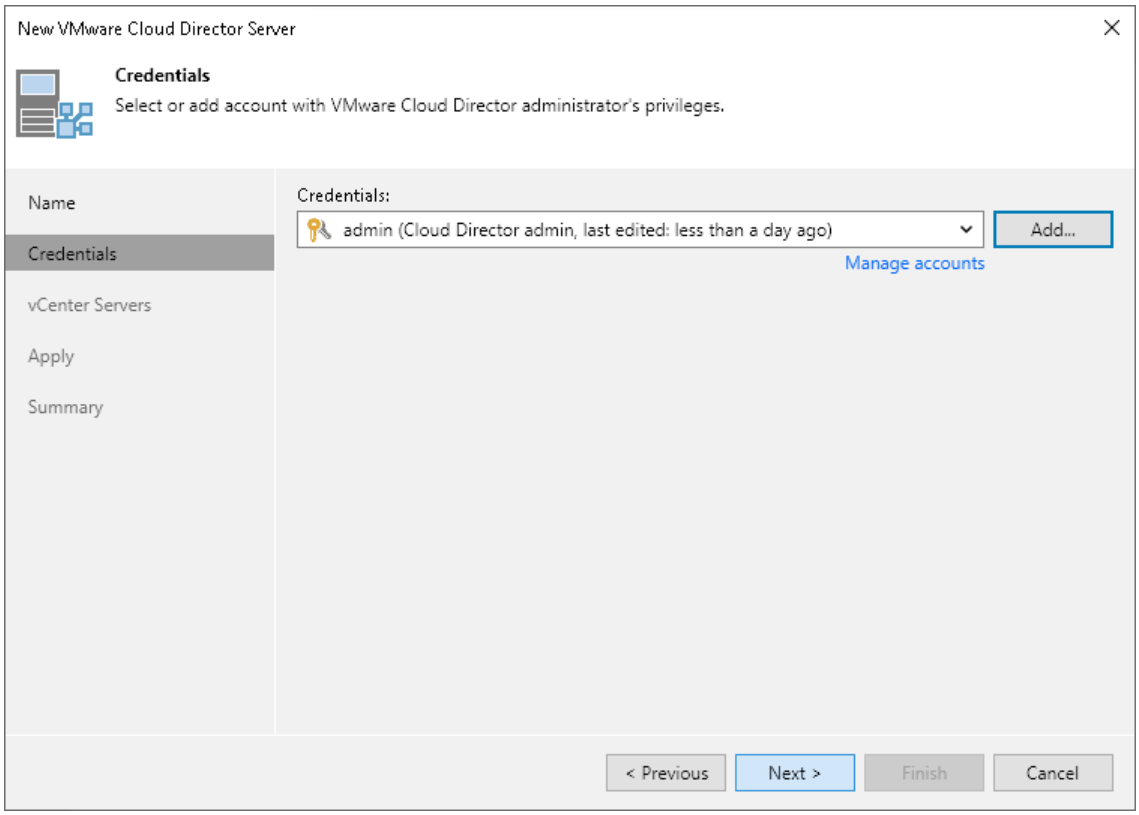

# Step 4. Specify Credentials for Underlying vCenter Servers

At the **vCenter Servers** step of the wizard, specify credentials for every vCenter Server added to VMware Cloud Director. If the vCenter Server is already added to the backup infrastructure, you do not need to specify credentials for it once again. Veeam Backup & Replication will automatically detect the credentials you provided when adding this vCenter Server and use them.

1. From the **vCenter servers** list, select a vCenter Server.
2. Click **Account** on the right and select credentials to connect to the vCenter Server. By default, Veeam Backup & Replication uses the same credentials that you have specified for VMware Cloud Director at the previous step of the wizard.

   If you have not set up the credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Credentials Manager.
3. Veeam Backup & Replication automatically detects a port used to communicate with the vCenter Server. If necessary, you can change the connection port for the vCenter Server. Click **vCenter** on the right and adjust the port number.

When you add a vCenter Server, Veeam Backup & Replication saves a thumbprint of the TLS certificate installed on the vCenter Server to the configuration database. During every subsequent connection to the server, Veeam Backup & Replication uses the saved thumbprint to verify the server identity and avoid the man-in-the-middle attack. For details on managing TLS Certificates, see Backup Server Certificate.

If the certificate installed on the server is not trusted, Veeam Backup & Replication displays a warning.

- To view detailed information about the certificate, click **View**.
- If you trust the server, click **Continue**.
- If you do not trust the server, click **Cancel**. Veeam Backup & Replication will display an error message, and you will not be able to connect to the server.

> **NOTE**
>
> When you update a certificate on a server, this server becomes unavailable in the Veeam Backup & Replication console. To make the server available again, acknowledge the new certificate at the **Credentials** step of the Edit Server wizard.

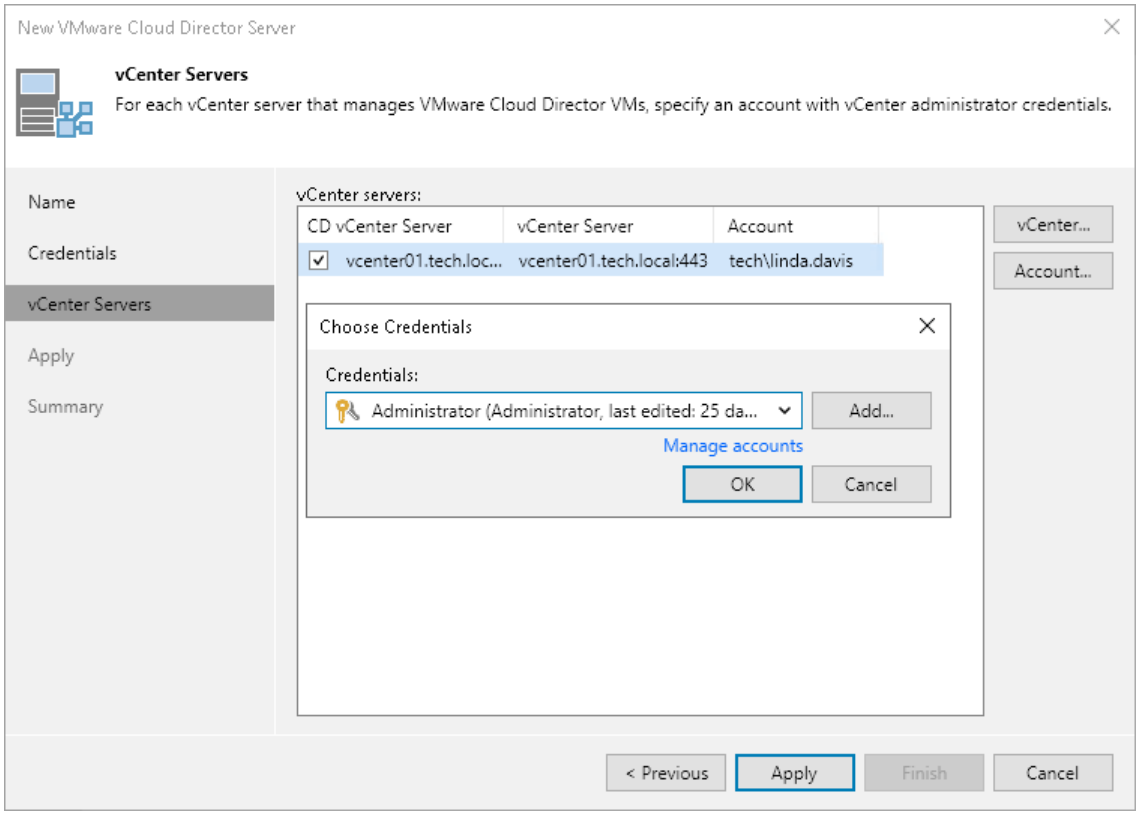
New VMware Cloud Director Server
vCenter Servers
For each vCenter server that manages VMware Cloud Director VMs, specify an account with vCenter administrator credentials.
Name
Credentials
vCenter Servers
Apply
Summary
vCenter servers:
CD vCenter Server
vCenter Server
Account
vcenter01.tech.loc...
vcenter01.tech.local:443
tech\linda.davis
vCenter...
Account...
Choose Credentials
Credentials:
Administrator (Administrator, last edited: 25 da...
Add...
Manage accounts
OK
Cancel
< Previous
Apply
Finish
Cancel

# Step 5. Finish Working with Wizard

At the **Apply** step of the wizard, complete the procedure of VMware Cloud Director adding.

1. Review details of VMware Cloud Director.
2. Click **Next**, and then click **Finish** to exit the wizard.

If vCenter Servers underlying VMware Cloud Director are already added to the backup infrastructure, they will not be added for the second time. Veeam Backup & Replication will create associations with the vCenter Servers and display them in the VMware Cloud Director hierarchy.

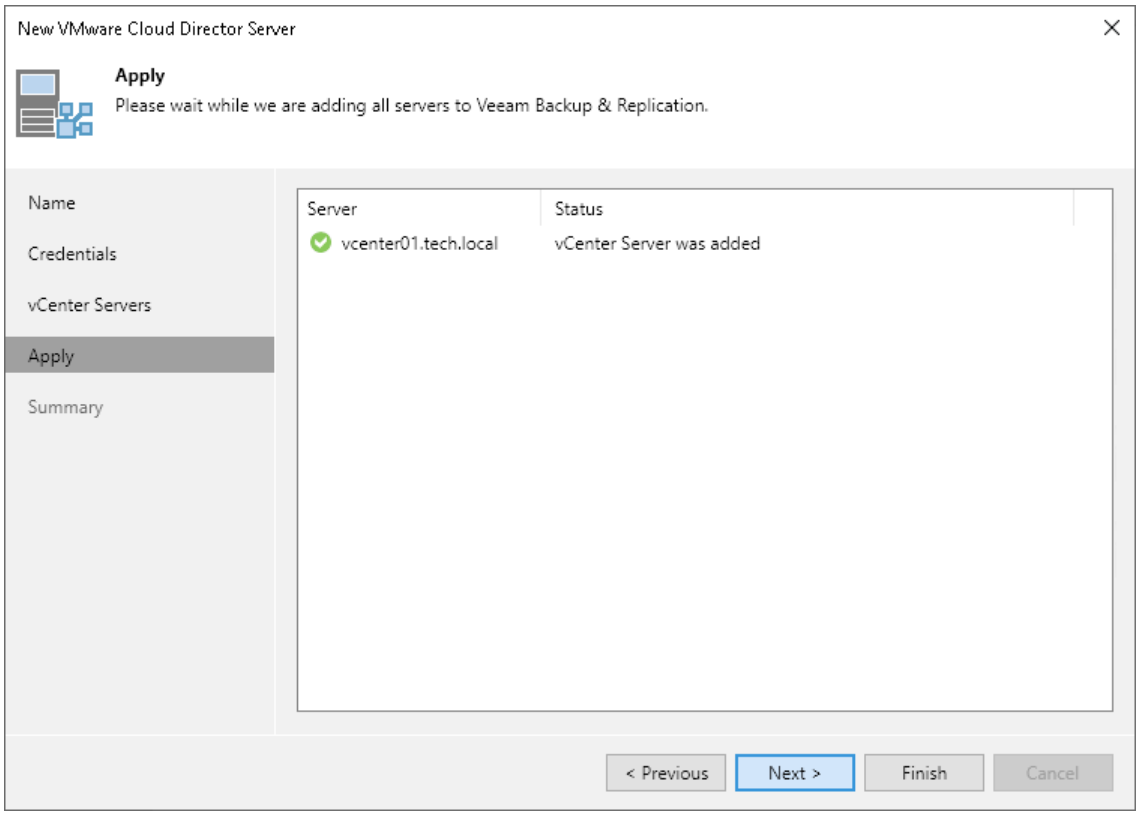

# Adding Microsoft Windows Servers

You must add to the backup infrastructure Microsoft Windows servers that you plan to use as backup infrastructure components and servers that you plan to use for various types of restore operations.

Before you add a Microsoft Windows server, check prerequisites. Then use the **New Windows Server** wizard to add the server.

## Before You Begin

Before you add a Microsoft Windows server to the backup infrastructure, check network connection settings of this server.

- Check permissions required to add the server. For more information, see Permissions.
- File and printer sharing must be enabled in network connection settings of the added Microsoft Windows server. On every connected Microsoft Windows server, Veeam Backup & Replication deploys two components:
  - Veeam Installer Service
  - Veeam Data Mover

  If file and printer sharing is not enabled, Veeam Backup & Replication will fail to deploy these components.

# Step 1. Launch New Windows Server Wizard

To launch the **New Windows Server** wizard, do one of the following:

- Open the **Backup Infrastructure** or **Files** view, in the inventory pane select the **Microsoft Windows** node and click **Add Server** on the ribbon.
- Open the **Backup Infrastructure** view. In the inventory pane, select the **Managed Servers** node and click **Add Server** on the ribbon or right-click the **Managed Servers** node and select **Add Server.** In the **Add Server** window, select **Microsoft Windows**.

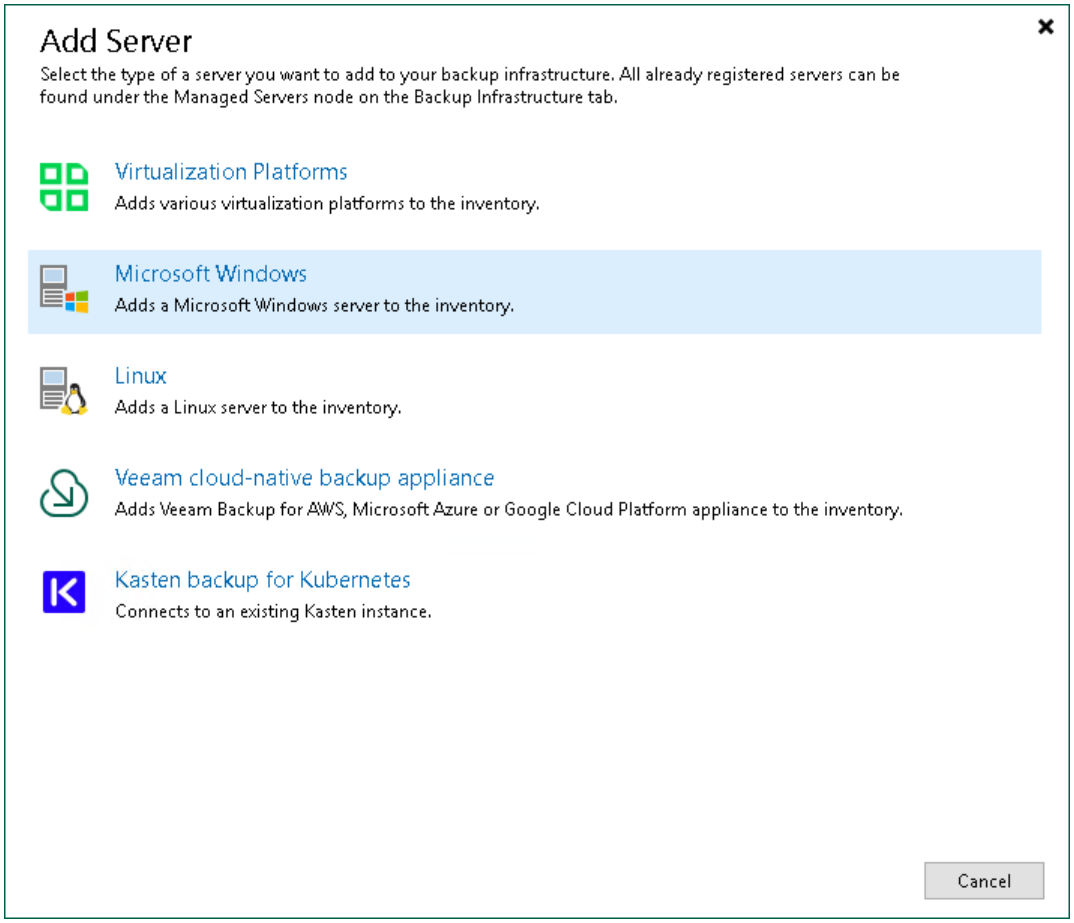

# Step 2. Specify Server Name or Address

At the **Name** step of the wizard, specify an address and description for the Microsoft Windows server.

1. Enter a full DNS name, or IPv4 or IPv6 address of the Microsoft Windows server. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
2. Provide a description for future reference. The default description contains information about the user who added the server, date and time when the server was added.

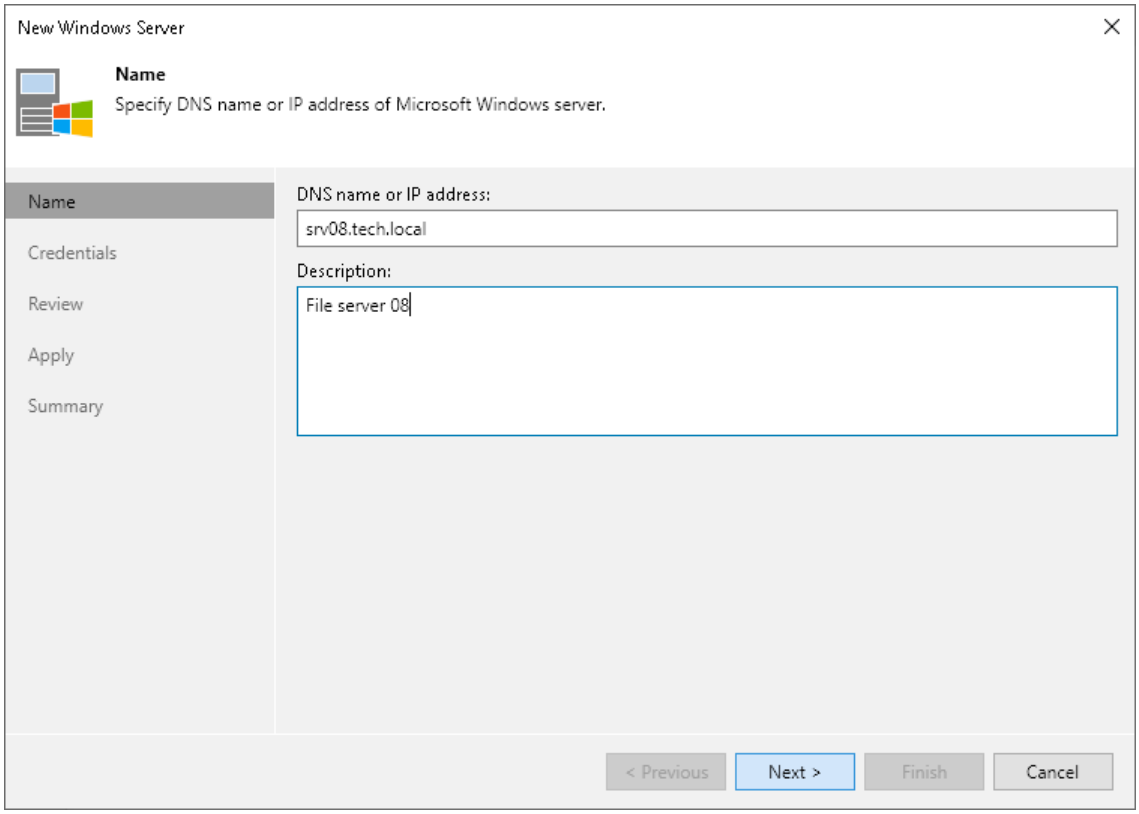

# Step 3. Specify Credentials

At the **Credentials** step of the wizard, specify credentials for the Microsoft Windows server.

1. From the **Credentials** list, select credentials for the account that has privileges described in section Permissions.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add the credentials. For more information, see Credentials Manager.

   Veeam Backup & Replication will use the provided credentials to deploy its components on the added server.

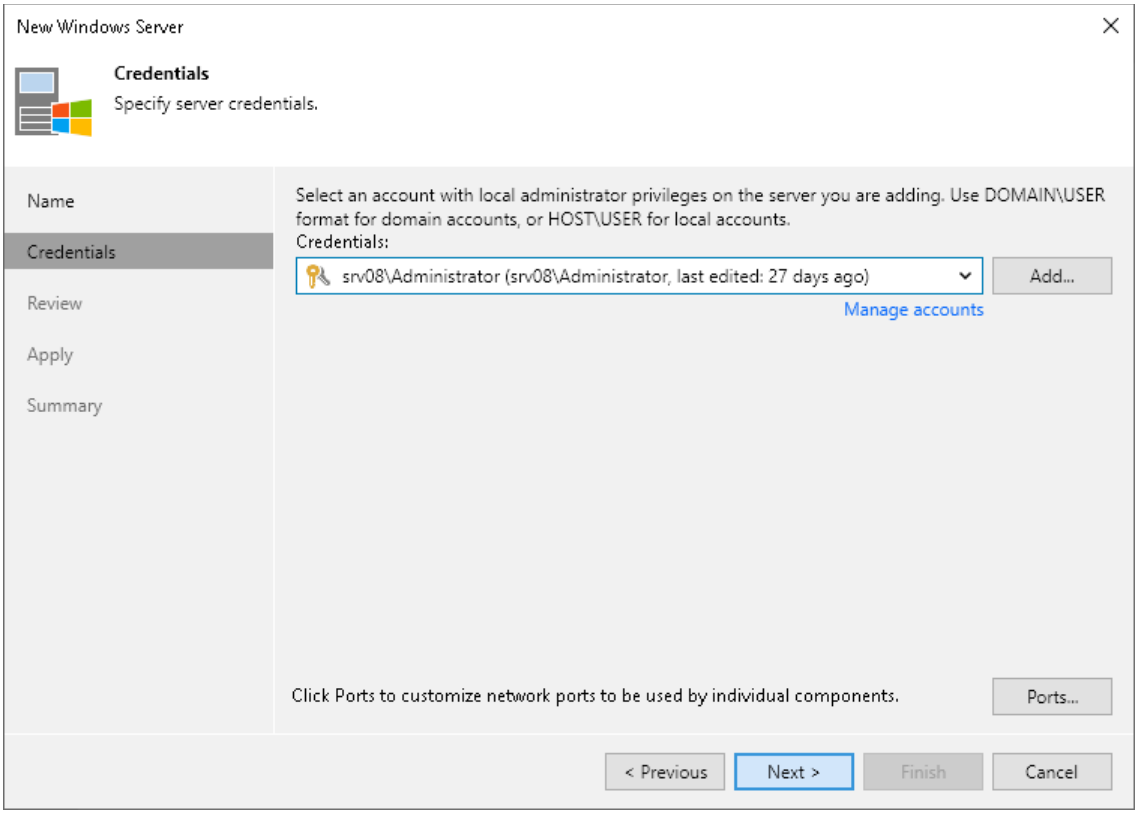


2. To customize network ports used by Veeam Backup & Replication components, click **Ports**. For default ports used by the Veeam Backup & Replication components, see Ports.

   If necessary, adjust the port numbers.

3. In the **Data transfer options** section of the **Network Settings** window, specify connection settings for file copy operations. Provide a range of ports that will be used as transmission channels between the source server and target server (one port per task). By default, Veeam Backup & Replication uses port range 2500-3300.

   If the virtual environment is not large and data traffic will not be significant, you can specify a smaller range of ports, for example, 2500-2509 to run 10 concurrent tasks at the same time. Note that Veeam Backup & Replication processes each VM disk as a separate task.

4. [For Microsoft Windows server deployed outside NAT] In the **Preferred TCP connection role** section select the **Run server on this side** check box. In the NAT scenario, the outside client cannot initiate a connection to the server on the NAT network. As a result, services that require initiation of the connection from outside can be disrupted. With this option selected, you will be able to overcome this limitation and initiate a "client-server" connection — that is, a connection in the direction of the Microsoft Windows server.

   The option applies if one of the following roles is assigned to the server: source VMware backup proxy in backup or replication scenarios, source repository in the backup copy scenario.

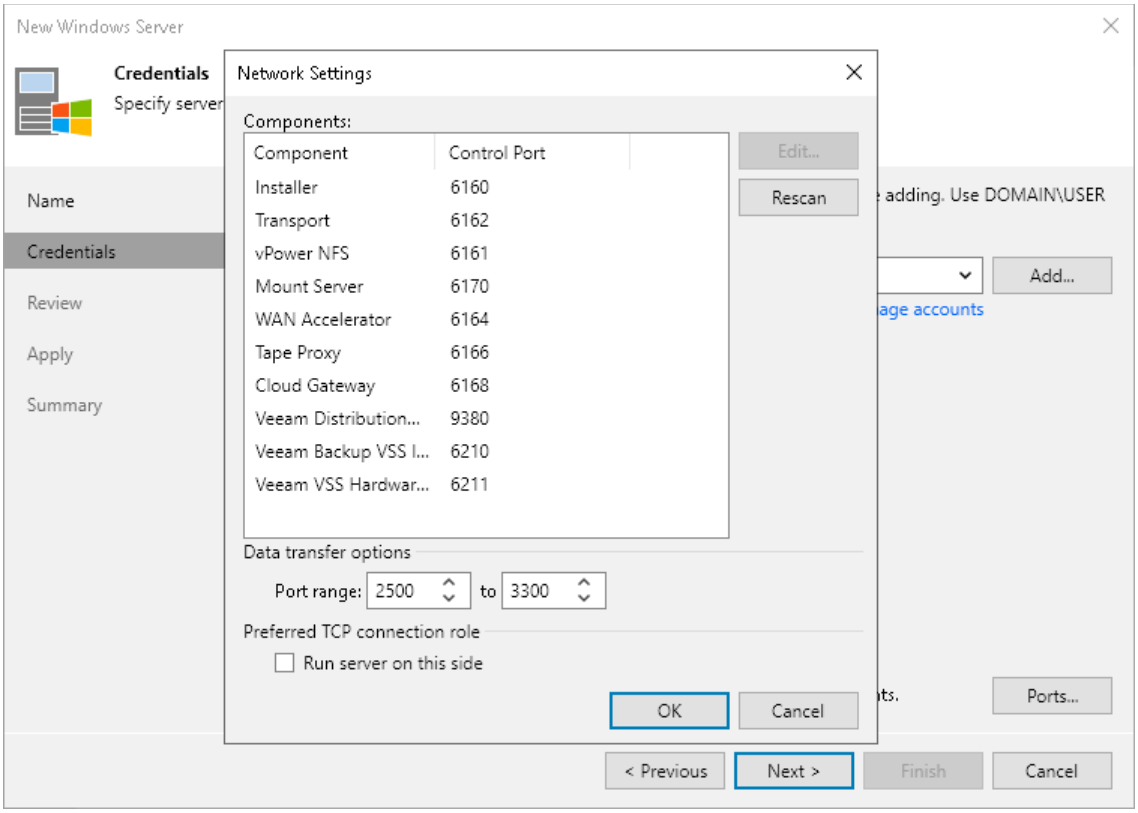

# Step 4. Review Components

At the **Review** step of the wizard, review what Veeam Backup & Replication components are already installed on the server and what components will be installed.

1. Review the components.
2. Click **Apply** to add the Microsoft Windows server to the backup infrastructure.

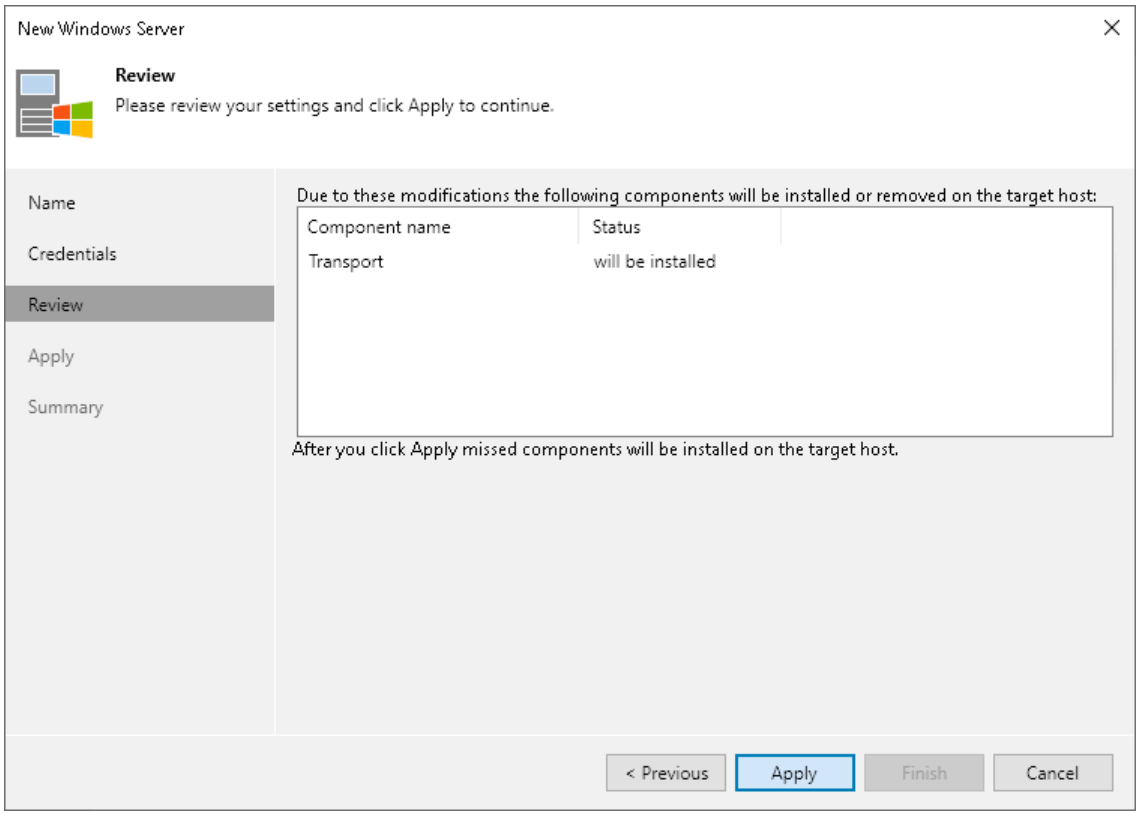

# Step 5. Apply Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all the required components. Click **Next** to complete the adding of the server.

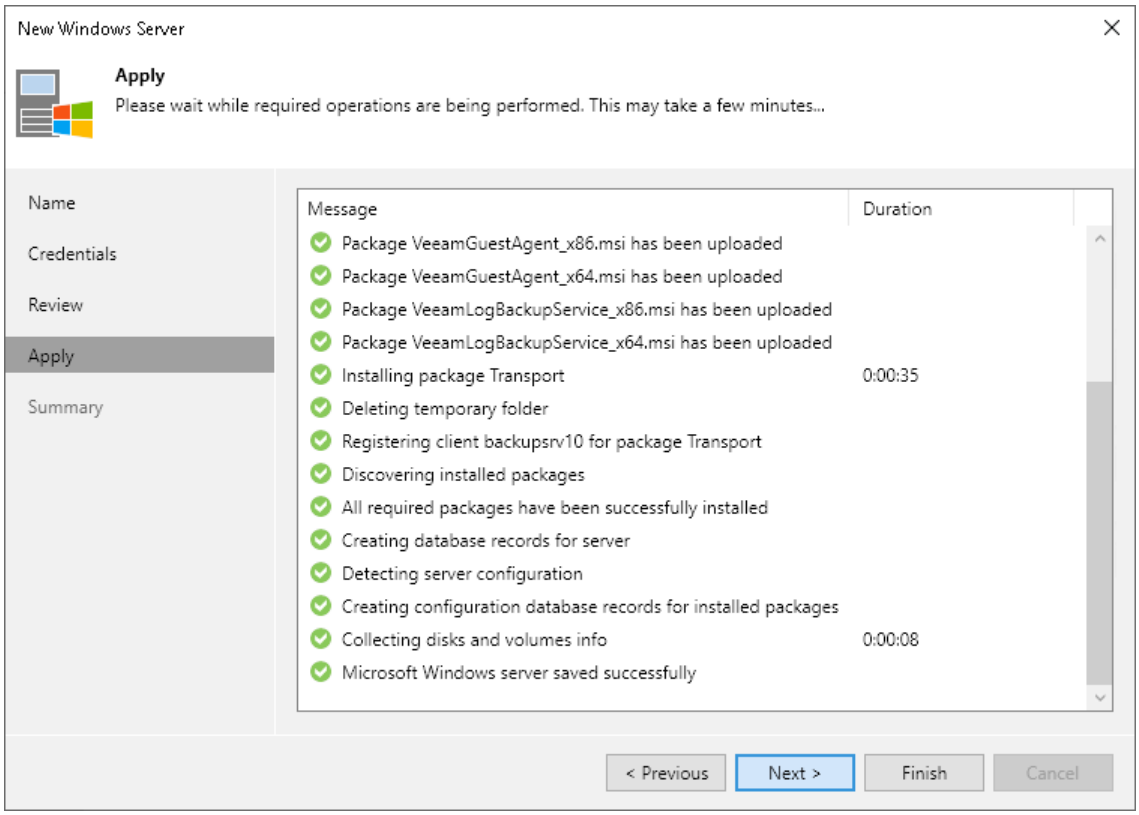

# Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of Microsoft Windows server adding.

1. Review details of the Microsoft Windows server.
2. Click **Next**, then click **Finish** to exit the wizard.

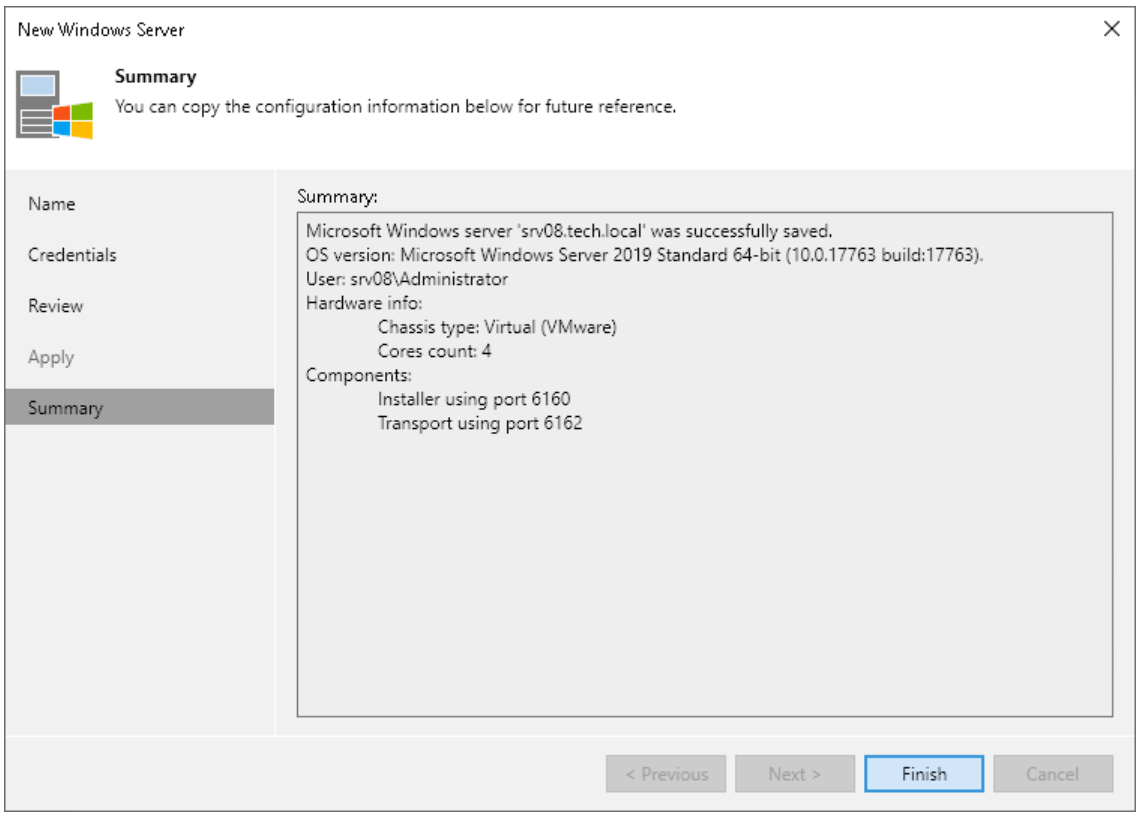

# Adding Linux Servers

You must add to the backup infrastructure Linux servers that you plan to use as backup infrastructure components and servers that you plan to use for various types of restore operations.

Before you add a Linux server, check prerequisites. Then use the **New Linux Server** wizard to add the server.

## Before You Begin

Before you add a Linux server to the Veeam Backup & Replication infrastructure, check the required permissions and the following prerequisites.

## Linux Firewalls

When you add a Linux server to the backup infrastructure, Veeam Backup & Replication automatically opens ports used by the Veeam Data Mover on the Linux server. Generally, Veeam Backup & Replication automatically open ports for most of popular firewalls (iptables, ufw, firewall-cmd). However, if for some reason the ports are not opened, you can open the ports manually. You can also specify these ports at the SSH Connection step of the **New Linux Server** wizard. Note that ports are opened dynamically: if 10 concurrent jobs are running, Veeam Backup & Replication opens ports 2500-2509.

If you use the `firewalld` tool, you can configure firewall rules to open ports only in necessary zones. By default, Veeam Backup & Replication opens ports in all active `firewalld` zones. If your firewall is configured for different zones, and you want to minimize security holes, you can configure Veeam Backup & Replication to open the ports only for certain zones. To do this, perform the following:

1. On the helper host or target Linux host, create the `/etc/VeeamNetConfig` file and define the following parameter:

   ```
   FirewalldZones=zone_name_1, zone_name_2
   ```

   where `zone_name_1, zone_name_2` is a list of zone names where the ports must be open. Veeam Backup & Replication will skip the zones that are not in this list.

2. [Only for helper host] If you select a Linux host that is already added to the Veeam Backup & Replication infrastructure, you should also add required zones to the `/opt/veeam/transport/VeeamTransportConfig` file.

   ```
   FirewalldZones=zone_name_1, zone_name_2
   ```

   **NOTE**

   Veeam Backup & Replication opens the port 2500 in all zones even if you have specified the required zones in configuration files.

## TLS Connection

Linux servers use the TLS connection. You can disable the TLS connection with a registry value for the servers that do not support the TLS connection. For more information, contact Veeam Customer Support.

# Step 1. Launch New Linux Server Wizard

To launch the **New Linux Server** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Managed Servers** node and select **Add Server**. Alternatively, you can click **Add Server** on the ribbon.
3. In the **Add Server** window, select **Linux**.

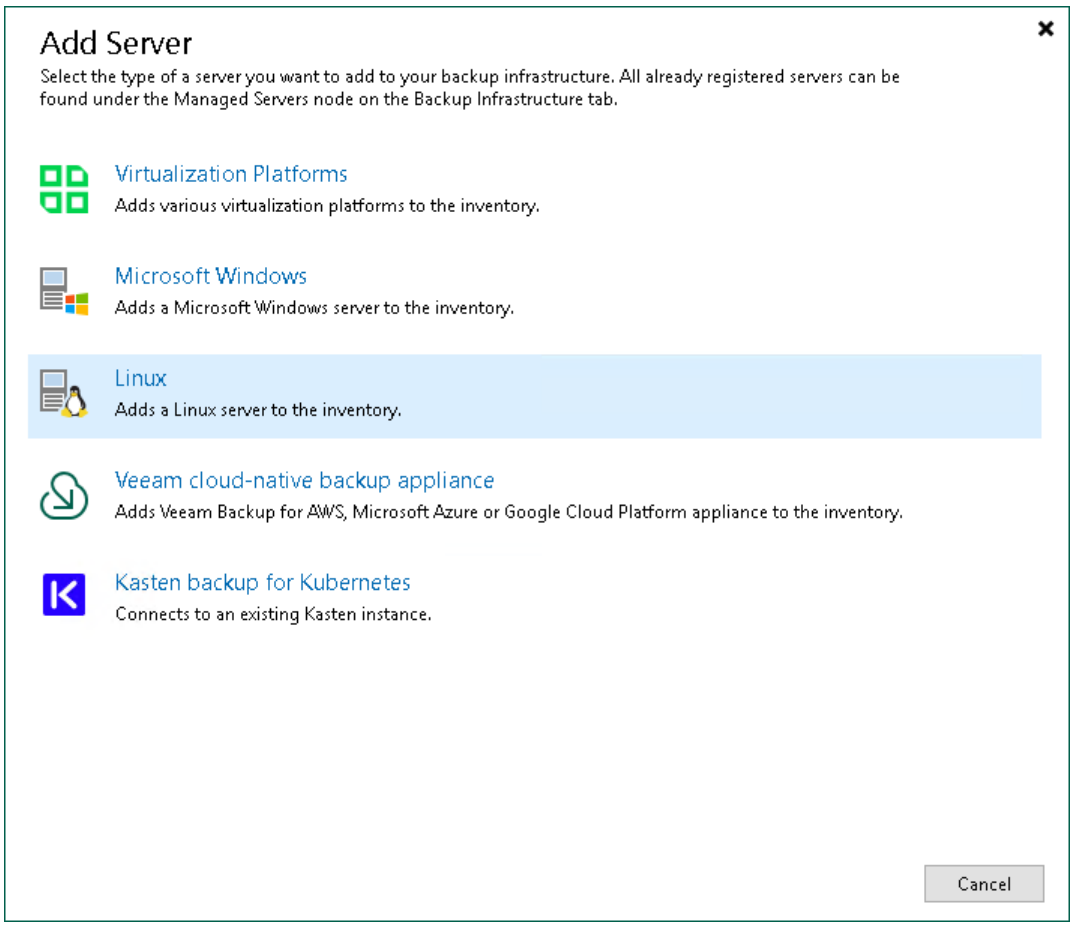

# Step 2. Specify Server Name or Address

At the **Name** step of the wizard, specify an address and description for the Linux server.

1. Enter a full DNS name,or IPv4 or IPv6 address of the Linux server. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
2. Provide a description for future reference. The default description contains information about the user who added the server, date and time when the server was added.

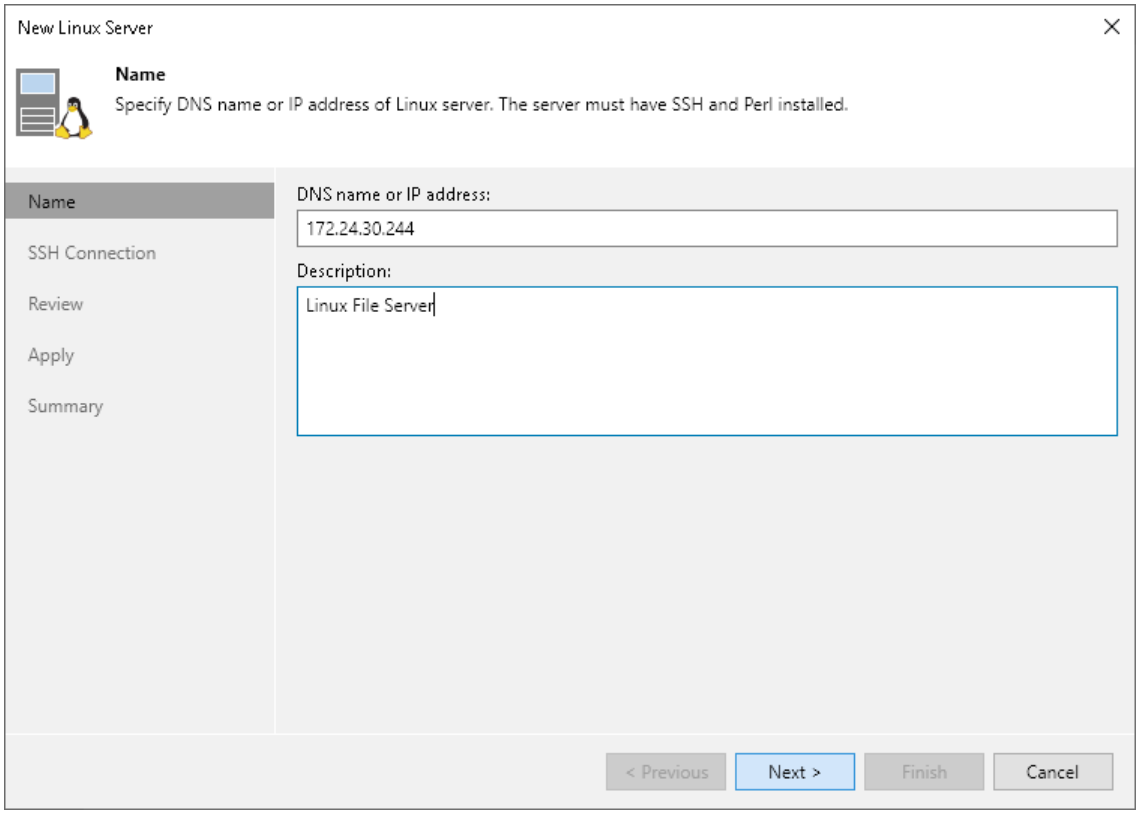

# Step 3. Specify Credentials and SSH Settings

At the **SSH Connection** step of the wizard, specify credentials for the Linux server and additional SSH connection settings.

1. From the **Credentials** list, select credentials for the account that has permissions described in section Permissions. You can select a credentials record that uses the password authentication method or credentials record that uses the Identity/Pubkey authentication method.

   > **NOTE**
   >
   > The account you selected must have the `home` directory created on the Linux server.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add the credentials. For more information, see Credentials Manager.

   To add a Linux server that you want to use as a hardened repository, click **Add** and select **Single-use credentials for hardened repository**. For more information about preparing a Linux server and setting up credentials, see Adding Hardened Repositories.

   > **NOTE**
   >
   > If you add a Linux server with single-use credentials, consider the following:
   >
   > - The folder with the repository must be accessible for accounts with user permissions (and not only root).
   > - SSH connection is necessary only for the deployment of Veeam Data Mover, or transport service. The transport service will be used to communicate with backup infrastructure components without the SSH connection. For security purposes, after you added the Linux server, you can disable SSH connection for the user account you use to connect to the Linux server. If you can work with the server from the console, disable SSH connection for the server itself.

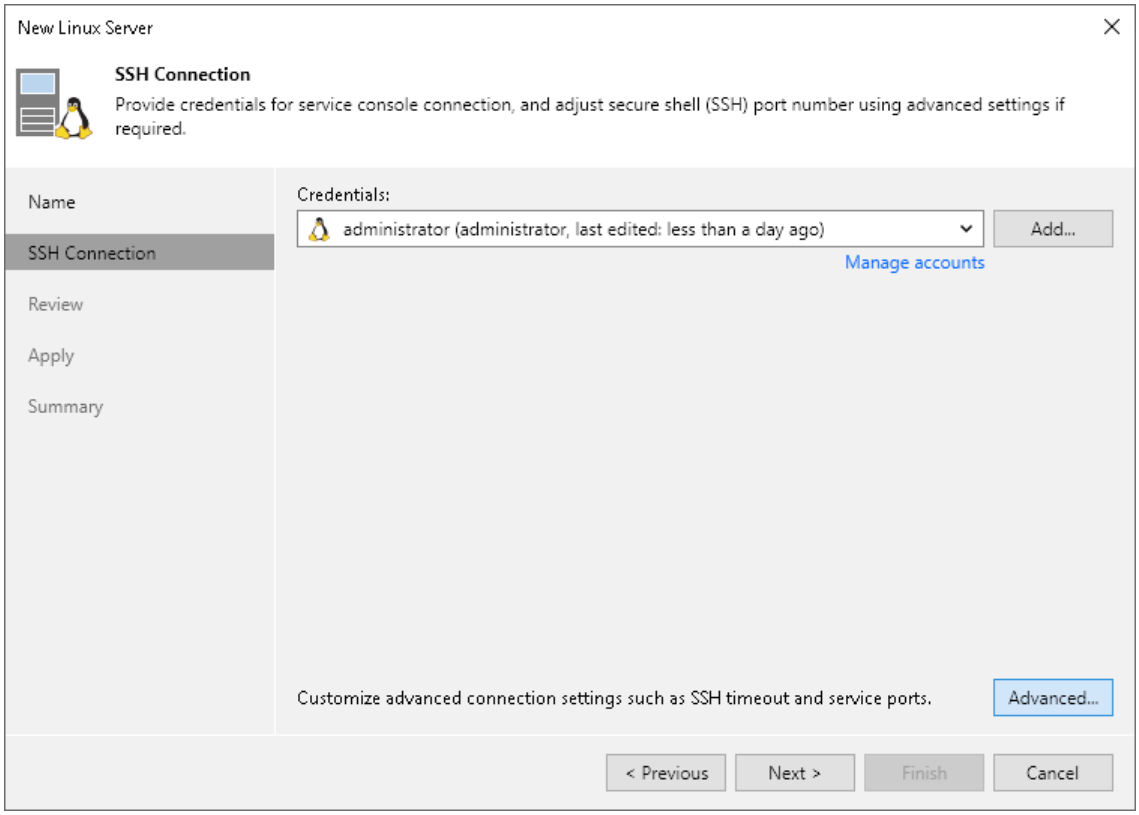

2. To configure SSH settings, click **Advanced**. This option becomes available after you have entered your credentials. In the **SSH Settings** window:

   a. In the **Service console connection** section, specify an SSH timeout. By default, the SSH timeout is set to 20000 ms. If a task targeted at the Linux server is inactive after the specified timeout, Veeam Backup & Replication will automatically terminate the task.

   b. In the **Data transfer options** section, specify connection settings for file copy operations. Provide a range of ports that will be used as transmission channels between the source host and target host (one port per task). By default, Veeam Backup & Replication uses port range 2500-3300. If the virtual environment is not large and data traffic will not be significant, you can specify a smaller range of ports, for example, 2500-2509 to run 10 concurrent tasks at the same time.

      Port 6162 is opened by default. It is a port used by Veeam Data Mover.

      > **NOTE**
      >
      > If you want to open these ports only for certain `firewalld` zones, you can specify the required zones in the configuration files. For instructions, see the Before You Begin section.

   c. [For Linux server deployed outside NAT] In the **Preferred TCP connection role** section, select the **Run server on this side** check box. In the NAT scenario, the outside client cannot initiate a connection to the server on the NAT network. As a result, services that require initiation of the connection from outside can be disrupted. With this option selected, you will be able to overcome this limitation and initiate a "client-server" connection — that is, a connection in the direction of the Linux server.

      The option applies if one of the following roles is assigned to the server: source VMware backup proxy in backup or replication scenarios, source repository in the backup copy scenario.

You can also change the SSH port over which you want to connect to the Linux server. For this, click the **Manage accounts** link and edit the account used to connect to the Linux server.

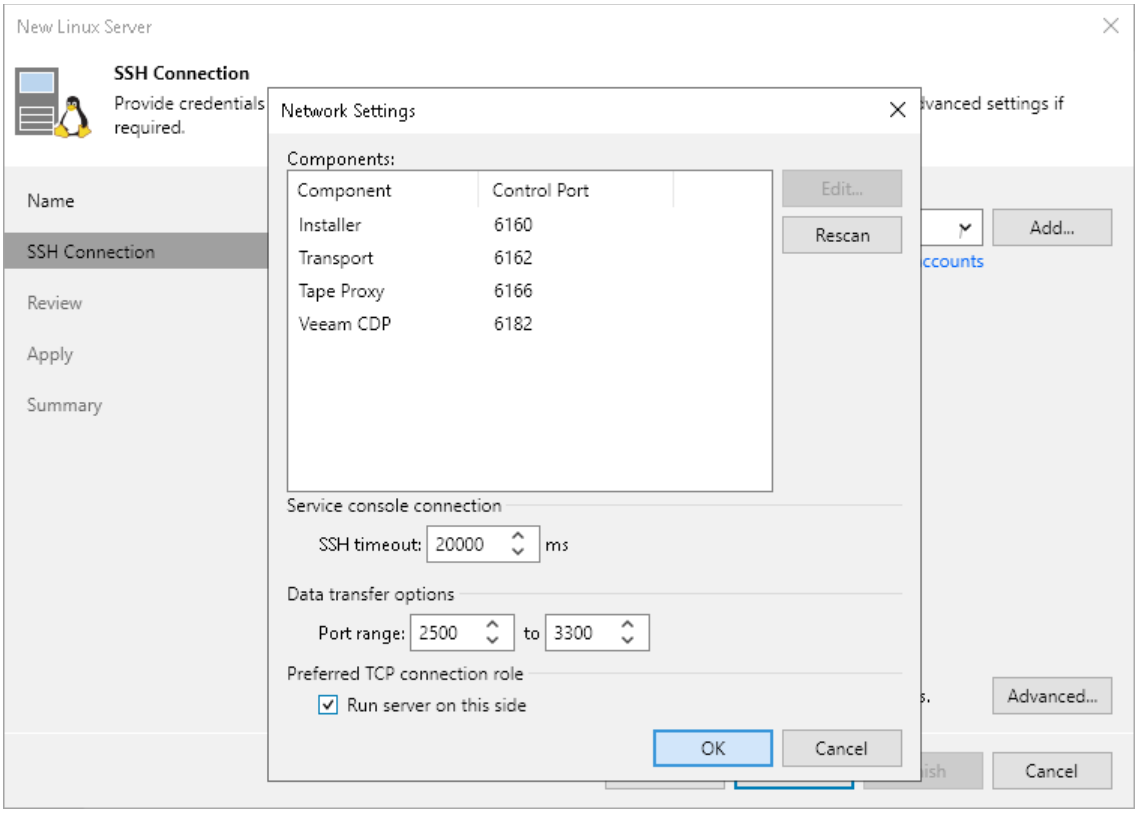

3. When you add a Linux server, Veeam Backup & Replication saves a fingerprint of the Linux host SSH key to the configuration database. During every subsequent connection to the server, Veeam Backup & Replication uses the saved fingerprint to verify the server identity and avoid the man-in-the-middle attack.

   To let you identify the server, Veeam Backup & Replication displays the SSH key fingerprint:

   - If you trust the server and want to connect to it, click **Yes**.
   - If you do not trust the server, click **No**. Veeam Backup & Replication will display an error message, and you will not be able to connect to the server.

> **NOTE**
>
> When you update an SSH key on a server, this server becomes unavailable in the Veeam Backup & Replication console. To make the server available again, acknowledge the new SSH key at the **SSH Connection** step of the Edit Server wizard.

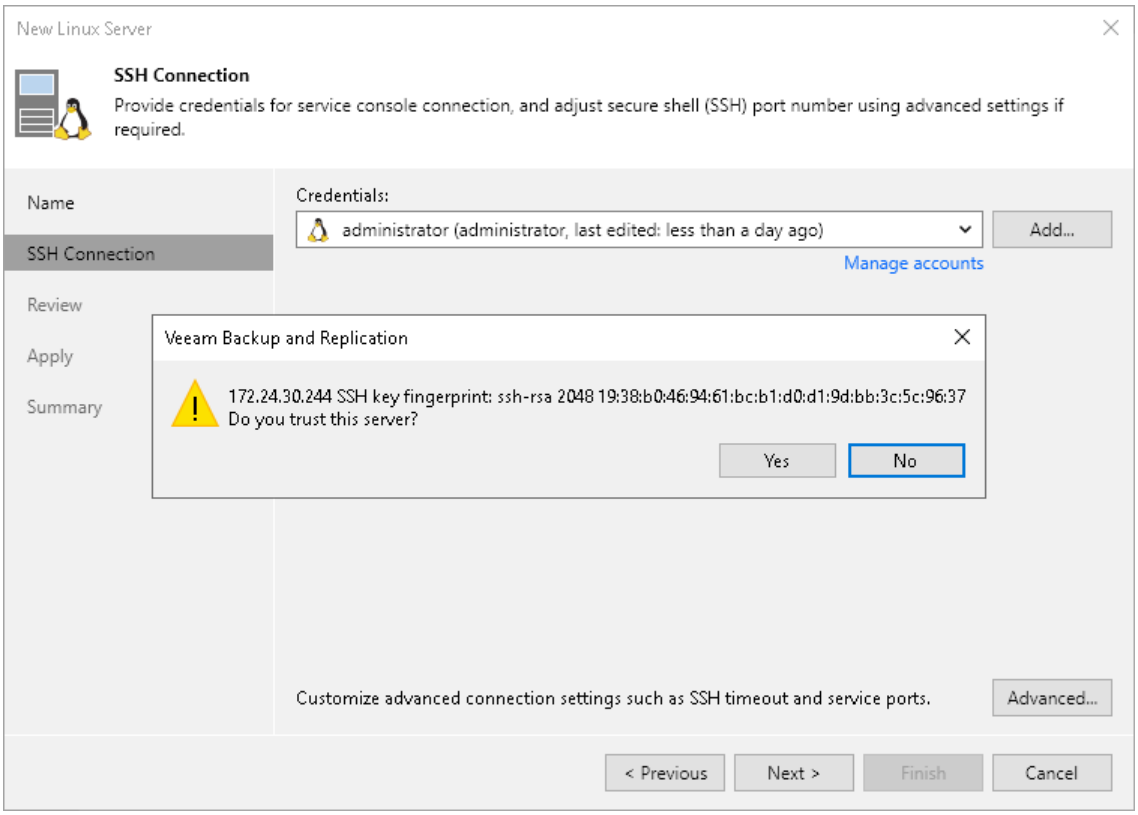

# Step 4. Review Components

At the **Review** step of the wizard, review what Veeam Backup & Replication components are already installed on the server and what components will be installed.

1. Review the components.
2. Click **Apply** to add the Linux server to the backup infrastructure.

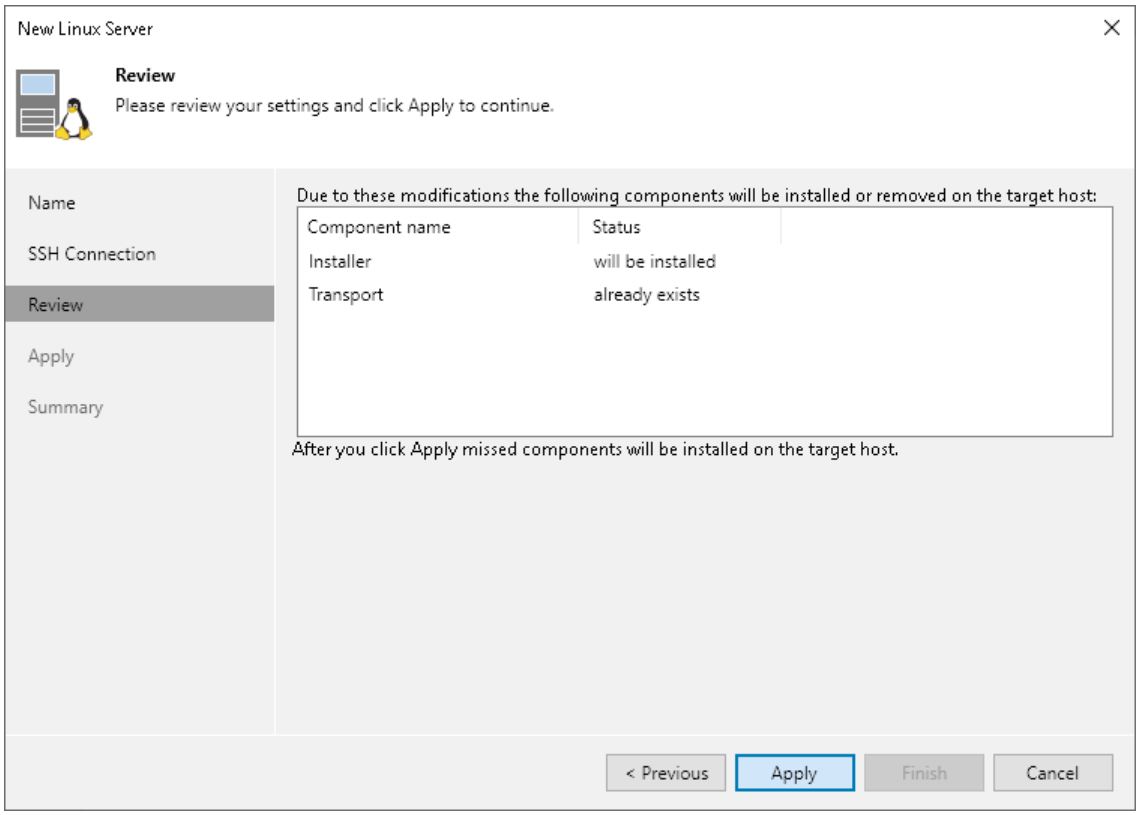

# Step 5. Apply Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all the required components. Click **Next** to complete the adding of the server.

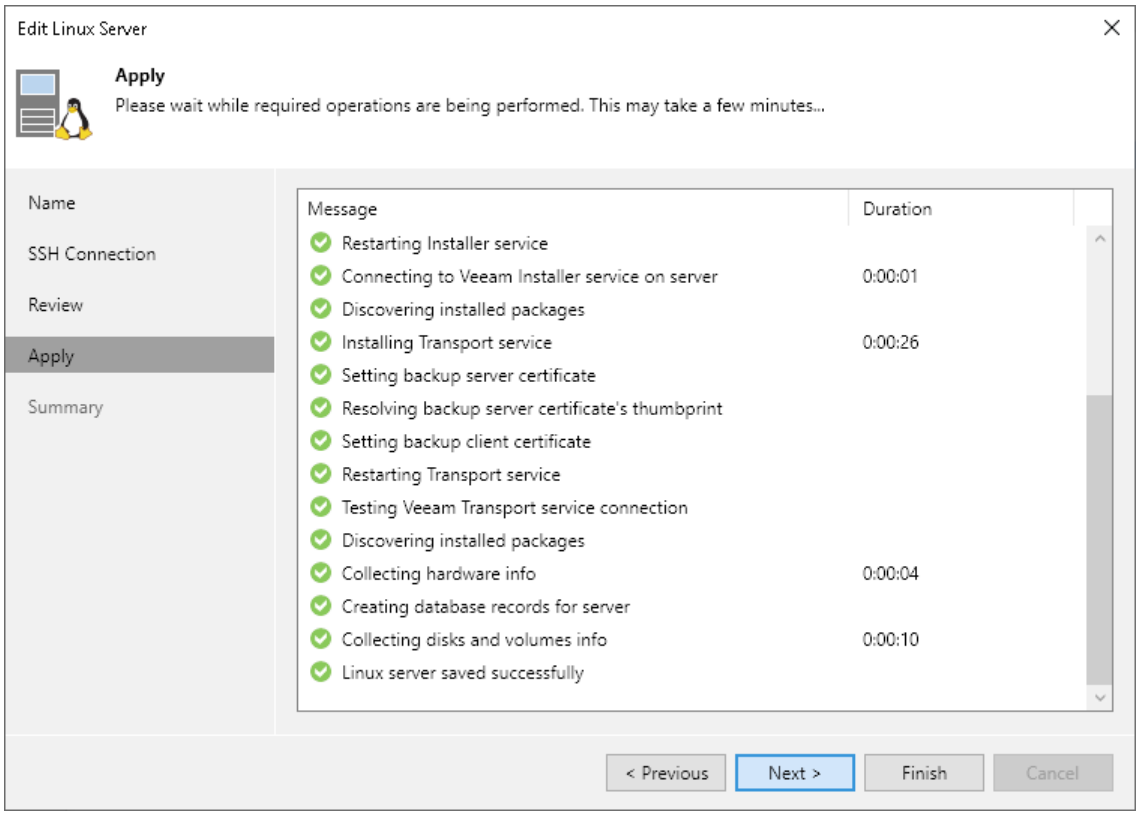

# Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of Linux server adding.

1. Review details of the Linux server.
2. Click **Next**, then click **Finish** to exit the wizard.

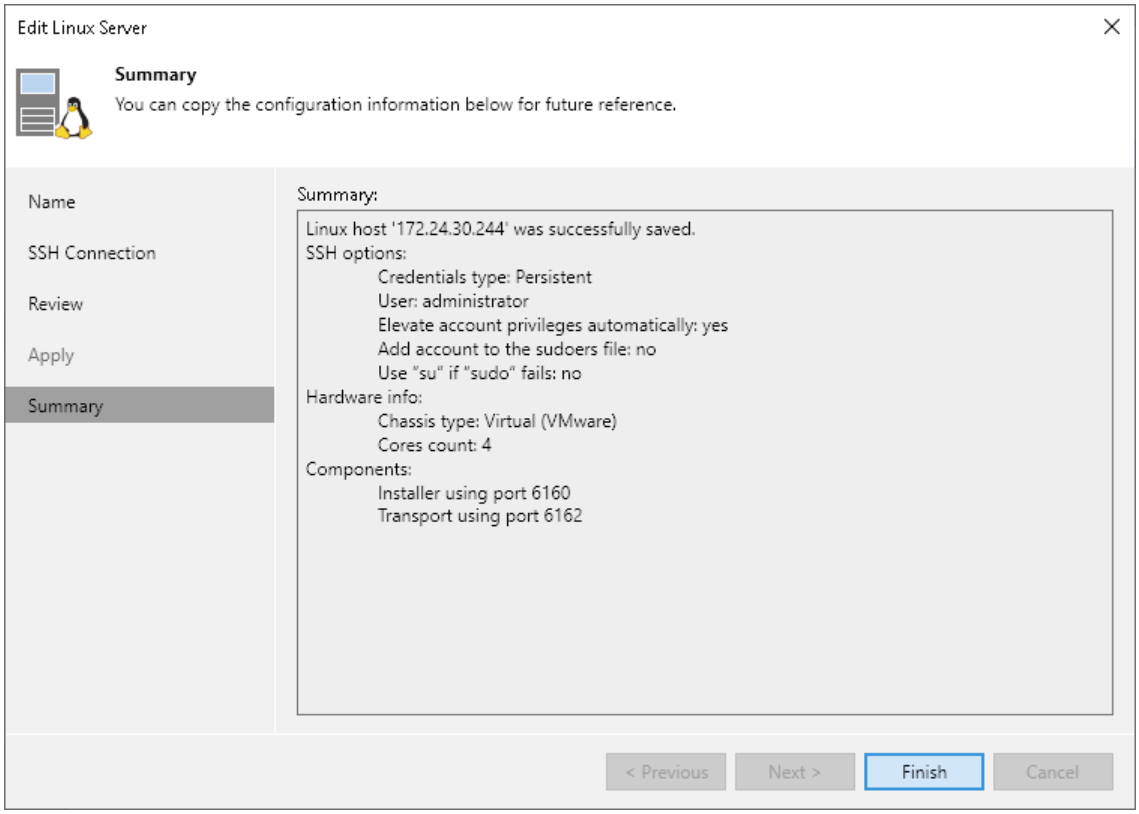

# Rescanning Servers

In some cases, you may need to rescan hosts or servers in the backup infrastructure. The rescan operation may be required if you have added or removed new disks and volumes to/from the host or server and want to display actual information in Veeam Backup & Replication. During the rescan operation, Veeam Backup & Replication retrieves information about disks and volumes that are currently connected to a host or server and stores this information to the configuration database.

Veeam Backup & Replication automatically performs a rescan operation every 4 hours. You can also start the rescan operation manually:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed servers**.
3. In the working area, select the server or host and click **Rescan** on the ribbon. Alternatively, you can right-click the server or host and select **Rescan**.

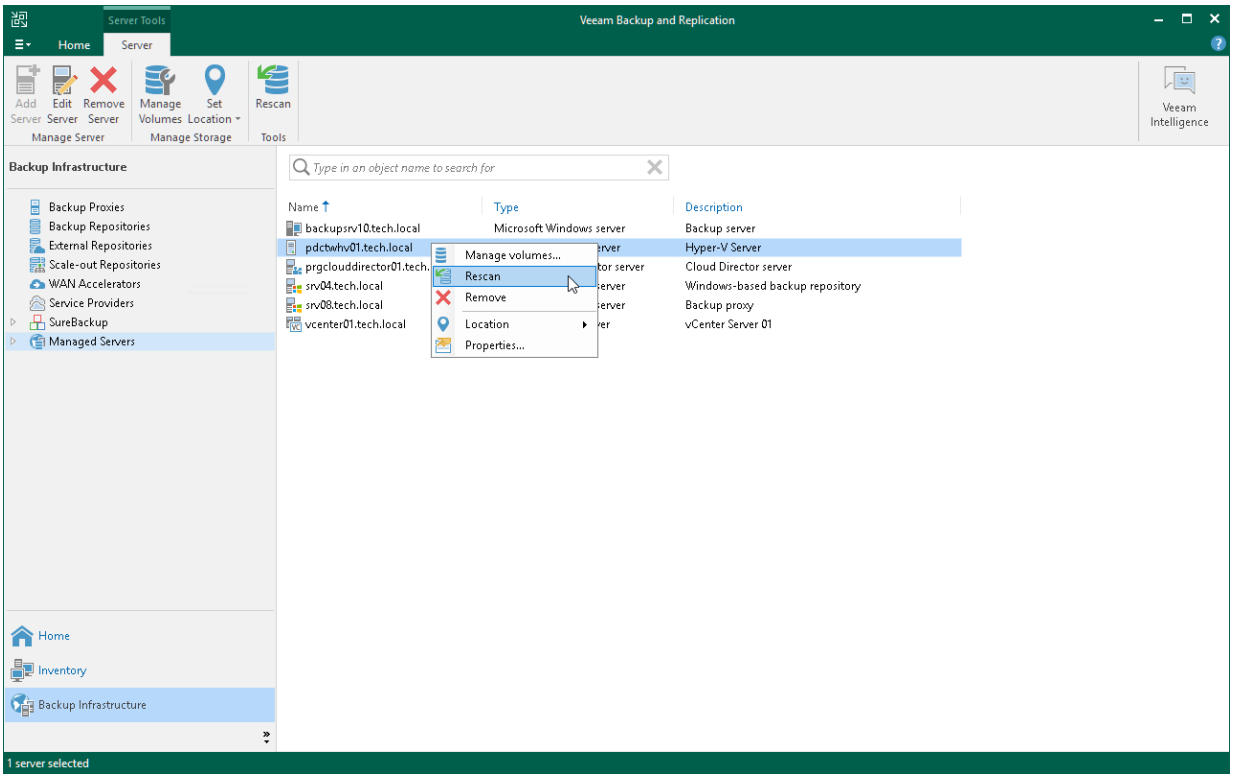

# Editing Server Settings

To edit settings of a server in the backup infrastructure:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed servers**.
3. In the working area, select the server and click **Edit Server** on the ribbon or right-click the server and select **Properties**.
4. You will follow the same steps as you have followed when adding the server. Edit server settings as required.

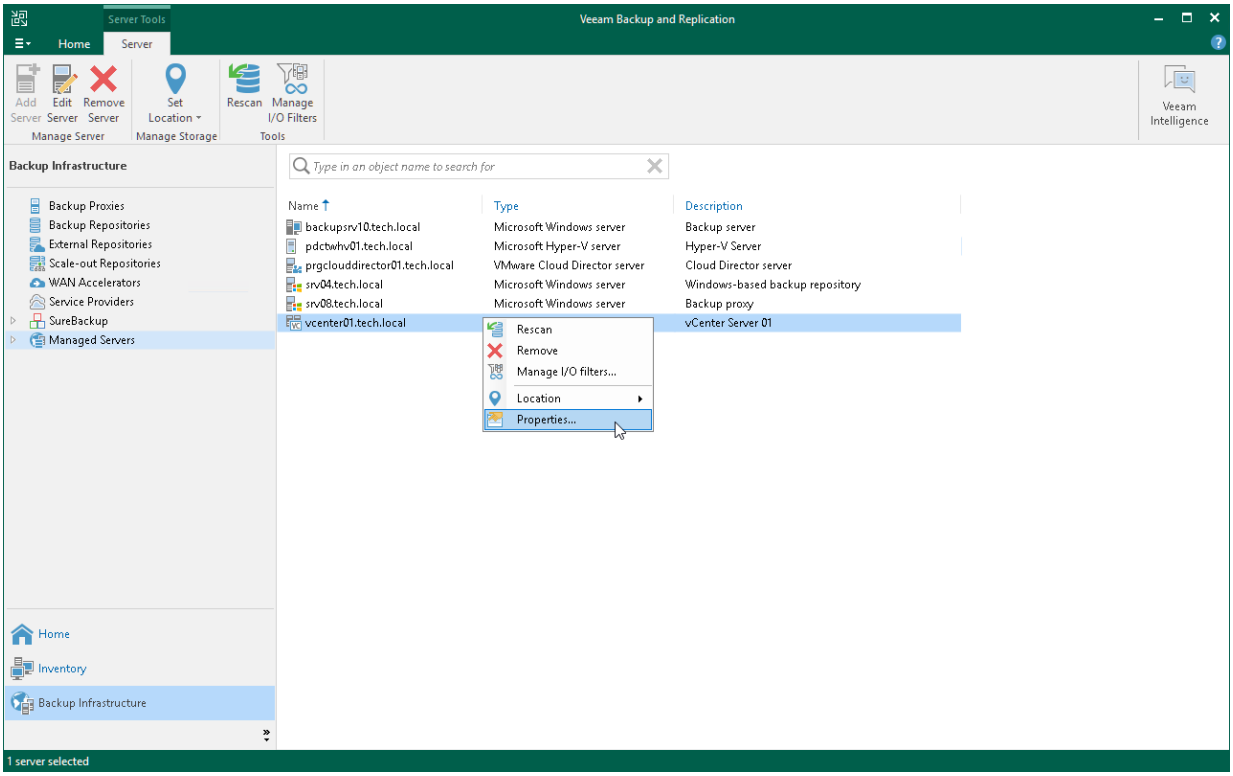

# Removing Servers

If you do not plan to use some server anymore, you can remove it from the backup infrastructure.

You cannot remove a server that has any dependencies. For example, you cannot remove a server that is referenced by a backup or replication job, performs the role of a VMware backup proxy or backup repository. To remove such server, you will need to delete all referencing jobs and roles first.

When you remove a server that is used as a target host or backup repository, backup files and replica files are not removed from disk. You can easily import these files later to Veeam Backup & Replication if needed.

> **NOTE**
>
> When you remove VMware Cloud Director from the backup infrastructure, vCenter Servers added to VMware Cloud Director are not removed. To remove the vCenter Server, in the inventory pane expand the **vCenter Servers** node, right-click the vCenter Server and select **Remove**.
>
> You cannot remove vCenter Servers added to VMware Cloud Director until the VMware Cloud Director server is removed from the backup infrastructure.

To remove a server from the backup infrastructure:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed servers**.
3. In the working area, select the server and click **Remove Server** on the ribbon or right-click the server and select **Remove**.

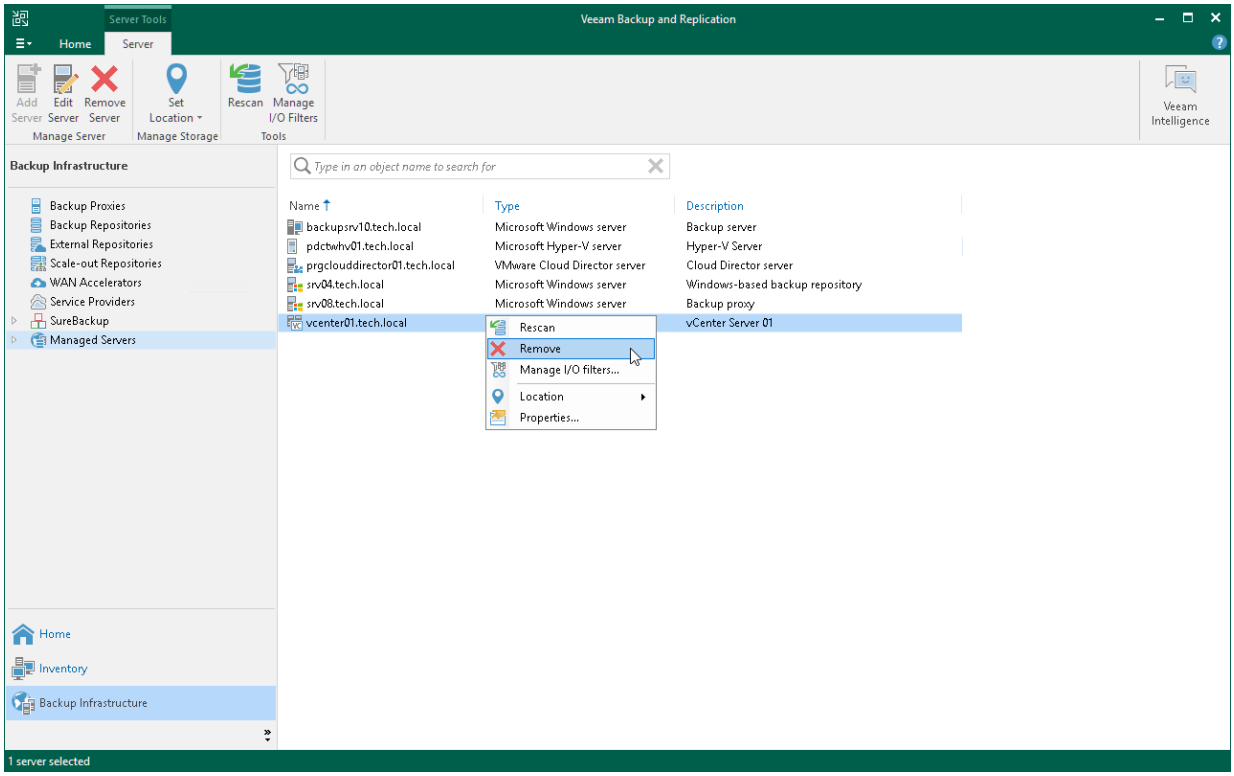

# General-Purpose Backup Proxies

A general-purpose backup proxy is a component that operates as a data mover. The backup proxy processes jobs and delivers backup and restore traffic.

## Usage Scenarios

General-purpose backup proxies can be used for the following operations:

- Unstructured Data Backup. In this case, proxies transfer data between an unstructured data source and a backup repository.
- Veeam Agent and storage system snapshot integration. In this case, proxies transfer data between a storage system and a backup repository. For more information on Veeam Agent Integration, see the Storage Snapshots Support section in the Veeam Agent Management Guide.

For more information on the backup infrastructure components required for different features, see feature descriptions.

## General-Purpose Backup Proxy Deployment

By default, the role of the general-purpose backup proxy is assigned to the backup server itself. However, this option is sufficient only for small installations where all components are located in the same network segment. For larger installations with larger workload, assign the role of a backup proxy to a dedicated server, as described in the Adding General-Purpose Backup Proxies section.

To optimize performance of several concurrent tasks, you can use several backup proxies. In this case, Veeam Backup & Replication will distribute the backup or restore workload between available backup proxies on per-task basis, taking into account proxy connectivity and their current load.

To minimize the network load during backup, locate the backup proxy closer to the source file share or storage in the computer network: at the best they should be one hop away from each other.

## Backup Proxy Services and Components

General-purpose backup proxies run light-weight services that take a few seconds to deploy. Deployment is fully automated. Veeam Backup & Replication installs the following components and services:

- **Veeam Backup VSS Integration** is a service that manages Microsoft VSS snapshots. Used for file backup.
- **Veeam VSS Hardware Snapshot Provider** is a service that extends Microsoft VSS and enables backups from storage snapshots. Used for Veeam Agents.
- **Veeam Installer Service** is an auxiliary service that is installed and started on any Windows server once it is added to the list of managed servers in the Veeam Backup & Replication console. This service analyzes the system, installs and upgrades necessary components and services depending on the role selected for the server.

## Requirements for General-Purpose Backup Proxy

Before you add a backup proxy to the inventory of the virtual infrastructure, check the following prerequisites and limitations:

- The backup proxy must meet the system requirements. For more information, see System Requirements.

- The role of a backup proxy for file backup and Veeam Agent backup from storage system snapshots must be assigned to a Microsoft Windows-managed server. For object storage backup, you can also use Linux proxy servers.
- [For Veeam Agent and Storage System Snapshot Integration] Check limitations listed in the Storage Snapshots Support section in the Veeam Agent Management Guide.

# Adding General-Purpose Backup Proxies

You must add to the backup infrastructure one or more backup proxies that you plan to use for moving backup data in the unstructured data backup or Veeam Agent and storage system snapshot integration.

# Step 1. Launch New Backup Proxy Wizard

To launch the **New Backup Proxy** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Proxies** node and select **Add proxy**. Alternatively, you can click **Add Proxy** on the ribbon.
3. In the **Add Backup Proxy** window, select **General-purpose backup proxy**.

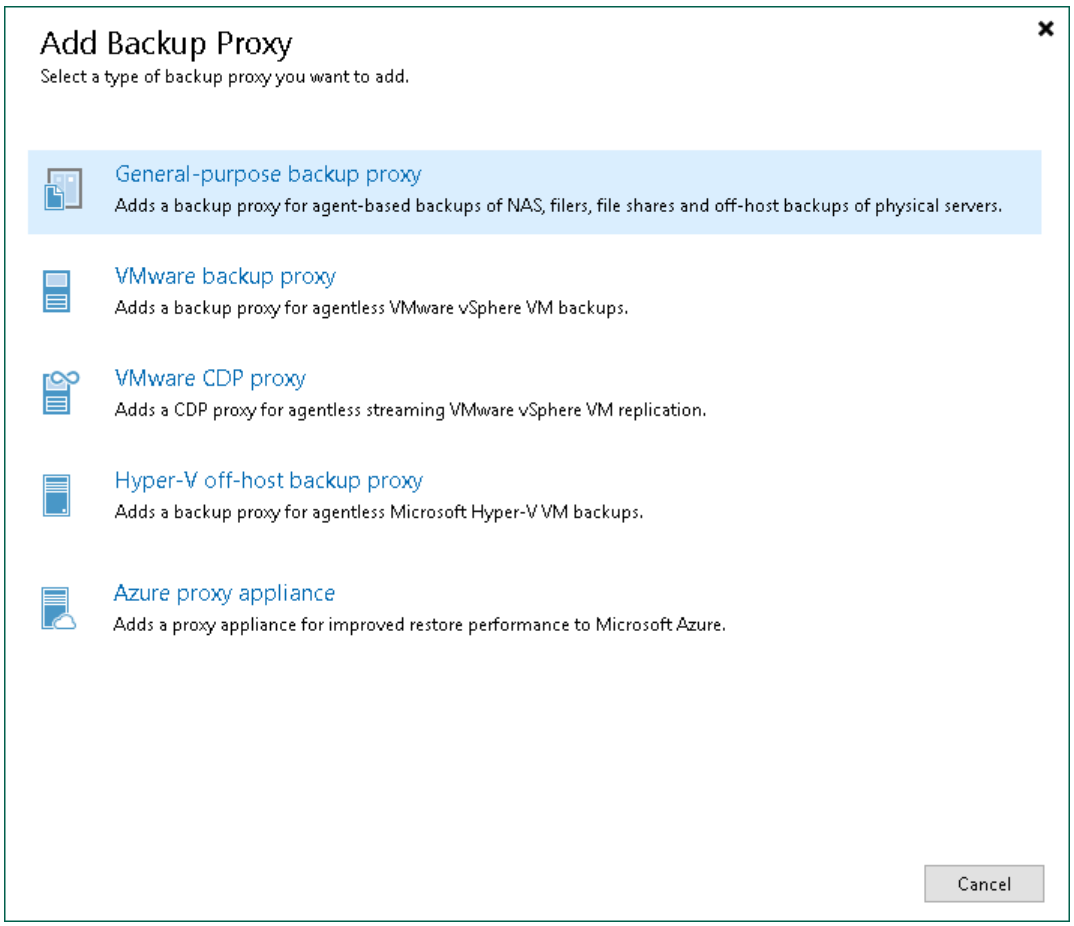

# Step 2. Choose Server

At the **Server** step of the wizard, specify server settings for the backup proxy.

1. From the **Choose server** list, select a Linux (for object storage backup only) or Microsoft Windows server that you want to use as a backup proxy.

   The list of servers contains only those managed servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, you can click **Add New** to open the **New Windows Server** or **New Linux Server** wizard. For more information, see the Adding Microsoft Windows Servers and Adding Linux Servers sections.

2. In the **Proxy description** field, provide a description for future reference. The default description contains information about the user who added the backup proxy, date and time when the backup proxy was added.

3. In the **Max concurrent tasks** field, specify the number of tasks that the backup proxy can process in parallel.

   If the number of parallel tasks reaches this value, the backup proxy will not start a new task until one of current tasks completes. Veeam Backup & Replication creates one task per every source file share. The recommended number of concurrent tasks is calculated automatically based on the amount of available resources. Backup proxies with multi-core CPUs can handle more concurrent tasks.

   For example, for a 4-core CPU, it is recommended that you specify a maximum of 8 concurrent tasks, for an 8-core CPU — 16 concurrent tasks. When defining the number of concurrent tasks, keep in mind network traffic throughput in the infrastructure.

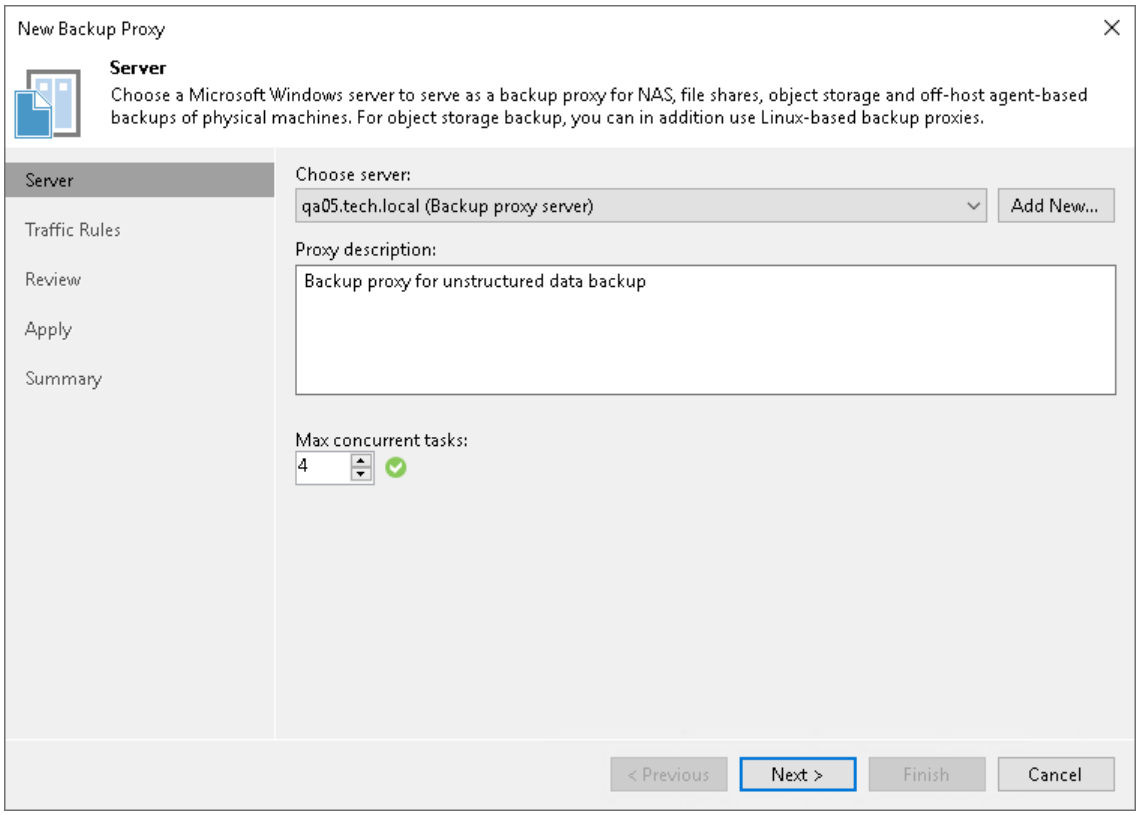

# Step 3. Configure Traffic Rules

At the **Traffic Rules** step of the wizard, configure network traffic rules. These rules help you throttle and encrypt traffic transferred between backup infrastructure components. For more information, see Configuring Network Traffic Rules.

The list of network traffic rules contains only rules applied to the backup proxy: its IP address falls into the IP range configured for the rule.

To view settings configured for the rule:

1. Select the rule in the list.
2. Click **View**. The **View Network Traffic Rule** window will display settings configured for the rule.

To modify network traffic settings:

1. Click the **Manage network traffic rules** link.
2. The **Global Network Traffic Rules** window will display the full list of all existing global network traffic rules.
3. Select the rule that you want to modify and click **Edit**. For more information on how to configure network traffic rules, see Configuring Network Traffic Rules.

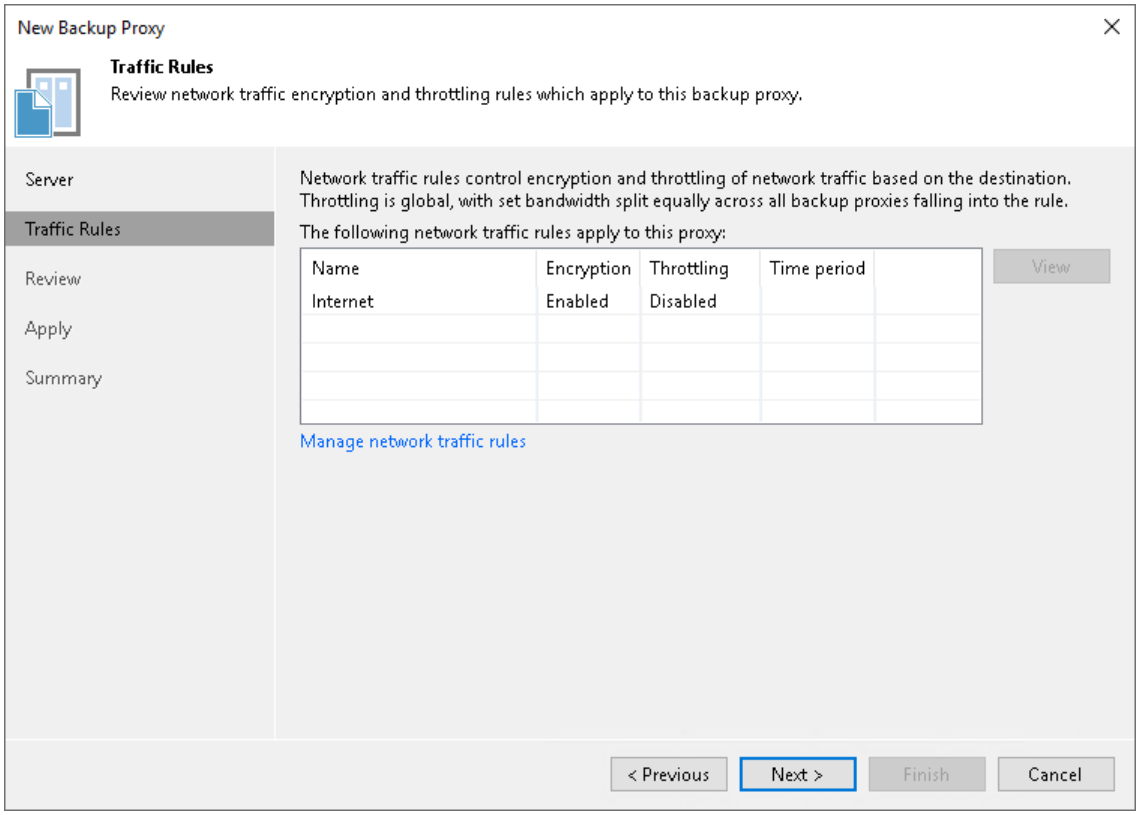

# Step 4. Review Components to Install

At the **Review** step of the wizard, review what Veeam Backup & Replication components are already installed on the server and click **Apply** to start installation of missing components.

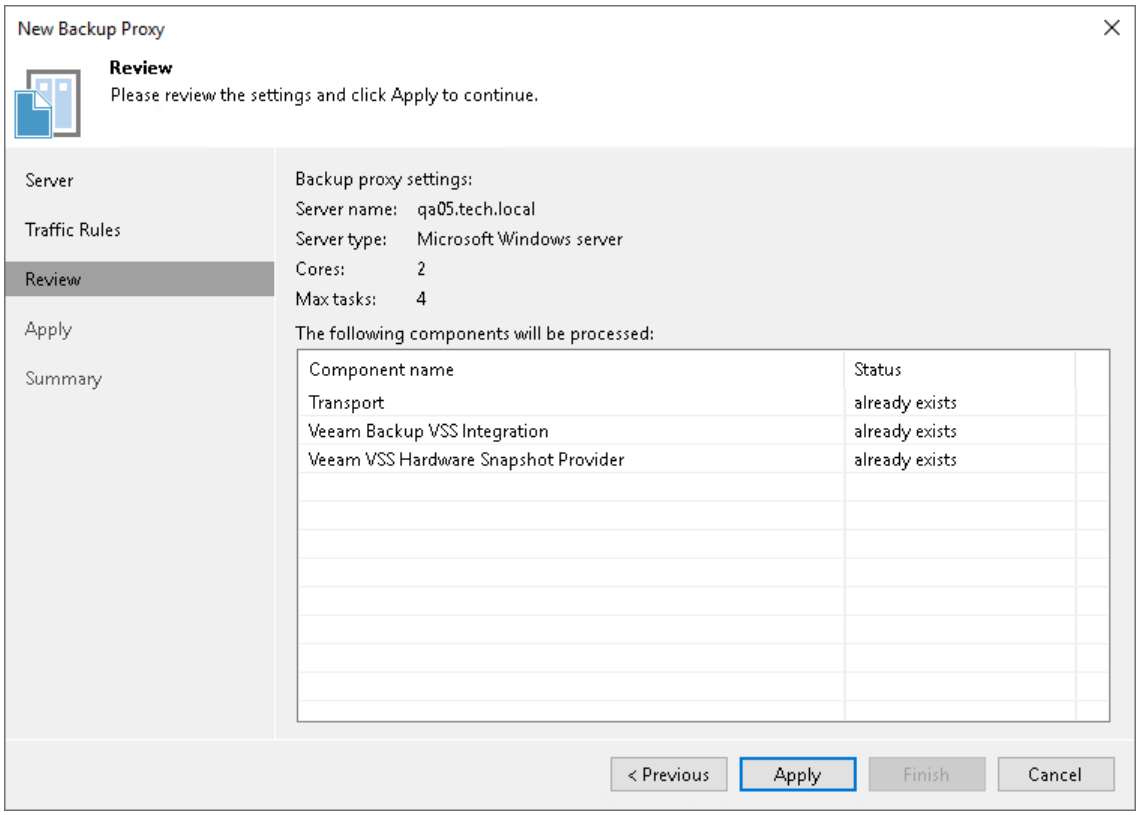

# Step 5. Apply Backup Proxy Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components. Click **Next** to complete the procedure of the backup proxy role assignment to the server.

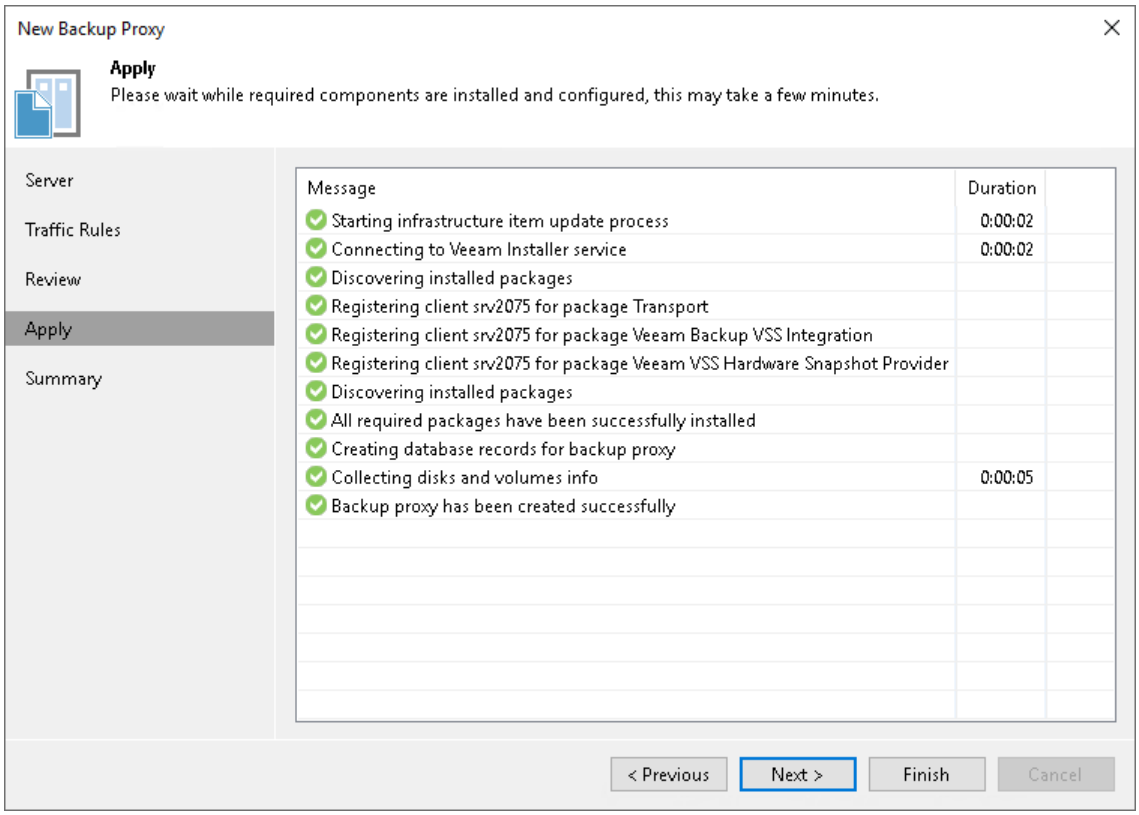

# Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added backup proxy and click **Finish** to exit the wizard.

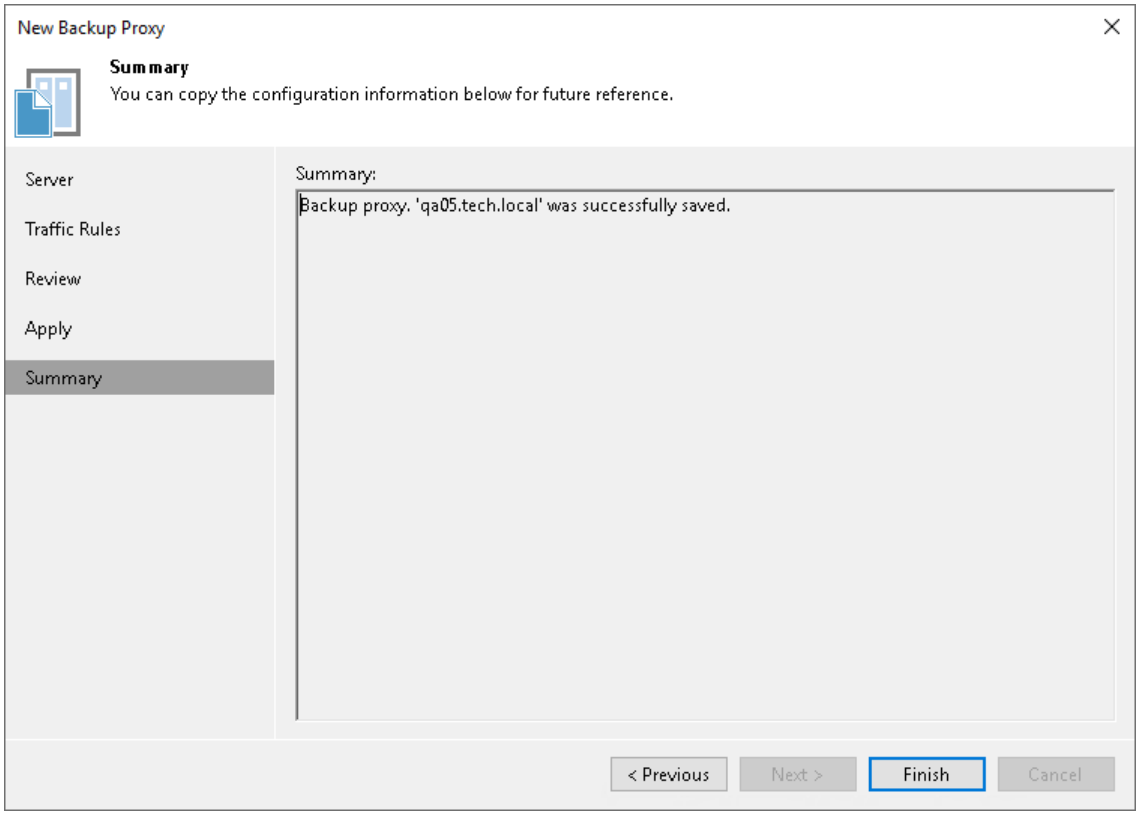

# VMware Backup Proxies

A VMware backup proxy is an architecture component that sits between the backup server and other components of the backup infrastructure. While the backup server administers tasks, the proxy processes jobs and delivers backup traffic.

Basic VMware backup proxy tasks include the following:

- Retrieving VM data from the production storage
- Compressing
- Deduplicating
- Encrypting
- Sending it to the backup repository (for example, if you run a backup job) or another VMware backup proxy (for example, if you run a replication job)

## Usage Scenarios

A VMware backup proxy is used for backup, replication, Quick Migration and other features. For more information on the backup infrastructure components required for different features, see the relevant feature description.

## VMware Backup Proxy Transport Modes

Depending on your backup architecture, a VMware backup proxy can use one of the following data transport modes:

- Direct storage access
- Virtual appliance
- Network

If the VM disks are located on the storage system and the storage system is added to the Veeam Backup & Replication console, the VMware backup proxy can also use the Backup from Storage Snapshots mode.

You can explicitly select the transport mode or let Veeam Backup & Replication automatically choose the mode. For details, see the Transport Modes in the Veeam Backup & Replication Guide and Configuring Backup Proxy for Storage Snapshots sections in the Storage System Snapshot Integration Guide.

## VMware Backup Proxy Deployment

By default, the role of the proxy is assigned to the backup server itself. However, this is sufficient only for small installations with low traffic load. For large installations, it is recommended to deploy dedicated backup proxies.

To optimize performance of several concurrent jobs, you can use several backup proxies. In this case, Veeam Backup & Replication will distribute the backup workload between available backup proxies. You can deploy backup proxies both in the primary site and in remote sites.

To deploy a proxy, you need to add a Windows-based or Linux-based server to Veeam Backup & Replication and assign the role of the VMware backup proxy to the added server. For requirements and limitations that backup proxies have, see Requirements and Limitations for VMware Backup Proxies.

# VMware Backup Proxy Services and Components

Backup proxies run light-weight services that take a few seconds to deploy. Deployment is fully automated. Veeam Backup & Replication installs the following components and services:

- **Veeam Installer Service** is an auxiliary service that is installed and started on any Windows server once it is added to the list of managed servers in the Veeam Backup & Replication console. This service analyzes the system, installs and upgrades necessary components and services depending on the role selected for the server.
- **Veeam Data Mover** is a component that performs data processing tasks on behalf of Veeam Backup & Replication, such as retrieving source VM data, performing data deduplication and compression, and storing backed-up data on the target storage.

# In This Section


- Requirements and Limitations for VMware Backup Proxies
- Transport Modes
- Adding VMware Backup Proxies
- Editing VMware Backup Proxy Settings
- Disabling and Removing VMware Backup Proxies

# Requirements and Limitations for VMware Backup Proxies

Before you assign the role of a VMware backup proxy, check the following prerequisites and limitations.

## Connection to Storage

The following list shows possible connections between the machine and storage that keeps backups of this machine. The first connection is the most efficient, the last one is the least efficient.

- A machine used as a VMware backup proxy should have direct access to the storage on which VMs reside or the storage where VM data is written. This way, the VMware backup proxy will retrieve data directly from the datastore bypassing LAN.
- The VMware backup proxy can be a VM with HotAdd access to VM disks on the datastore. This type of proxy also enables LAN-free data transfer between the host and the VMware backup proxy.
- If neither of the above scenarios is possible, you can assign the role of the VMware backup proxy to a machine whose network is located closer to the source or the target storage to which the proxy will connect. In this case, VM data will be transported over LAN using the NBD protocol.

## General Requirements and Limitations

- The role of a VMware backup proxy can be assigned to the following machines:
  - Physical or virtual Microsoft Windows machine
  - Physical or virtual Linux machine
- Before assigning the role of the VMware backup proxy to a machine, you must first add a vCenter server or ESXi host to the backup infrastructure.
- The machine must meet the system requirements. For more information, see System Requirements.
- The account that you specify when adding a server must have permissions described in section Permissions.
- It is recommended that you balance the number of tasks on backup proxies and backup repository to avoid the situation where some backup infrastructure resources remain idle while others are overloaded.
- You must add the machine to the Veeam Backup & Replication console as a managed server.
- If you back up proxies that use the Virtual appliance (HotAdd) mode to process VM data, the change block tracking mechanism (CBT) will be disabled. For more information on CBT, see Changed Block Tracking.
- If you back up encrypted VMs that use the Virtual appliance (HotAdd) mode, make sure backup proxies are also encrypted. For more information, see Encrypted VMs.
- A VMware backup proxy should be as close to the source data as possible with a high bandwidth connection. Consider a good connection between proxy and repository.
- VMware backup proxies that use VMware VDDK-based transport modes do not support processing virtual disks with the @ symbol in the disk name.

## Requirements and Limitations for VMware Backup Proxy on Linux

In addition to the general requirements and limitations, the following ones apply to Linux backup proxies:

- Linux backup proxies use the transport service for connection with backup infrastructure components. If the transport service cannot be installed, Linux backup proxy require SSH connection.
- You can assign the role of a VMware backup proxy to a Linux server added with single-use credentials, for example, a Linux server used as a hardened repository. For this configuration, only the Network mode (NBD) is supported, other transport modes will not be available for selection.
- Linux backup proxies cannot be used with VMware Cloud on AWS. This is because VDDK settings required by VMware cannot be enabled on Linux backup proxies.
- Linux backup proxies that use virtual appliance (HotAdd) transport mode do not support the VM copy scenario.
- Linux backup proxies cannot act as guest interaction proxies. For more information, see Guest Interaction Proxies.
- For Direct SAN with iSCSI access, note that Linux backup proxies must have the Open-iSCSI initiator enabled.
- For Direct NFS access, consider the following:
  - Linux backup proxies must have NFS client package installed.
  - Debian-based backup proxies must have the `nfs-common` package installed.
  - RHEL-based backup proxies must have the `nfs-utils` package installed.
- See recommendations for VMware backup proxy parameters in the Veeam Backup & Replication Best Practice Guide.

# Transport Modes

A transport mode is a method that is used by the Veeam Data Mover to retrieve VM data from the source and write VM data to the target. Job efficiency and time required for job completion greatly depend on the transport mode.

For data retrieval, Veeam Backup & Replication offers the following modes (starting from the most efficient):

- Direct storage access
- Virtual Appliance (HotAdd)
- Network

The Veeam Data Mover responsible for data retrieval runs on a VMware backup proxy. The transport mode can be defined in the settings of the VMware backup proxy that performs the job.

When you configure VMware backup proxy settings, you can manually select a transport mode, or let Veeam Backup & Replication select the most appropriate mode automatically. If you use automatic mode selection, Veeam Backup & Replication will scan VMware backup proxy configuration and its connection to the VMware vSphere infrastructure to choose the optimal transport mode. If several transport modes are available for the same VMware backup proxy, Veeam Backup & Replication will choose the mode in the following order: Direct storage access > Virtual appliance > Network.

The selected transport mode is used for data retrieval. For writing data to the target, Veeam Backup & Replication picks the transport mode automatically, based on the configuration of the VMware backup proxy and transport mode limitations.

Veeam Backup & Replication leverages VMware vStorage APIs for Data Protection (VADP) for all transport modes except for backup from storage snapshots, the direct NFS transport mode and the virtual appliance transport mode. VADP can be used for VMware vSphere starting from version 4.

Applicability and efficiency of each transport mode primarily depends on the type of datastore used by the source host — local or shared, and on the VMware backup proxy type — physical or virtual. The following table shows recommendations for installing the VMware backup proxy, depending on the storage type and desired transport mode.

| Production Storage Type | Direct Storage Access | Virtual Appliance | Network Mode |
|---|---|---|---|
| **Fiber Channel (FC) SAN** | Install a VMware backup proxy on a physical server with a direct FC access to the SAN. | Install a VMware backup proxy on a VM running on an ESXi host connected to the storage device. | This mode is *not recommended* on 1 Gb Ethernet but works well with 10 Gb Ethernet.<br>Install a VMware backup proxy on any machine on the storage network. |
| **iSCSI SAN** | Install a VMware backup proxy on a physical or virtual machine. | | |
| **NFS Storage** | | | |
| **Shared SAS** | Install a VMware backup proxy on a physical server with a direct SAS access to the SAN. | | |

| Production Storage Type | Direct Storage Access | Virtual Appliance | Network Mode |
|---|---|---|---|
| **vSAN** | Not supported. | Install a VMware backup proxy on a VM running on an ESXi host connected to the VSAN storage device. | |
| **vVol** | | Install a VMware backup proxy on a VM running on an ESXi host connected to the vVol storage. | |
| **Local Storage** | | Install a VMware backup proxy on a VM on every ESXi host. | |

**NOTE**

If you use VMware Cloud on AWS, the only available transport mode is Virtual appliance. We recommend you install a VMware backup proxy on a VM running on an ESXi host connected to the VSAN storage device.

# Direct Storage Access

In the Direct storage access mode, Veeam Backup & Replication reads/writes data directly from/to the storage system where VM data or backups are located. This mode comprises two transport modes:

- Direct SAN access
- Direct NFS access

## Direct SAN Access

The Direct SAN access transport mode is recommended for VMs whose disks are located on shared VMFS SAN LUNs that are connected to ESXi hosts over FC, FCoE, iSCSI, and on shared SAS storage.

In the Direct SAN access transport mode, Veeam Backup & Replication leverages VMware VADP to transport VM data directly from and to FC, FCoE and iSCSI storage over the SAN. VM data travels over the SAN, bypassing ESXi hosts and the LAN. The Direct SAN access transport method provides the fastest data transfer speed and produces no load on the production network.

The Direct SAN access transport mode can be used for all operations where the VMware backup proxy is engaged:

- Backup
- Replication
- VM copy
- Quick migration

- Entire VM restore
- VM disk restore
- Replica failback

# Requirements for the Direct SAN Access Mode

To use the Direct SAN access transport mode, make sure that the following requirements are met:

- It is strongly recommended that you assign the role of a VMware backup proxy working in the Direct SAN access mode to a physical machine. If you assign this role to a VM, the VMware backup proxy performance may not be optimal.
- A VMware backup proxy using the Direct SAN access transport mode must have a direct access to the production storage using a hardware or software HBA. If a direct SAN connection is not configured or not available when a job or task starts, the job or task will fail.
- SAN storage volumes presented as VMware datastores must be exposed to the OS of the VMware backup proxy that works in the Direct SAN access transport mode.

  The volumes must be visible in Disk Management but must not be initialized by the OS. Otherwise, the VMFS filesystem will be overwritten with NTFS, and volumes will become unrecognizable by ESXi hosts. To prevent volumes initialization, Veeam Backup & Replication automatically sets the SAN Policy within each proxy to Offline Shared and also sets the SAN LUNs to the Offline state.
- [For restore operations] A VMware backup proxy must have write access to LUNs where VM disks are located.

# Limitations for the Direct SAN Access Mode

- The Direct SAN access transport mode can be used to restore only thick VM disks.
- The transport mode is used for the entire VM, not for the virtual disk. It means that one VM can be processed in one transport mode only. If there are VM disks that cannot be processed in the Direct SAN access transport mode, then the rest of the disks also cannot be processed in this transport mode.
- You cannot use the Direct SAN access mode in the following cases:
  - For VMs residing on vSAN. You can use Virtual appliance and Network transport modes to process such VMs. For details on vSAN restrictions, see VDDK release notes. For example, release notes for VDDK 7.0.3.
  - If at least one VM disk is located on a vVol.
  - For Veeam Cloud Connect Replication because in this scenario Veeam Backup & Replication always creates VM replicas with thin disks.
  - For incremental restore due to VMware limitations. Either disable CBT for VM virtual disks for the duration of the restore process or select another transport mode for incremental restore.
  - For backing up VM templates.

- Veeam Backup & Replication uses the Direct SAN access transport mode to read and write VM data only during the first session of the replication job. During subsequent replication job sessions, Veeam Backup & Replication will use the Virtual appliance or Network transport mode on the target side. The source side proxy will keep reading VM data from the source datastore in the Direct SAN access transport mode.

  Veeam Backup & Replication writes VM data to the target datastore in the Direct SAN access transport mode only if disks of a VM replica are thick-provisioned. If disks are thin-provisioned, Veeam Backup & Replication will write VM data in the Network or Virtual appliance mode. By default, Veeam Backup & Replication replicates VM disks in the thin format.

  To write VM data to the target datastore in the Direct SAN access transport mode, select to convert VM disks to the thick format at the **Destination** step of the replication job wizard.

- [For restore operations] If you chose the **Manual selection** option in the backup proxy settings and specified the datastores manually, make sure that the `ReadOnly` value for each datastore is set to `false` on the backup proxy. You can use the diskpart command interpreter to verify the value. For more information, see Microsoft Docs.

- IDE and SATA disks can be processed in the Direct SAN access transport mode.

- Veeam Backup & Replication supports multipathing (MPIO) for Windows-based backup proxies in the Direct SAN access transport mode with the following conditions:

  - The Multipath I/O feature must be enabled in the Windows Server Manager console. For more information, see Microsoft Docs.
  - Zoning and masking on the FC switches and storage must be configured.
  - All proxy ports must have access to shared VMFS SAN LUNs where VMs disks are located on.

- Multipathing (MPIO) for Linux-based backup proxies in the Direct SAN access transport mode leverages only path failovers and not load balancing. These are limitations of the VMware VDDK, and the distributions supported for MPIO in the Direct SAN access transport mode are listed in the Virtual Disk Development Kit release notes for your vSphere version.

## Data Backup in Direct SAN Access Mode

To retrieve VM data blocks from a SAN LUN during backup, the VMware backup proxy uses metadata about the layout of VM disks on the SAN.

Data backup in the Direct SAN access transport mode includes the following steps:

1. The VMware backup proxy sends a request to the ESXi host to locate the necessary VM on the datastore.
2. The ESXi host locates the VM.
3. Veeam Backup & Replication triggers VMware vSphere to create a VM snapshot.
4. The ESXi host retrieves metadata about the layout of VM disks on the storage (physical addresses of data blocks).
5. The ESXi host sends metadata to the VMware backup proxy.
6. The VMware backup proxy uses metadata to copy VM data blocks directly from the source storage over the SAN.

7. The VMware backup proxy processes copied data blocks and sends them to the target.

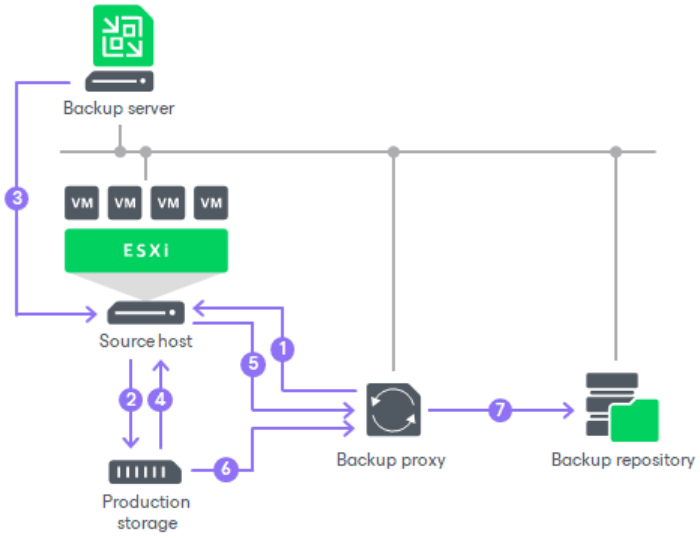


## Data Restore in Direct SAN Access Mode

The Direct SAN access transport mode can be used to restore VMs with thick disks. To restore VMs with thin disks, you can use the Direct NFS access, Virtual appliance or Network mode.

> **TIP**
>
> If you plan to process VMs that have both thin and thick disks, you can enable the Direct SAN access mode. However, Veeam Backup & Replication will use the Network transport mode to restore disks of these VMs. If you want to use Direct SAN access mode, restore all VM disks as thick.

Data restore in the Direct SAN access transport mode includes the following steps:

1. The VMware backup proxy retrieves data blocks from the backup repository or a datastore in the target site.
2. The VMware backup proxy sends a request to the ESXi host in the source site to restore data to a necessary datastore.
3. The ESXi host in the source site allocates space on the datastore.

4. Data blocks are written to the datastore.

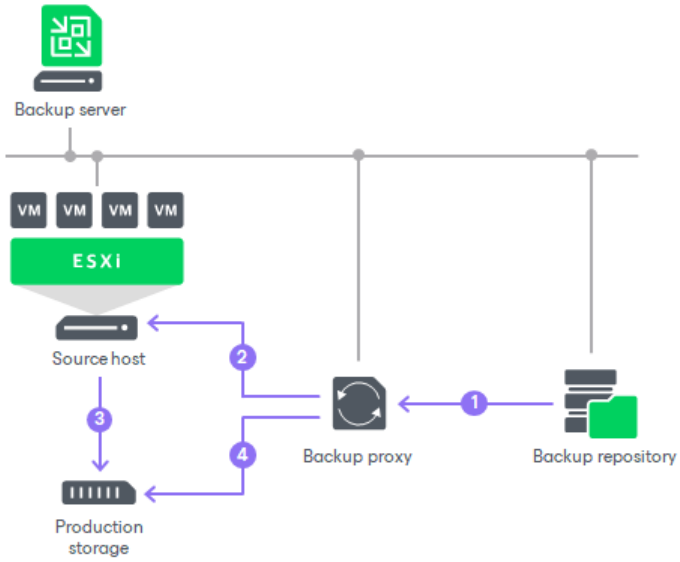


## Direct NFS Access

The Direct NFS access is a recommended transport mode for VMs whose disks are located on NFS datastores.

The Direct NFS access mode provides an alternative to the Network mode. When Veeam Backup & Replication processes VM data in the Network mode, it uses VMware VDDK to communicate with the ESXi host. This produces additional load on the ESXi host.

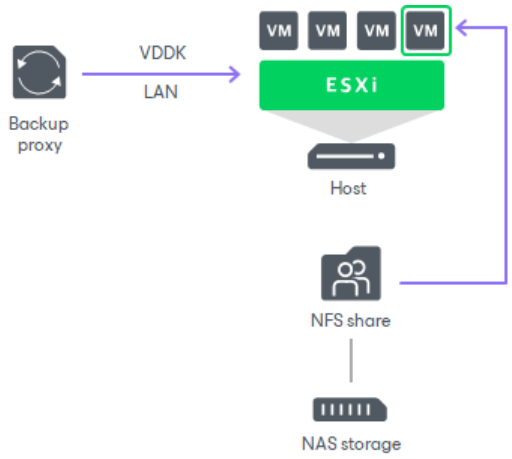

In the Direct NFS access mode, Veeam Backup & Replication bypasses the ESXi host and reads/writes data directly from/to NFS datastores. To do this, Veeam Backup & Replication deploys its native NFS client on the VMware backup proxy and uses it for VM data transport. VM data still travels over LAN but there is no load on the ESXi host.

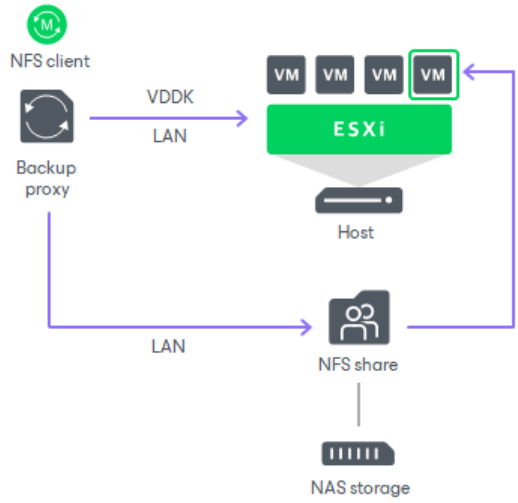


The Direct NFS access mode can be used for all operations where the VMware backup proxy is engaged:

- Backup
- Replication
- Quick Migration
- VM copy
- Entire VM restore
- VM disk restore
- Replica failback

# Requirements for the Direct NFS Access Mode

- Direct NFS access mode can be used in VMware vSphere environments running NFS version 3 and 4.1.
- The VMware backup proxy used for VM data processing must have access to the NFS datastores where VM disks are located. For more information, see VMware Backup Proxy for Direct NFS Access Mode.
- If NFS volumes are mounted on the ESXi host under names, not IP addresses, the volume names must be resolved by DNS from the VMware backup proxy.

# Limitations for Direct NFS Access Mode

- Veeam Backup & Replication cannot parse delta disks in the Direct NFS access mode. For this reason, the Direct NFS access mode has the following limitations:
  - The Direct NFS access mode cannot be used for VMs that have at least one snapshot.

- Veeam Backup & Replication uses the Direct NFS transport mode to read and write VM data only during the first session of the replication job. During subsequent replication job sessions, the VM replica will already have one or more snapshots. For this reason, Veeam Backup & Replication will use another transport mode to write VM data to the datastore on the target side. The source side proxy will keep reading VM data from the source datastore in the Direct NFS transport mode.

- If you enable the **Enable VMware tools quiescence** option in the job settings, Veeam Backup & Replication will not use the Direct NFS transport mode to process running Microsoft Windows VMs that have VMware Tools installed. The Direct NFS transport mode is not used because during VM quiescence VMware creates a snapshot with two delta disks per virtual disk.

- If a VM has some disks that cannot be processed in the Direct NFS access mode, Veeam Backup & Replication processes these VM disks in the Network transport mode.

## VMware Backup Proxy for Direct NFS Access Mode

To instruct the VMware backup proxy to use the Direct NFS access mode, you must choose the **Automatic selection** or **Direct storage access** option in the VMware backup proxy settings.

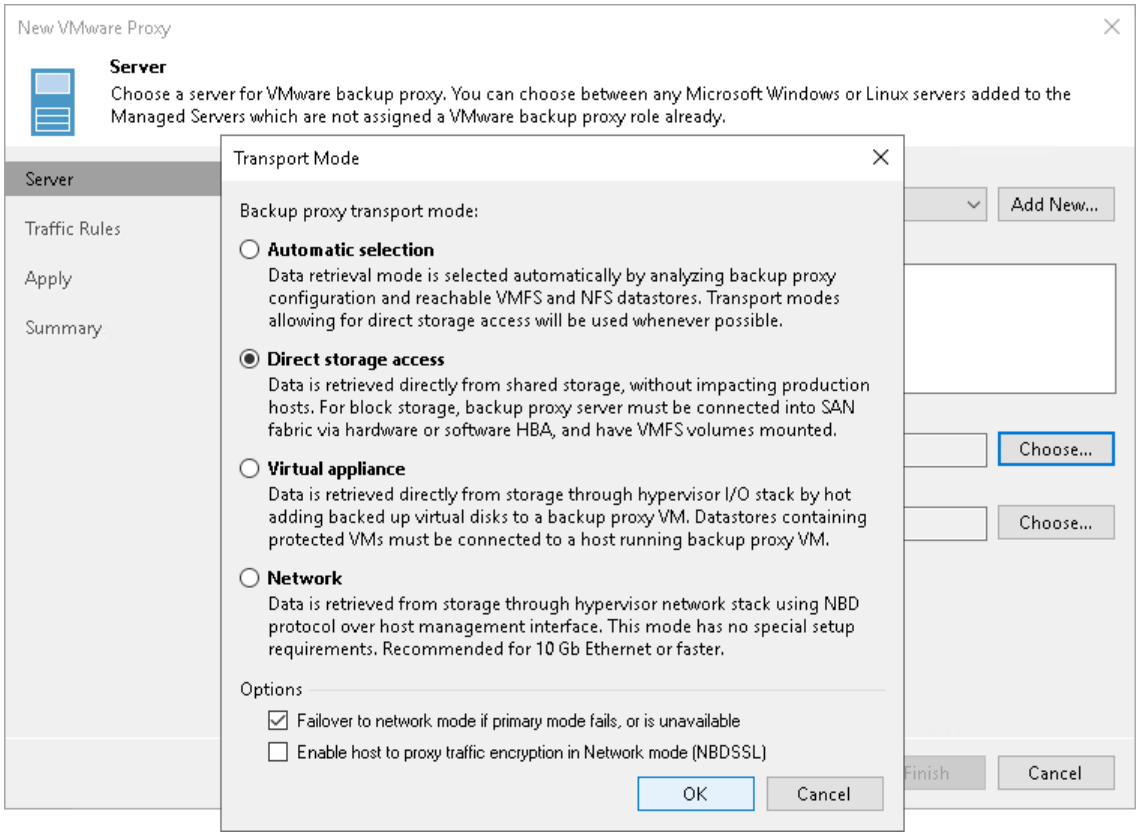


To read and write data in the Direct NFS transport mode, the VMware backup proxy must meet the following requirements:

1. The VMware backup proxy must have access to the NFS datastore.
2. The VMware backup proxy must have *ReadOnly/Write* permissions and root access to the NFS datastore.

Veeam Backup & Replication deploys its NFS agent on every VMware backup proxy when you assign the VMware backup proxy role to a Microsoft Windows server (physical or virtual). Linux backup proxies must have NFS client package installed. For more information, see Requirements and Limitations for VMware Backup Proxy on Linux.

# VMware Backup Proxy Selection

Veeam Backup & Replication selects backup proxies working in the Direct NFS access transport mode by the following rules:

- If you instruct Veeam Backup & Replication to select a VMware backup proxy automatically for a job or task, Veeam Backup & Replication picks a VMware backup proxy with the minimum number of hops to the NFS datastore. If there are several backup proxies with the equal number of hops in the backup infrastructure, Veeam Backup & Replication picks the least busy VMware backup proxy in the backup infrastructure.

  If all backup proxies with the minimum number of hops are busy at the moment, Veeam Backup & Replication waits until these backup proxies are free. Veeam Backup & Replication does not pick a VMware backup proxy that has a greater number of hops to the NFS datastore and works in the Direct NFS access or Virtual appliance transport mode.

- If you select one or more backup proxies explicitly for a job or task, Veeam Backup & Replication does not regard the number of hops to the NFS datastore. Veeam Backup & Replication picks the least busy VMware backup proxy working in the Direct NFS access transport mode.

  If all backup proxies working in the Direct NFS access transport mode are busy, Veeam Backup & Replication waits until these backup proxies are free. Veeam Backup & Replication does not pick a VMware backup proxy working in the Virtual appliance transport mode.

To detect the number of hops from a VMware backup proxy to the NFS datastore, Veeam Backup & Replication uses the host discovery process. During host discovery, Veeam Backup & Replication obtains information about the number of hops, checks to which NFS datastores the VMware backup proxy has access and what permissions the VMware backup proxy has on NFS datastores.

The host discovery process rescans all machines to which the VMware backup proxy role is assigned. The process starts automatically every 4 hours. Host discovery is also triggered when you change the transport mode settings and choose to use the Direct storage access for the VMware backup proxy.

If necessary, you can start the host discovery process manually. To do this, perform the **Rescan** operation for a machine to which the VMware backup proxy role is assigned.

## Data Backup in Direct NFS Access Mode

Data backup in the Direct NFS access transport mode is performed in the following way:

1. The VMware backup proxy sends a request to the ESXi host to locate a VM on the NFS datastore.
2. The ESXi host locates the VM.
3. Veeam Backup & Replication triggers VMware vSphere to create a VM snapshot.
4. The ESXi host retrieves metadata about the layout of VM disks on the storage (physical addresses of data blocks).
5. The ESXi host sends metadata to the VMware backup proxy.
6. The VMware backup proxy uses metadata to copy VM data blocks directly from the NFS datastore over LAN.

7. The VMware backup proxy processes copied data blocks and sends them to the target over LAN.

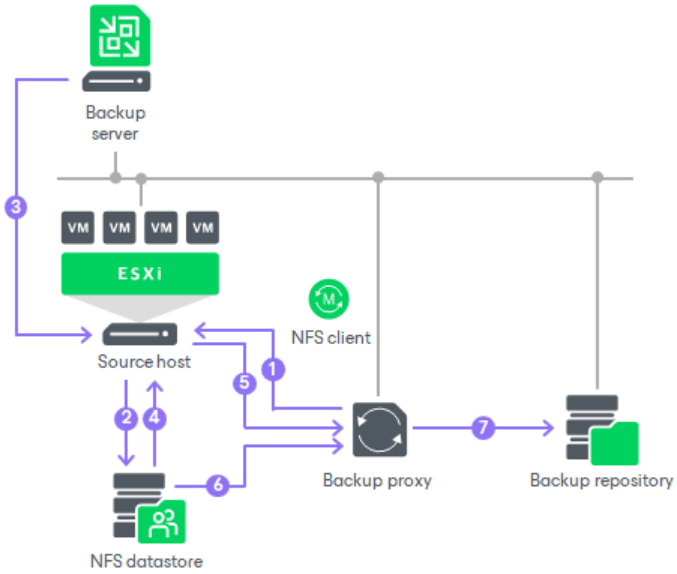


## Data Restore in Direct NFS Access Mode

The Direct NFS access transport mode can be used to restore VMs with thick and thin disks.

Data restore in the Direct NFS access transport mode is performed in the following way:

1. The VMware backup proxy retrieves data blocks from the backup repository or a datastore in the target site.
2. The VMware backup proxy sends a request to the ESXi host to restore data to an NFS datastore.
3. The ESXi host allocates space on the NFS datastore.
4. Data blocks obtained from the VMware backup proxy are written to the NFS datastore over LAN.

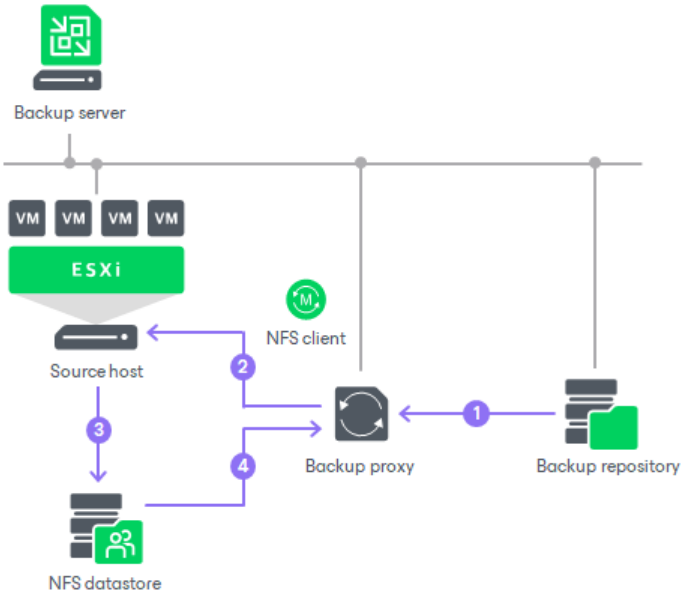

With the selected option **Same as source** from the Restore disk type drop-down list, Veeam Backup & Replication restores the disk with the same type as in the backup. However, during the restore in the Direct NFS access mode, the disk type can be affected by settings specified for your NFS server, VAAI, and so on. Thin disks can be restored as thin disks, thick eager disks as thick eager disks. Thus, Veeam Backup & Replication converts thick lazy zeroed disks to thick eager disks, if you use VAAI, otherwise to thin disks. For more information about disk types, see VMware Docs.

# Virtual Appliance (HotAdd)

The Virtual appliance mode is not so efficient as the Direct storage access mode but provides better performance than the Network mode. The Virtual appliance mode is recommended if the role of a VMware backup proxy is assigned to a VM.

In the Virtual appliance mode, Veeam Backup & Replication uses the VMware SCSI HotAdd capability that allows attaching devices to a VM while the VM is running. During backup, replication or restore disks of the processed VM are attached to the VMware backup proxy. VM data is retrieved or written directly from/to the datastore, instead of going through the network.

The Virtual appliance transport mode can be used for all operations where the VMware backup proxy is engaged:

- Backup
- Replication
- VM copy
- Quick Migration
- Entire VM restore
- VM disk restore
- Replica failback

# Requirements for the Virtual Appliance mode

To use the Virtual appliance transport mode, make sure that the following requirements are met:

- The role of a VMware backup proxy must be assigned to a VM.
- The VMware backup proxy and processed VMs must reside in the same datacenter.
- The VMware backup proxy must have access to disks of the VM that this proxy processes. For example, in a replication job, the source VMware backup proxy must have access to the disks of the source VM, the target proxy — to the disks of the replica. If a VMware backup proxy acts as both source and target proxy, it must have access to the disks of the source VM and replica. In restore operations, the VMware backup proxy must have access to disks of the restored VMs.
- [For NFS 3.0] If you plan to process VMs that store disks on the NFS datastore, you must configure Veeam Backup & Replication to use the proxy on the same host as VMs. This is required due to an issue described in this VMware KB article. For more information on how to configure the proxy, see this Veeam KB article.

  As an alternative, you can use ESXi 6.0 or later and NFS 4.1.

- The VMware backup proxy must have the latest version of VMware Tools installed. Note that the backup server installed on a VM can also perform the role of the VMware backup proxy that uses Virtual appliance transport mode. In this case, make sure the backup server has the latest version of VMware Tools installed.
- SCSI 0:X controller must be present on a VMware backup proxy. In the opposite case, VM data processing in the Virtual appliance transport mode will fail.

# Limitations for the Virtual Appliance mode

- [For vSphere 6.5 and later] If a source VM has vSAN disks and a VMware backup proxy used to process this VM has non-vSAN disks, backup and restore in the Virtual appliance mode is not supported.
- If a VMware backup proxy used to process a source VM resides on a VMFS 3 datastore, it must be formatted with proper block size to be able to mount the largest virtual disk of hot-added VMs:
  - 1 MB block size — 256 GB maximum file size
  - 2 MB block size — 512 GB maximum file size
  - 4 MB block size — 1024 GB maximum file size
  - 8 MB block size — 2048 GB maximum file size

  This limitation does not apply to VMFS-5 volumes that always have 1 MB file block size.
- For vSphere 5.5 and later the maximum supported VMDK size is 62 TB.
- [For Microsoft Windows proxy] Before running a data protection task, Veeam Backup & Replication disables the volume automount feature, and it remains disabled after the data protection task is completed.
- Backup and restore of IDE disks in the Virtual appliance mode is not supported.
- Backup and restore of SATA disks in the Virtual appliance mode is supported if you use VMware vSphere 6.0 and later.
- [For Quick Migration during Instant Recovery] Virtual appliance (HotAdd) transport mode cannot be used if the role of the backup proxy and mount server or backup repository where the backup file is stored are assigned to the same VM.

## Data Backup and Restore in Virtual Appliance Mode

The process of data retrieval in the Virtual appliance transport mode includes the following steps:

1. The backup server sends a request to the ESXi host to locate the necessary VM on the datastore.
2. The ESXi host locates the VM.
3. Veeam Backup & Replication triggers VMware vSphere to create a VM snapshot.
4. VM disks are attached (hot-added) to the VMware backup proxy.
5. Veeam Backup & Replication reads data directly from disks attached to the VMware backup proxy.

6. When the VM processing is complete, VM disks are detached from the VMware backup proxy and the VM snapshot is deleted.

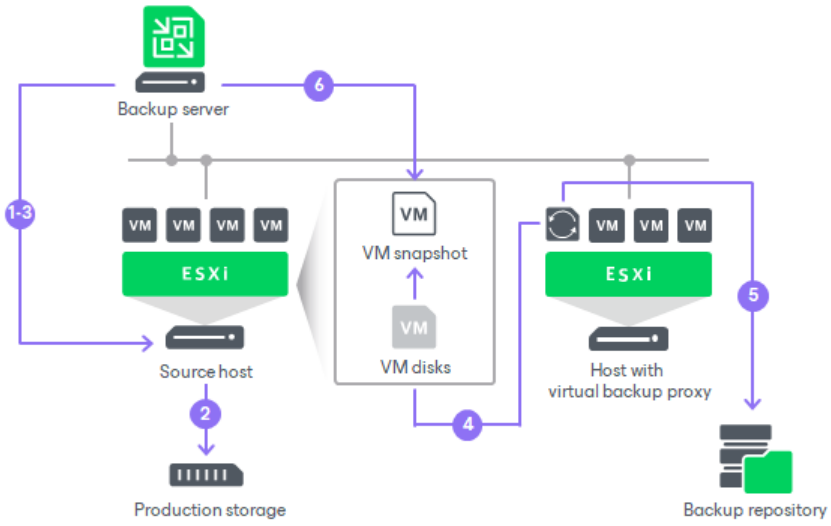


The process of data restore in the Virtual appliance mode works in a similar manner. VM disks from the backup are attached to the VMware backup proxy and Veeam Backup & Replication transports VM data to the target datastore. After the restore process is finished, VM disks are detached from the VMware backup proxy.

ESXi host interacts with VMware Cloud on AWS through VMware vCenter. Veeam Backup & Replication performs backup through the networkless Virtual appliance (HotAdd) mode.

## Virtual Appliance Mode for VMs on VSAN

To transport data of VMs residing on VSAN in the Virtual appliance mode, you must assign the VMware backup proxy role to a VM.

The VMware backup proxy VM must meet the following requirements:

- The VMware backup proxy VM must reside on an ESXi host connected to a VSAN cluster.

  Veeam Backup & Replication will retrieve data of processed VMs over the I/O stack of the ESXi host on which the VMware backup proxy is deployed.

- Disks of the VMware backup proxy VM must reside on the VSAN datastore.

If you have several backup proxies on ESXi hosts in the VSAN cluster, Veeam Backup & Replication chooses the most appropriate VMware backup proxy to reduce the backup traffic on the VSAN cluster network. To choose a VMware backup proxy, Veeam Backup & Replication checks HDDs directly attached to every ESXi host and calculates the amount of VM data on these HDDs. The preference is given to the ESXi host that has a direct access to an HDD with the maximum amount of VM data. This approach helps reduce workload on the ESXi I/O stack during data transport.

> **NOTE**
>
> Even if disks of a VM are located on a host where the VMware backup proxy is deployed, VSAN traffic may still be observed between hosts in the cluster. This behavior depends on the VSAN cluster itself and cannot be modified in Veeam Backup & Replication.

## Virtual Appliance Mode for VMs on vVol

To transport data of VMs residing on VVol in the Virtual appliance mode, you must assign the VMware backup proxy role to a VM.

The VMware backup proxy VM must meet the following requirement: for VM restore to a VVol in the Virtual appliance mode, the VMware backup proxy VM must be located on the same VVol datastore as the target VM. If this is not possible, the restore proxy must use a different transport mode.

> **NOTE**
>
> Due to the VMware vSphere mechanisms, data backup and restore for VMs on VVol in the Virtual appliance mode triggers the creation of auxiliary volumes in the storage array:
>
> - A temporary "snapshot" LUN is created when a VMware vSphere snapshot is triggered. The LUN is removed shortly afterward.
> - An additional "cloned" LUN is created during the data transport. When the VM processing is complete, the LUN is removed.
>
> The maximum amount of volumes depends on the storage array and VMware vSphere limits. For more information, see the documentation of your storage system vendor.

# Network Mode

The Network mode can be used with any infrastructure configuration. In this mode, data is retrieved through the ESXi host over LAN using the Network Block Device protocol (NBD).

The Network mode has low data transfer speed over LAN. To take the load off the LAN, Veeam Backup & Replication provides two alternative modes: Direct Storage Access and Virtual Appliance. However, the Network mode is the only applicable mode when the VMware backup proxy role is assigned to a physical machine and the host uses local storage. Also, the Network mode can be the best choice if you have a large virtual environment with hundreds of small VMs, with 10 Gb Ethernet networks and with a small change rate.

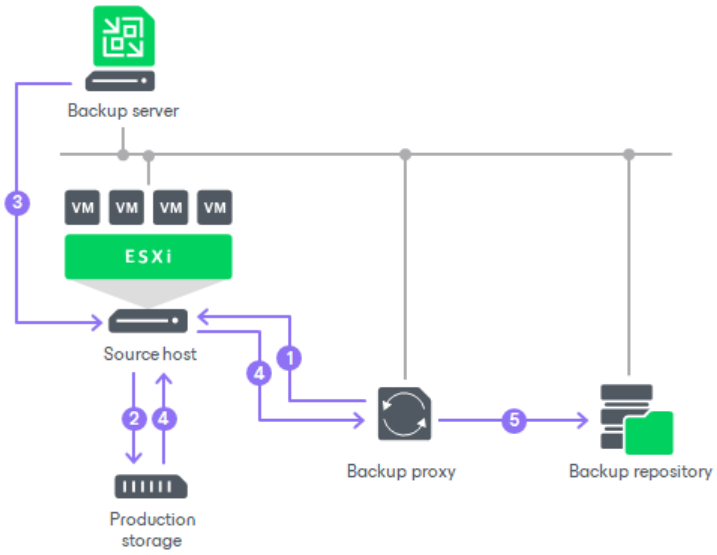


The process of data retrieval in Network mode includes the following steps:

1. The VMware backup proxy sends a request to the ESXi host on which the processed VM is registered to locate the VM on the datastore.

2. The ESXi host locates the processed VM on the datastore.
3. Veeam Backup & Replication instructs VMware vSphere to create a VMware vSphere VM snapshot.
4. ESXi host copies VM data blocks from the source storage and sends them to the VMware backup proxy over LAN.

   Note that the real data transfer speed may be significantly less than the available speed. This is because the VMware backup proxy and the ESXi host communicate over the ESXi management network.
5. The VMware backup proxy sends the data to target.

Veeam Backup & Replication processes VM disks in parallel. If VM disks are located on different storage types (for example, on the SAN and local storage), Veeam Backup & Replication uses different transport modes to process VM disks. In such scenario, it is strongly recommended that you select the **Failover to network mode if primary mode fails, or is unavailable** option when configuring the mode settings for the VMware backup proxy.

## Failover to Network Mode

You can instruct Veeam Backup & Replication to switch to the Network transport mode and transfer VM data over LAN if the primary transport mode — Direct storage access or Virtual appliance — fails during the job session. This option is enabled by default to ensure that jobs and tasks can be completed successfully in any situation. In scenarios when the data cannot be processed by other transport modes, failover to the Network transport mode applies automatically even if the failover option is disabled.

Note that data transport over LAN puts additional load on your production network and may potentially affect performance if you accomplish data protection and disaster recovery tasks in business hours.

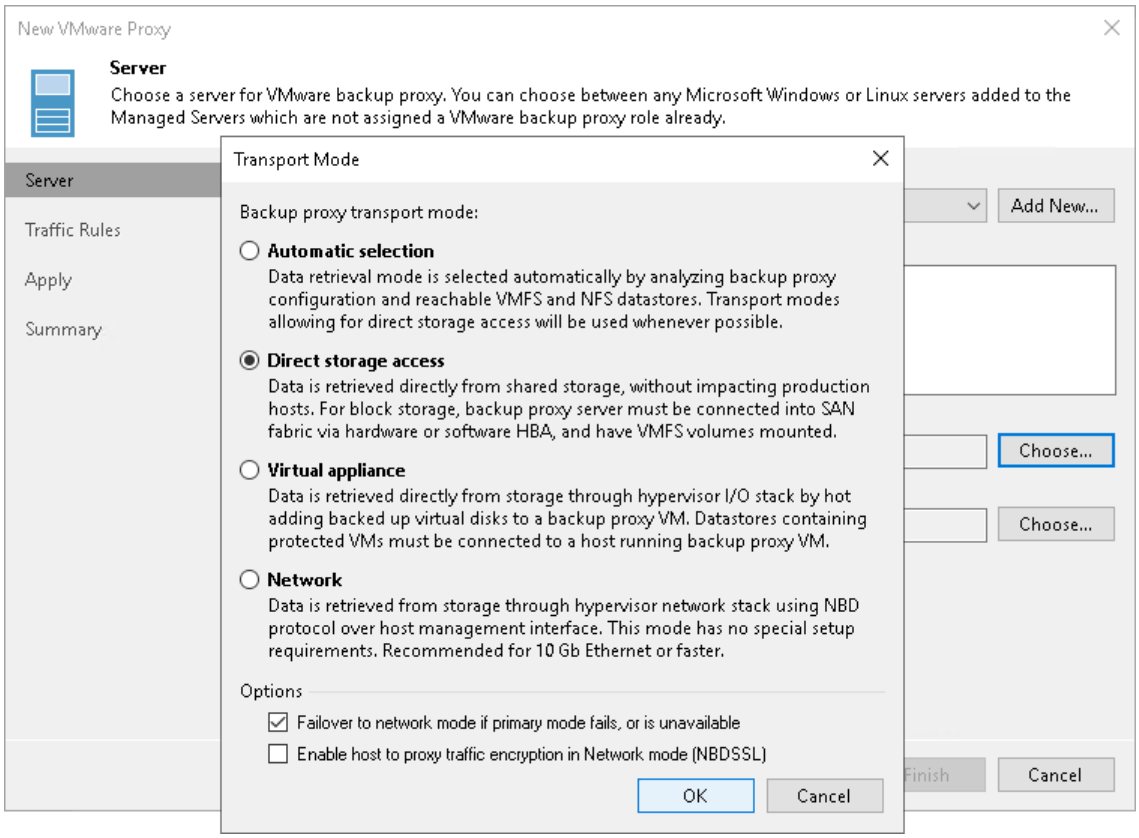

# Adding VMware Backup Proxies

You can configure one or more backup proxies in the backup infrastructure.

Before adding a VMware backup proxy, check prerequisites. Then use the **New VMware Proxy** wizard.

# Step 1. Launch New VMware Proxy Wizard

To launch the **New VMware Proxy** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Proxies** node and select **Add Proxy**. Alternatively, you can click **Add Proxy** on the ribbon.
3. In the **Add Backup Proxy** window, select **VMware backup proxy**.

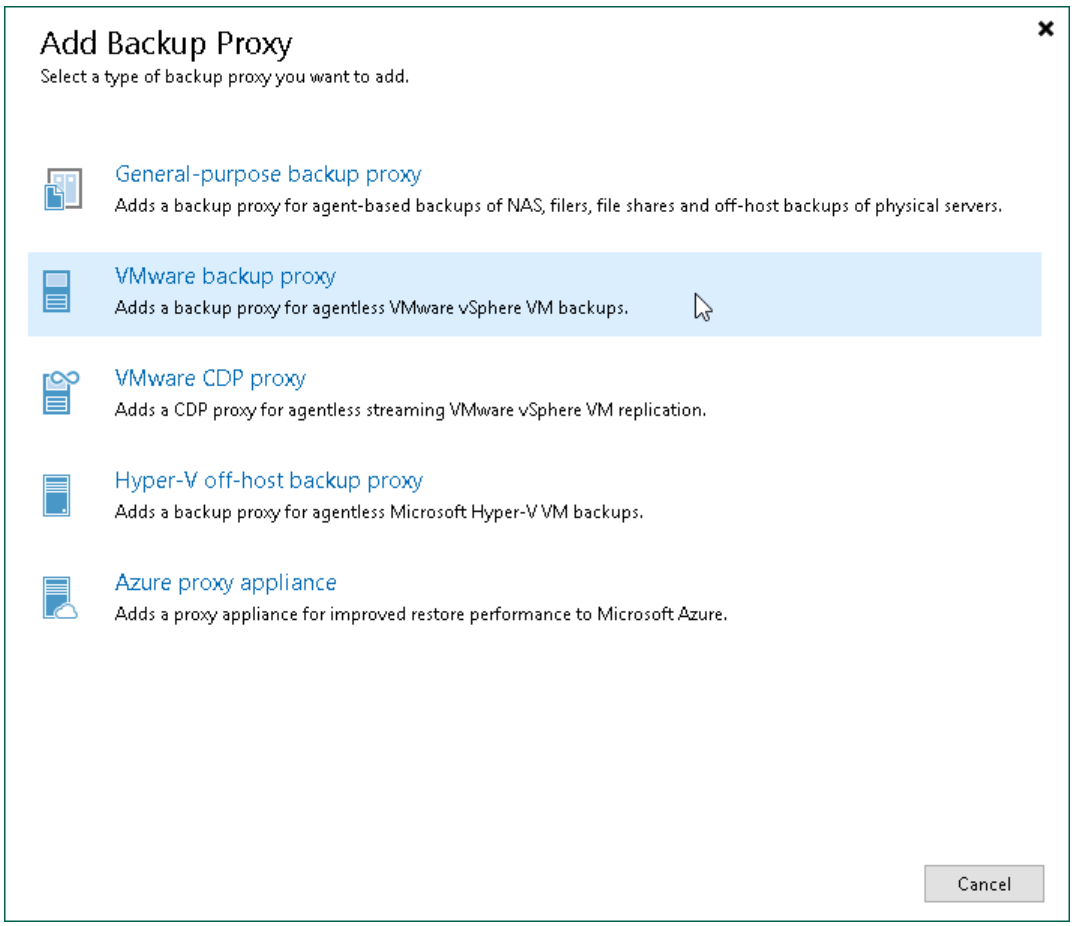

# Step 2. Choose Server

At the **Server** step of the wizard, specify server settings for the backup proxy:

1. From the **Choose server** list, select a Microsoft Windows or Linux server to which you want to assign the backup proxy role.

   If the server is not added to the backup infrastructure, click **Add New** to open the **Add Server** wizard. For more information, see Adding Microsoft Windows Servers or Adding Linux Servers.

   > **NOTE**
   >
   > If a user account specified for the Linux server does not have root or elevated to root permission, Veeam Backup & Replication will direct you to the **Edit Linux Server** wizard. In this wizard, you can change the user account.

2. In the **Proxy description** field, provide a description. The default description contains information about the user who added the backup proxy, date and time when the backup proxy was added.

3. By default, Veeam Backup & Replication analyzes the backup proxy configuration, defines to which datastores it has access and automatically selects the best transport mode depending on the type of connection between the backup proxy and datastores.

   You can select the data transport mode manually. Click **Choose** on the right of the **Transport mode** field. In the opened window, select one of the available modes. For more information, see Transport Modes.

4. In the **Options** section of the **Transport Mode** window, specify additional options for the selected transport mode:

   - [For the Direct storage access and Virtual appliance transport modes] If the primary transport mode fails during the job session, Veeam Backup & Replication will automatically fail over to the Network transport mode. To disable failover, clear the **Failover to network mode if primary mode fails, or is unavailable** check box. However, note that failover to the Network transport mode applies automatically if data cannot be processed by other transport modes even if the check box is disabled.
   - [For the Network mode] You can choose to transfer VM data over an encrypted TLS connection. To do this, select the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box. Traffic encryption puts more stress on the CPU of an ESXi host but ensures secure data transfer.

   > **NOTE**
   >
   > In some cases, the backup proxy may not be able to use some transport modes due to known limitations. For more information, see Transport Modes. If you assign the backup proxy role to a hardened repository, only the **Network** mode will be available. Other transport modes will be grayed out.

5. In the **Connected datastores** field, specify datastores to which the backup proxy has a direct SAN or NFS connection. By default, Veeam Backup & Replication automatically detects all datastores that the backup proxy can access.

   You can set up the list of datastores if you want the backup proxy to work with specific datastores. Click **Choose** on the right of the **Connected datastores** field, choose **Manual selection** and add datastores with which the backup proxy must work in the Direct storage access mode.

> **IMPORTANT**
>
> If the backup proxy is located on the Veeam Cloud Connect service provider side, the list of selected datastores is overridden. In this case, Veeam Backup & Replication automatically detects datastores that the backup proxy can access.

6. In the **Max concurrent** tasks field, specify the number of tasks that the backup proxy must handle in parallel. If this value is exceeded, the backup proxy will not start a new task until one of current tasks finishes.

   Veeam Backup & Replication creates one task per every VM disk. The recommended number of concurrent tasks is calculated automatically based on available resources. Backup proxies with multi-core CPUs can handle more concurrent tasks. For example, for a 4-core CPU, it is recommended that you specify maximum 8 concurrent tasks, for an 8-core CPU — 16 concurrent tasks. When defining the number of concurrent tasks, consider network traffic throughput in the virtual infrastructure.

   > **IMPORTANT**
   >
   > Limitation of concurrent tasks is ignored if the backup proxy acts as a target proxy for a Veeam Cloud Connect job.

7. Click **Next**.

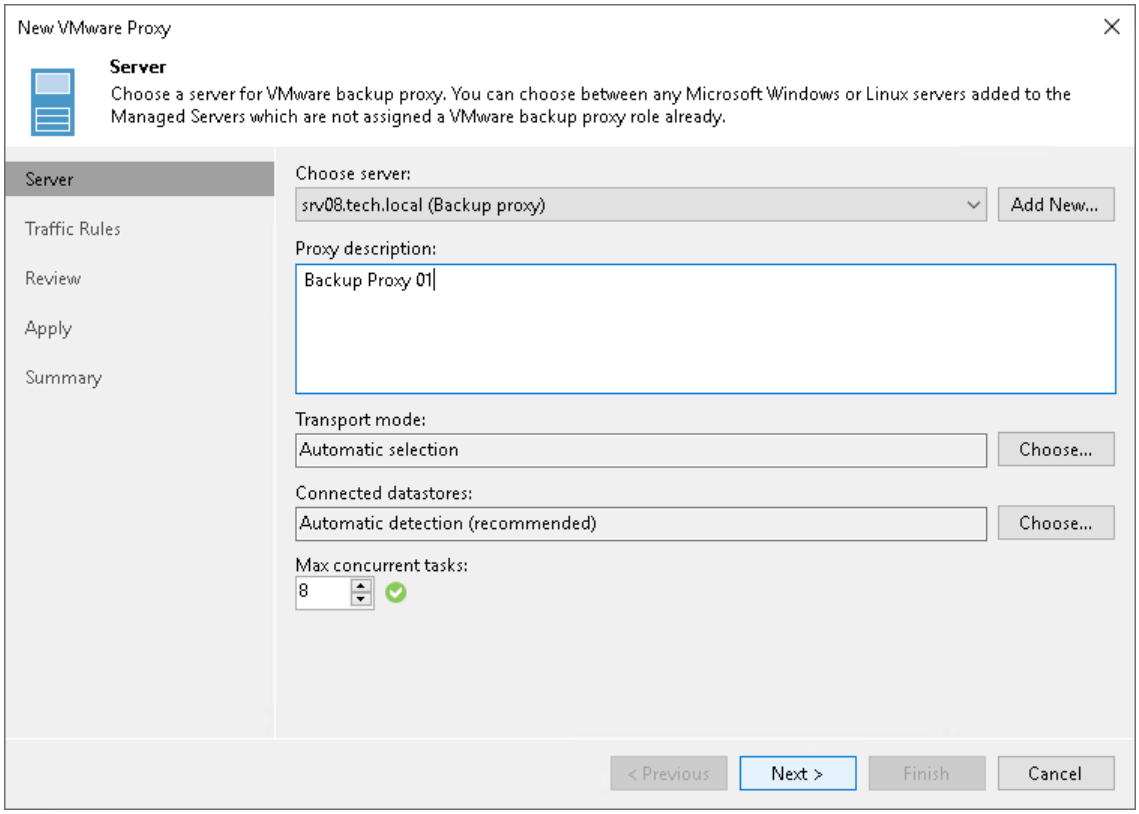

# Step 3. Select Proxy VM

The **Proxy VM** step will appear if the following conditions are met:

- You have selected the **Virtual Appliance** transport mode without the **Failover to network mode if primary mode fails, or is unavailable** check box at the Choose Server step of the wizard.
- Veeam Backup & Replication cannot identify the selected VM: for example, there are two VMs with the same BIOS UUID or BIOS UUID is not specified.

To select a VM from the virtual infrastructure:

1. At the **Proxy VM** step of the wizard, click the **Browse** button.
2. In the **Select Virtual Machine** window, select a VM.

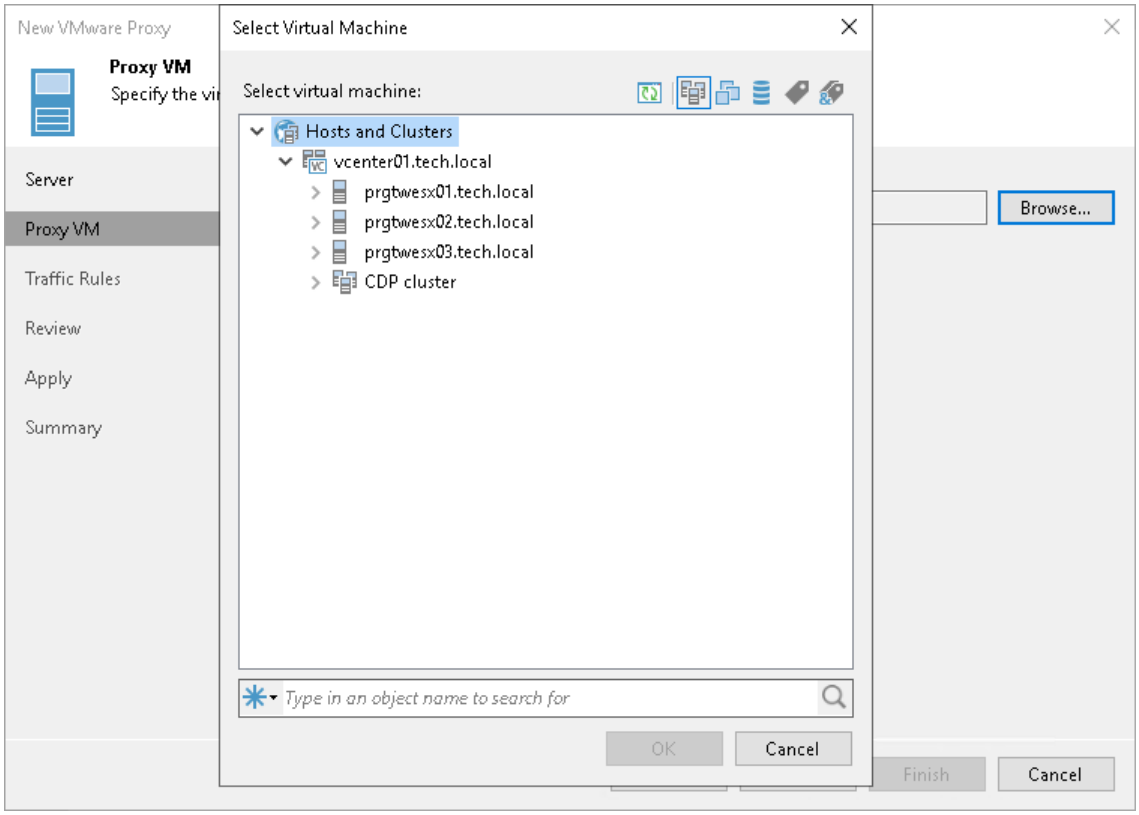

# Step 4. Configure Traffic Rules

At the **Traffic Rules** step of the wizard, configure network traffic rules. These rules help you throttle and encrypt traffic transferred between backup infrastructure components. For more information, see Configuring Network Traffic Rules.

The list of network traffic rules contains only the rules that are applicable to the VMware backup proxy: its IP address falls into the IP range of the rule.

To view rule settings:

1. Select a rule in the list.
2. Click **View** on the right of the rule list.

You can also modify network traffic settings:

1. Click **Manage network traffic rules** link at the bottom of the wizard.
2. In the opened window, you will see all global network traffic rules.
3. Select the rule that you want to modify and click **Edit**.

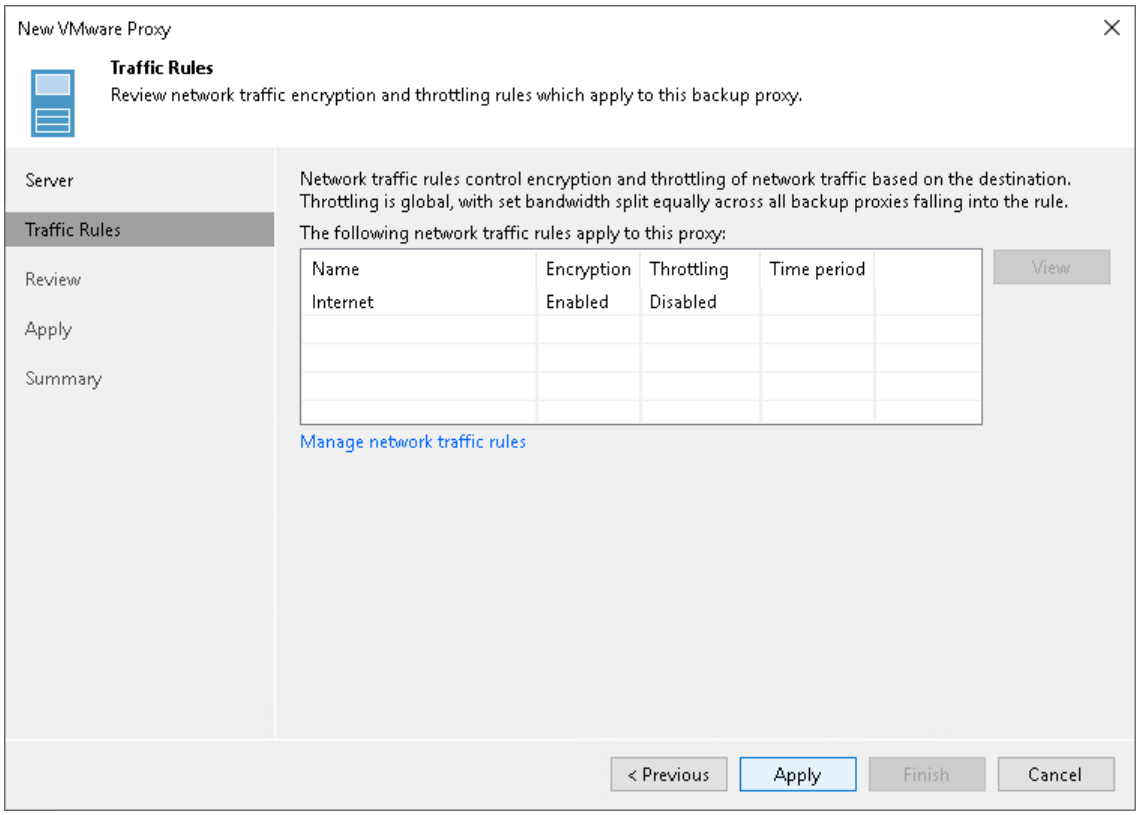

# Step 5. Review Components to Install

At the **Review** step of the wizard, review what Veeam Backup & Replication components are already installed on the server and click **Apply** to start installation of missing components.

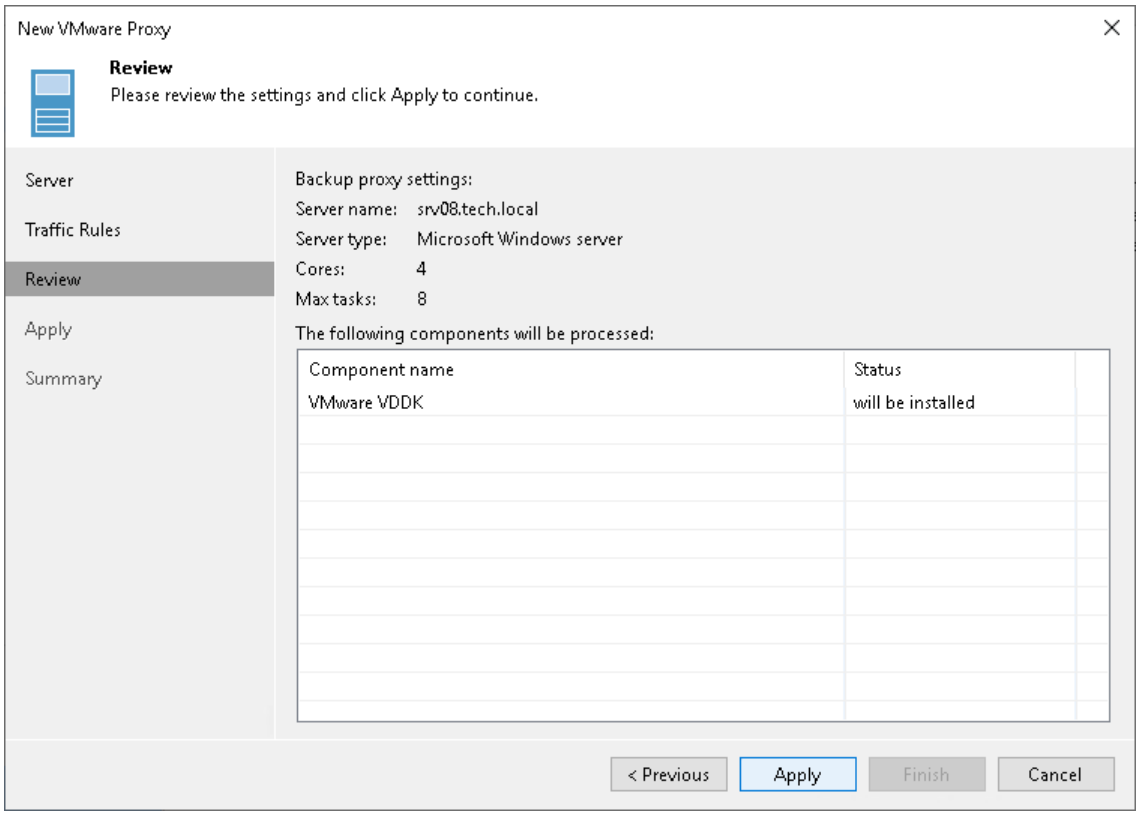

# Step 6. Apply VMware Backup Proxy Settings

At the **Apply** step of the wizard, wait for the VMware backup proxy to be added to the backup infrastructure. Then click **Next**.

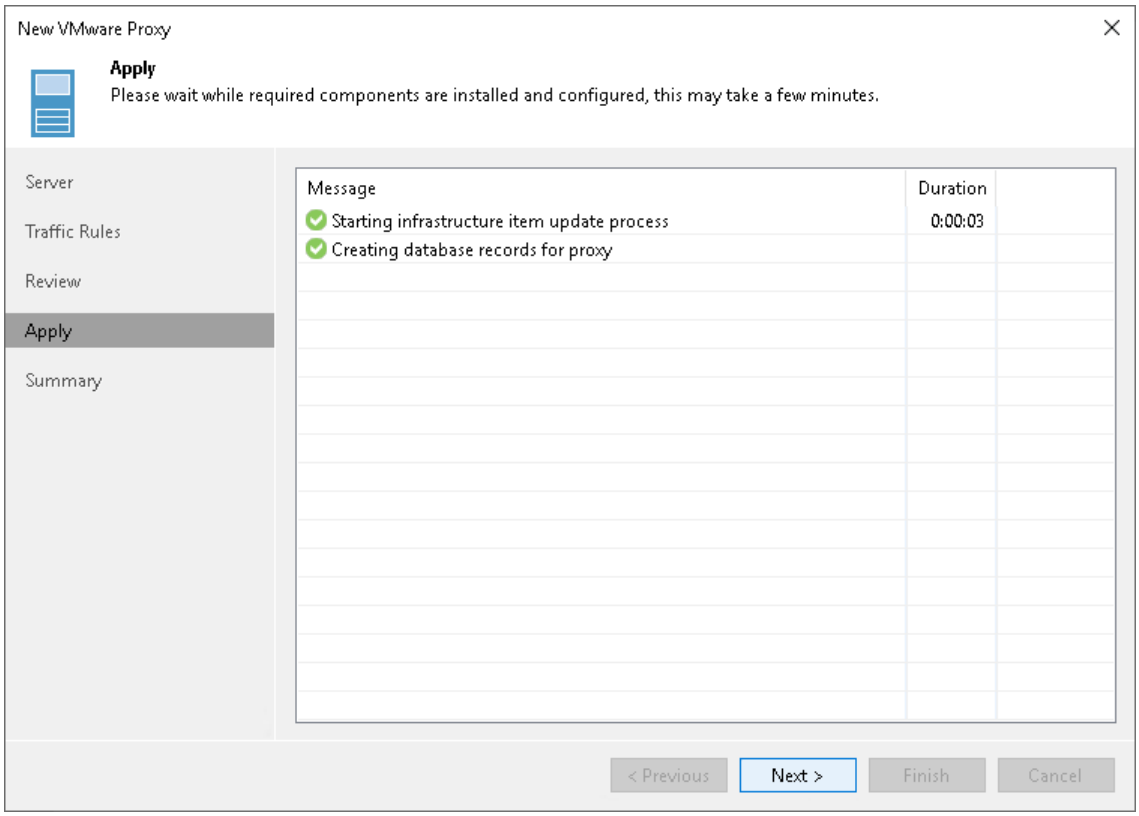

# Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, check that the VMware backup proxy is added. Then click **Finish** to exit the wizard.

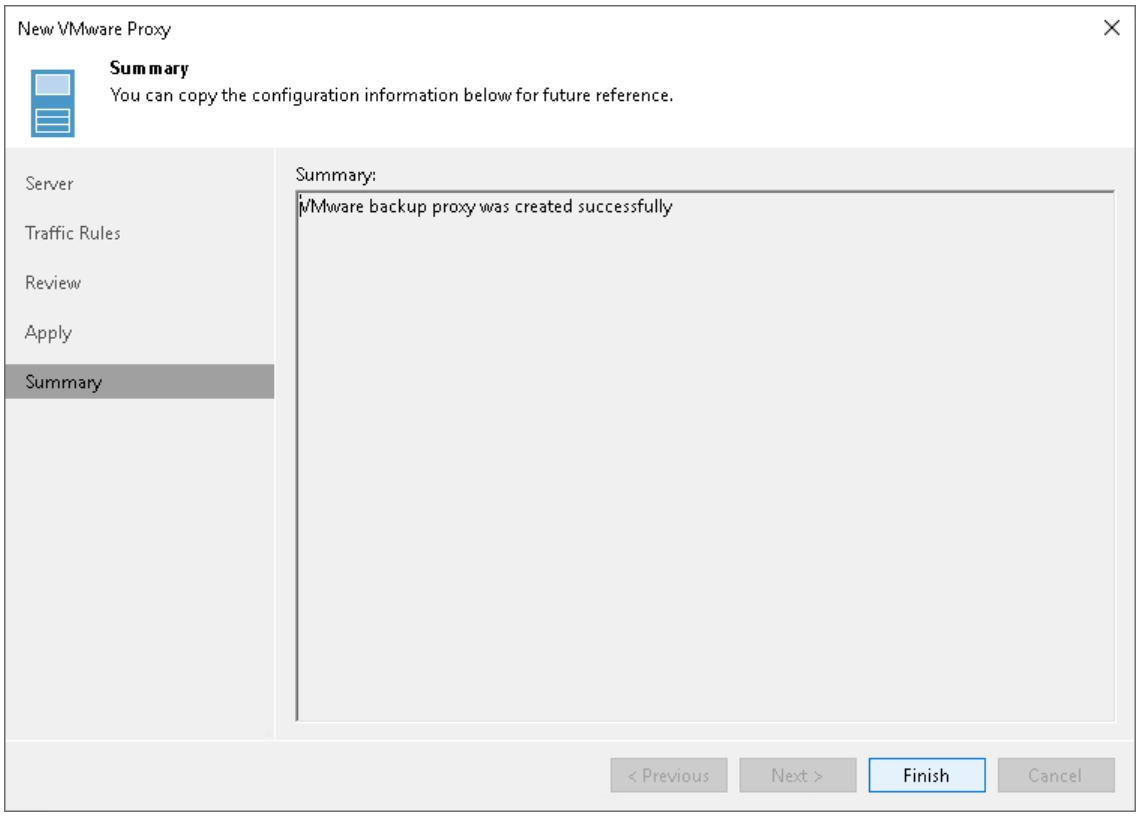

# Editing VMware Backup Proxy Settings

You can edit settings of backup proxies you have configured.

To edit VMware backup proxy settings:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the VMware backup proxy and click **Edit Proxy** on the ribbon or right-click the VMware backup proxy and select **Properties**.
4. Edit VMware backup proxy settings as required.

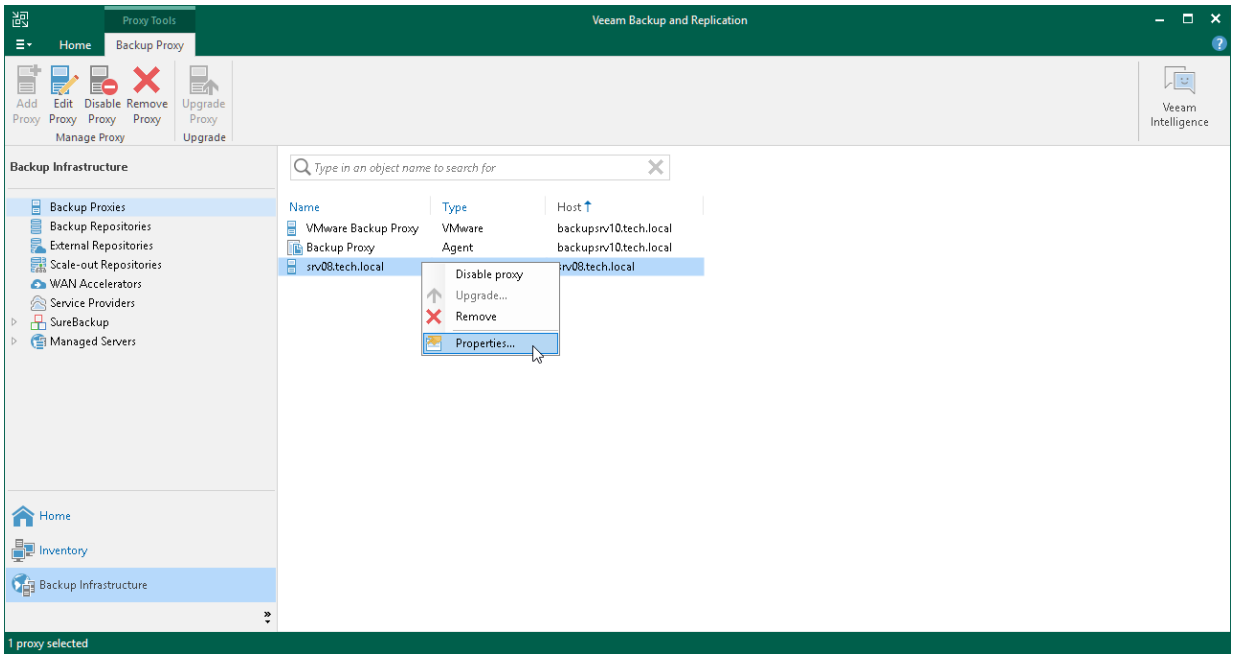

# Disabling and Removing VMware Backup Proxies

You can temporarily disable a VMware backup proxy or remove it from the backup infrastructure.

## Disabling Backup Proxies

When you disable a VMware backup proxy, Veeam Backup & Replication does not use this backup proxy for any jobs configured on the backup server. VMware backup proxy disabling can be helpful if you instruct Veeam Backup & Replication to automatically select backup proxies for jobs and do not want Veeam Backup & Replication to use specific backup proxies.

You can disable all VMware backup proxies, including the default backup proxy installed on the backup server. Do not disable all VMware backup proxies at once. Otherwise, Veeam Backup & Replication will not be able to perform backup, replication and restore operations that use VMware backup proxies.

To disable a VMware backup proxy:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the VMware backup proxy and click **Disable Proxy** on the ribbon or right-click the backup proxy and select **Disable proxy**.

You can enable a disabled VMware backup proxy at any time:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the VMware backup proxy and click **Disable Proxy** on the ribbon once again or right-click the VMware backup proxy and select **Disable proxy**.

## Removing VMware Backup Proxies

You can permanently remove a VMware backup proxy from the backup infrastructure. When you remove a VMware backup proxy, Veeam Backup & Replication unassigns the VMware backup proxy role from the server, and this server is no longer used as a VMware backup proxy. The actual server remains in the backup infrastructure.

You can remove all VMware backup proxies, including the default backup proxy installed on the backup server. Do not remove all VMware backup proxies at once. Otherwise, Veeam Backup & Replication will not be able to perform backup, replication and restore operations that use VMware backup proxies.

You cannot remove a VMware backup proxy that is explicitly selected in any backup, replication or VM copy job. To remove such VMware backup proxy, you first need to delete a reference to this VMware backup proxy in the job settings.

To remove a VMware backup proxy:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.

3. In the working area, select the VMware backup proxy and click **Remove Proxy** on the ribbon or right-click the VMware backup proxy and select **Remove**.

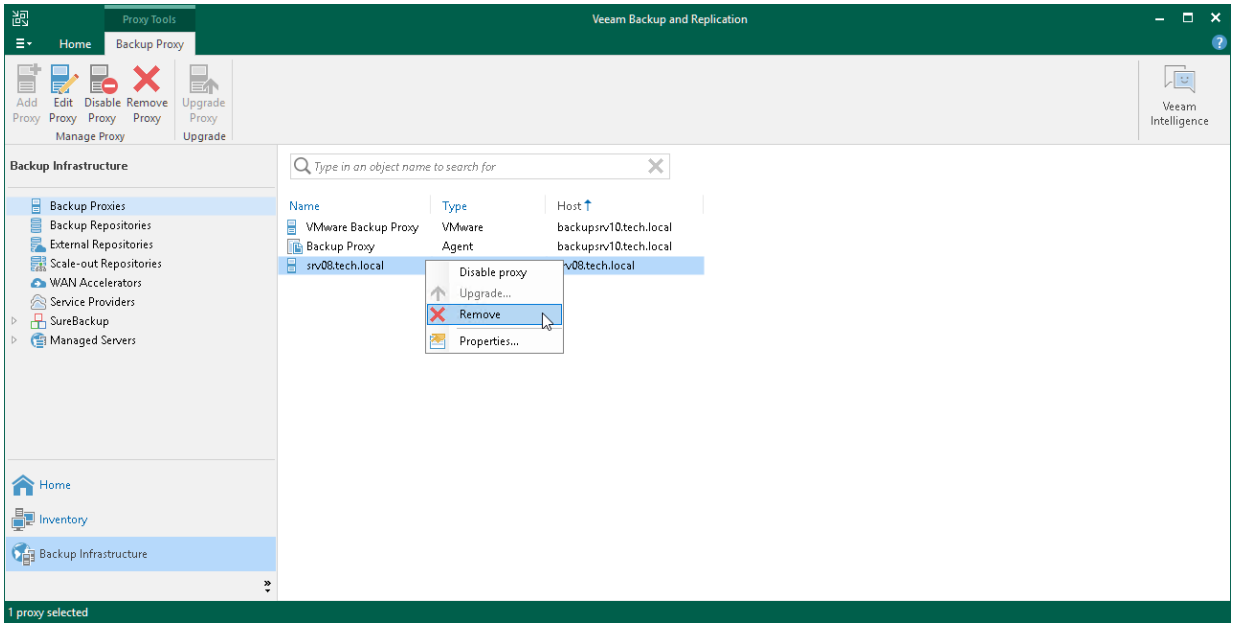

# VMware CDP Proxies

A VMware CDP proxy is a component that operates as a data mover and transfers data between the source and target hosts. Basically, VMware CDP proxy performs the following tasks:

- Receives VM data from the production storage
- Aggregates changed data
- Prepares data for short-term restore points
- Compresses and deduplicates data
- Encrypts and decrypts data
- Sends data to the storage in the disaster recovery site or another VMware CDP proxy

## Usage Scenarios

A VMware CDP proxy is required for continuous data protection. For more information on the backup infrastructure components required for CDP, see Continuous Data Protection (CDP).

## VMware CDP Proxy Deployment

You can assign the role of a VMware CDP proxy to any Windows-based or Linux-based virtual or physical server added to your Veeam Backup & Replication infrastructure. For information on how to add a server, see Adding Microsoft Windows Servers and Adding Linux Servers. For information on how to assign the VMware CDP proxy role, see Adding VMware CDP Proxies.

You need to configure at least two VMware CDP proxies:

- Source proxy in the production site.
- Target proxy in the disaster recovery site.

To optimize performance of several concurrent tasks, you can use several VMware CDP proxies in each site. In this case, Veeam Backup & Replication will distribute the restore workload between available proxies on per-task basis, taking into account proxy connectivity and their current load. For more information on which proxies are considered the most appropriate for continuous data protection, see How CDP Works.

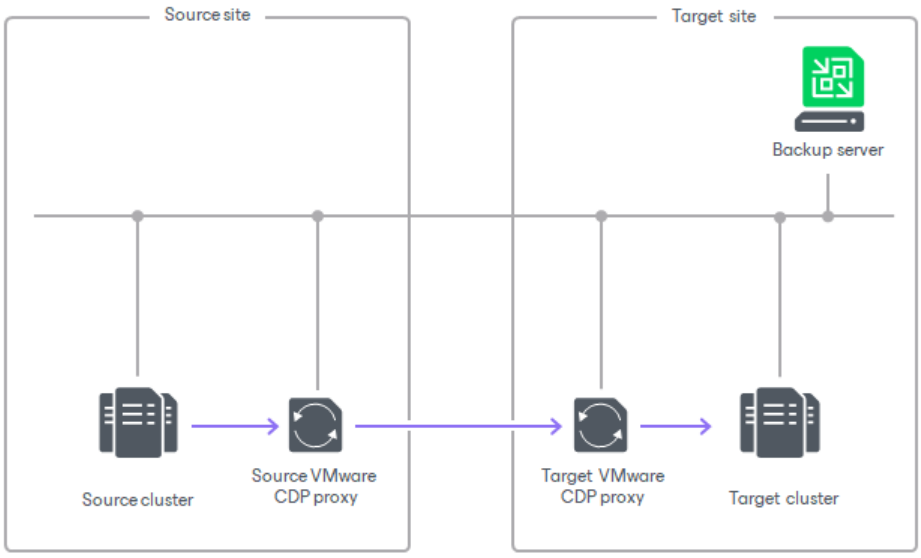

For better performance, use one VMware CDP proxy only as a source or as a target proxy. For example, if you have cross cluster or cross host replication (from ESXi 1 to ESXi 2, and from ESXi 2 to ESXi 1), it is better to have four VMware CDP proxies: one source proxy and one target proxy for data flow from ESXi 1 to ESXi 2, and one source proxy and one target proxy for data flow from ESXi 2 to ESXi 1.

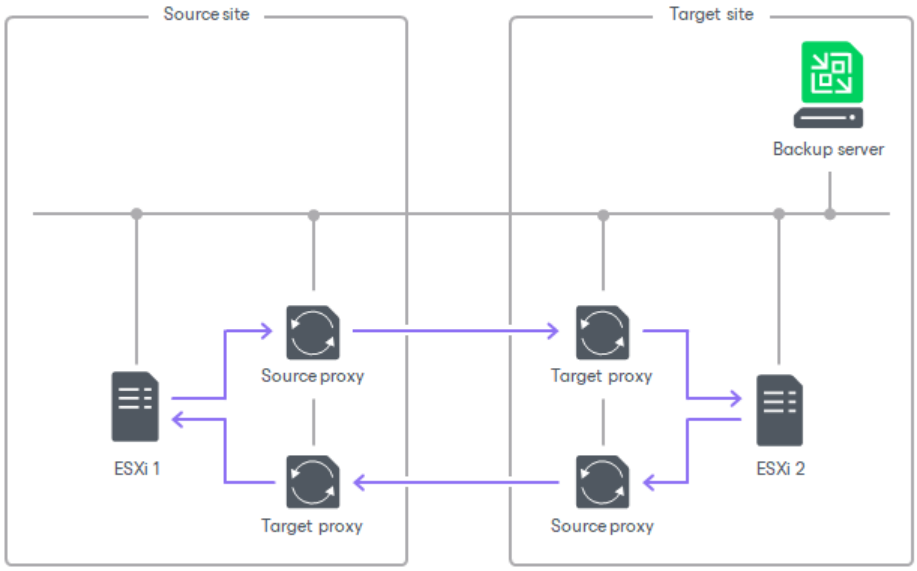


**NOTE**

If you deploy VMware CDP proxies on virtual machines, locate source proxies on the source host and target proxies on the target host.

# VMware CDP Proxy Services and Components

VMware CDP proxies run light-weight services that take a few seconds to deploy. Deployment is fully automated. Veeam Backup & Replication installs the following components and services:

- **Veeam CDP Proxy Service** manages all CDP activities such as data aggregation, data compression and decompression, data transfer and other.
- **Veeam Installer Service** is an auxiliary service that is installed and started on any Windows server once it is added to the list of managed servers in the Veeam Backup & Replication console. This service analyzes the system, installs and upgrades necessary components and services depending on the role selected for the server.
- **Veeam Data Mover** handles traffic sent during failback.

# VMware CDP Proxy RAM and Cache

By default, a VMware CDP proxy stores the received data into RAM. If RAM is less or equal to 16 GB, the VMware CDP proxy uses 50% of the memory for the OS and 50% for data processing. If RAM is larger than 16 GB, the VMware CDP proxy uses 8 GB for the OS and the rest of RAM for data processing. The VMware CDP proxy allocates at least 1 MB of RAM for each processed disk. This protective mechanism guarantees that disk processing will not stop even if some disks produce too much data or cannot be processed.

If a VMware CDP proxy runs out of the memory or cannot allocate space in the memory, the proxy starts storing data into the cache. If the cache and RAM gets full and there is no "free" VMware CDP proxy, Veeam Backup & Replication will use the protective mechanism. However, disk processing performance will be low.

Data is deleted from the cache or memory only after the proxy gets a notification that the target host has successfully saved data sent by the proxy.

# Requirements for VMware CDP Proxy

Before you assign the role of a backup proxy, check the following requirements:

- For system requirements, see System Requirements for VMware CDP Proxy.
- A VMware CDP proxy must be a Windows-based or Linux-based virtual or physical server.
- Before assigning the role of the VMware CDP proxy to a server, you must first add a vCenter server or VMware Cloud Director server to the backup infrastructure.
- [For VMware CDP proxies deployed on physical servers] Fast network between hosts and VMware CDP proxies is required.

# Adding VMware CDP Proxies

You must add to the backup infrastructure VMware CDP proxies that you plan to use for continuous data protection of VMs.

To add a VMware CDP proxy, use the **New VMware CDP Proxy** wizard.

# Step 1. Launch New VMware Proxy Wizard

To launch the **New VMware CDP Proxy** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Proxies** node and select **Add Proxy**. Alternatively, you can click **Add Proxy** on the ribbon.
3. In the **Add Backup Proxy** window, select **VMware CDP proxy**.

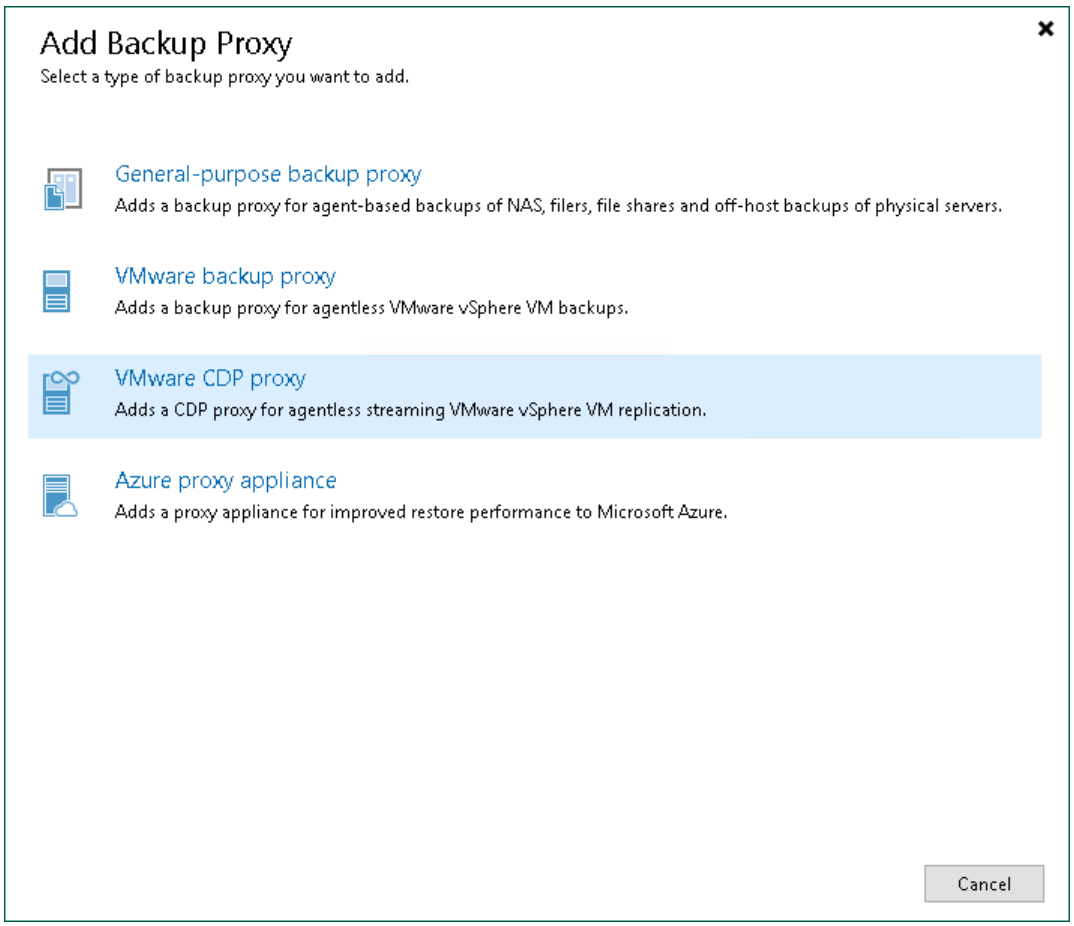

# Step 2. Select Server and Traffic Ports

At the **Server** step of the wizard, select a server which you want to use as the VMware CDP proxy, specify description and ports that will be used for communication:

1. From the **Choose server** drop-down list, select a physical or virtual Microsoft Windows-based or Linux-based server to which you want to assign the VMware CDP proxy role.

   If you have not added a server to the backup infrastructure, click **Add New** to open the **Add Server** wizard, and follow the instructions from the Adding Microsoft Windows Servers or Adding Linux Servers section.

2. In the **Proxy description** field, provide a description for future reference.

3. In the **CDP host traffic port** field, specify a port that the source and target host will use to communicate with the VMware CDP proxy.

   The range of available ports is 33032-33039.

4. In the **CDP proxy traffic port** field, specify a port that the proxy will use to communicate with other VMware CDP proxies.

   The range of available ports is 33032-33039.

   > **IMPORTANT**
   >
   > You must specify different values in the **CDP host traffic port** and **CDP proxy traffic port** fields.

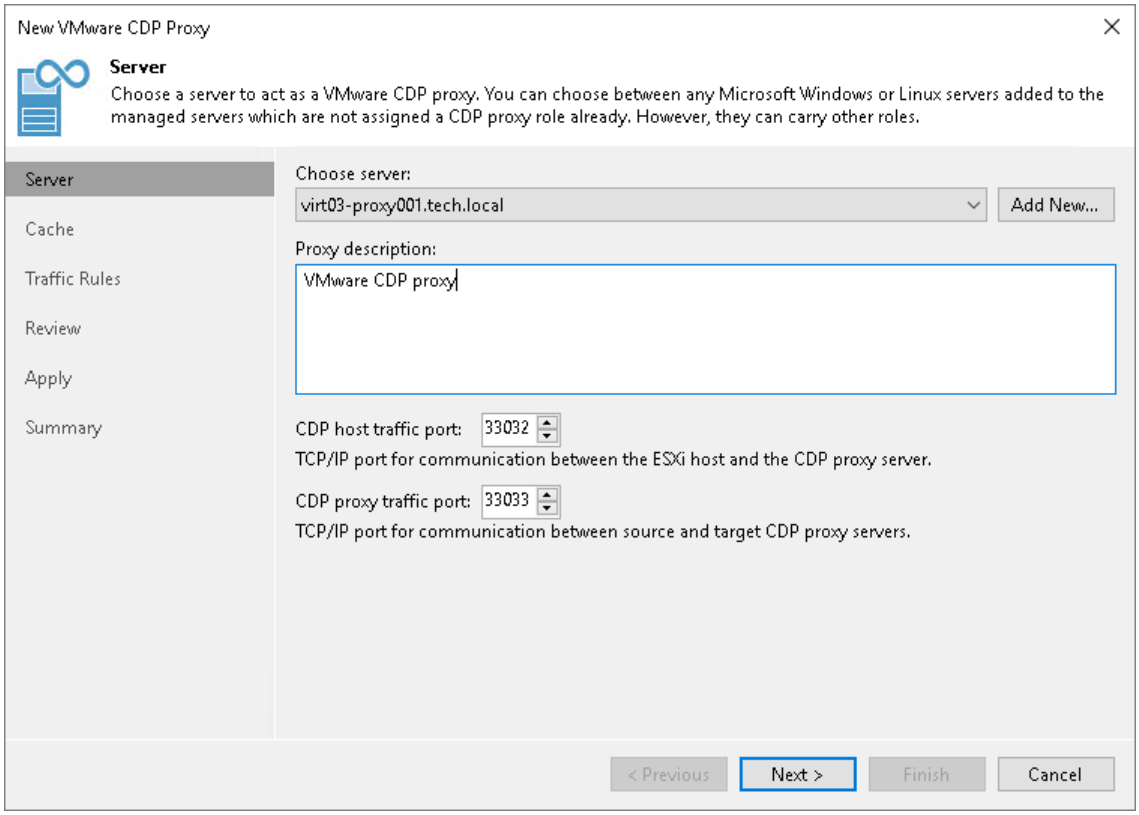

# Step 3. Configure Cache

At the **Cache** step of the wizard, specify the path to a folder where cached data will be stored and how much space can be used for storing the cache.

For more information on the cache, see VMware CDP Proxy Cache.

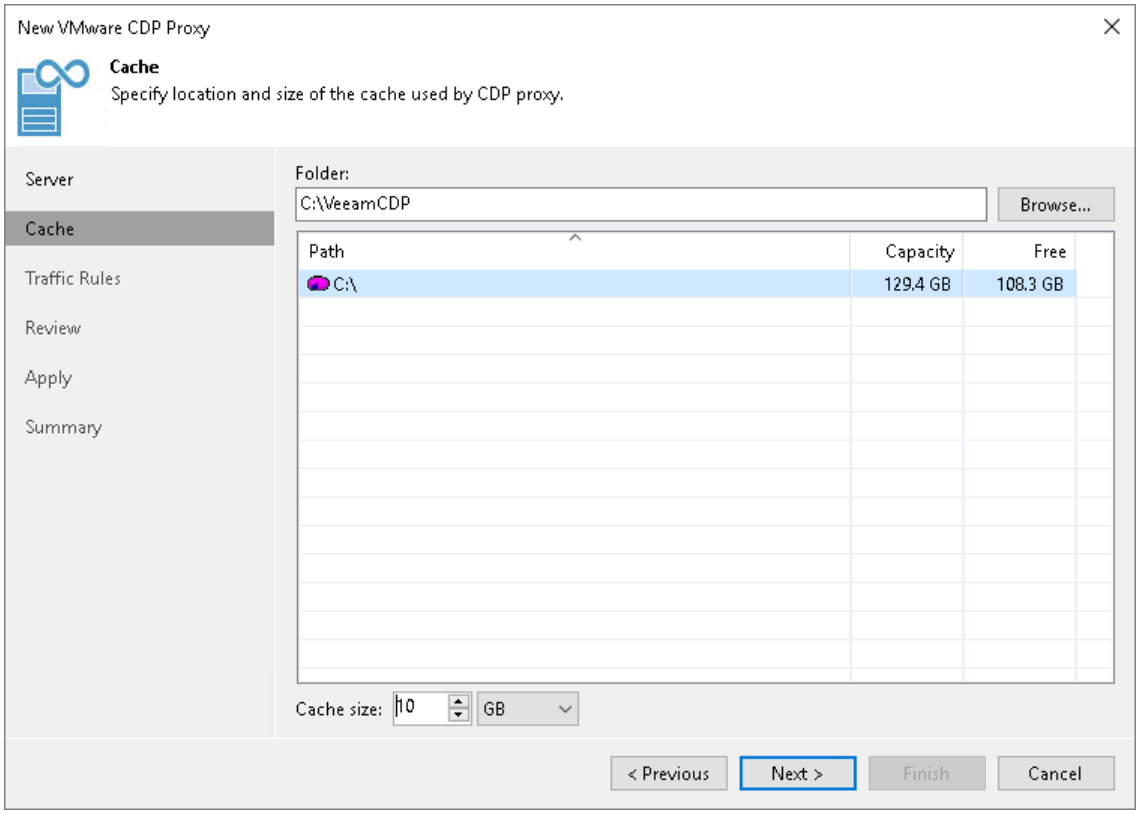

# Step 4. Configure Network Traffic Rules

At the **Traffic Rules** step of the wizard, configure network traffic rules. These rules help you reduce, throttle and encrypt traffic sent between backup infrastructure components. For more information, see Managing Network Traffic.

The list of network traffic rules contains only the rules that are applicable to the VMware CDP proxy — this means that the proxy IP address falls into the IP range of a rule.

To change network traffic settings:

1. Click the **Manage network traffic rules** link.
2. To edit or remove a rule, select the required rule and click **Edit** or **Remove**.
3. To add a rule, click **Add** and follow the instructions from the Configuring Network Traffic Rules section.
4. To choose networks over which Veeam Backup & Replication will send data between backup infrastructure components, click **Networks**. Then follow the instructions in steps 3-6 in the Specifying Preferred Networks section.

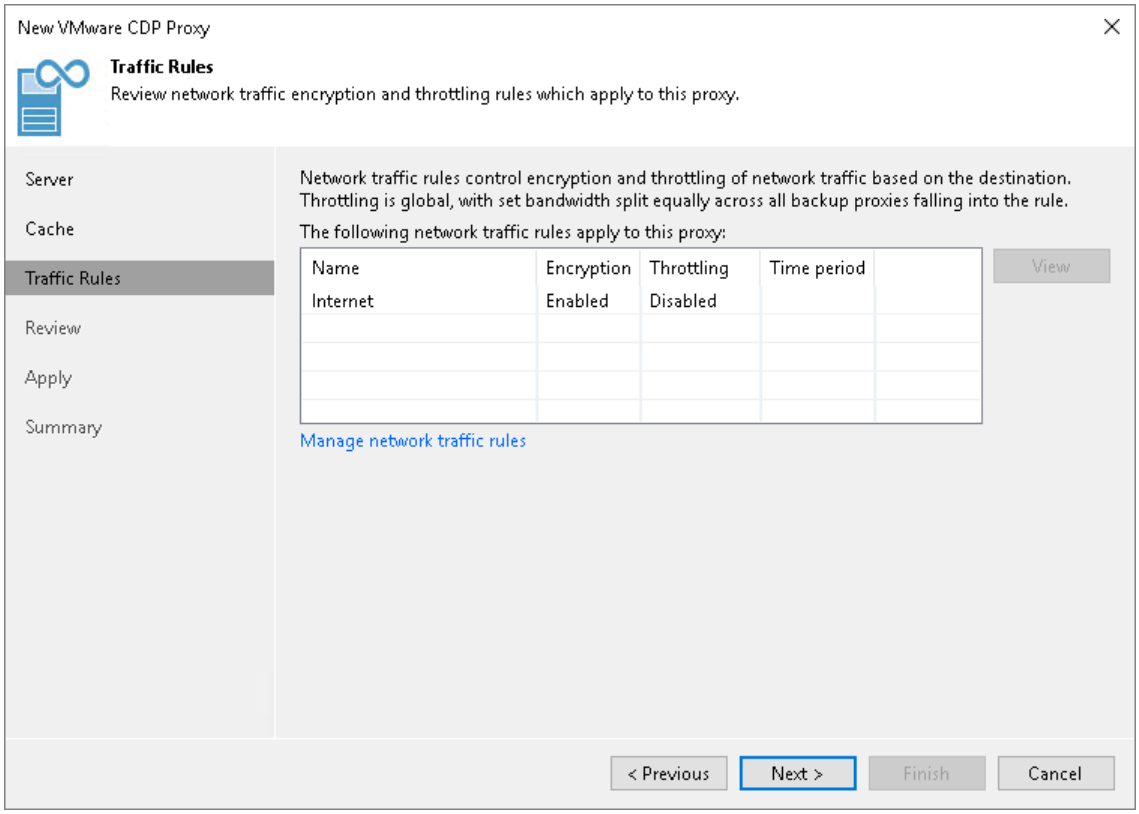

# Step 5. Review Settings and Install Components

At the **Review** step of the wizard, review components which are already installed on the server and which will be installed. Click **Apply** to start installation of missing components.

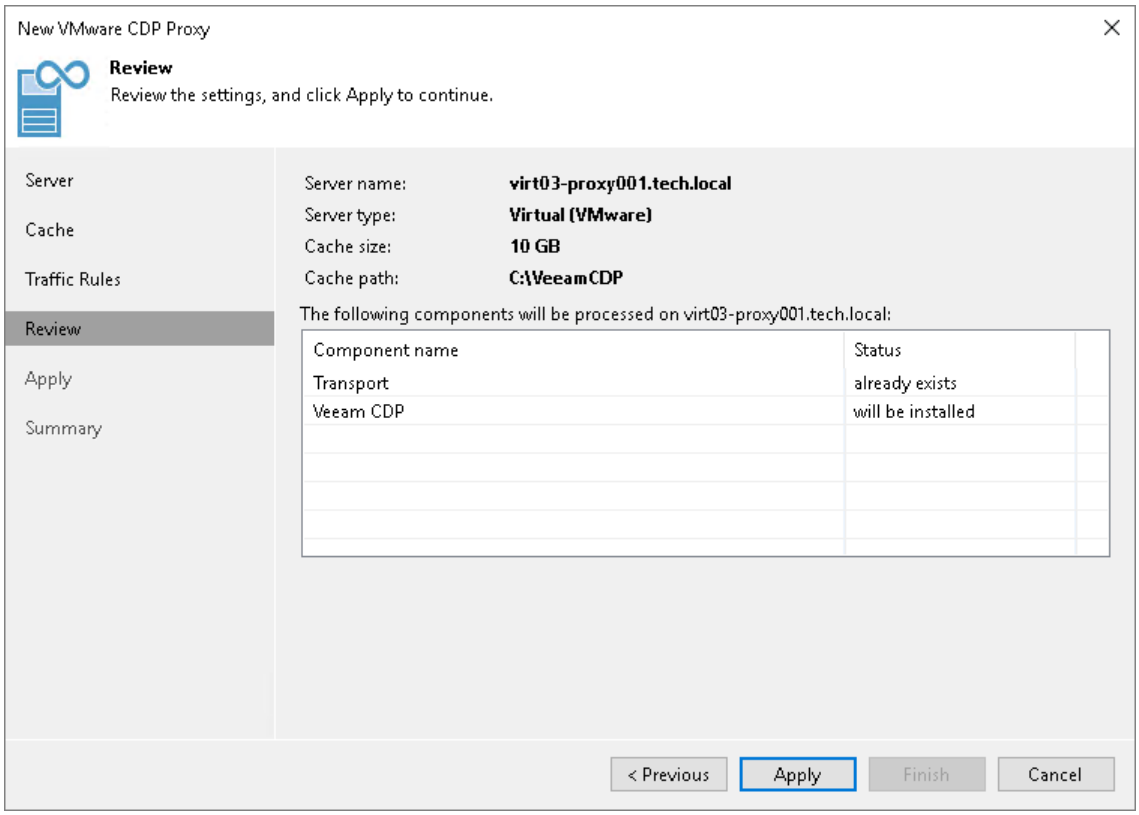

# Step 6. Apply Proxy Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all the required components. Click **Next** to complete the procedure of the VMware CDP proxy role assignment.

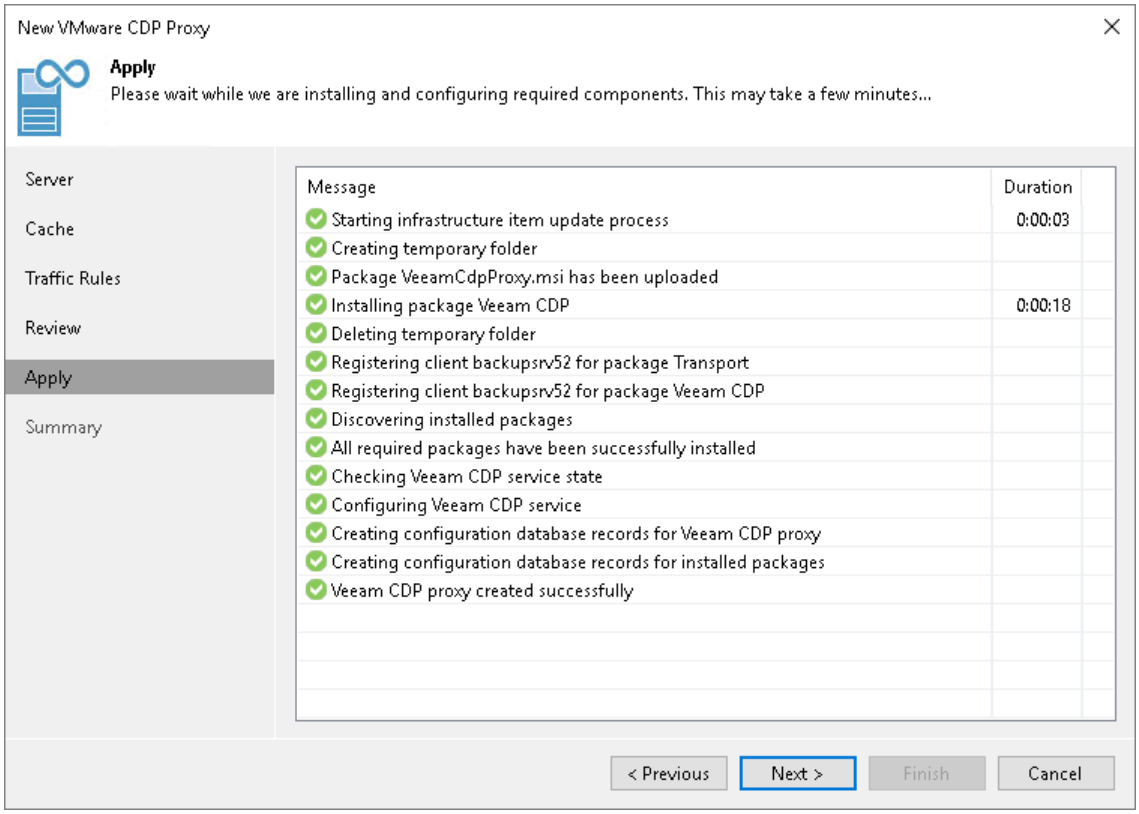

# Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the server added as the VMware CDP proxy and click **Finish** to exit the wizard.

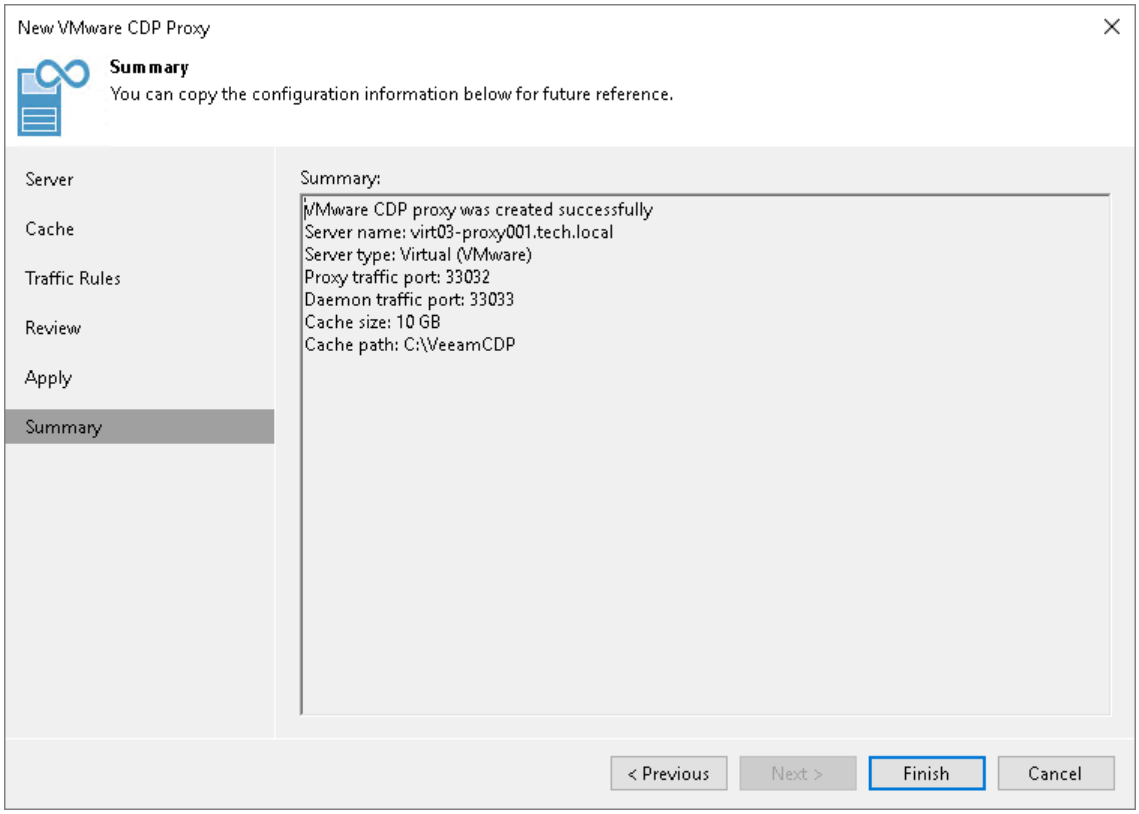

# Cache Repositories

A cache repository is a storage location where Veeam Backup & Replication keeps temporary cached metadata for the data backed up by the file backup jobs and object storage backup jobs. For more information about the cache repository, see the Backup Infrastructure for Unstructured Data Backup section.

# Backup Repositories

A backup repository is a storage location where Veeam keeps backup files, VM copies and metadata for replicated VMs.

> **NOTE**
>
> Consider the following:
>
> - Do not configure multiple backup repositories pointing to the same location or using the same path.
> - Do not configure multiple backup repositories with "nested" paths (when one repository path is a sub-path of another repository), for example:
>
> ```
> /mnt/repo/backups/
> /mnt/repo/backups/production/
> ```

You can use the following storage types as backup repositories:

- **Direct attached storage**. You can add virtual and physical servers as backup repositories:
  - Microsoft Windows server
  - Linux server
  - Hardened Repository
- **Network attached storage**. You can add the following network shares as backup repositories:
  - SMB (CIFS) share
  - NFS share
- **Deduplicating storage appliances**. You can add the following deduplicating storage appliances as backup repositories:
  - Dell Data Domain
  - ExaGrid
  - Fujitsu ETERNUS CS800
  - HPE StoreOnce
  - Infinidat InfiniGuard
  - Quantum DXi
- **Object storage**. You can use cloud storage services as backup repositories. For more information, see Object Storage Repository.

## Related Topics

- External Repositories
- Scale-Out Backup Repositories

# External Repositories

An external repository is a read-only repository. You can use Veeam Backup & Replication to copy, import and restore backups created by Veeam Backup for AWS, Veeam Backup for Microsoft Azure and Veeam Backup for Google Cloud from external repositories to the cloud or on-premises repositories. This way, you can migrate workloads between cloud, on-premises and virtual infrastructures.

Veeam Backup & Replication supports the following types of external repositories:

- Amazon S3 (with Standard storage class assigned)
- Azure Blob (Hot and Cool access tier)
- Google Cloud (with Standard storage class assigned)

## Deployment

To start working with backups created by Veeam Backup for AWS, Veeam Backup for Microsoft Azure and Veeam Backup for Google Cloud, you must add a repository that contains backups of Amazon EC2 instances, Azure VMs or Google Cloud VMs to the Veeam Backup & Replication infrastructure as an external repository. For more information, see Adding External Repositories.

## Usage Scenarios

You can perform the following operations:

- Copy backups to on-premises repositories.
- Restore EC2 instances, Azure VMs and Google Cloud VMs to AWS.
- Restore Azure VMs, EC2 instances and Google Cloud VMs to Microsoft Azure.
- Restore Google Cloud instances, EC2 instances and Azure VMs to Google Cloud.
- Restore guest OS files and folders.
- Export disks of EC2 instances, Azure VMs and Google Cloud VMs.
- Govern retention policies.
- Perform Instant Recovery to VMware vSphere.
- Perform Instant Recovery to Microsoft Hyper-V.

> **NOTE**
>
> Consider the following:
>
> - During the process of copying backups, or restore to Amazon EC2 or Microsoft Azure, data of EC2 instances and Azure VMs may migrate from one geographic location to another. In this case, Veeam Backup & Replication displays a warning and stores a record about data migration to job or task session details. For more information, see Locations.
> - You cannot use an external repository as a target for backup or backup copy jobs.

# How External Repository Works

Continue with this section to learn more on how external repositories work.

## Ownership

Ownership defines what entity can own an Amazon S3, Azure Blob or Google Cloud storage repository at a time.

# How Does Taking Ownership Occur

After Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud has finished its initial backup job session, it becomes the rightful owner of both a storage repository and backup files in that repository.

Taking ownership of such a repository along with its backup files by the Veeam Backup & Replication client consists of the following consecutive steps:

- Step 1. Taking ownership of a repository.

  Reclaiming ownership of a repository occurs every time a client adds object storage as an external repository to the Veeam Backup & Replication console.

- Step 2. Taking ownership of backup files in the repository.

  Becoming an owner of backup files in object storage is only possible after Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud launches the backup job session which is referring to backups you are trying to take ownership of (for example, backup files that are located in the repository you have added at the step one).

  During its session, Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud verifies the owner of a repository and if it finds out that the owner has been changed, it changes the owner of each backup file in that repository by creating a new checkpoint that refers to a new rightful owner. Such a checkpoint will be used during subsequent sessions of a backup job to repeat owner verification.

It is possible, however, that after you add an external repository, you never launch the associated backup job again. In such a scenario, Veeam Backup & Replication will not be able to manage retention policies, but you will still be able to restore external repository data, remove backups from external repositories and perform backup copy.

# Taking Ownership by Another Veeam Backup & Replication Client

Ownership of a repository along with its backup files can only be granted to one Veeam Backup & Replication client at a time.

Therefore, if a client *A* adds an external repository that has previously been added by the client *B*, the client *B* completely loses its ownership privileges.

Losing privileges means that the client *B* will no longer be able to manage retention policies. All the previously created backup copy jobs and restore sessions will be failing.

Ownership, however, can easily be reclaimed by re-adding the same object storage anew.

## Cache

Veeam Backup & Replication caches data that is being retrieved from external repositories every time a backup copy job or restore session is performed.

Such an approach helps not only to reduce the number of cost-expensive operations incurred by AWS, Microsoft Azure or Google Cloud, but also decrease the amount of traffic being sent over the network.

Consider the following:

- Cache is created on a gateway server while the following activities are being processed:
  - Backup copy jobs.
  - Restore sessions.
- Cache is not created upon the addition of an external repository to the Veeam Backup & Replication console.
- Cache consists of metadata of blocks being retrieved from external repositories.
- Cache is written to:
  - On a Windows-based gateway server: `C:\ProgramData\Veeam\ExternalCache`
  - On a Linux-based gateway server: `/var/veeam/ExternalCache`
- Cache is reused and updated during each subsequent execution of a backup copy job or restore session.
- Cache size limit is 10GB. Once reached, Veeam will purge obsolete cache by 20% preserving most frequently used parts. Purging is done by the maintenance job.
- Cache is removed in the following cases:
  - An external repository has been removed from the Veeam Backup & Replication console.
  - The gateway server has been changed in the settings of the external repository configuration.
  - The backup file has been removed from the external repository.
  - During the maintenance job session.
- Cache can be removed manually without affecting the backup infrastructure in any negative way.

## Encryption

Backups that reside in Amazon S3 buckets, Azure Blob storage and Google Cloud storage can be encrypted by Veeam Backup for AWS, Veeam Backup for Microsoft Azure and Veeam Backup for Google Cloud. Moreover, password for such encrypted backups may change on a daily basis. For example, there is a backup chain in Amazon S3 bucket that consists of 10 restore points, each of which was encrypted with different password. Therefore, there are 10 different passwords in total that have been used.

To be able to decrypt each restore point in such a backup chain without having to provide each previously used password separately, Veeam Backup & Replication implements the ability of backward hierarchical decryption.

Backward hierarchical decryption requires you to provide only the latest password so that all the previously created restore points can be decrypted as well. For example, there are three restore points: A, B, and C. The point A was encrypted with password 1, B with password 2, and C with password 3. Therefore, you will only need to know the password of the C point to decrypt points C, B, and A.

If you plan to perform data recovery operations with encrypted backups, you must provide a password for these backups in the **External Repository** wizard:

- [For Veeam Backup for AWS] At the Encryption step of the wizard.
- [For Veeam Backup for Microsoft Azure] At the Encryption step of the wizard.
- [For Google Cloud] At the Bucket step of the wizard.

## Managing Retention Policy

A retention policy is set in the backup policy settings of Veeam Backup for AWS, Veeam Backup for Microsoft Azure and Veeam Backup for Google Cloud and defines the number of restore points to keep in repositories.

Retention policies are initially managed by Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud until a Veeam Backup & Replication client reclaims ownership of a repository and all the backup files in such a repository.

Once ownership is reclaimed, Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud ceases to govern retention policies and the Veeam Backup & Replication client becomes responsible for removing obsolete restore points from repositories.

The restore points that fall under the retention policy will be removed upon the next successful session of the maintenance job.

When a Veeam Backup & Replication client removes an external repository from its scope, it relinquishes its ownership which then goes directly to Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud until another Veeam Backup & Replication client reclaims it anew and so forth.

> **IMPORTANT**
>
> A retention policy can only be applied by the Veeam Backup & Replication client that is the rightful owner of an Amazon S3 object storage repository and its backup files.

## Maintenance Job

The maintenance job is a system job that is executed automatically every 24 hours.

The maintenance job does the following:

- Purges obsolete restore points that fall under the retention policy.

  To be able to purge obsolete restore points from external repositories due to the retention policy threshold, a Veeam Backup & Replication client must be the owner of a repository and its backup files.
- Purges cache by 20% of the size limit. By default, the size limit is 10GB.
- Saves its session results to the configuration database.

  The session results can be found in the **History** view under the **System** node.

# Adding External Repositories

You can add the following types of external repositories:

- External Amazon S3 Storage
- External Azure Blob Storage

- External Google Cloud Storage

## Adding External Amazon S3 Storage

To add an Amazon S3 storage as an external repository, do the following:

1. Launch the New External Repository wizard.
2. Specify the repository name and description.
3. Specify the cloud account.
4. Specify cloud storage details.
5. Configure encryption.
6. Wait for the repository to be added to the infrastructure.
7. Finish working with the wizard.

## Step 1. Launch New External Repository Wizard

To launch the **New Backup Repository** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane select the **External Repositories** node and click **Connect to Repository** on the ribbon. At the **Connect to External Repository** window, select **Veeam Backup for AWS**.
- Open the **Backup Infrastructure** view. In the inventory pane right-click the **External Repositories** node and select **Connect to...** At the **Connect to External Repository** window, select **Veeam Backup for AWS**.

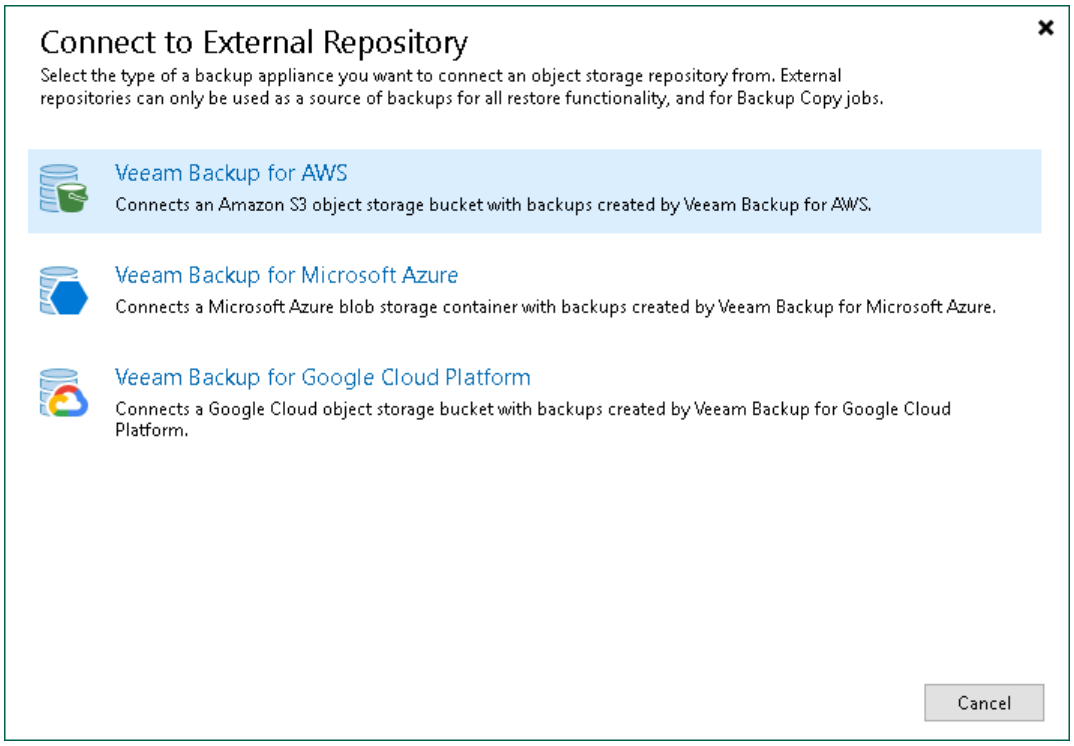

## Step 2. Specify External Repository Name

At the **Name** step of the wizard, specify a name and description for the external repository.

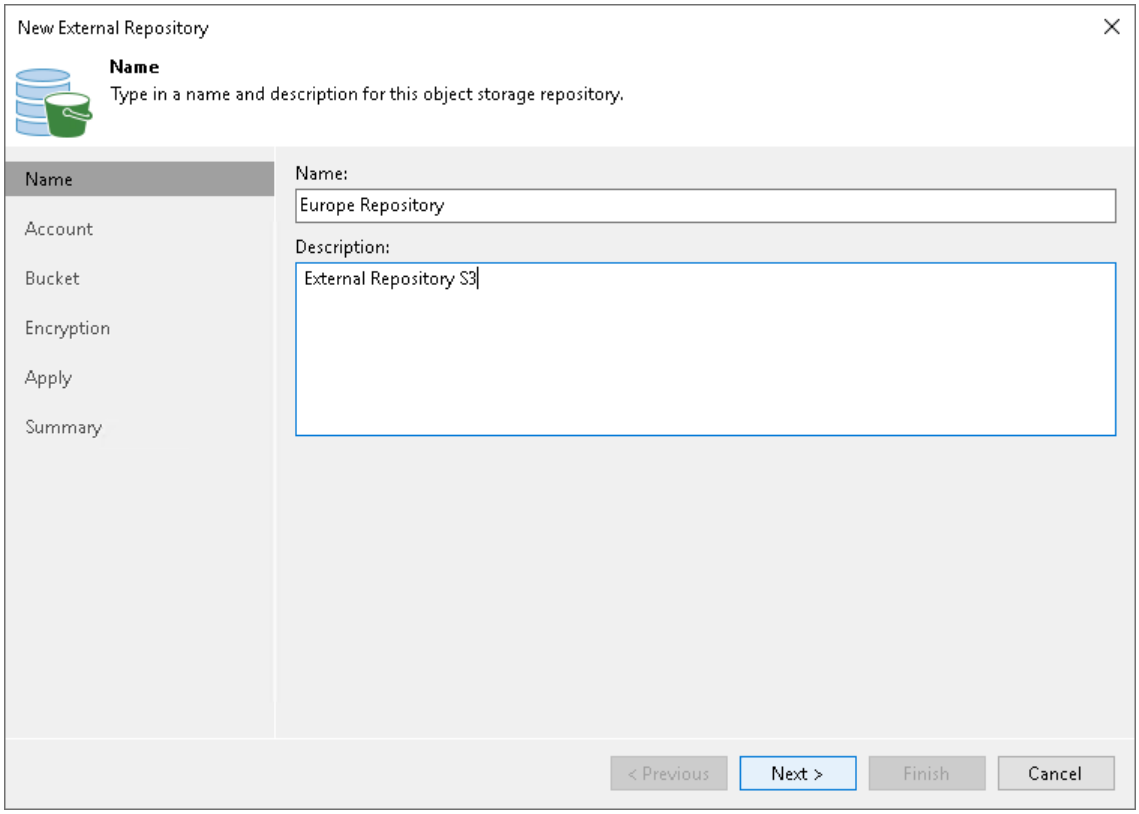

## Step 3. Specify Cloud Storage Account

At the **Account** step of the wizard, specify Amazon S3 connection settings:

1. From the **Credentials** drop-down list, select user credentials to access an Amazon S3 bucket with Amazon EC2 instance backups.

   If you have not set up credentials beforehand in the Cloud Credentials Manager, click the **Manage cloud accounts** link or click **Add** on the right to add the necessary credentials.

2. From the **Data center region** drop-down list, select the AWS region where the Amazon S3 bucket is located: *Global*, *GovCloud (US)*, or *China*.

3. From the **Gateway server** drop-down list, select a gateway server that will be used to access the Amazon S3 bucket. We recommend that you use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited.

   The gateway server caches data when you copy backups or perform restore operations. The gateway server helps you decrease the amount of traffic being sent over the network and reduce data transfer costs. For more information on caching data, see Cache.

   By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. If the Veeam Backup & Replication server resides in a region that differs from the Amazon region where your Amazon S3 bucket resides, choose a server that is located close to the bucket. You can choose any Microsoft Windows or Linux server that is added to your Veeam Backup & Replication infrastructure and has internet connection. Note that the server must be added to the infrastructure beforehand. For more information on how to add a server, see the Adding Microsoft Windows Servers and Adding Linux Servers sections.

   > **NOTE**
   >
   > Consider the following:
   >
   > - On the gateway server, Veeam Backup & Replication deploys Veeam Data Mover that handles traffic sent when you work with external repository data. If Veeam Data Mover becomes outdated, you must upgrade it as described in section Upgrading External Repositories.
   > - If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.

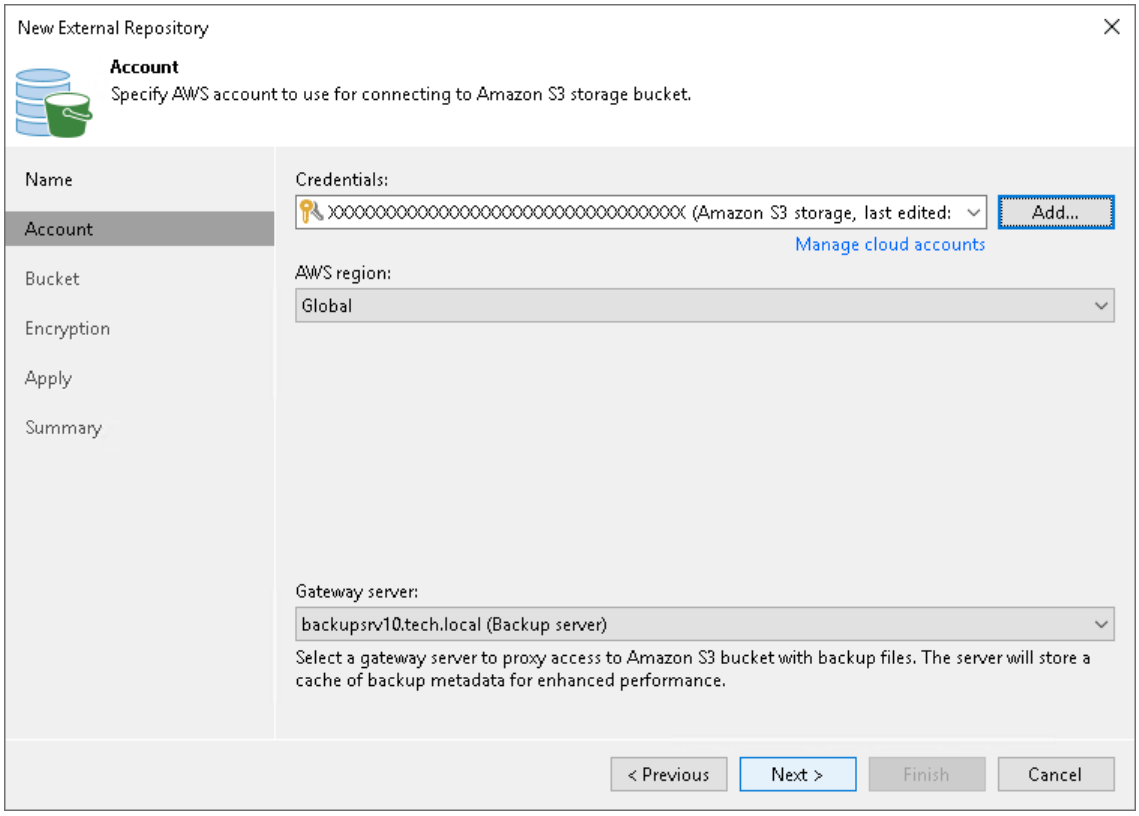
New External Repository
Account
Specify AWS account to use for connecting to Amazon S3 storage bucket.
Name
Account
Bucket
Encryption
Apply
Summary
Credentials:
(Amazon S3 storage, last edited:
Add...
Manage cloud accounts
AWS region:
Global
Gateway server:
backupsrv10.tech.local (Backup server)
Select a gateway server to proxy access to Amazon S3 bucket with backup files. The server will store a cache of backup metadata for enhanced performance.
< Previous
Next >
Finish
Cancel

## Step 4. Specify Cloud Storage Details

At the **Bucket** step of the wizard, specify an Amazon S3 bucket and folder where Amazon EC2 instance backups reside:

1. From the **Data center** drop-down list, select an AWS region where the Amazon S3 bucket is located.
2. From the **Bucket** drop-down list, select the necessary Amazon S3 bucket where EC2 instance backups reside.
3. Click **Browse** to select a folder in the Amazon S3 bucket where EC2 instance backups reside.

   If you do not see the required folder, make sure that the repository you are trying to add is created on the Veeam Backup for AWS server.

**NOTE**

Consider the following:

- Only Standard storage class is supported.
- If another Veeam Backup & Replication client has already added the same folder, you will be prompted whether to reclaim ownership of such a folder. For more information about ownership, see Ownership.

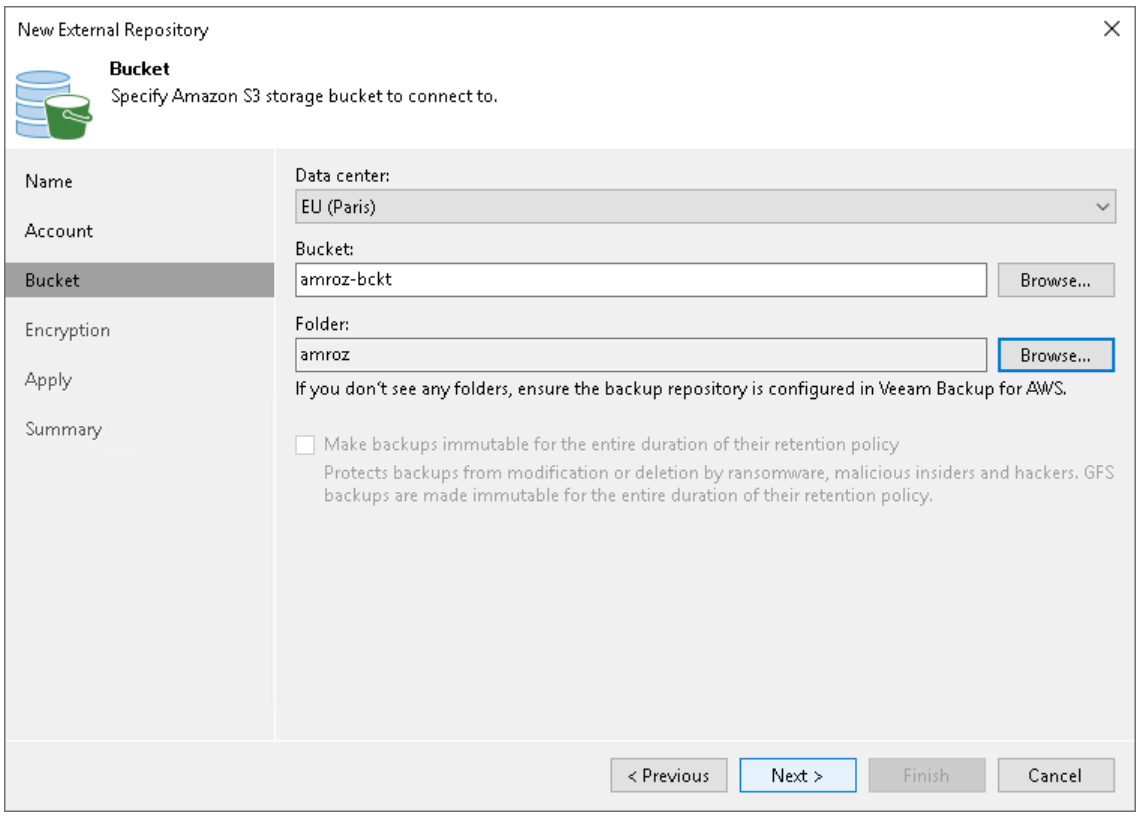

## Step 5. Configure Encryption

At the **Encryption** step of the wizard, provide a key for decryption:

- If data in the external repository is encrypted with a key management service (KMS) customer managed key (CMK), Veeam Backup & Replication shows the used key. In this case, Veeam Backup & Replication will automatically decrypt the backups.

  > **NOTE**
  >
  > The cloud account you specified at the **Account** step of the wizard, must have access to KMS customer managed keys. For more information about KMS key permissions, see the AWS KMS Encryption section in the Veeam Backup for AWS Guide.

- If data in the external repository is encrypted with a password, select **Enable backup file encryption** and then click **Perform Veeam encryption with the following password**. From the drop-down list, select the password that must be used to decrypt the data. If the password is correct, Veeam Backup & Replication will automatically decrypt the backups.

  If you have not added the password beforehand, click the **Manage passwords** link or the **Add** button to add the necessary password. For more information on adding passwords, see Creating Passwords.

  If you do not specify the decryption password, you can do it later. For more information, see Viewing External Repository Data.

For more information on how backups are encrypted, see the S3 Repository Encryption section in the Veeam Backup for AWS User Guide.

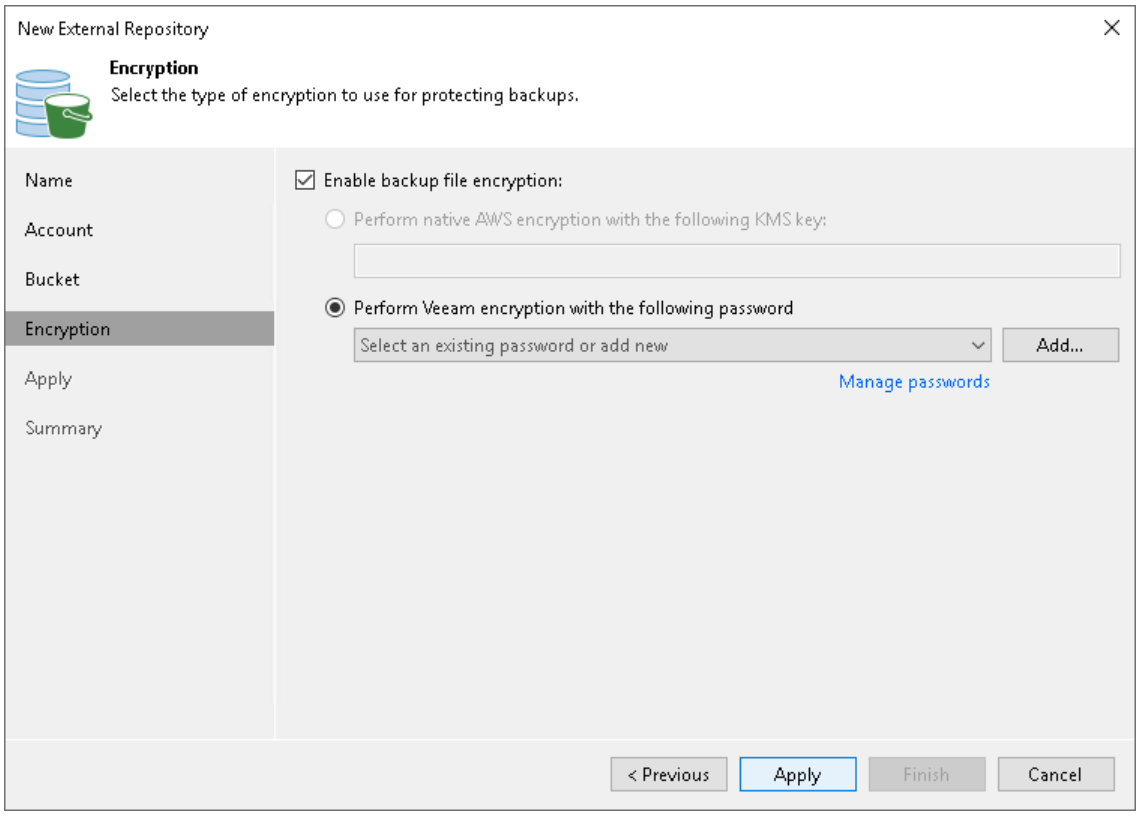

## Step 6. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to install and configure all required components. Then click **Next** to complete the procedure of adding the external repository to the backup infrastructure.

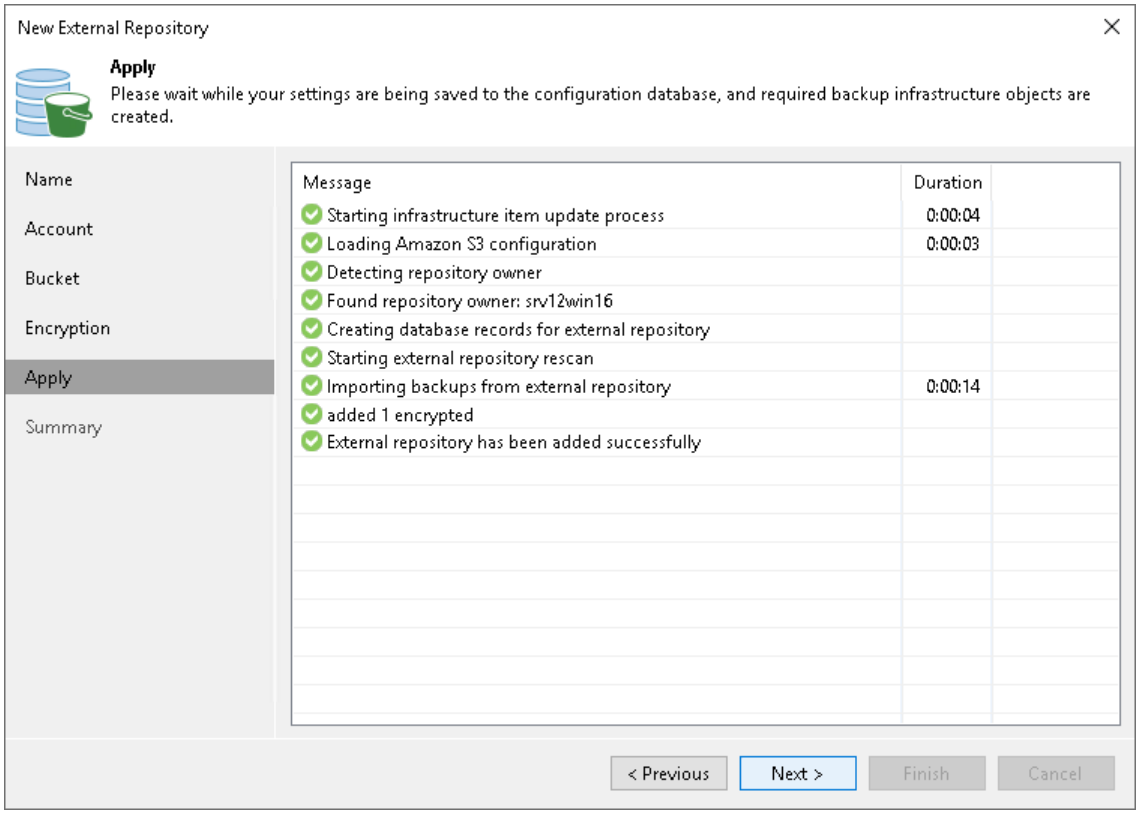

### Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review the external repository settings and click **Finish**.

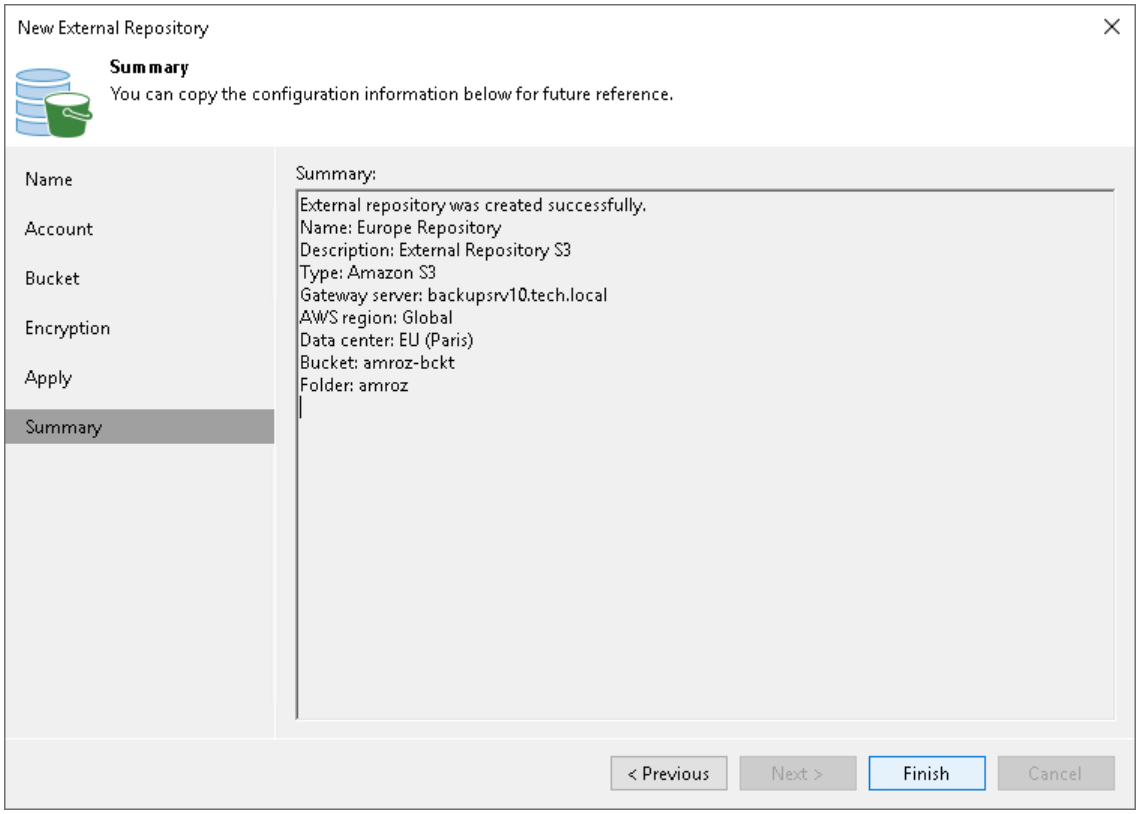


## Adding External Azure Blob Storage

To add an Microsoft Azure Blob storage as an external repository, do the following:

1. Launch the New External Repository wizard.
2. Specify the repository name.
3. Specify the cloud account.
4. Select the Azure Blob container.
5. Configure encryption.
6. Wait for the repository to be added to the infrastructure.
7. Finish working with the wizard.

## Step 1. Launch New External Repository Wizard

To launch the **New External Repository** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane select the **External Repositories** node and click **Connect to Repository** on the ribbon. At the **Connect to External Repository** window, select **Veeam Backup for Microsoft Azure**.
- Open the **Backup Infrastructure** view. In the inventory pane right-click the **External Repositories** node and select **Connect to...** At the **Connect to External Repository** window, select **Veeam Backup for Microsoft Azure**.

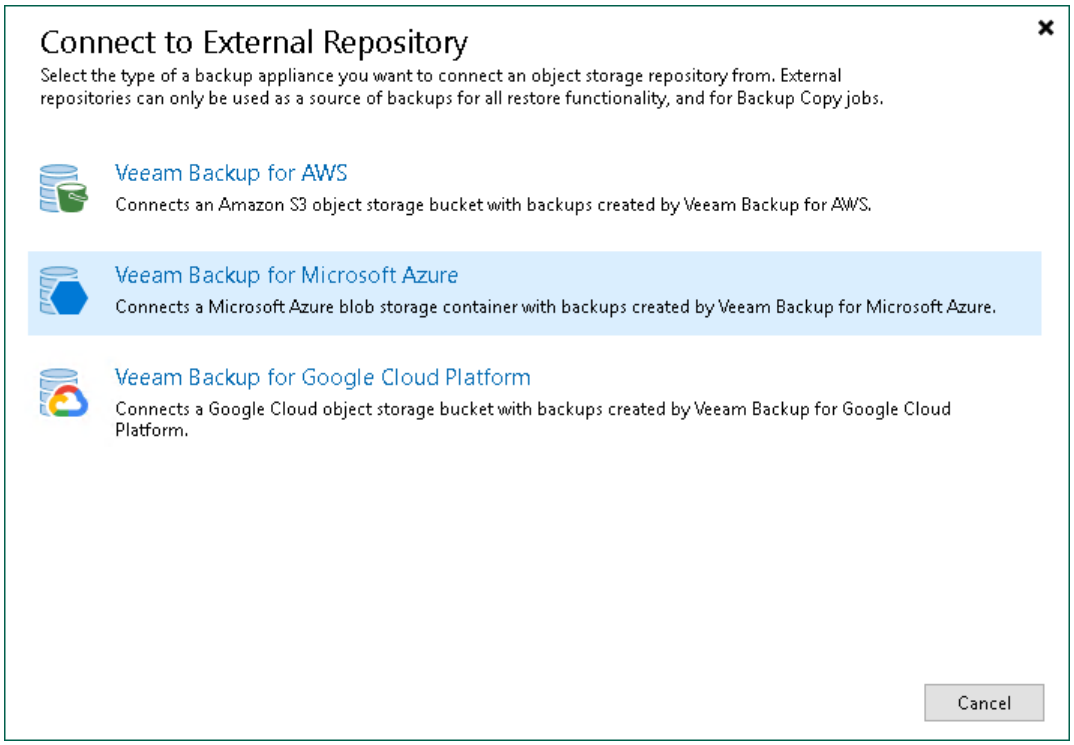

## Step 2. Specify External Repository Name

At the **Name** step of the wizard, specify a name and description for the external repository.

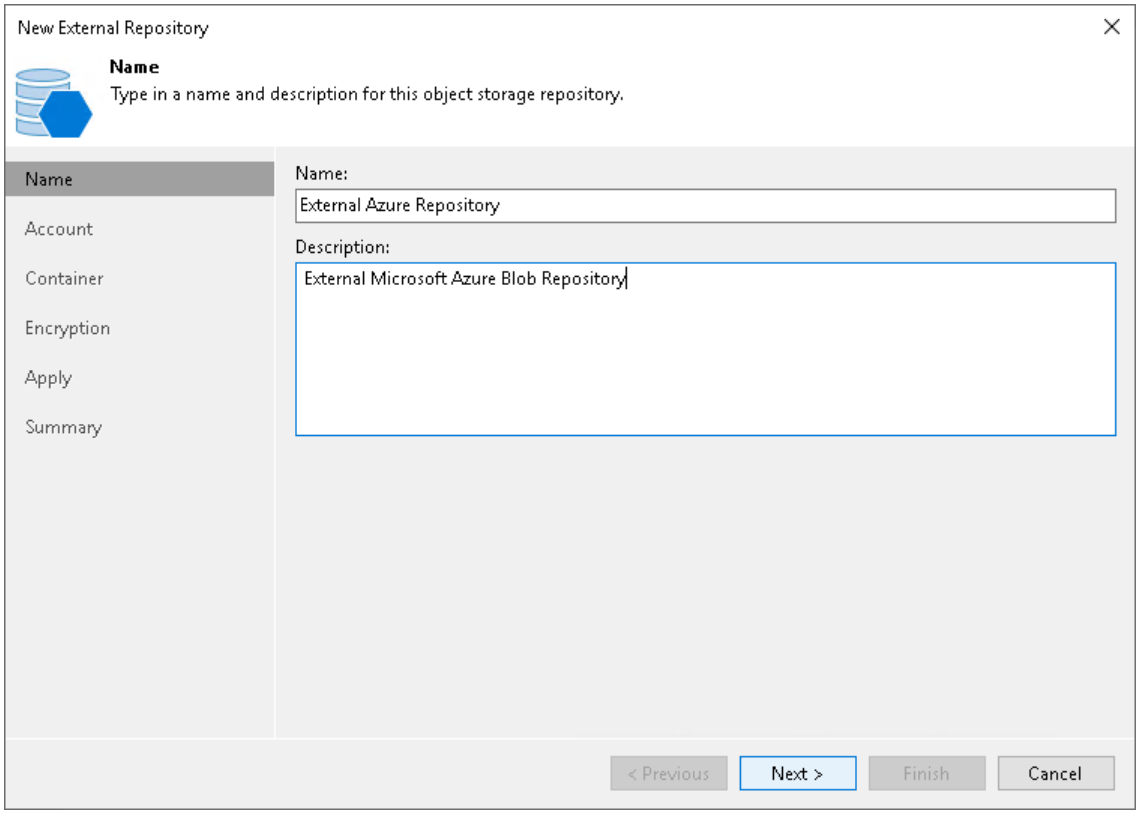

## Step 3. Specify Cloud Storage Account

At the **Account** step of the wizard, specify settings for an account which will be used to connect to Azure Blob storage:

1. From the **Credentials** drop-down list, select user credentials to access your Azure Blob storage.

   If you have not set up credentials beforehand in the Cloud Credentials Manager, click the **Manage cloud accounts** link or click **Add** on the right to add the necessary credentials.

2. From the **Region** drop-down list, select the region type.

3. From the **Gateway server** drop-down list, select a gateway server that will be used to access Azure Blob storage. We recommend that you use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited.

   The gateway server caches data when you copy backups or perform restore operations. The gateway server helps you decrease the amount of traffic being sent over the network and reduce data transfer costs. For more information on caching data, see Cache.

   By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. If the Veeam Backup & Replication server resides in a region that differs from the Azure region where your Blob storage resides, choose a server that is located close to the storage. You can choose any Microsoft Windows or Linux server that is added to your Veeam Backup & Replication infrastructure and has internet connection. Note that the server must be added to the infrastructure beforehand. For more information on how to add a server, see the Adding Microsoft Windows Servers and Adding Linux Servers sections.

   > **NOTE**
   >
   > Consider the following:
   >
   > - On the gateway server, Veeam Backup & Replication deploys Veeam Data Mover that handles traffic sent when you work with external repository data. If Veeam Data Mover becomes outdated, you must upgrade it as described in section Upgrading External Repositories.
   > - If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.

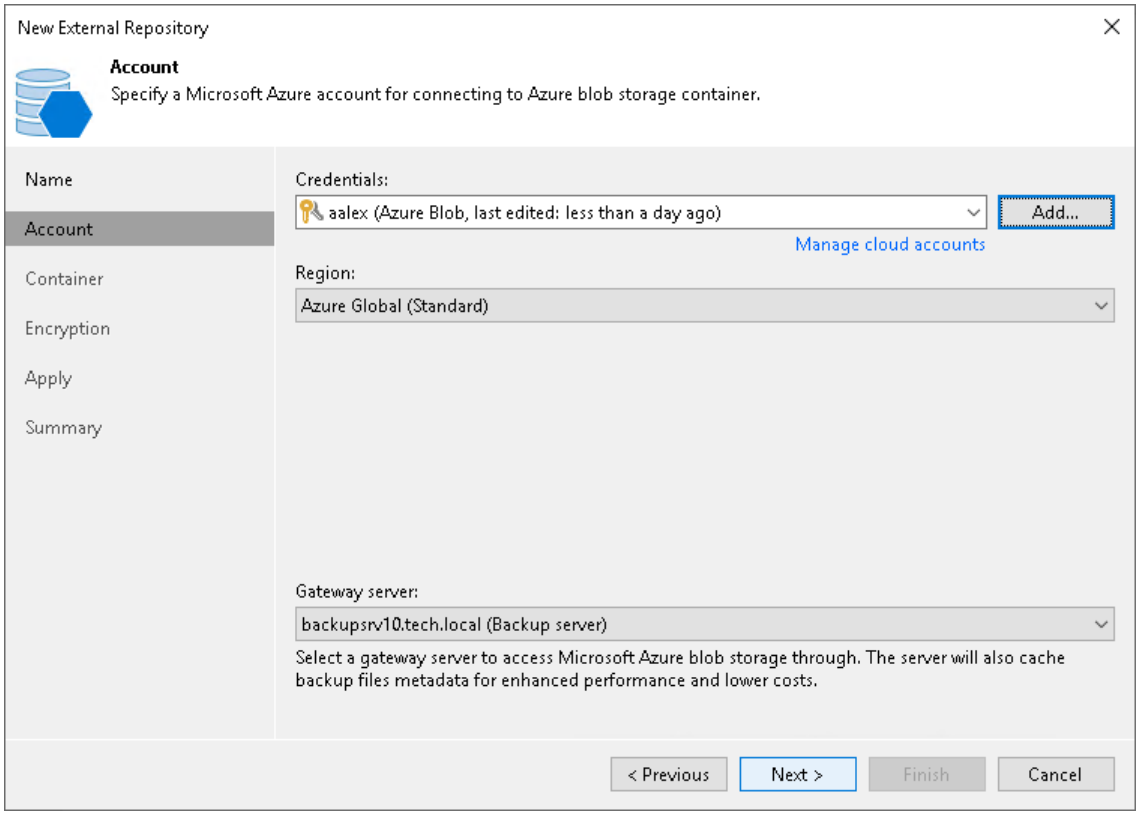
New External Repository
Account
Specify a Microsoft Azure account for connecting to Azure blob storage container.
Name
Account
Container
Encryption
Apply
Summary
Credentials:
aalex (Azure Blob, last edited: less than a day ago)
Add...
Manage cloud accounts
Region:
Azure Global (Standard)
Gateway server:
backupsrv10.tech.local (Backup server)
Select a gateway server to access Microsoft Azure blob storage through. The server will also cache backup files metadata for enhanced performance and lower costs.
< Previous
Next >
Finish
Cancel

## Step 4. Select Azure Blob Container

At the **Container** step of the wizard, specify Azure Blob storage container settings:

1. From the **Container** drop-down list, select a container that contains backups created by Veeam Backup for Microsoft Azure.
2. Click **Browse** to select a folder that contains backups created by Veeam Backup for Microsoft Azure.

   If you do not see the required folder, make sure that the repository you are trying to add is created on the Veeam Backup for Microsoft Azure server.

**NOTE**

Consider the following:

- Only Hot and Cool access tiers are supported.
- If another Veeam Backup & Replication client has already added the same folder, you will be prompted whether to reclaim ownership of such a folder. For more information about ownership, see Ownership.

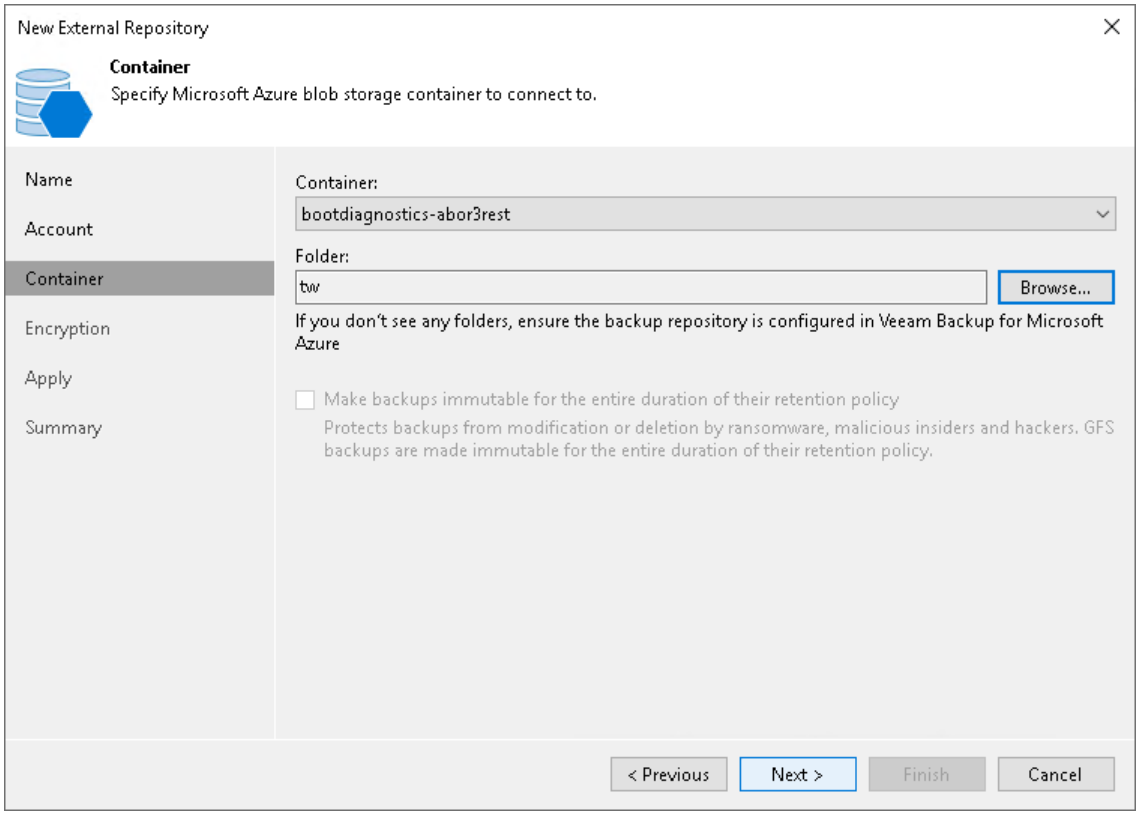

## Step 5. Configure Encryption

At the **Encryption** step of the wizard, provide a key for decryption:

- If data in the external repository is encrypted with a Key Vault key, Veeam Backup & Replication shows the used Key Vault and encryption key. In this case, Veeam Backup & Replication will automatically decrypt the backups.

  > **NOTE**
  >
  > The cloud account you specified at the **Account** step of the wizard, must have access to Key Vault keys. For more information about Key Vault key permissions, see the Service Account Permissions section in the Veeam Backup for Microsoft Azure Guide.

- If data in the external repository is encrypted with a password, select **Enable backup file encryption** and then click **Perform Veeam encryption with the following password**. From the drop-down list, select the password that must be used to decrypt the data. If the password is correct, Veeam Backup & Replication will automatically decrypt the backups.

  If you have not added the password beforehand, click the **Manage passwords** link or the **Add** button to add the necessary password. For more information on adding passwords, see Creating Passwords.

  If you do not specify the decryption password, you can do it later. For more information, see Viewing External Repository Data.

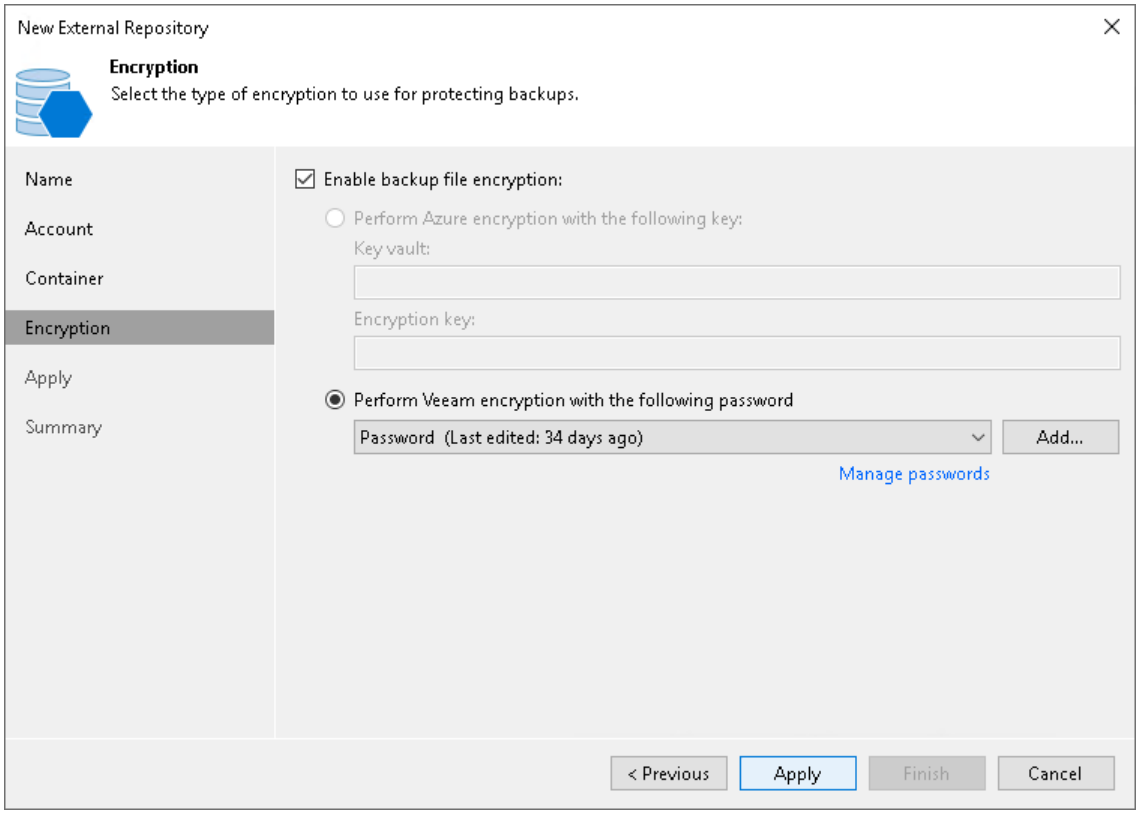

## Step 6. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to install and configure all required components. Then click **Next** to complete the procedure of adding the external repository to the backup infrastructure.

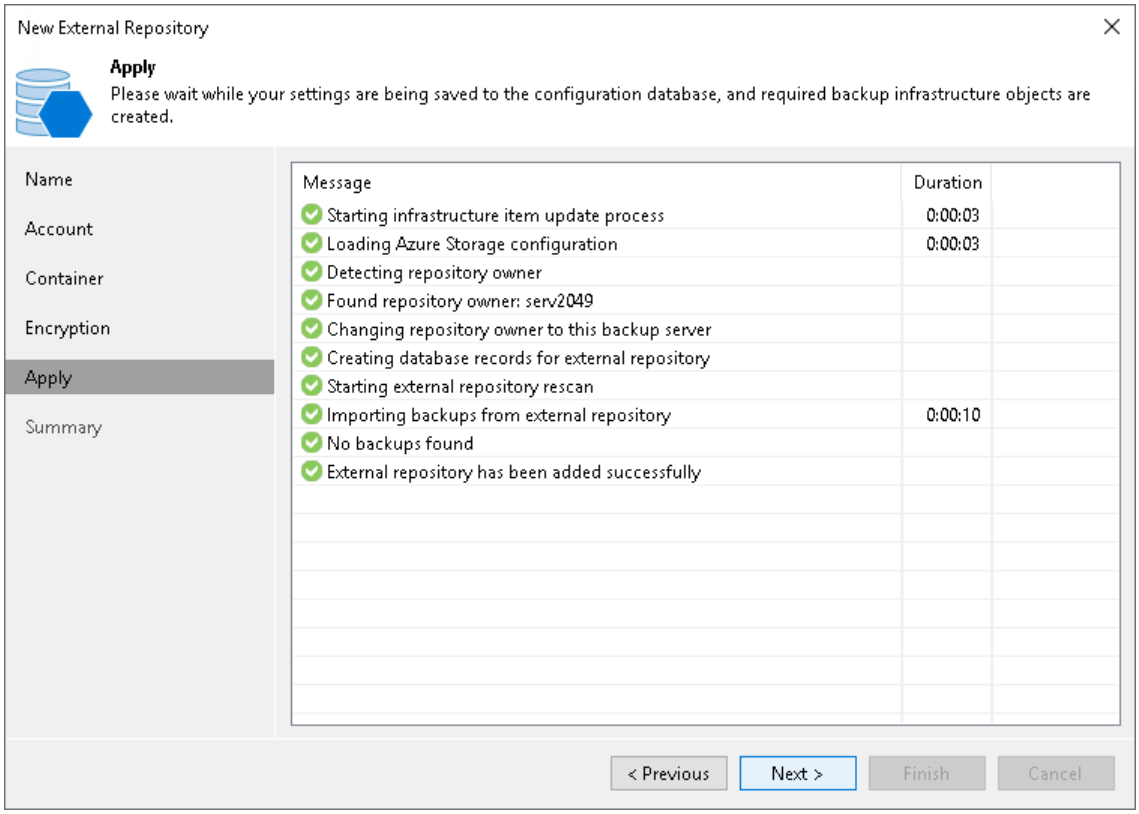

### Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review the repository configuration settings and click **Finish**.

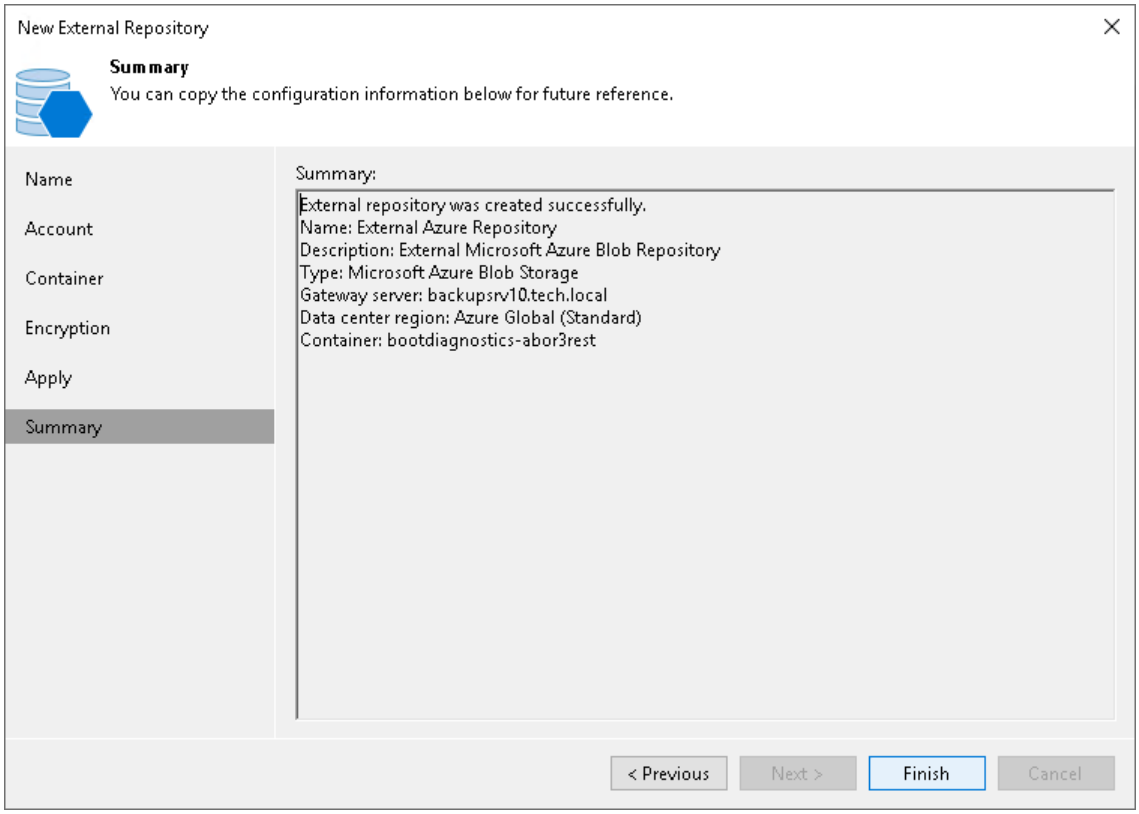


## Adding External Google Cloud Storage

To add an external Google Cloud storage repository to the backup infrastructure, do the following:

1. Launch the New External Repository wizard.
2. Specify the repository name.
3. Specify the cloud account.
4. Specify cloud storage details.
5. Apply settings.
6. Finish working with the wizard.

## Step 1. Launch New External Repository Wizard

To launch the **New External Repository** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane select the **External Repositories** node and click **Connect to Repository** on the ribbon. At the **Connect to External Repository** window, select **Veeam Backup for Google Cloud**.
- Open the **Backup Infrastructure** view. In the inventory pane right-click the **External Repositories** node and select **Connect to...** At the **Connect to External Repository** window, select **Veeam Backup for Google Cloud**.

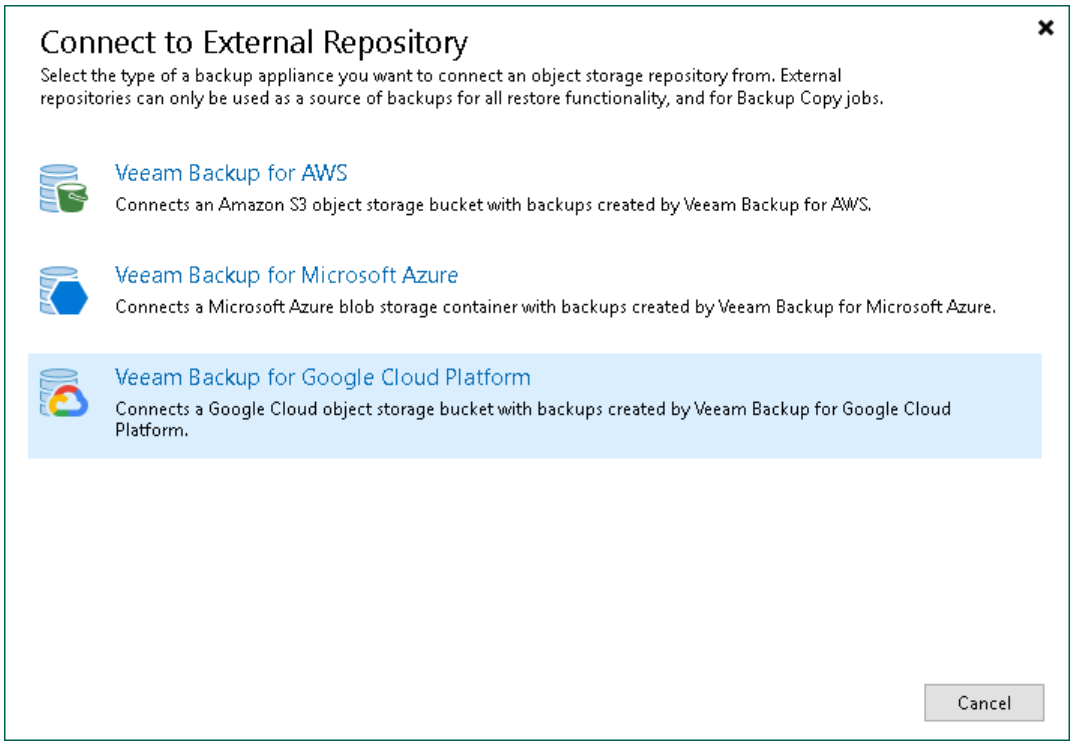

## Step 2. Specify External Repository Name

At the **Name** step of the wizard, specify a name and description for the external repository.

1. In the **Name** field, enter a name for the external repository.
2. In the **Description** field, enter an optional description.

   The default description contains information about the user who added the external repository, date and time when the external repository was added.

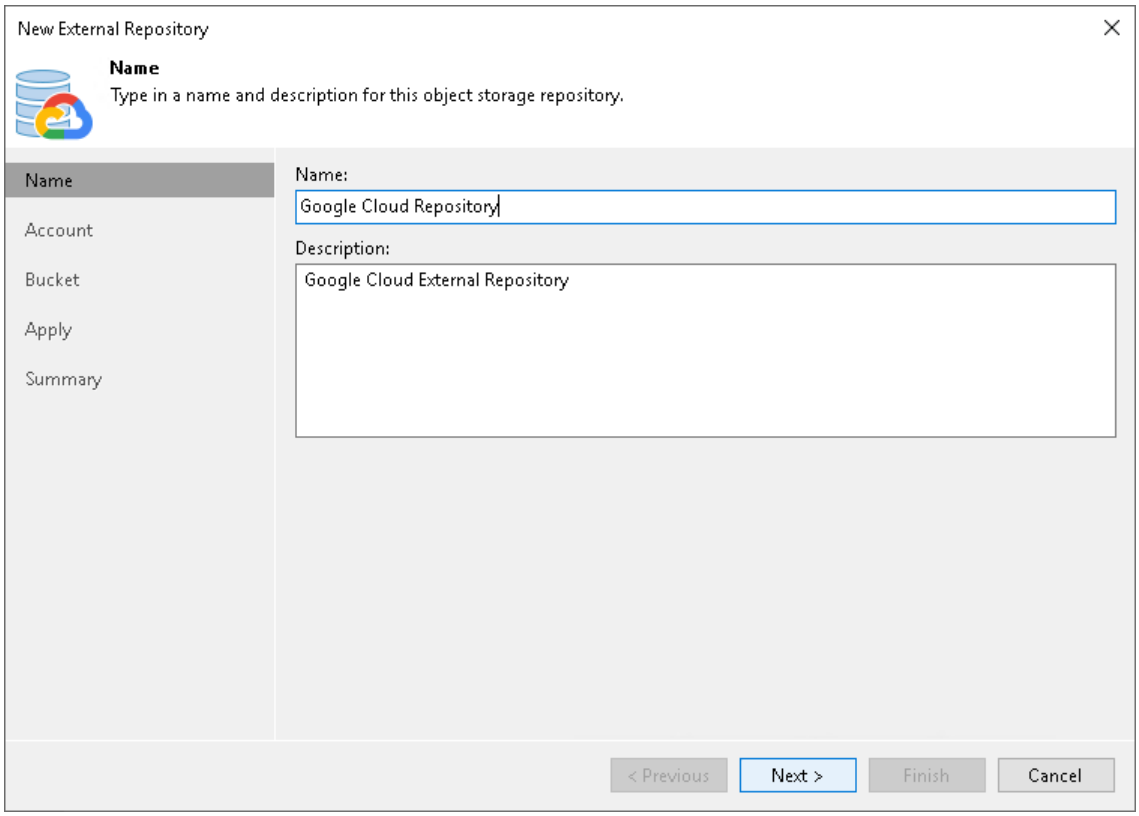

## Step 3. Specify Cloud Storage Account

At the **Account** step of the wizard, specify Google Cloud connection settings:

1. From the **Credentials** drop-down list, select user credentials to access a storage bucket with the backups.

   If you have not set up credentials beforehand, click the **Manage cloud accounts** link or click **Add** on the right and add the necessary credentials, as described in section Google Cloud Accounts .

2. From the **Gateway server** drop-down list, select a server that will be used to access the Google Cloud storage.

   You can select any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. You may want to use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers and Adding Linux Servers. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

   On the gateway server, Veeam Backup & Replication deploys Veeam Data Mover. Veeam Data Mover is responsible for handling ingress/egress requests that are sent to/from the gateway server when you work with the external repository data in Veeam Backup & Replication. If Veeam Data Mover becomes outdated, you must upgrade it as described in section Upgrading External Repositories.

   Gateway servers store cached data. For more information, see Cache.

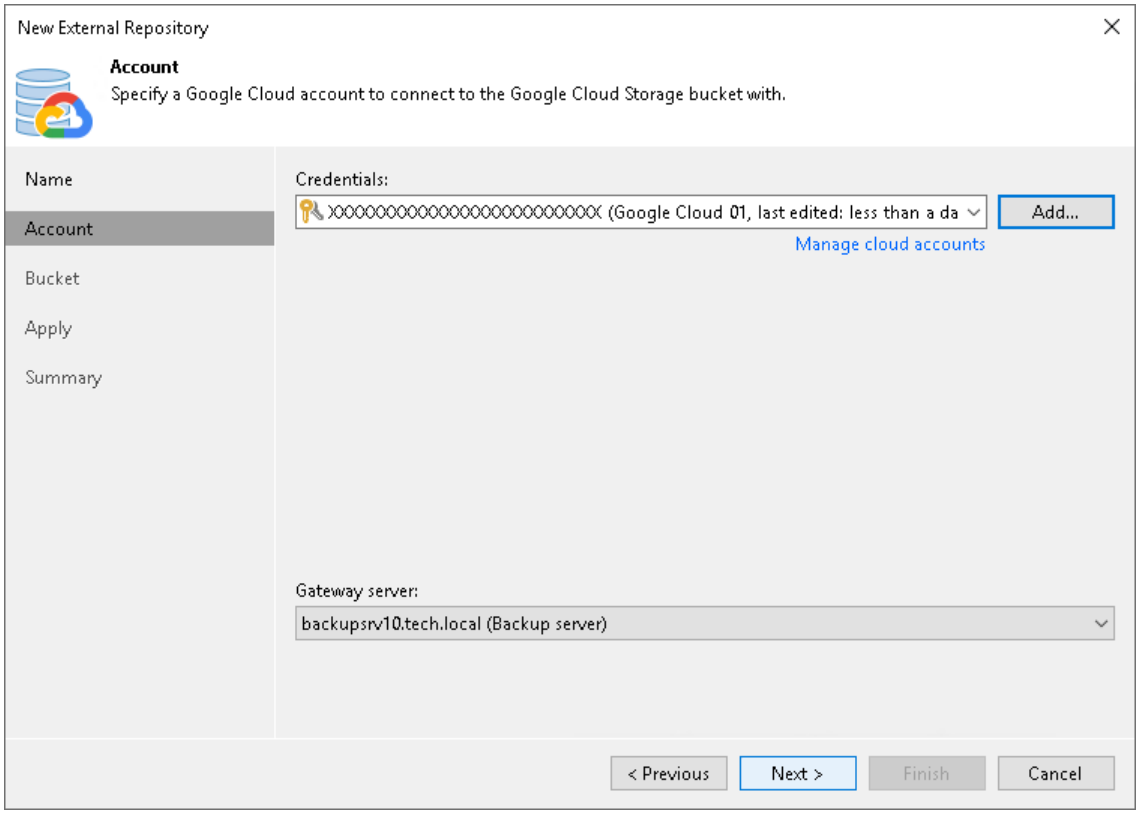

## Step 4. Specify Cloud Storage Details

At the **Bucket** step of the wizard, specify a Google Cloud bucket and folder where Google Cloud instance backups reside:

1. From the **Data center region** drop-down list, select a region.
2. Next to the **Bucket** field, click **Browse** and select a bucket.
3. In the **Select Folder** field, select a cloud folder where the data will be stored. To do it, click **Browse** and select an existing folder.
4. If the folder contains encrypted backups, select the **Use this password for encrypted backups** check box and provide a password. If you skip this step for encrypted backups, Veeam Backup & Replication will add such backups to the **External Repository (Encrypted)** node. For more information, see Viewing External Repository Data.

   For more information about encryption, see Encryption.

**NOTE**

Consider the following:

- Only Standard storage class is supported.
- If another Veeam Backup & Replication client has already added the same folder, you will be prompted whether to reclaim ownership of such a folder. For more information about ownership, see Ownership.

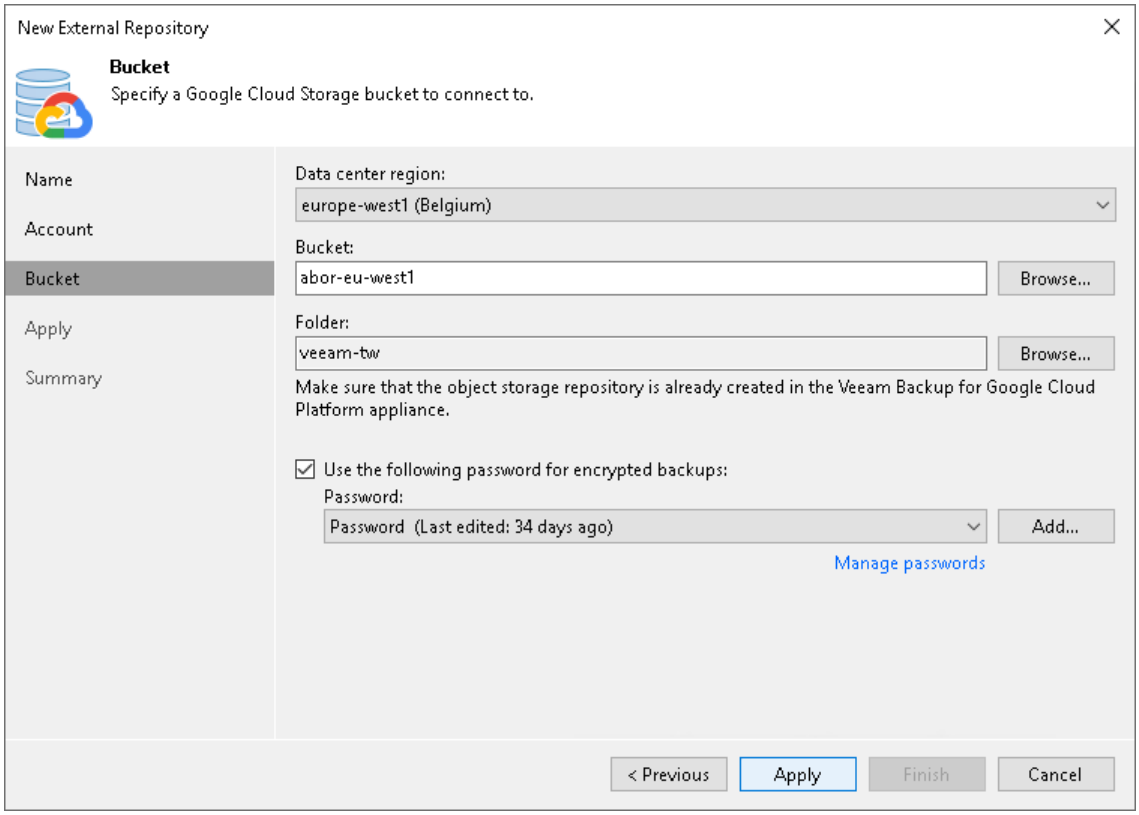

## Step 5. Apply Settings

At the **Apply** step of the wizard, wait until Veeam Backup & Replication applies settings and completes adding the external repository. Then click **Next**.

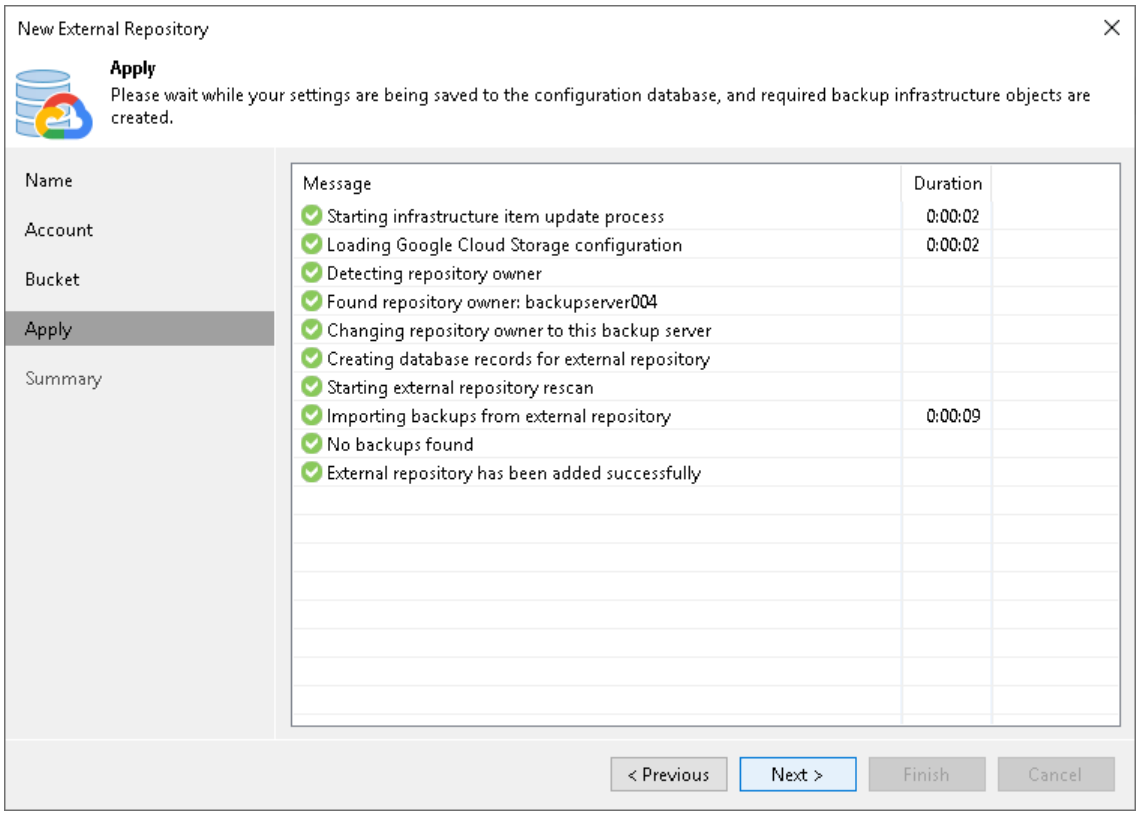

### Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review the external repository settings and click **Finish**.

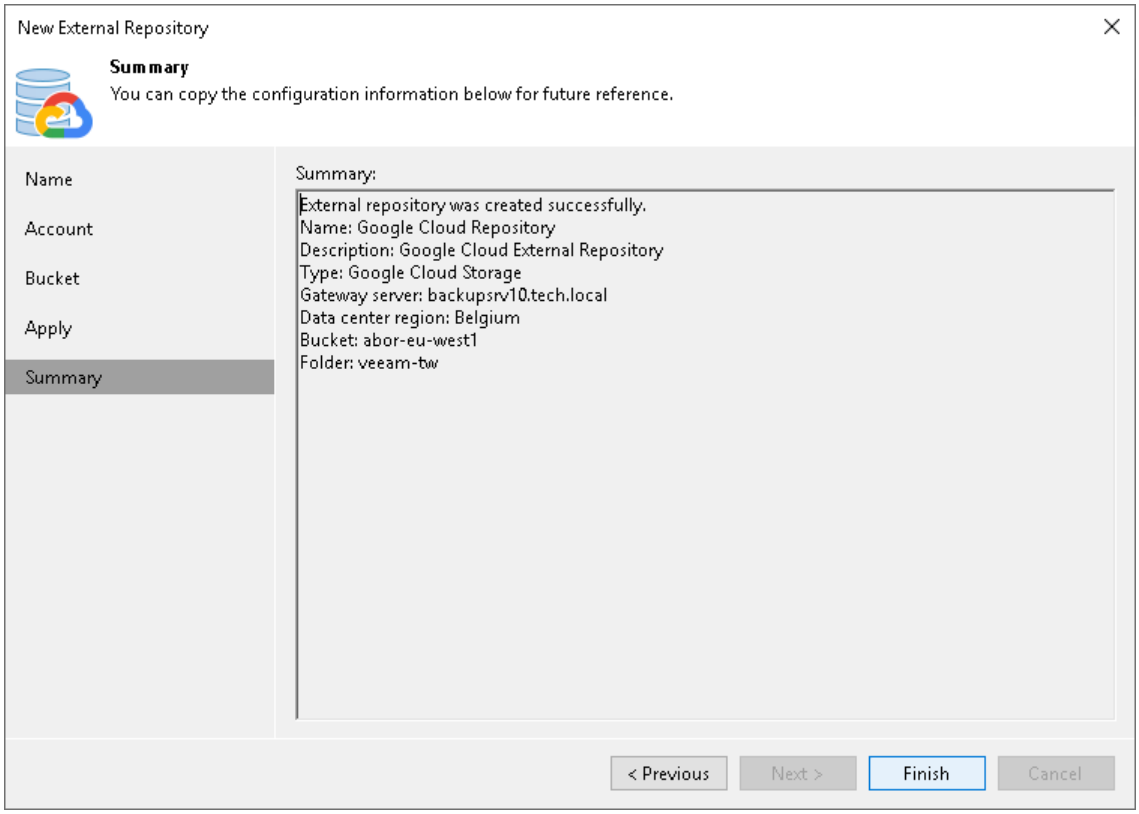


# Managing External Repositories

You can manage your external repositories in various ways: rescan external repositories to synchronize their state with the state of an object storage, upgrade external repositories by uploading a new version of Veeam Data Mover to the gateway server, edit settings of the external repositories, or remove the external repositories. For more information about these options, see the following sections.

## Viewing External Repository Data

After you add an external repository to the backup infrastructure, you can view backups in **External Repository** and **External Repository (Encrypted)** nodes of the **Backups** node in the **Home** view.

- In the **External Repository** node, Veeam Backup & Replication displays:
  - Amazon EC2 instance backups that were decrypted at the Bucket step of the **New External Repository** wizard.
  - Microsoft Azure VMs that were decrypted at the Container step of the **New External Repository** wizard.
  - Google Cloud VM instances that were decrypted at the Bucket step of the **New External Repository** wizard.
- In the **External Repository (Encrypted)** node, Veeam Backup & Replication displays:
  - EC2 instance backups that were encrypted by Veeam Backup for AWS.

- Microsoft Azure VMs that were encrypted by Veeam Backup for Microsoft Azure.
- Google Cloud VM instances that were encrypted by Veeam Backup for Google Cloud.

To decrypt backups, select a backup policy that created the backups you want to decrypt, click **Specify Password** on the ribbon, provide a password and click **OK**.

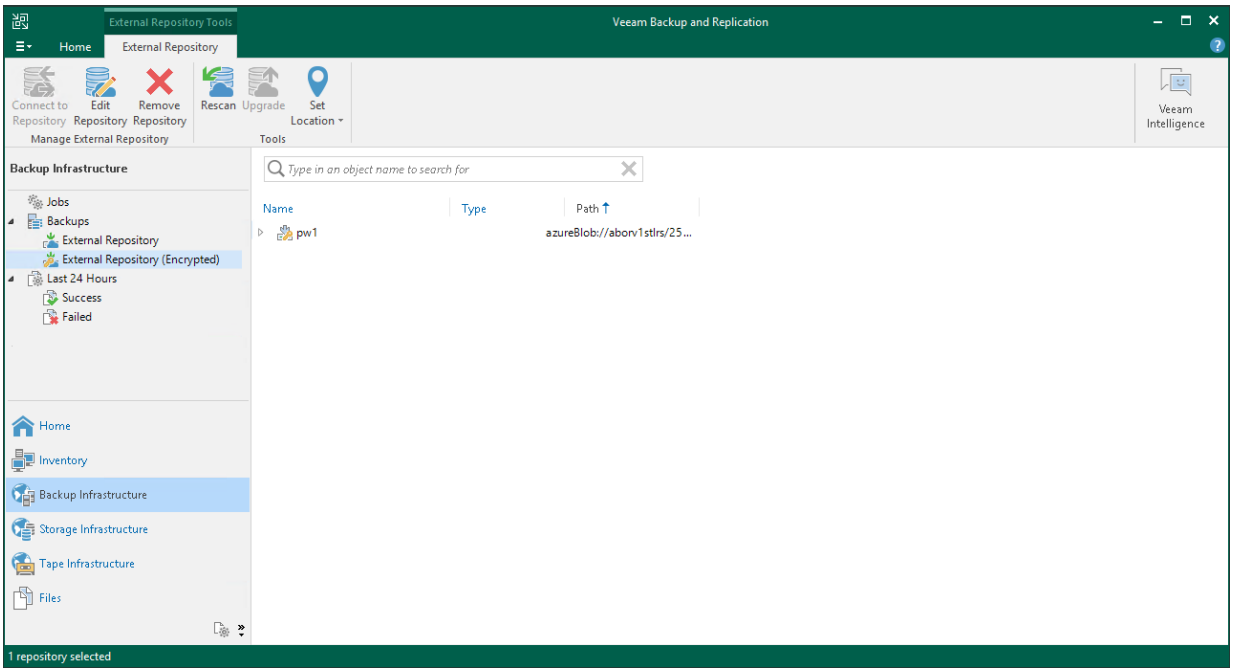


# Rescanning External Repositories

To synchronize the state of an external repository with the state of an object storage (Amazon S3, Azure Blob or Google Cloud Storage), you can use the rescan feature. During rescan, Veeam Backup & Replication fetches newly created restore points and other required metadata.

Consider the following:

- Rescan is done automatically in the following cases:
  - After you add an external repository to the backup infrastructure.
  - Every 24 hours.
  - After a backup chain is modified in the object storage. For example, if a restore point is added or deleted.
- Rescan session results are saved to the configuration database and can be found in the **History** view under the **System** node.

To rescan external repositories manually:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **External Repositories**.
3. Select a repository you want to rescan and click **Rescan** on the ribbon menu or right-click a repository and select **Rescan**.

If you have more than one external repository added to the scope, you may want to rescan all the repositories altogether. For that, right-click the root **External Repositories** node in the navigation pane and select **Rescan**.

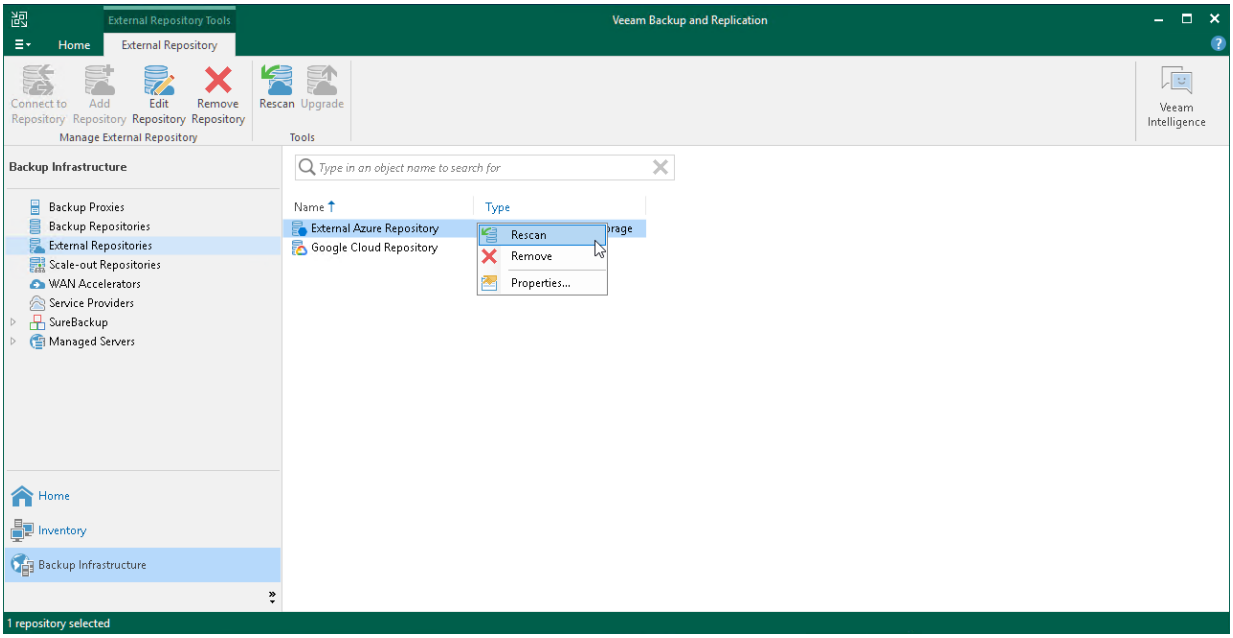


# Upgrading External Repositories

When you upgrade an external repository, a new version of Veeam Data Mover is uploaded to the gateway server. Veeam Data Mover is responsible for handling ingress/egress requests that are sent to/from the gateway server during working with the external repository data in Veeam Backup & Replication.

Upload of Veeam Data Mover is done directly to a gateway server which you specify at the Account step of the **New External Repository** wizard.

To upgrade an external repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **External Repositories**.

3. Select a repository you want to upgrade and click **Upgrade** on the ribbon menu or right-click a repository and select **Upgrade**.

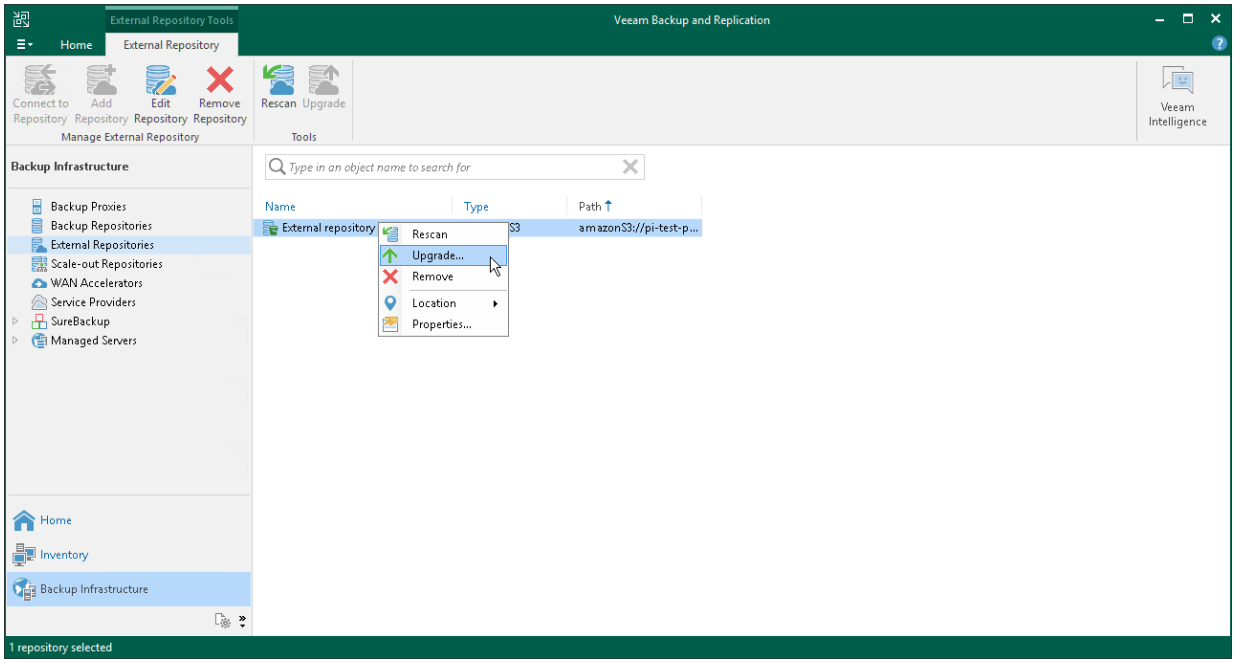


# Editing Settings of External Repositories

To edit settings of an external repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **External Repositories**.
3. In the working area, select an external repository and click **Edit Repository** on the ribbon or right-click the external repository and select **Properties**.
4. Follow the steps of the **Edit External Repository** wizard and edit settings as required.

   Note that some settings cannot be modified and will remain disabled during editing.

> **NOTE**
>
> Veeam Backup & Replication automatically determines and sets locations for external repositories based on the datacenter region. You can check the datacenter region for each external repository at the Bucket step of the **Edit External Repository** wizard. For more information on locations, see Managing Locations.

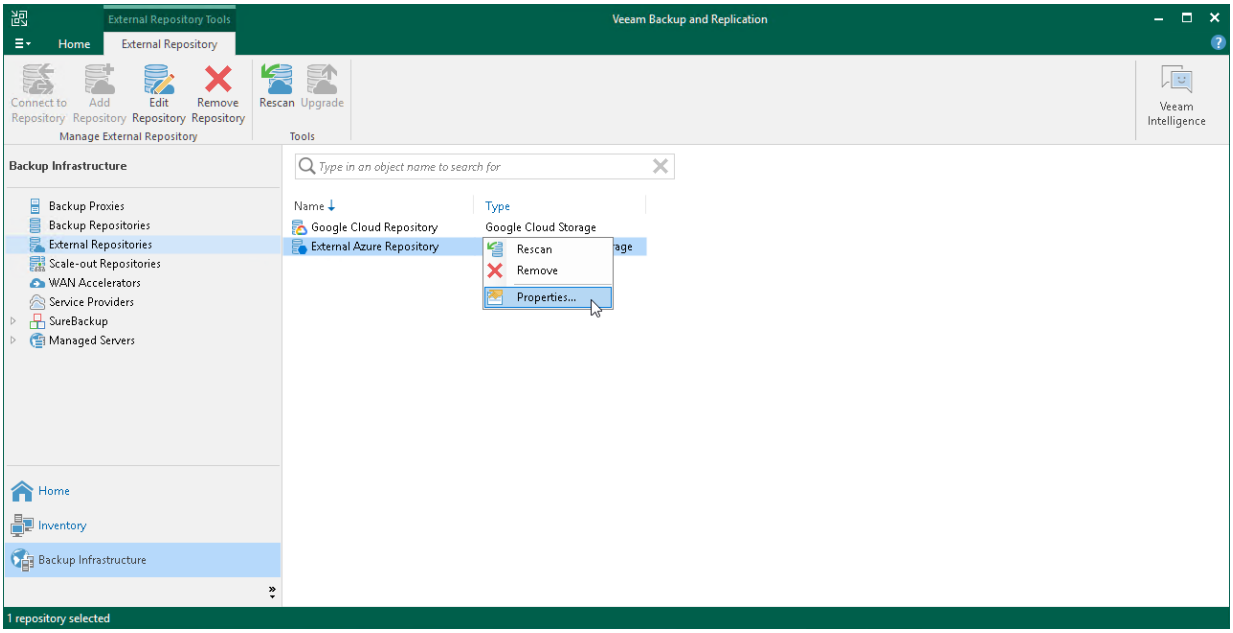


# Removing External Repositories

You can remove an external repository from the backup infrastructure.

When you remove an external repository, Veeam Backup & Replication does the following:

- Relinquishes ownership.
- Removes associated external repository records from the configuration database.
- Removes associated cache from the gateway server.

To remove an external repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **External Repositories**.

3. In the working area, select an external repository and click **Remove Repository** on the ribbon or right-click the external repository and select **Remove**.

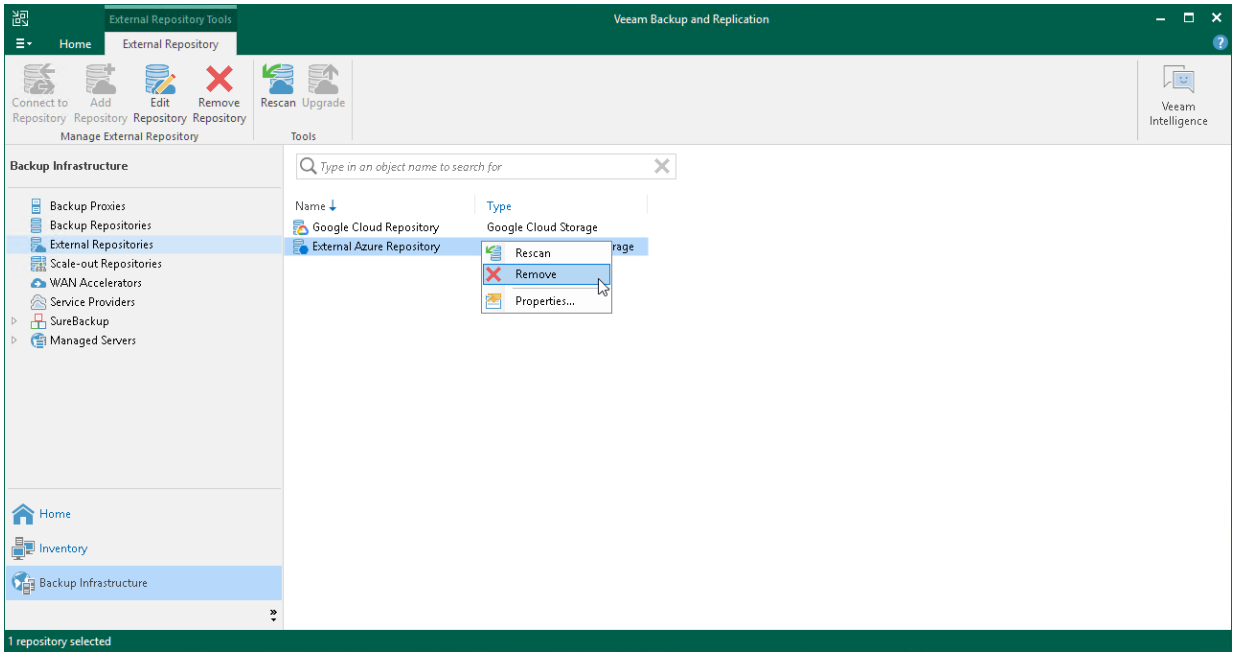

# Microsoft Windows Server

You can use a Microsoft Windows server with local or directly attached storage as a backup repository. The storage can be a local disk, directly attached disk-based storage (such as a USB hard drive), or iSCSI/FC SAN LUN in case the server is connected into the SAN fabric.

## Microsoft Windows Repository Deployment

To communicate with a Microsoft Windows-based repository, Veeam Backup & Replication uses two Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on a VMware backup proxy
- Veeam Data Mover on the Microsoft Windows repository

When any job addresses the backup repository, Veeam Data Mover on the VMware backup proxy establishes a connection with Veeam Data Mover on the backup repository, enabling efficient data transfer over LAN or WAN.

The Data Mover is installed automatically when you add a server to Veeam Backup & Replication as a managed server.

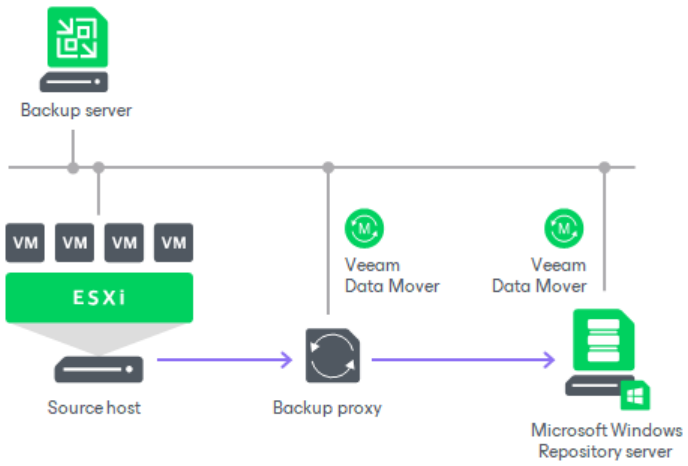


## vPower NFS Server

Windows repositories can be configured to function as vPower NFS Servers. In this case, Veeam Backup & Replication will run the Veeam vPower NFS Service directly in the backup repository (namely, on the managing Windows server to which storage is attached) and provide ESXi hosts with transparent access to backed-up VM images stored on the backup repository. For more information, see Veeam vPower NFS Service.

## Requirements for Microsoft Windows Server Based Repositories

A machine performing the role of a repository must meet the following requirements:

- The role of the repository can be assigned to a Microsoft Windows machine (physical or virtual). The machine must meet the system requirements. For more information, see System Requirements.
- You must add the machine to the Veeam Backup & Replication console as a managed server.

- If you want to use Fast Clone in the Microsoft Windows backup repository, the machine must also meet requirements listed in section Fast Clone.

# Adding Microsoft Windows Repositories

This section describes how to add a Microsoft Windows server as a backup repository.

To add a backup repository, use the **New Backup Repository** wizard.

## Step 1. Launch New Backup Repository Wizard

To launch the **New Backup Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add Backup Repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Direct Attached Storage** > **Microsoft Windows**.

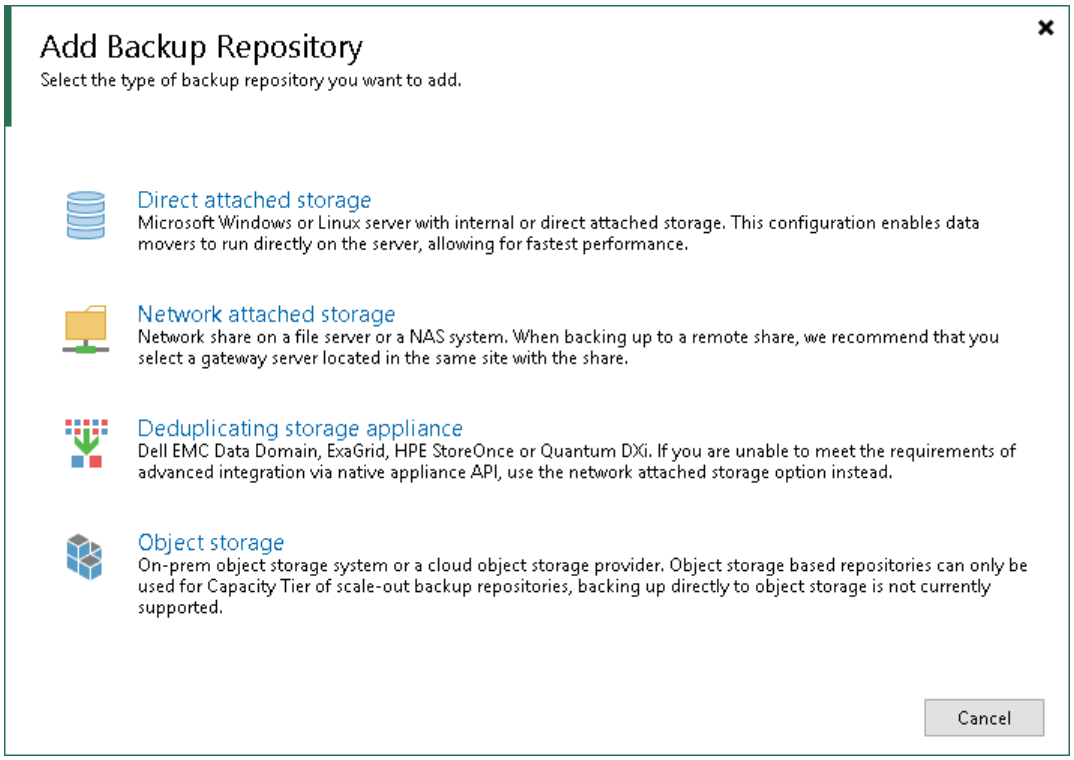

## Step 2. Specify Backup Repository Name and Description

At the **Name** step of the wizard, specify a name and description for the backup repository:

1. In the **Name** field, specify a name for the backup repository.
2. In the **Description** field, provide a description for future reference.

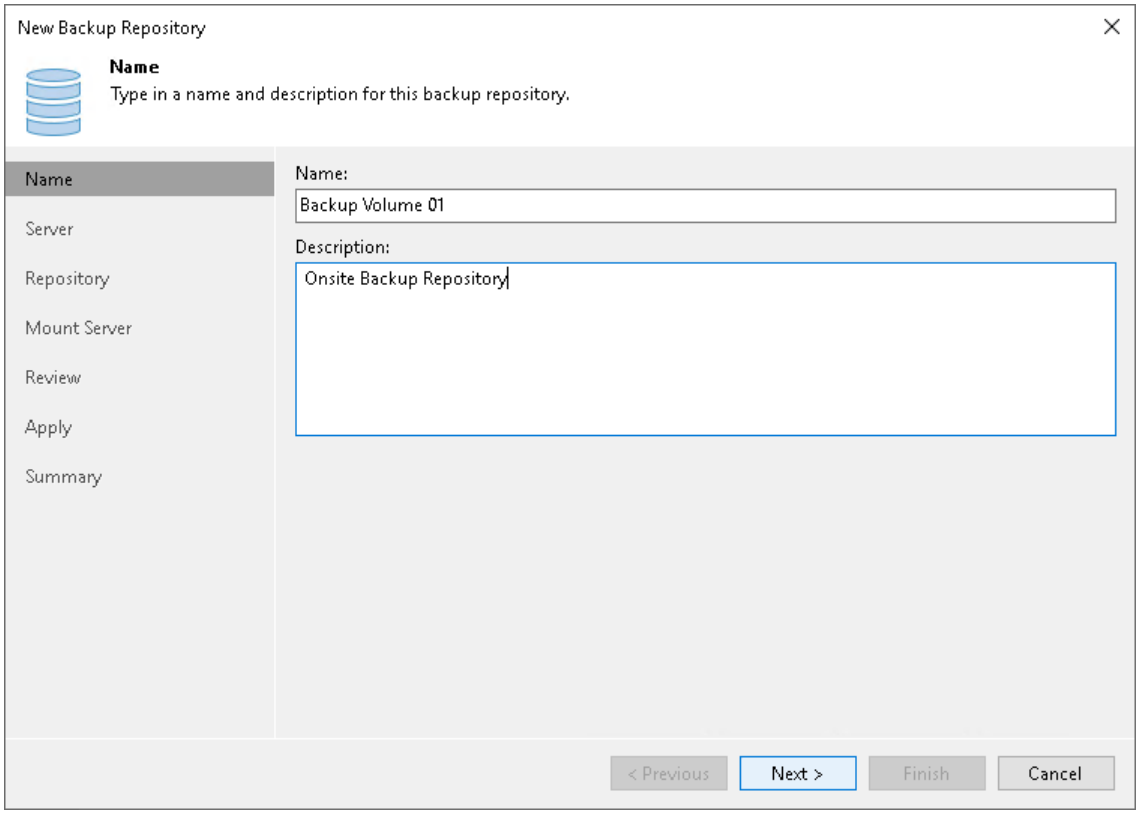

## Step 3. Specify Server or Shared Folder Settings

To configure settings for a Microsoft Windows server:

1. From the **Repository server** list, select a Microsoft Windows server that you want to use as a backup repository. The **Repository server** list contains only those servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, you can click **Add New** on the right to open the **New Windows Server** wizard.
2. Click **Populate** to see a list of disks connected to the server, their capacity and free space.

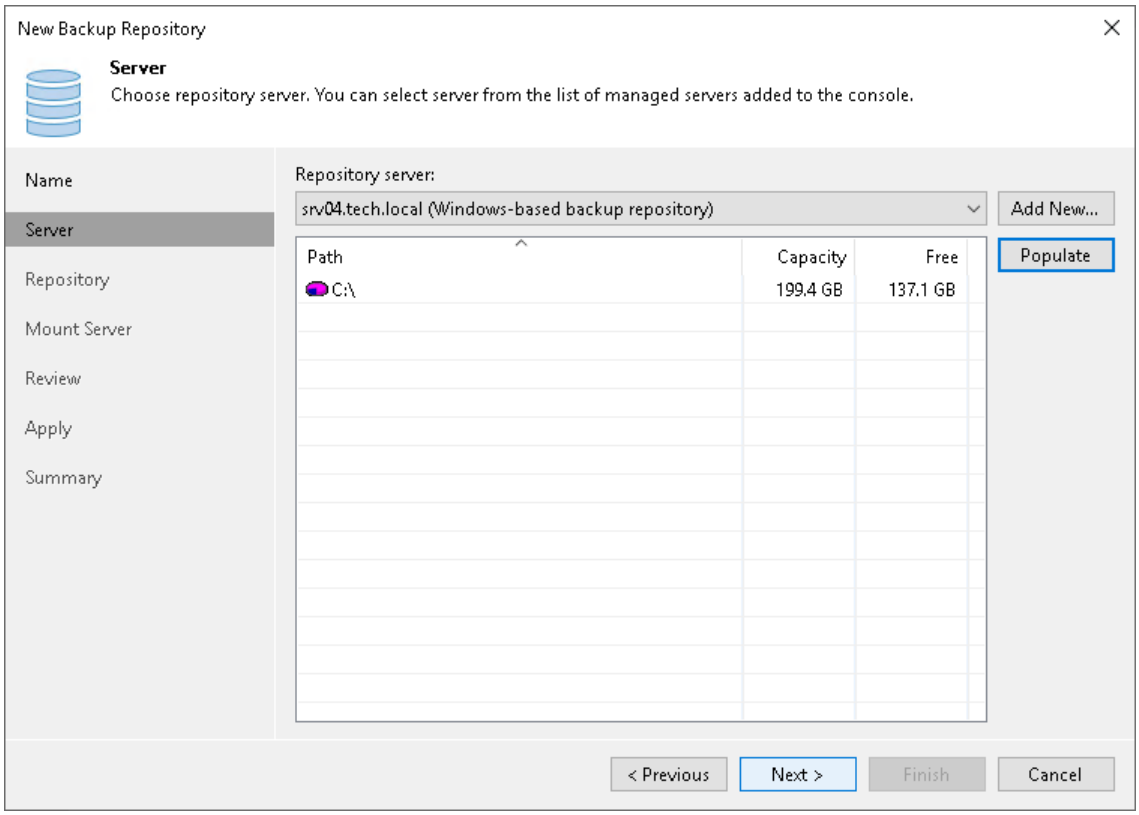

## Step 4. Configure Backup Repository Settings

At the **Repository** step of the wizard, configure general repository settings including path to the repository folder and load control, and also advanced repository settings.

# Configuring General Repository Settings

To configure general repository settings:

1. In the **Location** section, specify a path to the folder where backup files must be stored. Click **Populate** to check capacity and available free space in the selected location.
2. Use the **Load control** section to limit the number of concurrent tasks and data ingestion rate for the backup repository. These settings will help you control the load on the backup repository and prevent possible timeouts of storage I/O operations.
   - Select the **Limit maximum concurrent tasks** check box and specify the maximum allowed number of concurrent tasks for the backup repository. If this value is exceeded, Veeam Backup & Replication will not start a new task until one of current tasks finishes. For more information, see Limiting the Number of Concurrent Tasks.

     > **NOTE**
     >
     > Consider the following:
     >
     > - Limitation of concurrent tasks is ignored if the backup repository acts as a target storage for a Veeam Cloud Connect job.
     > - It is not recommended that you disable the **Limit maximum concurrent tasks to N** option for backup repositories with per-machine backup chains. In case of per-machine backup chains, synthetic operations (synthetic full backup, backup files merge and transformation) work in parallel for every workload in the backup. The number of parallel operations is limited by the number of concurrent tasks that can be performed on the backup repository. If you disable the **Limit maximum concurrent tasks to N** option (which results in using an unlimited number of slots), the load on the backup repository may be high.

   - Select the **Limit read and write data rates to** check box and specify the maximum rate to restrict the total speed of reading and writing data to the backup repository disk. For more information, see Limitation of Read and Write Data Rates for Backup Repositories.

     > **NOTE**
     >
     > The **Limit read and write data rates to** settings does not apply to health checks performed as part of backup and backup copy jobs. Even if you limit read/write rate for a backup repository, the health check will consume resources of the backup repository regardless of this setting. Consider this limitation when configuring basic and health check schedules for backup and backup copy jobs.

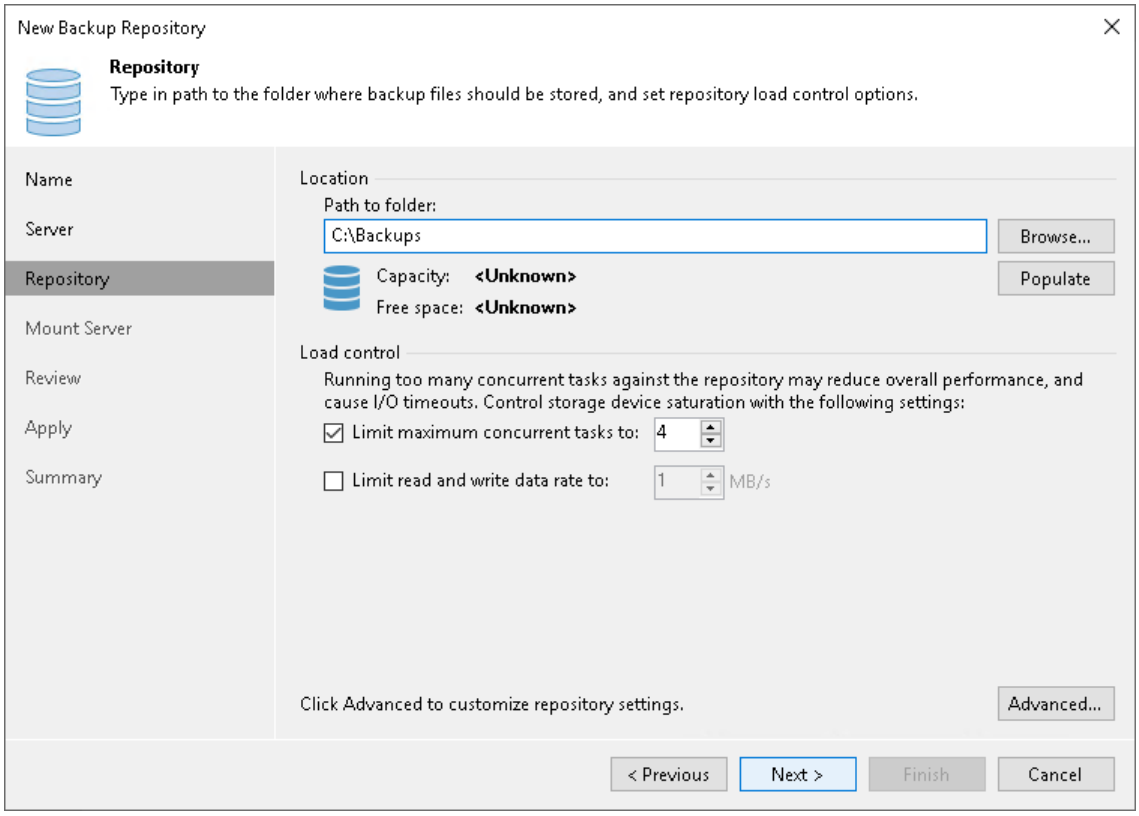


# Configuring Advanced Repository Settings

To configure advanced repository settings:

1. Click **Advanced**.
2. For storage systems using a fixed block size, select the **Align backup file data blocks** check box. Veeam Backup & Replication will align VM data saved to a backup file at a 4 KB (64 KB is also possible for ReFS) block boundary.
3. When you enable compression for a backup job, Veeam Backup & Replication compresses VM data at the source side and then transports it to the target side. Writing compressed data to a deduplicating storage appliance results in poor deduplication ratios as the number of matching blocks decreases. To overcome this situation, select the **Decompress backup data blocks before storing** check box. If data compression is enabled for a job, Veeam Backup & Replication will compress VM data on the source side, transport it to the target side, decompress VM data on the target side and write raw VM data to the storage device to achieve a higher deduplication ratio.

> **NOTE**
>
> Veeam Backup & Replication does not compress VM data if encryption is enabled for a job and the **Decompress backup data blocks before storing** check box is selected in the settings of the target backup repository. Therefore, in the job statistics, you may observe a higher amount of transferred data (the **Transferred** counter) as compared to a job for which encryption is disabled. For more information on job statistics, see Viewing Real-Time Statistics.
>
> In the properties of the backup created with encryption, you may also see that the backup size (the **Backup Size** column) is larger than the original VM size (the **Original Size** column). For more information on backup properties, see Viewing Backup Properties.

4. Select the **This repository is backed by rotated drives** check box if you plan to use a backup repository with rotated drives. For more information on how to configure rotated drives, see Deploying Backup Repositories with Rotated Drives.

5. To create a separate backup file for every machine in the job, make sure that the **Use per-machine backup files** check box is selected. If you clear the check box, Veeam Backup & Replication will create single-file backups. For more information on the backup chain formats and their limitations, see Backup Chain Formats.

> **NOTE**
>
> Changing of the **Use per-machine backup files** setting after the repository was already created does not take any effect. To change backup chain format, follow the instructions provided in Upgrading Backup Chain Formats.

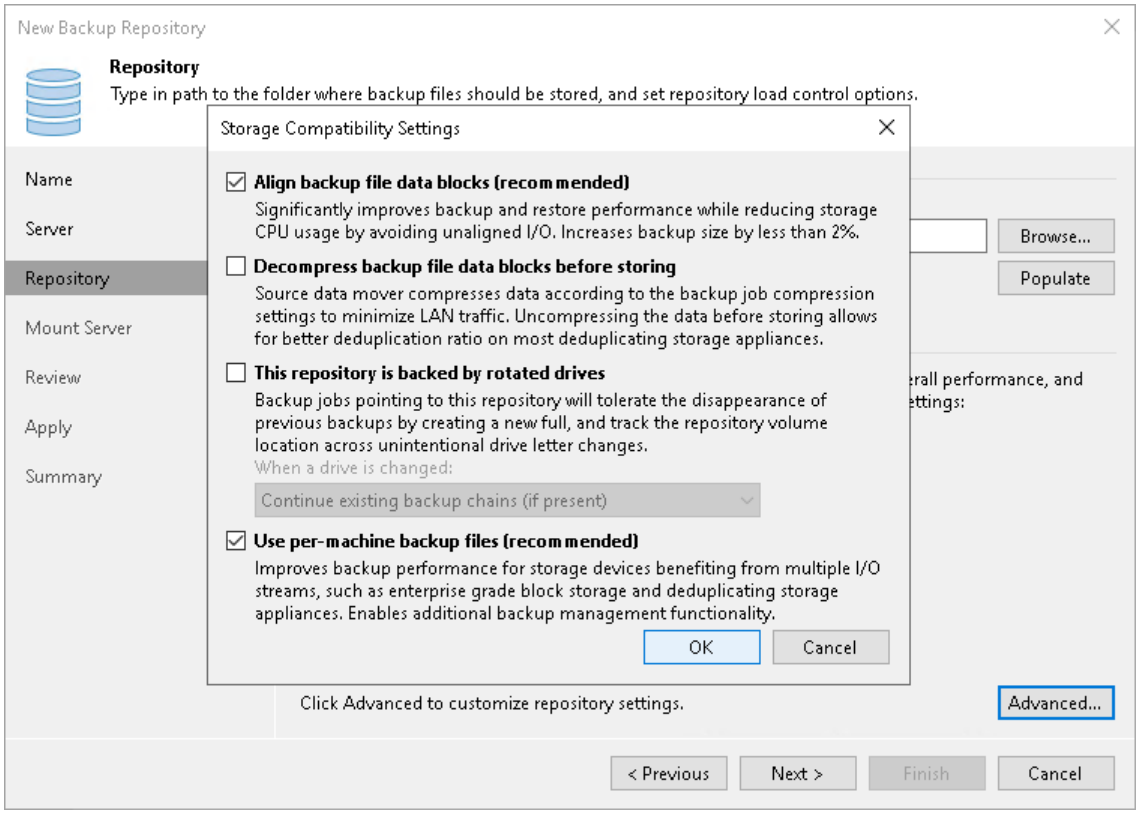

# Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for file-level and application items restore.

1. From the **Mount Server** list, select a server that you want to use as a mount server. The mount server is required for file-level and application items restore. During the restore process, Veeam Backup & Replication mounts the VM disks from the backup file residing in the backup repository to the mount server. As a result, VM data does not have to travel over the network, which reduces the load on the network and speed up the restore process. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers added to the backup infrastructure. If the server is not added to the backup infrastructure, click **Add New** on the right to open the **New Windows Server** wizard. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder that will be used for writing cache during mount operations.

3. To make the backup repository accessible by the Veeam vPower NFS Service, select the **Enable vPower NFS service on the mount server** check box. Veeam Backup & Replication will enable the vPower NFS Service on your selected mount server.

4. To customize network ports used by the vPower NFS Service, click **Ports**. For information on ports used by default, see Ports.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

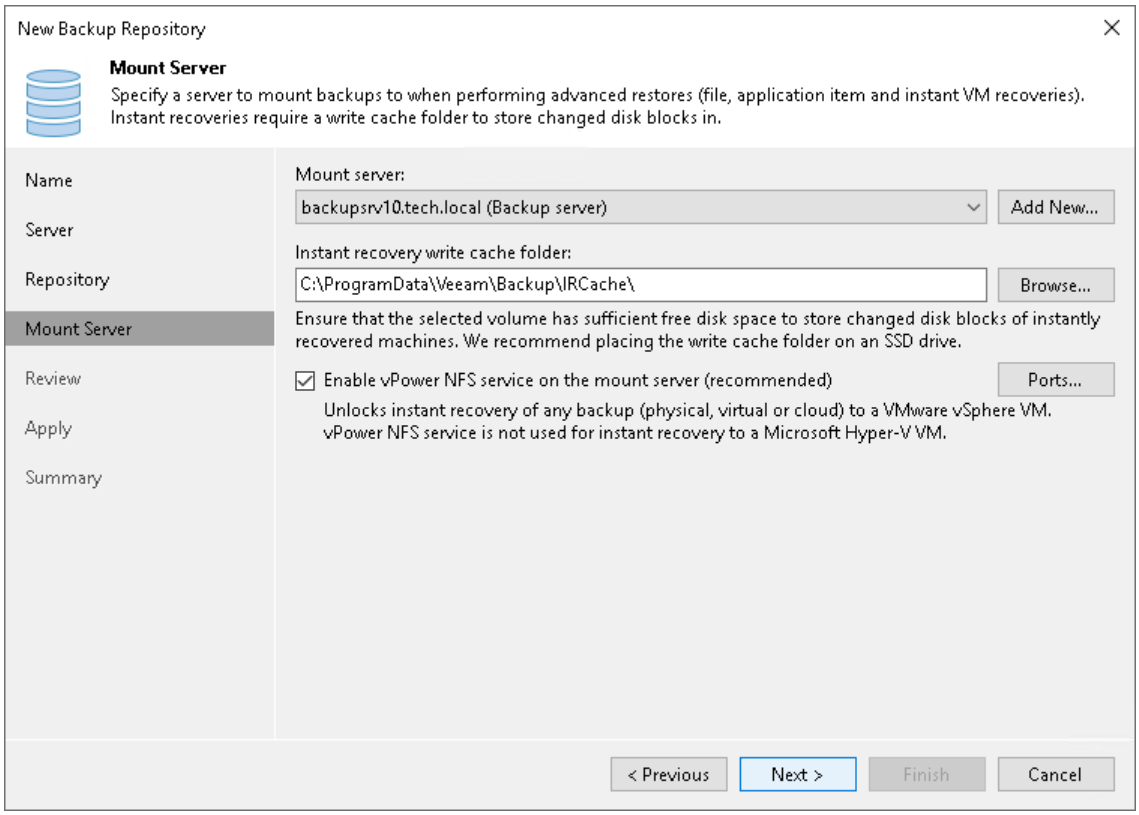

## Step 6. Review Properties and Components

At the **Review** step of the wizard, review details of the backup repository and specify importing settings.

1. Review the backup repository settings and list of components that will be installed on the backup repository server.
2. If the backup repository contains backups that were previously created by Veeam Backup & Replication, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Disk (Imported)** node.
3. If the backup repository contains guest file system index files that were previously created by Veeam Backup & Replication, select the **Import guest file system index** check box. Index files will be imported with backup files, and you will be able to search for guest OS files inside imported backups.

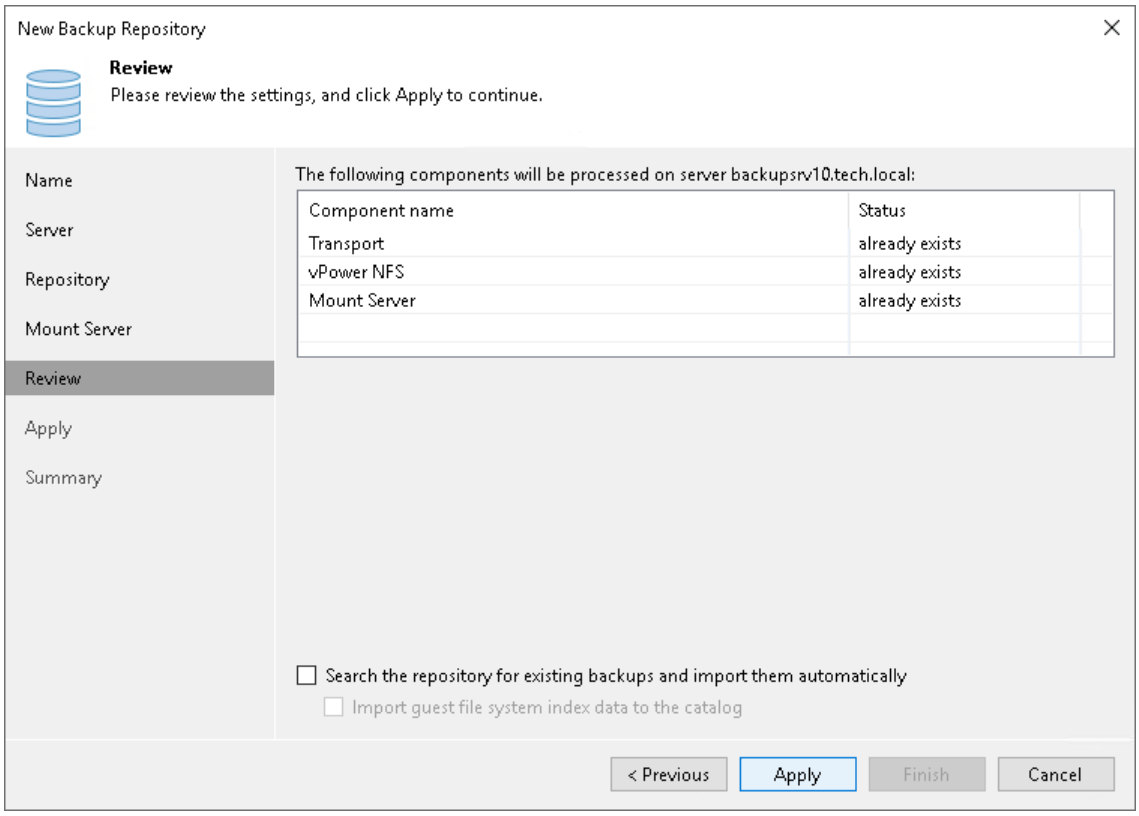

# Step 7. Apply Backup Repository Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to install and configure all required components. Then click **Next** to complete the procedure of adding the backup repository to the backup infrastructure.

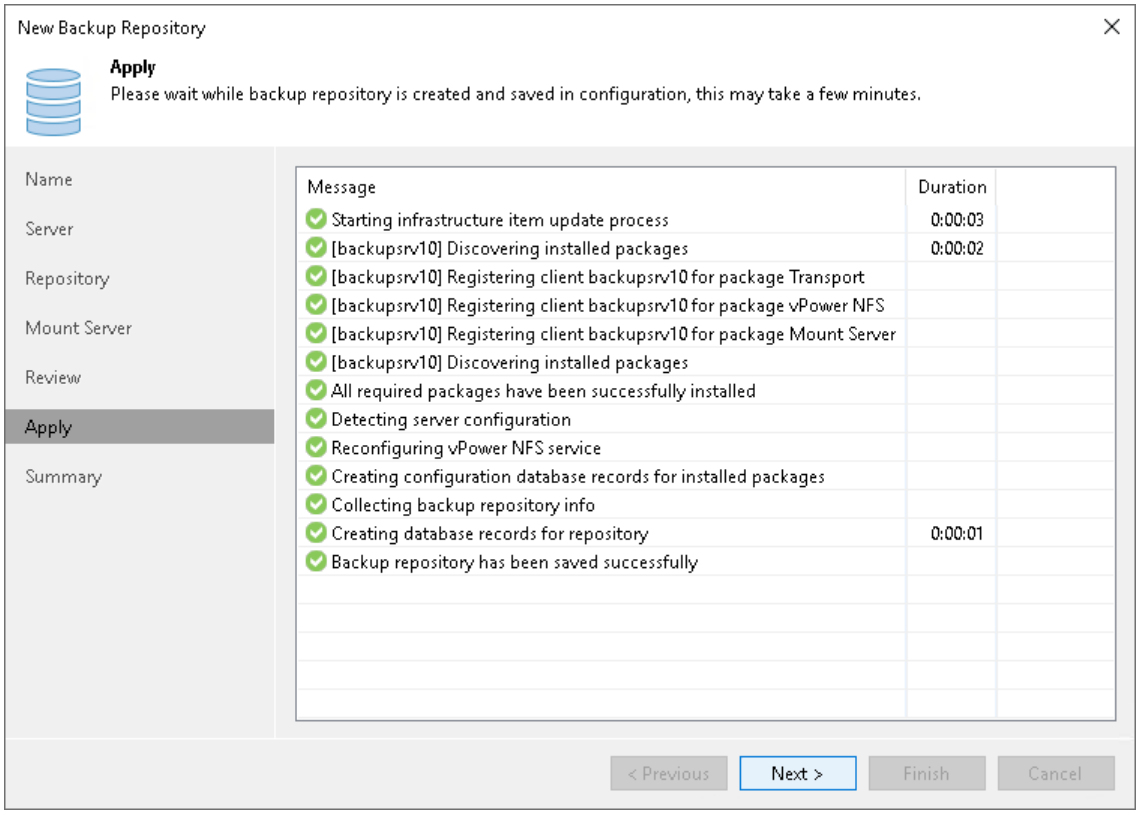

# Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added backup repository. Then click **Finish** to exit the wizard.

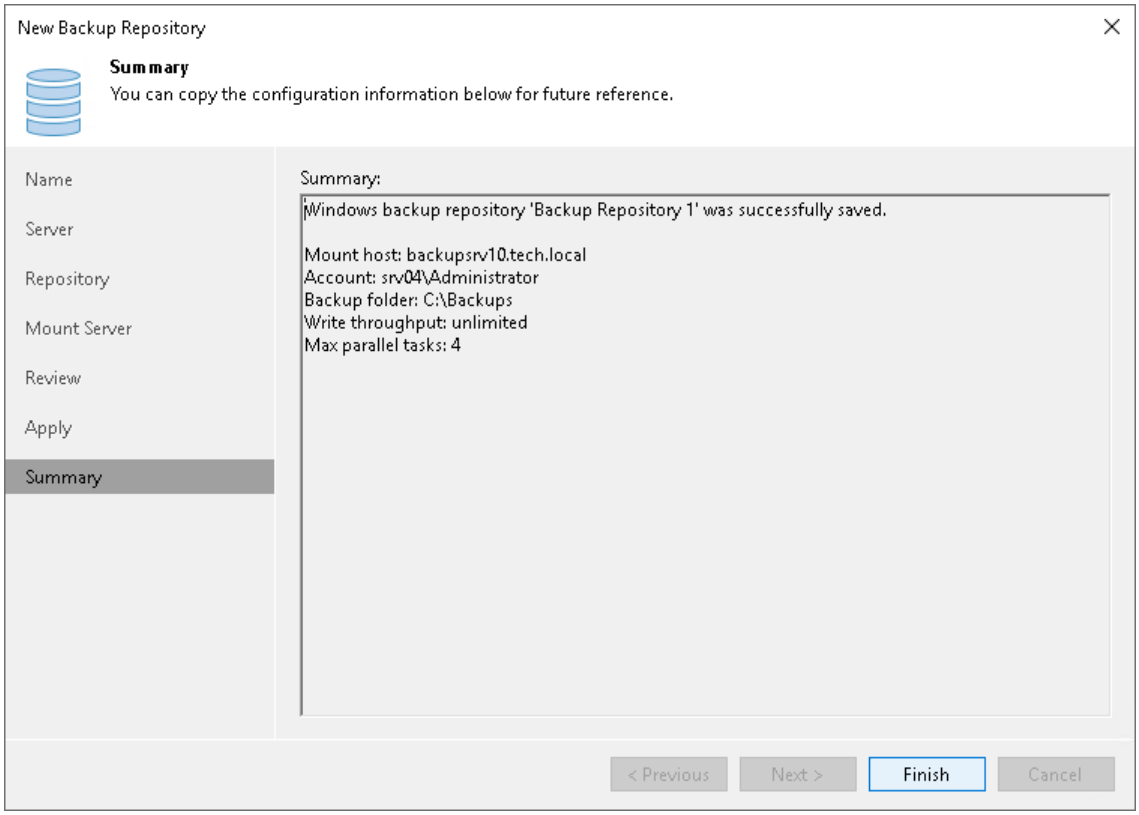

# Linux Server

You can add a Linux server with local, directly attached storage or mounted NFS as a backup repository. The storage can be a local disk, directly attached disk-based storage (such as a USB hard drive), NFS share, or iSCSI/FC SAN LUN in case the server is connected into the SAN fabric.

A Linux repository with single-use credentials and the immutability feature provides additional protection for your backup files. For more information, see Hardened Repository.

## Linux Backup Repository Deployment

To communicate with a Linux-based repository, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the backup proxy
- Veeam Data Mover on the Linux backup repository

For more information about Veeam Data Movers communication with a Linux-based server, see Veeam Data Mover Service.

Veeam Data Mover establishes a connection with the source-side Data Mover on the backup proxy, enabling efficient data transfer over LAN or WAN.

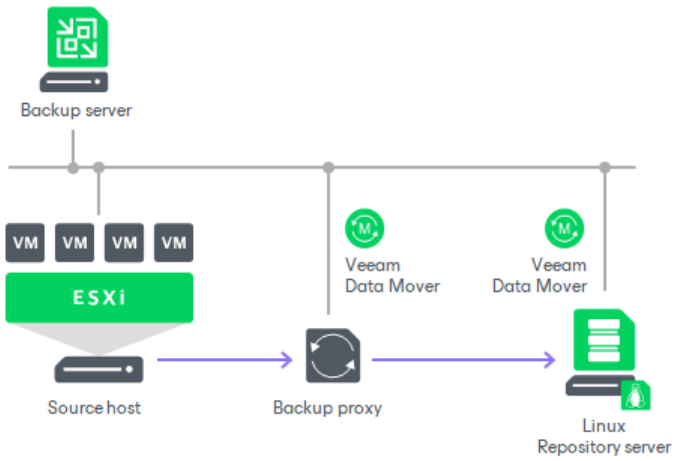


## Requirements for Linux Backup Repositories

A machine performing the role of a repository must meet the following requirements:

- The role of the repository can be assigned to a Linux machine (physical or virtual). The machine must meet the system requirements. For more information, see System Requirements.
- You must add the machine to the Veeam Backup & Replication console as a managed server.
- If Veeam Data Mover Service is non-persistent, Veeam Backup & Replication uses the SSH protocol to communicate with Linux backup repositories and requires the SCP utility in Linux repositories. Make sure that the SSH daemon is properly configured and SCP utility is available on the Linux host.
- If you want to use Fast Clone in the Linux backup repository, the machine must also meet requirements listed in section Fast Clone.

- Depending on the Linux distribution, Veeam services use one of the following Linux firewall managers to operate correctly:
    - `firewalld`
    - `ufw`
    - `iptables`
    - [For IPv6] `ip6tables`

    If none of these firewall managers are installed, make sure that you open all required ports manually. For more information, see Ports.

You can place both repositories (hardened and standard) on one Linux server only if you used single-use credentials when adding the host. Standard repository is a repository added with persistent credentials and disabled immutability. For more hardened repository limitations, see Requirements and Limitations.

# Adding Linux Repositories

This section describes how to add a Linux server as a backup repository.

To add a backup repository, use the **New Backup Repository** wizard.

# Step 1. Launch New Backup Repository Wizard

To launch the **New Backup Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add Backup Repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Direct Attached Storage** > **Linux**.

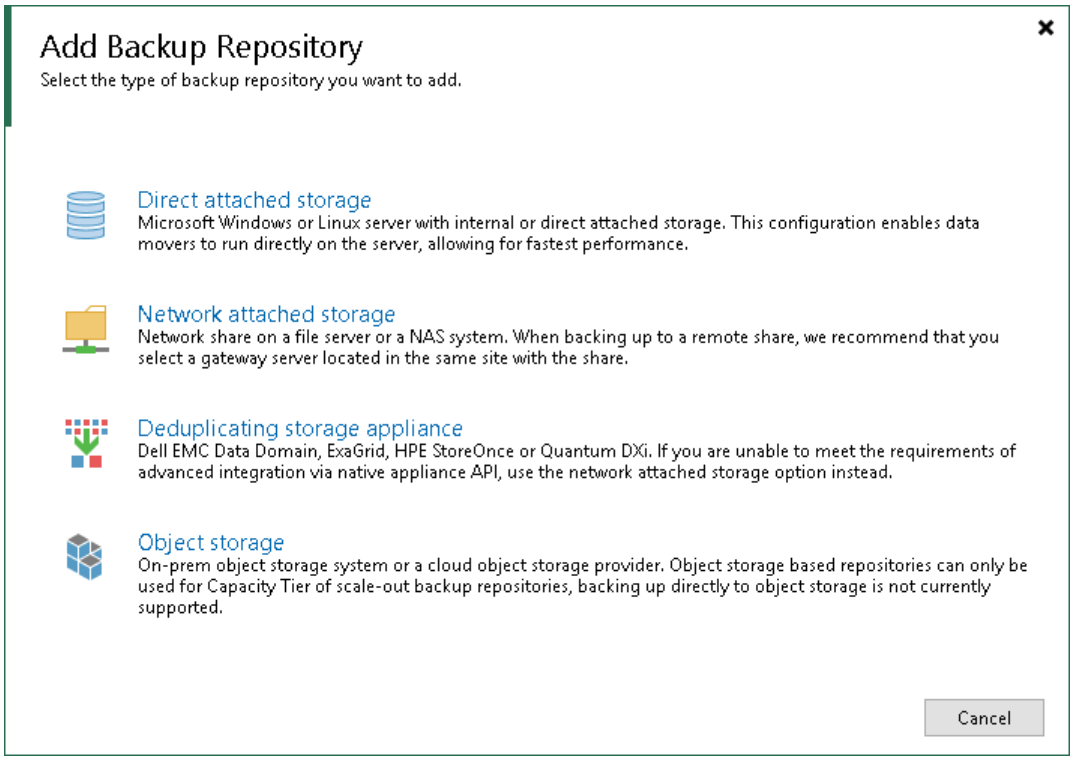

## Step 2. Specify Backup Repository Name and Description

At the **Name** step of the wizard, specify a name and description for the backup repository:

1. In the **Name** field, specify a name for the backup repository.
2. In the **Description** field, provide a description for future reference.

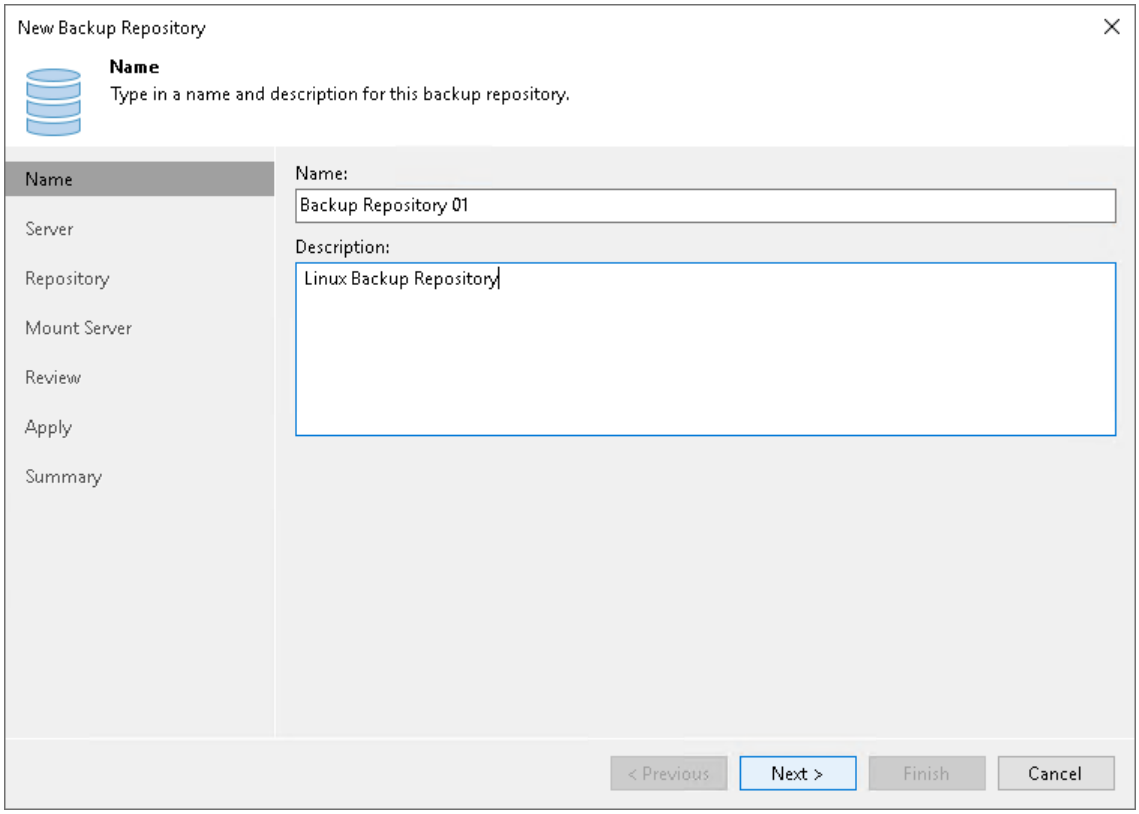

## Step 3. Specify Server Settings

To configure settings for a Linux server:

1. From the **Repository server** list, select the Linux server that you want to use as a backup repository.

   The **Repository server** list contains only those servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, you can click **Add New** on the right to open the **New Linux Server** wizard.

   > **NOTE**
   >
   > Note that you cannot add ExaGrid, Quantum DXi, Fujitsu ETERNUS CS800 and Infinidat InfiniGuard servers as Linux backup repositories. These servers are integrated with Veeam Backup & Replication, and thus must be added as deduplicating storage appliances.

2. Click **Populate** to see a list of disks connected to the server, their capacity and free space.

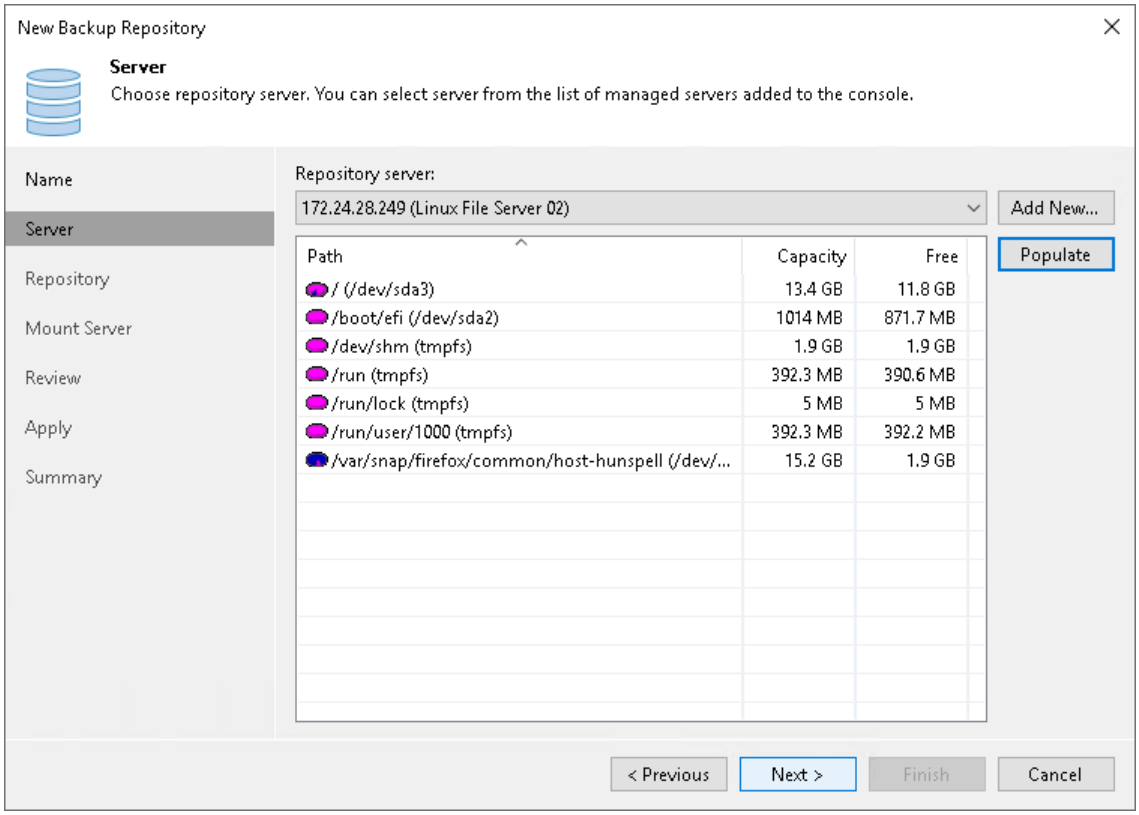

## Step 4. Configure Backup Repository Settings

At the **Repository** step of the wizard, configure general repository settings including path to the repository folder and load control, and also advanced repository settings.

# Configuring General Repository Settings

To configure general repository settings:

1. In the **Location** section, specify a path to the folder where backup files must be stored. Click **Populate** to check capacity and available free space in the selected location.
2. Select the **Use fast cloning on XFS volumes** check box to enable copy-on-write functionality. In terms of Veeam Backup & Replication, this functionality is known as Fast Clone. For more information, see Fast Clone.

3. Use the **Load control** section to limit the number of concurrent tasks and data ingestion rate for the backup repository. These settings will help you control the load on the backup repository and prevent possible timeouts of storage I/O operations.

- Select the **Limit maximum concurrent tasks** check box and specify the maximum allowed number of concurrent tasks for the backup repository. If this value is exceeded, Veeam Backup & Replication will not start a new task until one of current tasks finishes. For more information, see Limiting the Number of Concurrent Tasks.

  > **NOTE**
  >
  > Consider the following:
  >
  > - Limitation of concurrent tasks is ignored if the backup repository acts as a target storage for a Veeam Cloud Connect job.
  > - If you use backup repositories with per-machine backup chains, it is recommended to select the **Limit maximum concurrent tasks to N** check box. This option reduces the number of parallel operations performed by synthetic operations (synthetic full backup, backup files merge and transformation). Otherwise, the load on the backup repository may be high.

- Select the **Limit read and write data rate to** check box and specify the maximum rate to restrict the total speed of reading and writing data to the backup repository disk. For more information, see Limitation of Read and Write Data Rates for Backup Repositories.

  > **NOTE**
  >
  > The **Limit read and write data rate to** setting does not apply to health checks performed as part of backup and backup copy jobs. Even if you limit read/write rate for a backup repository, the health check will consume resources of the backup repository regardless of this setting. Consider this limitation when configuring basic and health check schedules for backup and backup copy jobs.

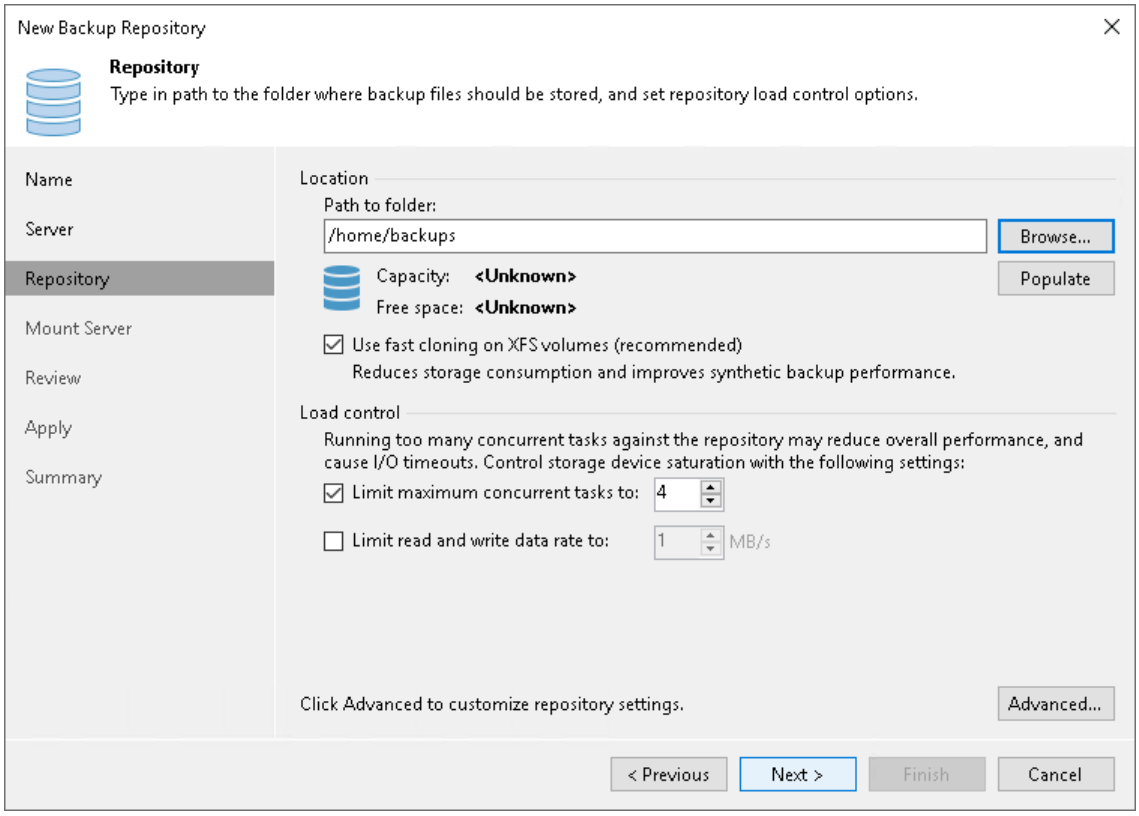


# Configuring Advanced Repository Settings

To configure advanced repository settings:

1. Click **Advanced**.
2. For storage systems using a fixed block size, select the **Align backup file data blocks** check box. Veeam Backup & Replication will align VM data saved to a backup file at a 4 KB block boundary.
3. When you enable compression for a backup job, Veeam Backup & Replication compresses VM data at the source side and then transports it to the target side. Writing compressed data to a deduplicating storage appliance results in poor deduplication ratios as the number of matching blocks decreases. To overcome this situation, select the **Decompress backup data blocks before storing** check box. If data compression is enabled for a job, Veeam Backup & Replication will compress VM data on the source side, transport it to the target side, decompress VM data on the target side and write raw VM data to the storage device to achieve a higher deduplication ratio.

> **NOTE**
>
> Veeam Backup & Replication does not compress VM data if encryption is enabled for a job and the **Decompress backup data blocks before storing** check box is selected in the settings of the target backup repository. Therefore, in the job statistics, you may observe a higher amount of transferred data (the **Transferred** counter) as compared to a job for which encryption is disabled. For more information on job statistics, see Viewing Real-Time Statistics.
>
> In the properties of the backup created with encryption, you may also see that the backup size (the **Backup Size** column) is larger than the original VM size (the **Original Size** column). For more information on backup properties, see Viewing Backup Properties.

4. Select the **This repository is backed by rotated drives** check box if you plan to use a backup repository with rotated drives. For more information on how to configure rotated drives, see Deploying Backup Repositories with Rotated Drives.

5. To create a separate backup file for every machine in the job, make sure that the **Use per-machine backup files** check box is selected. If you clear the check box, Veeam Backup & Replication will create single-file backups. For more information on the backup chain formats and their limitations, see Backup Chain Formats.

> **NOTE**
>
> Changing of the **Use per-machine backup files** setting after the repository was already created does not take any effect. To change backup chain format, follow the instructions provided in Upgrading Backup Chain Formats.

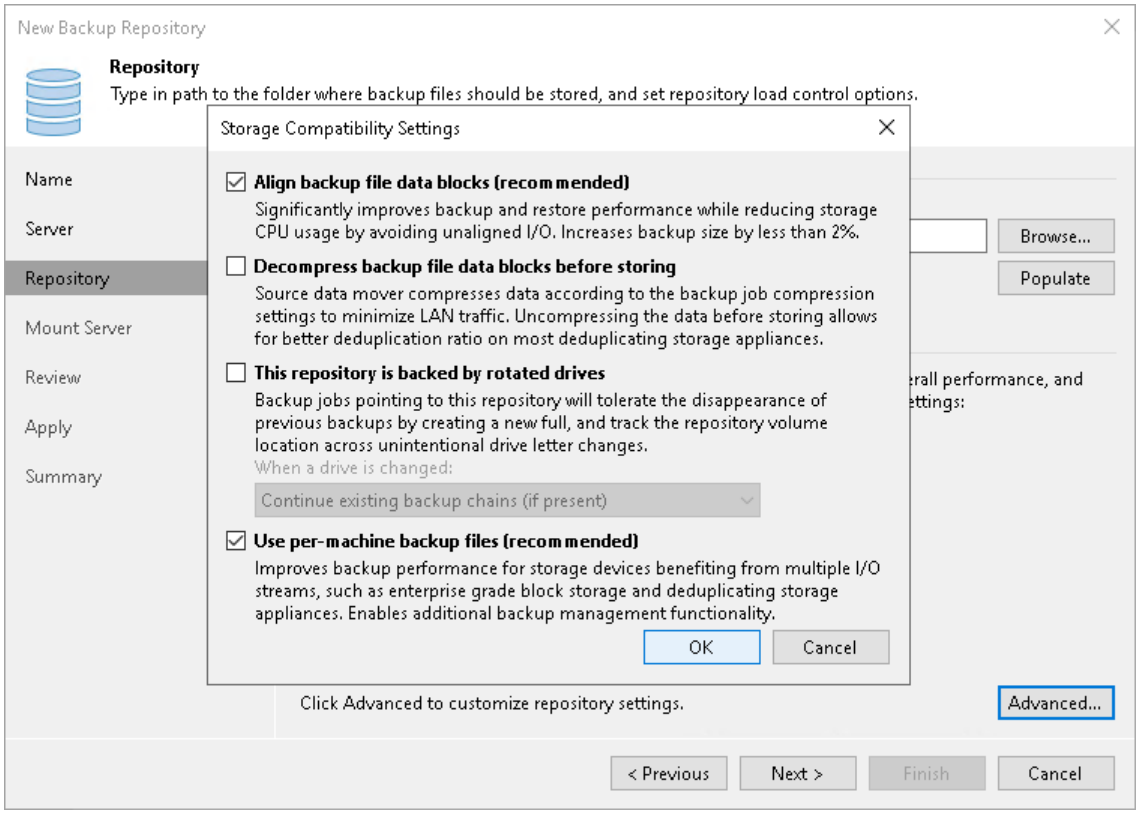

# Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for file-level and application items restore.

1. From the **Mount Server** list, select a server that you want to use as a mount server. The mount server is required for file-level and application items restore. During the restore process, Veeam Backup & Replication mounts the VM disks from the backup file residing in the backup repository to the mount server. As a result, VM data does not have to travel over the network, which reduces the load on the network and speed up the restore process. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers added to the backup infrastructure. If the server is not added to the backup infrastructure, click **Add New** on the right to open the **New Windows Server** wizard. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder that will be used for writing cache during mount operations.

3. To make the backup repository accessible by the Veeam vPower NFS Service, select the **Enable vPower NFS service on the mount server** check box. Veeam Backup & Replication will enable the vPower NFS Service on your selected mount server.

4. To customize network ports used by the vPower NFS Service, click **Ports**. For information on ports used by default, see Ports.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

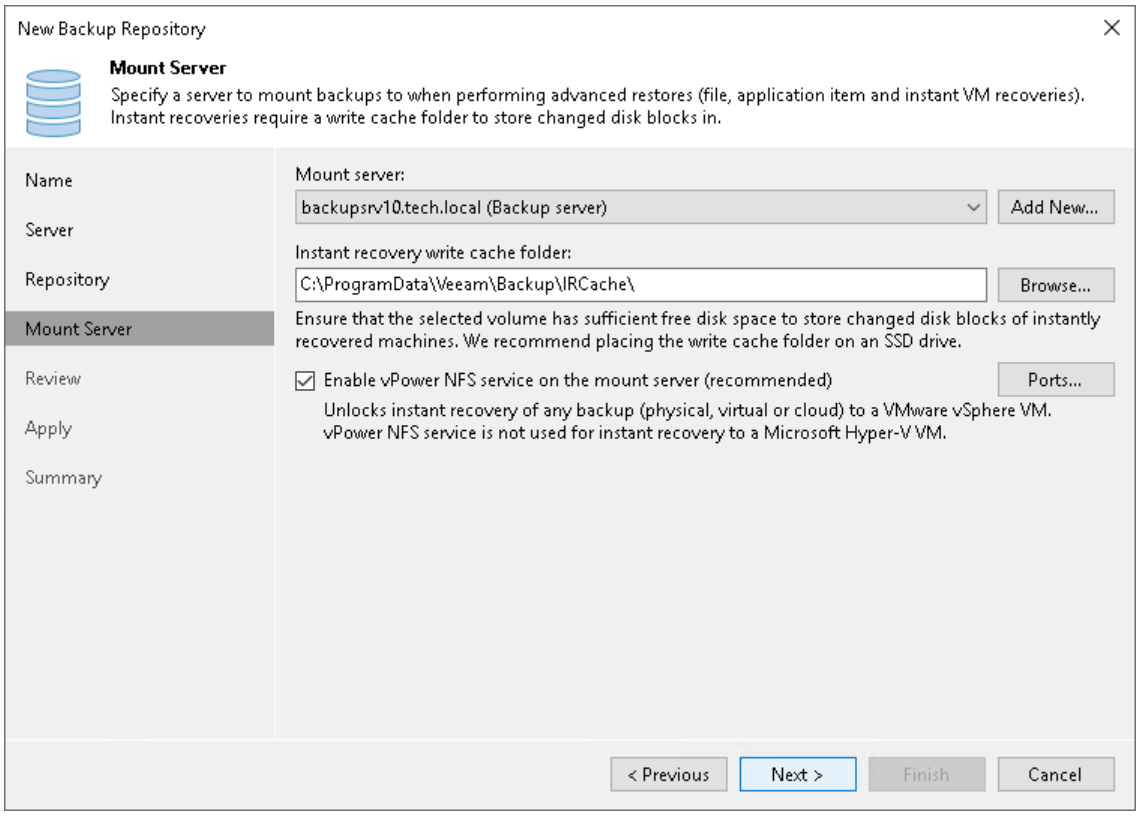

## Step 6. Review Properties and Components

At the **Review** step of the wizard, review details of the backup repository and specify importing settings.

1. Review the backup repository settings and list of components that will be installed on the backup repository server.
2. If the backup repository contains backups that were previously created by Veeam Backup & Replication, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Disk (Imported)** node.
3. If the backup repository contains guest file system index files that were previously created by Veeam Backup & Replication, select the **Import guest file system index** check box. Index files will be imported with backup files, and you will be able to search for guest OS files inside imported backups.

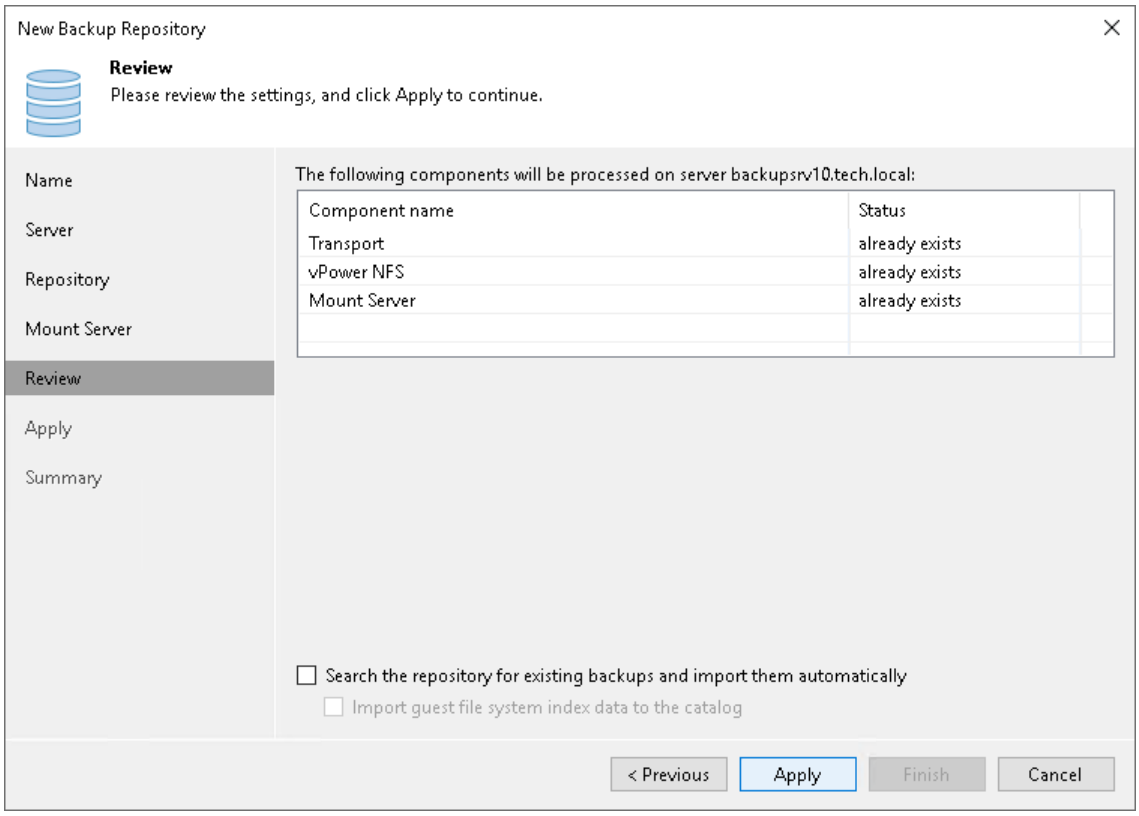

# Step 7. Apply Backup Repository Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to install and configure all required components. Then click **Next** to complete the procedure of adding the backup repository to the backup infrastructure.

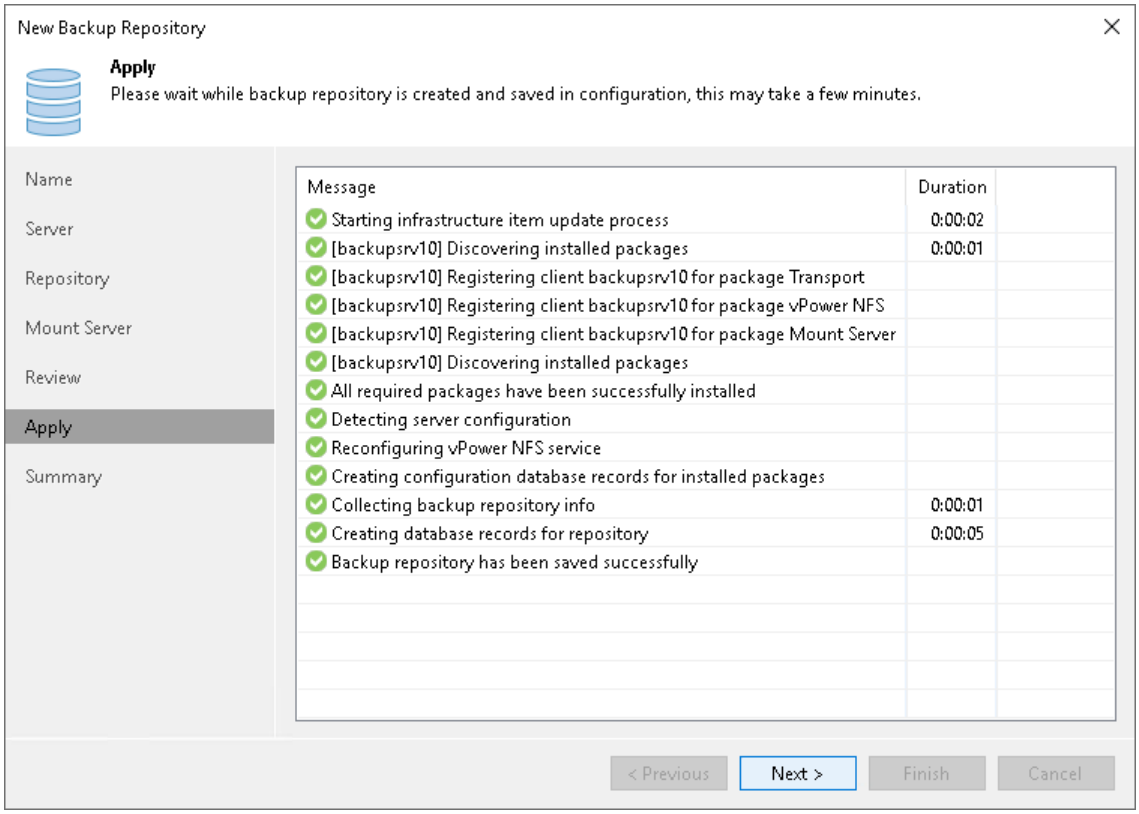

# Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added backup repository. Then click **Finish** to exit the wizard.

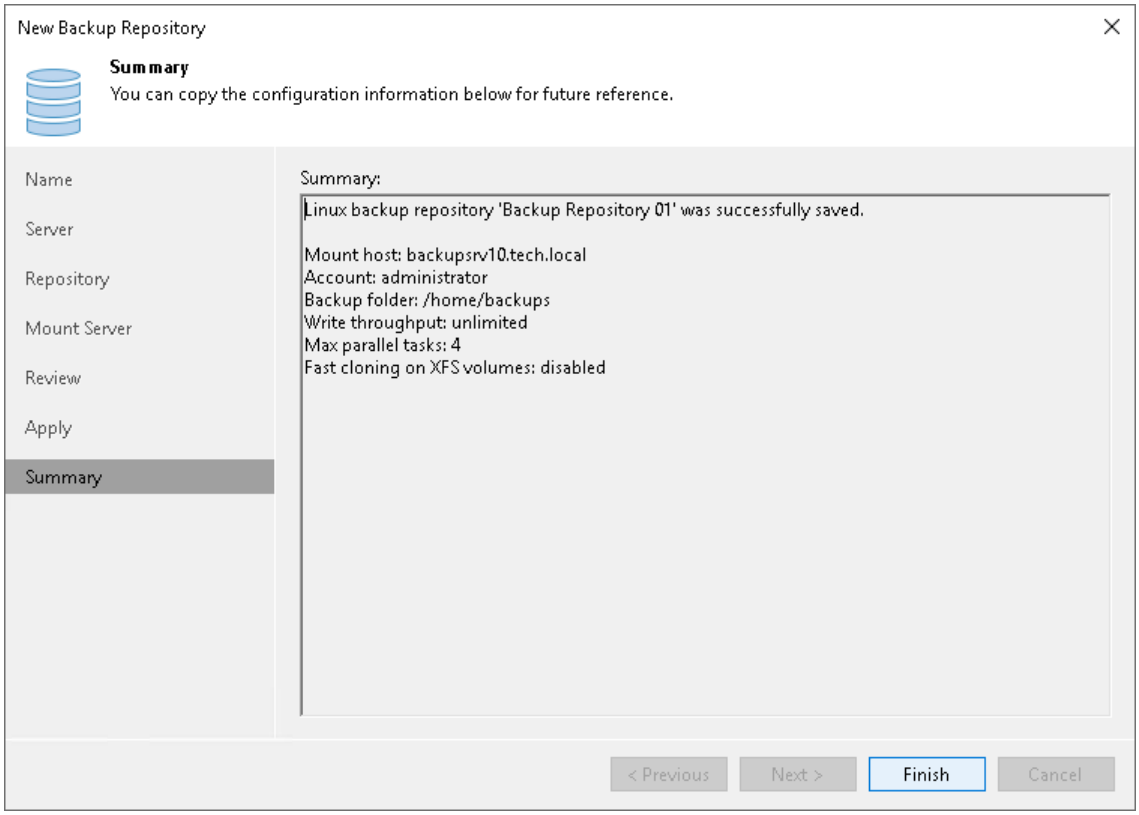

# Hardened Repository

To protect your backup files from loss as a result of malware activity or unplanned actions, you can add to your backup infrastructure a hardened repository based on a Linux server. The hardened repository supports the following features:

- **Immutability**: when you add a hardened repository, you specify the time period while backup files must be immutable. During this period, backup files stored in this repository cannot be moved, modified or deleted, but can be copied.
- **Single-use credentials**: credentials that are used only once to deploy Veeam Data Mover, or transport service, while adding the Linux server to the backup infrastructure. These credentials are not stored in the backup infrastructure. Even if the Veeam Backup & Replication server is compromised, the attacker cannot get the credentials and connect to the hardened repository.

> **NOTE**
>
> For security reasons, you cannot assign other roles to the hardened repository except for the VMware backup proxy working in the **Network** mode. For more information, see Hardened Repository as VMware Backup Proxy.

## About Hardened Repository

In the hardened repository, you can store the following backups created by backup and backup copy jobs:

- Image-level virtual machine backups and backup copies (VMware, Hyper-V, Cloud Director, Nutanix AHV, OLVM, Proxmox and RHV)
- Physical machine backups and backup copies (Microsoft Windows, Linux, MacOS, AIX, Solaris)
- Cloud machine backup copies (Microsoft Azure, AWS, Google Cloud)
- Unstructured data backups (file backups and object storage backups) and their backup copies
- Application-aware processing log backups and backup copies (Microsoft SQL transaction log files, Oracle archived log files, PostgreSQL WAL files)
- Enterprise application backups and backup copies (SAP HANA, Oracle RMAN, SAP on Oracle, Microsoft SQL Server)

Also, backups created by VeeamZIP, Copy Backup jobs, Move Backup jobs, and Export Backup jobs are supported.

## In This Section

- How Immutability Works
- Hardened Repository as Performance Extent
- Hardened Repository as VMware Backup Proxy

## How Immutability Works

Immutability is managed by the following services:

- The Veeam Data Mover Service (`veeamtransport`). It communicates with the backup server, gets information about the immutability time period, and forwards it to the Veeam Immutability Service that manages immutability attributes. The service uses port 6162 and runs as a non-root user.
- The Veeam Immutability Service (`veeamimmureposvc`). It checks file immutability attributes every 20 minutes, calculates the time until a file needs to be immutable, and sets or removes the immutable attribute. The service runs with root permissions as a child process of the Veeam Data Mover Service.

These services are deployed and updated by the Veeam Installer Service for Linux (`veeamdeploymentsvc`). For more information on how the service works, see Veeam Installer Service.

# Timeshift Detection

The Veeam Immutability Service also performs timeshift detection. The mechanism works in the following way:

1. After the Veeam Immutability Service starts, it creates the `/etc/veeam/immureposvc/timeLog` file. The file is updated every 10 minutes and contains the following information:
   - `systemTime` — current UTC value (Unix timestamp)
   - `moveTime` — timeshift value based on comparison between two metrics:
     - difference between current and previous UTC values
     - the value of the check interval (in seconds)

     For example, the value of the check interval is *600* seconds (10 minutes). The previous UTC value is *1699367670*, the current UTC value is *1699368270*. The difference between current and previous UTC values also equals *600* seconds. As two metrics have the same value, the current value of the `moveTime` parameter will not change.
   - `hwTime` — current RTC value (Unix timestamp)
   - `accelerationTime` — timeshift value based on comparison between two metrics:
     - difference between current and previous RTC values
     - difference between current and previous UTC values
2. If the value of the `moveTime` or `accelerationTime` parameter exceeds *86400* seconds (24 hours), retention operations will be blocked with the warning in the backup job session.

**NOTE**

Consider the following:

- To detect timeshift more precisely, RTC should be set to UTC.
- If RTC is disabled, the `accelerationTime` parameter value will be ignored. The Veeam Immutability Service will check only the `moveTime` parameter.
- If you shut down the repository for more than 24 hours, retention operations will be blocked.
- If retention operations are blocked, the Veeam Immutability Service will not set or remove the immutable attribute for existing and new backup files.

To get back to the normal operation mode, you will need to access the hardened repository under the user account with root privileges and perform the necessary operations. For more information, see this Veeam KB article.

The described hardened repository architecture prevents backup files from being deleted or modified by a potential attacker even if they exploit the Veeam Data Mover Service or compromise the NTP server. For more details, see Protect against Ransomware with Immutable Backups.

# Immutability for Image-Level VM and Physical Machine Backups

For image-level VM backups and physical machine backups, immutability works in the following way:

1. For each backup file, Veeam Backup & Replication creates a `.veeam.N.lock` file that contains information about the immutability time period. Also, the `xattr` attribute with immutable time period is set on each backup file. These files are stored on the hardened repository.
2. Backup files become immutable for the configured time period (minimum 7 days, maximum — 9999). The immutability period is set according to the following:
   - The count of the immutability period starts from the moment the last restore point in the active backup chain is created. For example:
     - The full backup file of the active chain was created on January 12. The first increment was created on January 13. The last increment was created on January 14.
     - The immutability period is set for 10 days and will be automatically extended for all backup files in the active chain. If there are several chains in the backup, Veeam Backup & Replication does not extend the immutability for inactive chains.
     - Full and incremental backup files will be immutable until January 24: the date of the last restore point creation (January 14) + 10 days.
   - The immutability flag is set on the file only when the current backup session is completed.

     **NOTE**

     For image-level VM backups, consider the following:

     - If you use per-machine backup with single metadata file or per-machine backup with separate metadata files format for storing backups and the restore point was incomplete, the immutability flag will be set only on successfully created backup files.
     - If you use single-file backup format for storing backups and the restore point was incomplete, the immutability flag will not be set on the backup file.

     For more information about backup chain formats, see Backup Chain Formats.
   - If you increase the immutability period in repository settings, a new value will be applied for the active and next chains.
   - If you decrease the immutability period and a new value for the next restore point is less than the previous one, it will be applied for the next incremental backups in the active chain and for the next chains. For example:
     - The immutability period is set for 20 days. The full backup file of the active chain was created on November 1. The first increment was created on November 2. The second increment was created on November 3. Full and incremental backup files will be immutable until November 23: the date of the last restore point creation (November 3) + 20 days.

        - If you decrease the immutability period to 7 days on November 4, all previous backup files in the active chain will be immutable until November 23. The next incremental backup in the active chain will be immutable until November 11: the date of the last restore point creation (November 4) + 7 days.
    - If you decrease the immutability period and a new value for the next restore point is greater than the previous one, it will be applied for all backup files in the active chain and in the next chains. For example:
        - The immutability period is set for 20 days. The full backup file of the active chain was created on November 1. The first increment was created on November 2. The second increment was created on November 3. Full and incremental backup files will be immutable until November 23: the date of the last restore point creation (November 3) + 20 days.
        - If you decrease the immutability period to 7 days on November 22, all backup files in the active chain will be immutable until November 29: the date of the last restore point creation (November 22) + 7 days.
    - If you use GFS retention policy, see Retention Scenarios.
- When the immutability time period expires, the immutability service makes backup files mutable again so they can be deleted or modified.

# Immutability for Application-Aware Processing Log Backups

For application-aware processing log backups, immutability works in the following way:

1. For each log backup file, Veeam Backup & Replication creates a `.veeam.N.lock` file that contain information about immutability time period. These files are stored on the hardened repository.
2. Log backup files become immutable for the configured time period (minimum 7 days, maximum — 9999). The immutability period is set according to the following:
    - Newly created log backup files are being updated several times according to the interval settings. Thus, the count of the immutability period starts and the immutability flag is set on the file only when the log backup job finished writing any data to the VLB file.
    - The immutability period is not extended for log backup files in the active chain. If there are several chains in the backup, Veeam Backup & Replication also does not extend the immutability for old chains.

      **IMPORTANT**

      If the immutability period is expired for log backup files in the active chain, these files become mutable and may be potentially removed. In this case, the application restore may fail as the log backup chain becomes incomplete. To mitigate risks, make sure that the immutability period covers all backups in the active backup chain.

    - If you increase the immutability period in repository settings, a new value will be applied for all log backup files created after the last successful image-level VM backup or physical machine backup. If you decrease the immutability period, a new value will be applied only for the next log backup files.
    - If you use GFS retention policy, see Retention Scenarios.
- When the immutability time period expires, the immutability service makes backup files mutable again so they can be deleted or modified.

## Immutability for Other Backups

For information on how immutability works for unstructured data backups and enterprise application backups, see the following sections:

- Unstructured Data Backups in Immutable Repositories
- Veeam Plug-in for Oracle RMAN
- Veeam Plug-in for SAP HANA
- Veeam Plug-in for SAP on Oracle
- Veeam Plug-in for Microsoft SQL Server

## Retention Scenarios

Consider the following retention scenarios:

- An immutability retention overrides a job retention: if the job retention period is shorter than the immutability period, Veeam Backup & Replication does not delete backup files when the retention period is over, but only when the immutability period expires.
- Veeam Backup & Replication compares the immutability period of the backup repository and the GFS backup file lifetime, and sets an immutability period for full backup files with GFS retention policy as equal to the longest of these periods. For example: the backup repository immutability period is 10 days; the GFS backup file lifetime is 3 years; the backup file will be immutable for 3 years; the increments from this full backup file will be immutable for 10 days from the moment of the last increment creation.
- The immutability period for backup files produced with VeeamZIP or Export Backup jobs is set according to the following:
  - [With enabled retention period] Veeam Backup & Replication compares the immutability period of the backup repository and the retention period, and sets an immutability period for backup files with retention period as equal to the longest of these periods. For example: the backup repository immutability period is 1 month; the VeeamZIP or Export Backup backup file lifetime is 7 years; the backup file will be immutable for 7 years.

    > **NOTE**
    >
    > If a hardened repository is a part of a scale-out backup repository with the capacity tier added and the move policy enabled and is used as a target for VeeamZIP or Export Backup jobs, Veeam Backup & Replication ignores the VeeamZIP or Export Backup retention period. The immutability time period for VeeamZIP or Export Backup backup files equals the period specified in the setting of a hardened repository.

  - [With disabled retention period] Veeam Backup & Replication ignores the VeeamZIP or Export Backup retention period. The immutability time period for VeeamZIP or Export Backup backup files equals the period specified in the setting of a hardened repository.

## Hardened Repository as Performance Extent

You can add a hardened repository to your scale-out backup repository as a performance extent. For more information, see Scale-Out Backup Repository and Add Performance Extents.

# Immutability Limitations

- Mutable and immutable extents must not be mixed within one scale-out backup repository.
- If your scale-out backup repository includes hardened repositories and deduplicating storage appliances with enabled immutability (StoreOnce, Dell Data Domain), consider the following:
  - The immutability time period must be the same on all extents added to this scale-out repository. If you want to change the immutability period on any extent, a new value will be applied to all extents.
  - If you want to upgrade to the latest Veeam Backup & Replication version, before starting the operation, you must specify the same immutability period for all extents added to this scale-out repository.

# Using Capacity Tiers and Hardened Repositories

If you use the capacity tier with move option, note that having a hardened repository as a performance extent will affect the capacity tier behavior. You will not be able to move immutable backup files, because they cannot be deleted from the performance extent. Veeam Backup & Replication will copy such backup files to the capacity tier. When the immutability time period is over, Veeam Backup & Replication will delete these files from the performance extent. For more information on copy and move policies, see Copying Backups to Capacity Tier and Moving Backups to Capacity Tier.

If a hardened repository is a part of a scale-out backup repository with the capacity tier added and the move policy enabled, Veeam Backup & Replication ignores the GFS retention policy. The immutability time period for full backup files equals the period specified in the setting of a hardened repository.

# Evacuating Immutable Backups

If you evacuate your backups from an immutable performance extent, Veeam Backup & Replication will copy them instead of moving. When the immutability time period is over, you will need to delete these files manually. If the target extent is also immutable, the immutability of the target extent will apply to copied backup files. For more information on evacuating backups, see Evacuating Backups from Extents.

When you evacuate backups to the hardened Linux extent, the immutability period of the full chain will be chosen as maximum between the following values:

- The immutability period determined for the original chain.
- The time of creation of the last restore point in the chain plus the immutability period determined for the target extent.

## Hardened Repository as VMware Backup Proxy

You can assign the VMware backup proxy role to the hardened repository. In this case, only the *Network* mode (NBD) is supported. Other transport modes will not be available for selection.

> **NOTE**
>
> A VMware backup proxy requires VMware VDDK components to be installed. This increases the risk of hardened repository attacks through VMware VDDK vulnerabilities. It is recommended to assign this role to another managed server when possible.

To add a hardened repository as a VMware backup proxy, use the **New VMware Proxy** wizard. For more information, see Adding VMware Backup Proxies.

# Requirements and Limitations

For the hardened repository, consider the following requirements and limitations.

# Linux Server

- The role of the hardened repository can be assigned to a Linux machine with local or remotely attached block storage. The machine must meet system requirements for backup repositories.

  > **NOTE**
  >
  > To reduce the attack surface, use a physical machine with local storage. For RAID configuration, recommendations are the following:
  >
  > - [For the operating system] RAID 1 on SSDs with at least 100 GB disk space should be used.
  > - [For backup data] RAID 6/60 with write-back cache should be used. At least one disk must be configured for the drive roaming.
  > - Internal disk cache must be disabled.
  > - RAID stripe size should be 128 or 256 KB.

- The Linux distribution must be 64-bit due to Veeam Data Mover requirements. If you use the following Linux distributions, you also need to upgrade Veeam Backup & Replication to version 12.1.2 (build 12.1.2.172):
  - RHEL 8 and 9 with DISA STIG profile enabled.
  - Rocky Linux 9 with DISA STIG profile enabled.
- The Linux machine file system must support immutable files and extended attributes modified by the chattr and setxattr commands. We recommend using XFS for performance and space efficiency reasons (block cloning support).
- As the hardened repository requires the block storage, you cannot use the following storage types:
  - NFS share or a Linux machine with the mounted NFS volume.
  - A Linux machine with the mounted SMB (CIFS) volume.
- Depending on the Linux distribution, Veeam services use one of the following Linux firewall managers to operate correctly:
  - `firewalld`
  - `ufw`
  - `iptables`
  - [For IPv6] `ip6tables`

  If none of these firewall managers are installed, make sure that you open all required ports manually. For more information, see Ports.
- You must add the Linux machine to the Veeam Backup & Replication console as a managed server. The hardened repository cannot be shared between different Veeam Backup & Replication servers.
- The Linux machine should have redundant network connection.
- Veeam Agent for Linux must not be installed on the Linux server.

### Repository

- For the separate directory that you created for the backup data, consider the following:
  - Both `owner` and `group` must be the user account you use to connect to the Linux server.
  - Directory permissions must be *0700*.
  - Directory must not have a sticky bit.
- To store backup files in a repository, use only a forward incremental backup method with enabled active full backup or synthetic full backup. Once a backup file becomes immutable, it can be merged or deleted only when the immutability time period expires. For this reason, you cannot select a reverse or a forever forward incremental backup method.
- For importing a backup, use VBK backup files. Metadata files of a backup chain (.VBM) cannot be immutable because they are updated on every job pass.
- Veeam Backup & Replication does not support symlinks in the path to the hardened repository.
- Only Veeam Data Mover Service child processes can run under the root account on a hardened repository.

### Immutability Feature

- To use the immutability feature for backup copy jobs, enable the GFS retention policy. For more information, see Long-Term Retention Policy (GFS).
- Do not use the immutability feature for a Nutanix Mine infrastructure. As Mine repositories contain thin-provisioned disks, there may be the case when Veeam Backup & Replication uses full storage capacity of a repository and cannot delete backup files from the file system.

## Preparing Ubuntu Linux Server as Hardened Repository

This section includes security considerations for installing and configuring the Linux server that will be used as a hardened repository. Recommendations are based on Security Technical Implementation Guides (STIGs) created and maintained by the Defense Information Systems Agency (DISA) for Ubuntu 20.04 LTS. For more information, see DISA STIGs Document Library.

### Installing Ubuntu Linux Server

To install Ubuntu 20.04 LTS, download the server install image from the Ubuntu Releases page. For more information on the installer and options of the installation wizard, see the official Canonical guide.

During installation process, consider the following Veeam recommendations:

1. Before you boot the installer, enable UEFI secure boot to prevent unsigned Linux kernel modules from being loaded.
2. In the GRUB menu, select the **Boot and Install with the HWE kernel** option to support the latest hardware.
3. At the welcome screen of the installation wizard, select the language for the installer and the default language for the installed system. For troubleshooting purposes, it is recommended to select the English language.
4. At the **Installer update available** step of the installation wizard, select the **Continue without updating** option.

5. At the **Keyboard configuration** step of the installation wizard, set up the keyboard layout used in your backup infrastructure.
6. At the **Network connections** step of the installation wizard, do the following:
   - If you have several network interface cards, create a bond to provide the network failover in case of connection issues. For the bond mode, select one of the following options:
     - *balance-rr* (if you use EtherChannel without LACP)
     - *802.3ad* (if you use EtherChannel with LACP)
     - *active-backup* (for other configurations)

     Example:

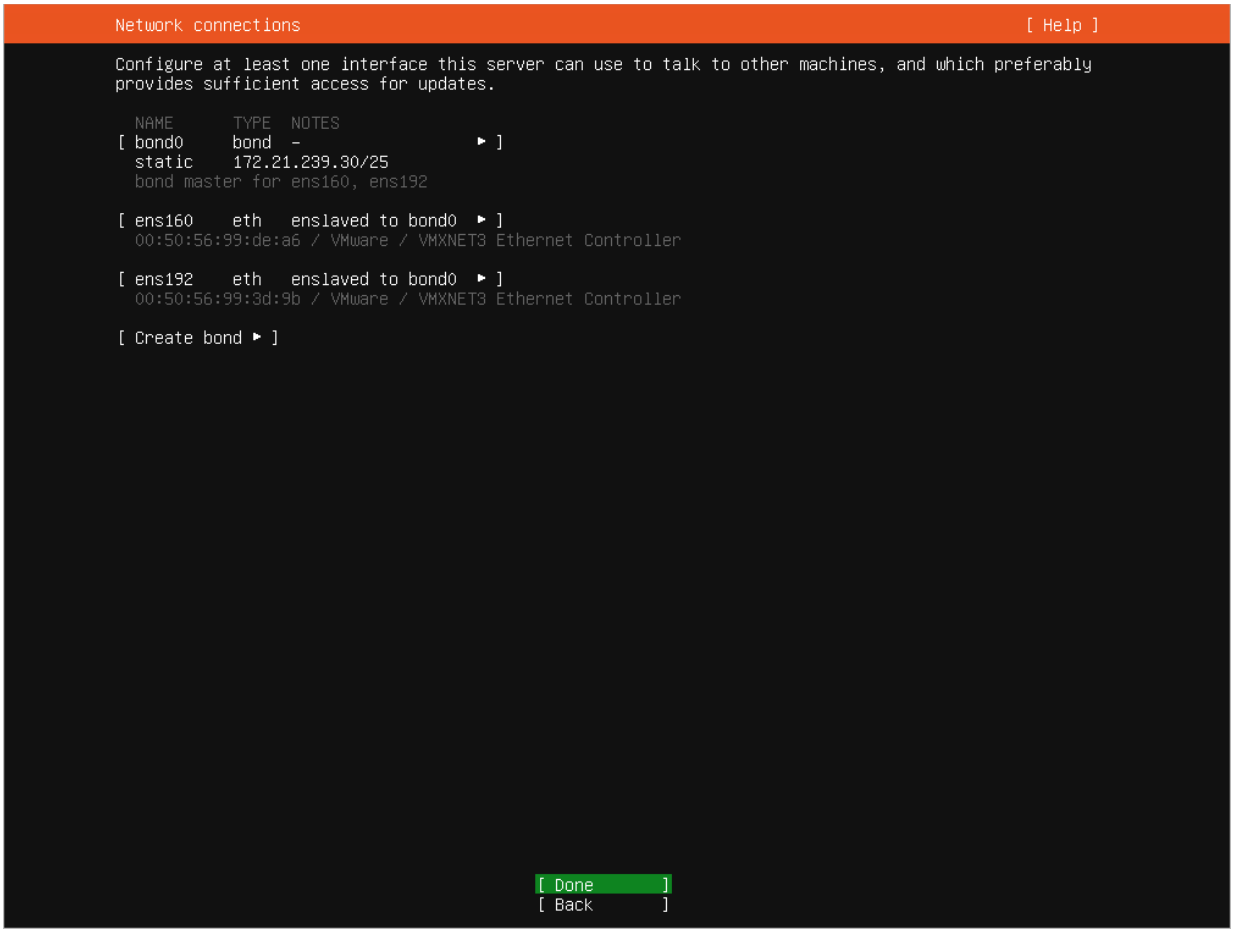


   - If you have only one network interface card and cannot create a bond, assign the static IP address to the network interface to reduce the risk of connection issues, for example, with the DHCP server.

- At the **Configure proxy** step of the installation wizard, specify the proxy server if required.
- At the **Configure Ubuntu archive mirror** step of the installation wizard, leave the default mirror address.

- At the **Storage configuration** step of the installation wizard, follow recommendations from CIS Benchmarks for Ubuntu Linux 20.04 LTS STIG for partitioning.

  For the operating system, use the ext4 file system. Example:

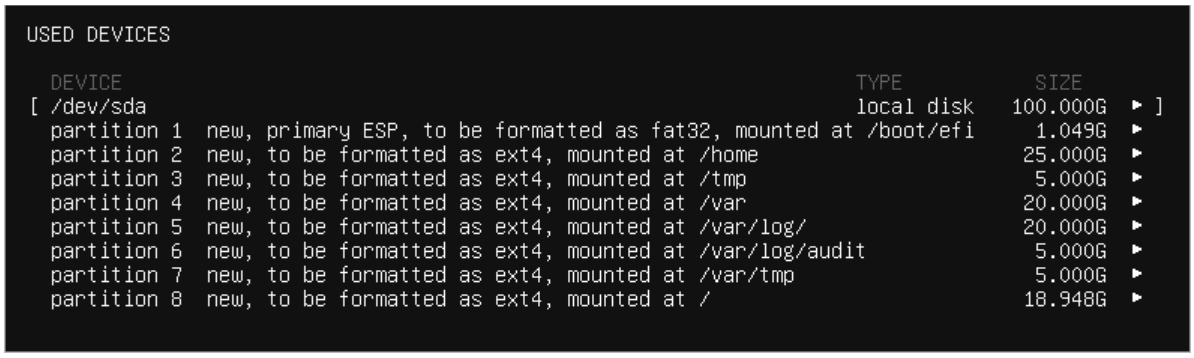


  For the backup data, use the XFS file system. Example:

  > **NOTE**
  >
  > To be compliant with DISA STIG UBTU-20-010414, you do not need to enable disk encryption for the operating system. To protect data in backups, use Veeam Backup & Replication built-in encryption instead. For more information, see Storage Settings.

  After you add partitions for all disks, click **Continue** in the **Confirm destructive action** window to apply changes. Note that all data on the disks will be deleted.

10. At the **Profile setup** step of the installation wizard, specify a hostname and a user account that you will use to connect to the Linux server. Consider that by default it will have sudo permissions. After you add a hardened repository to the backup infrastructure, you must remove this user account from the `sudo` group. For more information, see Post-Installation.
11. At the **SSH Setup** step of the installation wizard, select the **Install OpenSSH server** check box. The OpenSSH server is required to be compliant with DISA STIG UBTU-20-010042 and for deployment and upgrade of Veeam Data Mover.
12. At the **Featured Server Snaps** step of the installation wizard, do not install any additional packages. Click **Done** to start the installation process.

After the installation finishes, remove the installation media and reboot the system.

# Configuring Ubuntu Linux Server

When you configure system settings, consider the following Veeam recommendations:

- To be compliant with DISA STIG UBTU-20-010009, set a password for GRUB. To configure the setting manually, do the following:

  a. Create a password using the `grub-mkpasswd-pbkdf2` command:

  ```
  grub-mkpasswd-pbkdf2
  Enter password:
  Reenter password:
  PBKDF2 hash of your password is grub.pbkdf2.sha512.10000.C0F70D240A8BC5
  C1BC4E1303EC4F040957C1AF1BB8E99EED573133D3A017BE9B2BB48E52577A141B3A695
  2527A9D1BEF13E2BB29978DA71F2D867EBB03545021.C4E81CAE7B464E78B15DF0A578B
  63BAB3A0CB180C311AFA5A85F6245800D11D40B37B817C3F30348EE603AF725B7E09B98
  A291114B0206D[…]
  ```

  b. Add a user name and a password hash at the end of the `/etc/grub.d/40_custom` file:

  ```
  set superusers="root"
  password_pbkdf2 root grub.pbkdf2.sha512.10000.C0F70D240A8BC5F[…]
  ```

  c. To disable asking for credentials after rebooting the system and require them only when editing boot menu entries, open the `/etc/grub.d/10_linux` file and add the `--unrestricted` parameter to the `CLASS` variable:

  ```
  CLASS="--class gnu-linux --class gnu --class os --unrestricted"
  ```

  d. Update the GRUB configuration:

  ```
  sudo update-grub
  ```

- To be compliant with DISA STIG UBTU-20-010455, disable all wireless network adapters, as described in this DISA STIG article.
- If you do not use the proxy server and the Linux server has outgoing HTTP internet access allowed, limit outgoing HTTP traffic to the Ubuntu servers only or use an internal Ubuntu mirror. To receive Linux security updates, there must be the access to the Linux distribution security update servers.
- For the separate directory that you created for the backup data, allow access only for the user account you created during the installation. Use the following commands:
  - To assign the directory's owner:

    ```
    chown -R owner:group <dir_path>
    ```

    Both `owner` and `group` must be the user account you created during the installation.
  - To allow access to the directory only for its owner and the root account:

```
chmod 700 <dir_path>
```

- To be compliant with DISA STIG UBTU-20-010012, you must have only two users:
  - The root account. Note that by default the root account has a blank password and cannot be used for connection.
  - The user account you created during the installation. This account will be used to connect to the Linux server and deploy required Veeam Backup & Replication components including persistent Veeam Data Mover, or transport service. For more information about Veeam Data Movers, see this section.

  By default, the user account you created during the installation is the member of the sudo group and has enough privileges to deploy and install required Veeam Backup & Replication components. In that case, when you add a Linux server as a hardened repository to the backup infrastructure and specify single-use credentials, you do not need to enter the password for the root account. After the repository is added, you must remove the user account from the sudo group to make it a non-root account. To do this, perform the following steps:

  a. Allow the user account to reboot and shutdown the operating system:

  ```
  sudo bash -c "echo 'user1 ALL = (root) NOEXEC: /usr/sbin/reboot' >> /et
  c/sudoers"
  sudo bash -c "echo 'user1 ALL = (root) NOEXEC: /usr/sbin/shutdown' >> /
  etc/sudoers"
  ```

  b. Remove the user account from the `sudo` group:

  ```
  sudo deluser user1 sudo
  ```

  Note that the next time you log in with this user account, it will lose sudo permissions. if you need to execute commands as a privileged user, you must boot the operating system into the single user mode.

## Configuring DISA STIG Compliance for Ubuntu Linux Server

The Linux server based on Ubuntu 20.04 LTS can be configured according to the following DISA STIGs. Settings can be applied manually or using the automatic configuration script provided by Veeam. For more information about the script, see this Veeam page.

| DISA STIG ID | Description | How to apply |
|---|---|---|
| UBTU-20-010000 | The Ubuntu operating system must provision temporary user accounts with an expiration time of 72 hours or less. | Not applicable. A hardened repository does not use temporary user accounts. |

| DISA STIG ID | Description | How to apply |
|---|---|---|
| UBTU-20-010002 | The Ubuntu operating system must enable the graphical user logon banner to display the Standard Mandatory DoD Notice and Consent Banner before granting local access to the system through a graphical user logon. | Not applicable. The operating system is installed without GUI. |
| UBTU-20-010003 | The Ubuntu operating system must display the Standard Mandatory DoD Notice and Consent Banner before granting local access to the system through a graphical user logon. | Not applicable.The operating system is installed without GUI. |
| UBTU-20-010004 | The Ubuntu operating system must retain a user's session lock until that user reestablishes access using established identification and authentication procedures. | Not applicable.The operating system is installed without GUI. |
| UBTU-20-010005 | The Ubuntu operating system must allow users to directly initiate a session lock for all connection types. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010006 | The Ubuntu operating system must map the authenticated identity to the user or group account for PKI-based authentication. | Not applicable. A hardened repository uses only a root account and a non-root account with reduced permissions. |
| UBTU-20-010007 | The Ubuntu operating system must enforce 24 hours/1 day as the minimum password lifetime. Passwords for new users must have a 24 hours/1 day minimum password lifetime restriction. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010008 | The Ubuntu operating system must enforce a 60-day maximum password lifetime restriction. Passwords for new users must have a 60-day maximum password lifetime restriction. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010010 | The Ubuntu operating system must uniquely identify interactive users. | Not applicable. A hardened repository uses only a root account and a non-root account with reduced permissions. |
| UBTU-20-010013 | The Ubuntu operating system must automatically terminate a user session after inactivity timeouts have expired. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010014 | The Ubuntu operating system must require users to reauthenticate for privilege escalation or when changing roles. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010016 | The Ubuntu operating system default filesystem permissions must be defined in such a way that all authenticated users can read and modify only their own files. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010033 | The Ubuntu operating system must implement smart card logins for multifactor authentication for local and network access to privileged and non-privileged accounts. | Not applicable. A hardened repository uses only a root account and a non-root account with reduced permissions. |
| UBTU-20-010035 | The Ubuntu operating system must use strong authenticators in establishing nonlocal maintenance and diagnostic sessions. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010036 | The Ubuntu operating system must immediately terminate all network connections associated with SSH traffic after a period of inactivity. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010037 | The Ubuntu operating system must immediately terminate all network connections associated with SSH traffic at the end of the session or after 10 minutes of inactivity. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010038 | The Ubuntu operating system must display the Standard Mandatory DoD Notice and Consent Banner before granting any local or remote connection to the system. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. A consent banner includes the following text: "WARNING: Unauthorized access to this system is forbidden and will be prosecuted by law. By accessing this system, you agree that your actions may be monitored if unauthorized usage is suspected". |
| UBTU-20-010043 | The Ubuntu operating system must configure the SSH daemon to use Message Authentication Codes (MACs) employing FIPS 140-2 approved cryptographic hashes to prevent the unauthorized disclosure of information and detect changes to information during transmission. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010044 | The Ubuntu operating system must configure the SSH daemon to use FIPS 140-2 approved ciphers to prevent the unauthorized disclosure of information and detect changes to information during transmission. | Not applicable. SSH connection is necessary only for the deployment and upgrade of Veeam Data Mover and can be disabled after you add the hardened repository to the backup infrastructure. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010047 | The Ubuntu operating system must not allow unattended or automatic login through SSH. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010048 | The Ubuntu operating system must be configured so that remote X connections are disabled, unless to fulfill documented and validated mission requirements. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010049 | The Ubuntu operating system SSH daemon must prevent remote hosts from connecting to the proxy display. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010050 | The Ubuntu operating system must enforce password complexity by requiring that at least one upper-case character be used. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010051 | The Ubuntu operating system must enforce password complexity by requiring that at least one lower-case character be used. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010052 | The Ubuntu operating system must enforce password complexity by requiring that at least one numeric character be used. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010053 | The Ubuntu operating system must require the change of at least 8 characters when passwords are changed. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
|---|---|---|
| UBTU-20-010054 | The Ubuntu operating system must enforce a minimum 15-character password length. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010055 | The Ubuntu operating system must enforce password complexity by requiring that at least one special character be used. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010056 | The Ubuntu operating system must prevent the use of dictionary words for passwords. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010057 | The Ubuntu operating system must be configured so that when passwords are changed or new passwords are established, pwquality must be used. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010060 | The Ubuntu operating system, for PKI-based authentication, must validate certificates by constructing a certification path (which includes status information) to an accepted trust anchor. | Not applicable. A hardened repository does not use PKI-based authentication. |
| UBTU-20-010063 | The Ubuntu operating system must implement multifactor authentication for remote access to privileged accounts in such a way that one of the factors is provided by a device separate from the system gaining access. | Not applicable. A hardened repository does not use remote access to privileged accounts. |
| UBTU-20-010064 | The Ubuntu operating system must accept Personal Identity Verification (PIV) credentials. | Not applicable. A hardened repository does not use interactive users and smart cards. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010065 | The Ubuntu operating system must electronically verify Personal Identity Verification (PIV) credentials. | Not applicable. A hardened repository does not use interactive users and smart cards. |
| UBTU-20-010066 | The Ubuntu operating system for PKI-based authentication, must implement a local cache of revocation data in case of the inability to access revocation information through the network. | Not applicable. A hardened repository does not use interactive users and smart cards. |
| UBTU-20-010070 | The Ubuntu operating system must prohibit password reuse for a minimum of five generations. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010072 | The Ubuntu operating system must automatically lock an account until the locked account is released by an administrator when three unsuccessful logon attempts have been made. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010074 | The Ubuntu operating system must be configured so that the script which runs each 30 days or less to check file integrity is the default one. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010075 | The Ubuntu operating system must enforce a delay of at least 4 seconds between logon prompts following a failed logon attempt. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010100 | The Ubuntu operating system must generate audit records for all account creations, modifications, disabling, and termination events that affect /etc/passwd. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010101 | The Ubuntu operating system must generate audit records for all account creations, modifications, disabling, and termination events that affect /etc/group. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010102 | The Ubuntu operating system must generate audit records for all account creations, modifications, disabling, and termination events that affect /etc/shadow. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010103 | The Ubuntu operating system must generate audit records for all account creations, modifications, disabling, and termination events that affect /etc/gshadow. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010104 | The Ubuntu operating system must generate audit records for all account creations, modifications, disabling, and termination events that affect /etc/opasswd. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010117 | The Ubuntu operating system must alert the ISSO and SA (at a minimum) in the event of an audit processing failure. | To configure the setting manually, see this DISA STIG article. |
| UBTU-20-010118 | The Ubuntu operating system must shut down by default upon audit failure (unless availability is an overriding concern). | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010122 | The Ubuntu operating system must be configured so that audit log files are not read or write-accessible by unauthorized users. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
|---|---|---|
| UBTU-20-010123 | The Ubuntu operating system must be configured to permit only authorized users ownership of the audit log files. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010124 | The Ubuntu operating system must permit only authorized groups ownership of the audit log files. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010128 | The Ubuntu operating system must be configured so that the audit log directory is not write-accessible by unauthorized users. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010133 | The Ubuntu operating system must be configured so that audit configuration files are not write-accessible by unauthorized users. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010134 | The Ubuntu operating system must permit only authorized accounts to own the audit configuration files. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010135 | The Ubuntu operating system must permit only authorized groups to own the audit configuration files. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010136 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the su command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010137 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the chfn command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010138 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the mount command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010139 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the umount command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010140 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the ssh-agent command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010141 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the ssh-keysign command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010142 | The Ubuntu operating system must generate audit records for any use of the setxattr, fsetxattr, lsetxattr, removexattr, fremovexattr, and lremovexattr system calls. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010148 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the chown, fchown, fchownat, and lchown system calls. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010152 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the chmod, fchmod, and fchmodat system calls. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010155 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the creat, open, openat, open_by_handle_at, truncate, and ftruncate system calls. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010161 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the sudo command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010162 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the sudoedit command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010163 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the chsh command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010164 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the newgrp command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010165 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the chcon command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
|---|---|---|
| UBTU-20-010166 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the apparmor_parser command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010167 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the setfacl command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010168 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the chacl command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010169 | The Ubuntu operating system must generate audit records for the use and modification of the tallylog file. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010170 | The Ubuntu operating system must generate audit records for the use and modification of faillog file. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010171 | The Ubuntu operating system must generate audit records for the use and modification of the lastlog file. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010172 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the passwd command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010173 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the unix_update command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010174 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the gpasswd command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010175 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the chage command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010176 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the usermod command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010177 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the crontab command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010178 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the pam_timestamp_check command. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010179 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the init_module and finit_module syscalls. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010181 | The Ubuntu operating system must generate audit records for successful/unsuccessful uses of the delete_module syscall. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010182 | The Ubuntu operating system must produce audit records and reports containing information to establish when, where, what type, the source, and the outcome for all DoD-defined auditable events and actions in near real time. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010198 | The Ubuntu operating system must initiate session audits at system start-up. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010199 | The Ubuntu operating system must configure audit tools with a mode of 0755 or less permissive. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010200 | The Ubuntu operating system must configure audit tools to be owned by root. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010201 | The Ubuntu operating system must configure the audit tools to be group-owned by root. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010205 | The Ubuntu operating system must use cryptographic mechanisms to protect the integrity of audit tools. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010211 | The Ubuntu operating system must prevent all software from executing at higher privilege levels than users executing the software and the audit system must be configured to audit the execution of privileged functions. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010215 | The Ubuntu operating system must allocate audit record storage capacity to store at least one weeks' worth of audit records, when audit records are not immediately sent to a central audit record storage facility. | To configure the setting manually, see this DISA STIG article. |
| UBTU-20-010216 | The Ubuntu operating system audit event multiplexor must be configured to off-load audit logs onto a different system or storage media from the system being audited. | Applicable if the infrastructure have such components. To configure the setting manually, see this DISA STIG article. |
| UBTU-20-010217 | The Ubuntu operating system must immediately notify the SA and ISSO (at a minimum) when allocated audit record storage volume reaches 75% of the repository maximum audit record storage capacity. | To configure the setting manually, see this DISA STIG article. |
| UBTU-20-010230 | The Ubuntu operating system must record time stamps for audit records that can be mapped to Coordinated Universal Time (UTC) or Greenwich Mean Time (GMT). | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010244 | The Ubuntu operating system must generate audit records for privileged activities, nonlocal maintenance, diagnostic sessions and other system-level access. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010267 | The Ubuntu operating system must generate audit records for any successful/unsuccessful use of unlink, unlinkat, rename, renameat, and rmdir system calls. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010277 | The Ubuntu operating system must generate audit records for the /var/log/wtmp file. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010278 | The Ubuntu operating system must generate audit records for the /var/run/utmp file. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010279 | The Ubuntu operating system must generate audit records for the /var/log/btmp file. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010296 | The Ubuntu operating system must generate audit records when successful/unsuccessful attempts to use modprobe command. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010297 | The Ubuntu operating system must generate audit records when successful/unsuccessful attempts to use the kmod command. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010298 | The Ubuntu operating system must generate audit records when successful/unsuccessful attempts to use the fdisk command. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010300 | The Ubuntu operating system must have a crontab script running weekly to offload audit events of standalone systems. | To configure the setting manually, see this DISA STIG article. |
| UBTU-20-010400 | The Ubuntu operating system must limit the number of concurrent sessions to ten for all accounts or account types. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010403 | The Ubuntu operating system must monitor remote access methods. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010404 | The Ubuntu operating system must encrypt all stored passwords with a FIPS 140-2 approved cryptographic hashing algorithm. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010405 | The Ubuntu operating system must not have the telnet package installed. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010406 | The Ubuntu operating system must not have the rsh-server package installed. | To configure the setting manually, see this DISA STIG article<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010407 | The Ubuntu operating system must be configured to prohibit or restrict the use of functions, ports, protocols, and services, as defined in the PPSM CAL and vulnerability assessments. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010408 | The Ubuntu operating system must prevent direct login into the root account. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010409 | The Ubuntu operating system must disable account identifiers (individuals, groups, roles, and devices) after 35 days of inactivity. | Not applicable. A hardened repository uses only a root account and a non-root account with reduced permissions. |
| UBTU-20-010410 | The Ubuntu operating system must automatically remove or disable emergency accounts after 72 hours. | Not applicable. A hardened repository uses only a root account and a non-root account with reduced permissions. |
| UBTU-20-010411 | The Ubuntu operating system must set a sticky bit on all public directories to prevent unauthorized and unintended information transferred through shared system resources. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010412 | The Ubuntu operating system must be configured to use TCP syncookies. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010413 | The Ubuntu operating system must disable kernel core dumps so that it can fail to a secure state if system initialization fails, shutdown fails or aborts fail. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010415 | The Ubuntu operating system must deploy Endpoint Security for Linux Threat Prevention (ENSLTP). | Not applicable. A hardened repository must not have any third party software installed. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010416 | The Ubuntu operating system must generate error messages that provide information necessary for corrective actions without revealing information that could be exploited by adversaries. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010417 | The Ubuntu operating system must configure the /var/log directory to be group-owned by syslog. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010418 | The Ubuntu operating system must configure the /var/log directory to be owned by root. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010419 | The Ubuntu operating system must configure the /var/log directory to have mode "0755" or less permissive. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010420 | The Ubuntu operating system must configure the /var/log/syslog file to be group-owned by adm. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010421 | The Ubuntu operating system must configure /var/log/syslog file to be owned by syslog. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010422 | The Ubuntu operating system must configure /var/log/syslog file with mode 0640 or less permissive. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010423 | The Ubuntu operating system must have directories that contain system commands set to a mode of 0755 or less permissive. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010424 | The Ubuntu operating system must have directories that contain system commands owned by root. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010425 | The Ubuntu operating system must have directories that contain system commands group-owned by root. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010426 | The Ubuntu operating system library files must have mode 0755 or less permissive. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010427 | The Ubuntu operating system library directories must have mode 0755 or less permissive. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010428 | The Ubuntu operating system library files must be owned by root. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010429 | The Ubuntu operating system library directories must be owned by root. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010430 | The Ubuntu operating system library files must be group-owned by root or a system account. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010431 | The Ubuntu operating system library directories must be group-owned by root. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010432 | The Ubuntu operating system must be configured to preserve log records from failure events. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010433 | The Ubuntu operating system must have an installed application firewall to control remote access methods. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010434 | The Ubuntu operating system must enable and run the uncomplicated firewall (ufw). | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010435 | The Ubuntu operating system must, for networked systems, compare internal information system clocks at least every 24 hours with a server which is synchronized to one of the redundant United States Naval Observatory (USNO) time servers, or a time server designated for the appropriate DoD network (NIPRNet/SIPRNet), or the Global Positioning System (GPS). | To configure the setting manually, see this DISA STIG article. Note that you must set the "-R" flag in the `/etc/defaults/chrony` file.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010436 | The Ubuntu operating system must synchronize internal information system clocks to the authoritative time source when the time difference is greater than one second. | Not applicable. A hardened repository uses additional parameter to be compliant with STIG UBTU-20-010435. |
| UBTU-20-010438 | The Ubuntu operating system's Advance Package Tool (APT) must be configured to prevent the installation of patches, service packs, device drivers, or Ubuntu operating system components without verification they have been digitally signed using a certificate that is recognized and approved by the organization. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010439 | The Ubuntu operating system must be configured to use AppArmor. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010440 | The Ubuntu operating system must allow the use of a temporary password for system logons with an immediate change to a permanent password. | Not applicable. A hardened repository uses only a root account and a non-root account with reduced permissions. |
| UBTU-20-010441 | The Ubuntu operating system must be configured such that Pluggable Authentication Module (PAM) prohibits the use of cached authentications after one day. | Not applicable. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010442 | The Ubuntu operating system must implement NIST FIPS-validated cryptography to protect classified information and for the following: to provision digital signatures, to generate cryptographic hashes, and to protect unclassified information requiring confidentiality and cryptographic protection in accordance with applicable federal laws, Executive Orders, directives, policies, regulations, and standards. | Not applicable. |
| UBTU-20-010443 | The Ubuntu operating system must only allow the use of DoD PKI-established certificate authorities for verification of the establishment of protected sessions. | Not applicable. A hardened repository does not use smart card authentication. |
| UBTU-20-010444 | Ubuntu operating system must implement cryptographic mechanisms to prevent unauthorized modification of all information at rest. | To protect data in backups, use Veeam Backup & Replication built-in encryption. For more information, see Data Encryption. |
| UBTU-20-010445 | Ubuntu operating system must implement cryptographic mechanisms to prevent unauthorized disclosure of all information at rest. | To protect data in backups, use Veeam Backup & Replication built-in encryption. For more information, see Data Encryption. |
| UBTU-20-010446 | The Ubuntu operating system must configure the uncomplicated firewall to rate-limit impacted network interfaces. | Not applicable. |
| UBTU-20-010447 | The Ubuntu operating system must implement non-executable data to protect its memory from unauthorized code execution. | To configure the setting manually, see this DISA STIG article. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010448 | The Ubuntu operating system must implement address space layout randomization to protect its memory from unauthorized code execution. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010449 | The Ubuntu operating system must be configured so that Advance Package Tool (APT) removes all software components after updated versions have been installed. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010450 | The Ubuntu operating system must use a file integrity tool to verify correct operation of all security functions. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010451 | The Ubuntu operating system must notify designated personnel if baseline configurations are changed in an unauthorized manner. The file integrity tool must notify the System Administrator when changes to the baseline configuration or anomalies in the operation of any security functions are discovered. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010453 | The Ubuntu operating system must display the date and time of the last successful account logon upon logon. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010454 | The Ubuntu operating system must have an application firewall enabled. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
| --- | --- | --- |
| UBTU-20-010456 | The Ubuntu operating system must have system commands set to a mode of 0755 or less permissive. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010457 | The Ubuntu operating system must have system commands owned by root or a system account. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010458 | The Ubuntu operating system must have system commands group-owned by root or a system account. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010459 | The Ubuntu operating system must disable the x86 [Ctrl+Alt+Delete] key combination if a graphical user interface is installed. | Not applicable. The operating system is installed without GUI. |
| UBTU-20-010460 | The Ubuntu operating system must disable the x86 [Ctrl+Alt+Delete] key combination. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010462 | The Ubuntu operating system must not have accounts configured with blank or null passwords. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |
| UBTU-20-010463 | The Ubuntu operating system must not allow accounts configured with blank or null passwords. | To configure the setting manually, see this DISA STIG article.<br>This setting is also included to the automatic configuration script provided by Veeam. |

| DISA STIG ID | Description | How to apply |
|---|---|---|
| UBTU-20-010461 | The Ubuntu operating system must disable automatic mounting of Universal Serial Bus (USB) mass storage driver. | To configure the setting manually, see this DISA STIG article.<br><br>This setting is also included to the automatic configuration script provided by Veeam. |

# Preparing Red Hat Enterprise Linux Server as Hardened Repository

This section includes security considerations for installing and configuring the Linux server that will be used as a hardened repository. Recommendations are based on Security Technical Implementation Guides (STIGs) created and maintained by the Defense Information Systems Agency (DISA) for Red Hat Enterprise Linux (RHEL) 9.4. For more information, see DISA STIGs Document Library.

## Installing RHEL Server

To install Red Hat Enterprise Linux (RHEL), download the server install image from the Red Hat Product Downloads page. For more information on the installer and options of the installation wizard, see the official Red Hat guide. Before you start the installation process, see Requirements and Limitations.

During installation process, consider the following Veeam recommendations:

1. Before you boot the installer, enable UEFI secure boot to prevent unsigned Linux kernel modules from being loaded.
2. At the welcome screen of the installation wizard, select the language for the installer and the default language for the installed system. For troubleshooting purposes, it is recommended to select the English language.
3. On the **Installation Summary** step of the installation wizard, add any additional languages, keyboard layouts, and set your time zone as needed.
4. Go to the **Network & Host Name** screen and create a bond. To do this, perform the following steps:
   a. Select the plus icon and choose **Bond** in the drop down menu.
   b. Specify a name for the connection and the interface.
   c. Press **Add**.
   d. Select **Ethernet** as your connection type and add the device using the drop down menu.
   e. Set the bonding mode:
      - *Round Robin* (if you use EtherChannel without LACP).
      - *802.3ad* (if you use EtherChannel with LACP).
      - *Active-backup* (for other configurations).
   f. Repeat this process for the second network card.

5. Select **IPv4 Settings** or **IPv6 Settings**, depending on your connection. Configure your IP address and DNS. A static IP address should be used if possible. Click **Save**.
6. Specify a hostname.
7. Go to the **Connect to Red Hat** screen and log in using your Red Hat account details.
8. Go to the **Installation Destination** screen. Select all the disks on your server (excluding any external storage devices such as SD cards), set your **Storage Configuration** to **Custom** and click **Done**.
9. At the **Manual Partitioning** screen, add required partitions and mount points using the **Standard Partition** device type. For a disk with 100 GiB capacity, the following partitioning scheme is recommended:

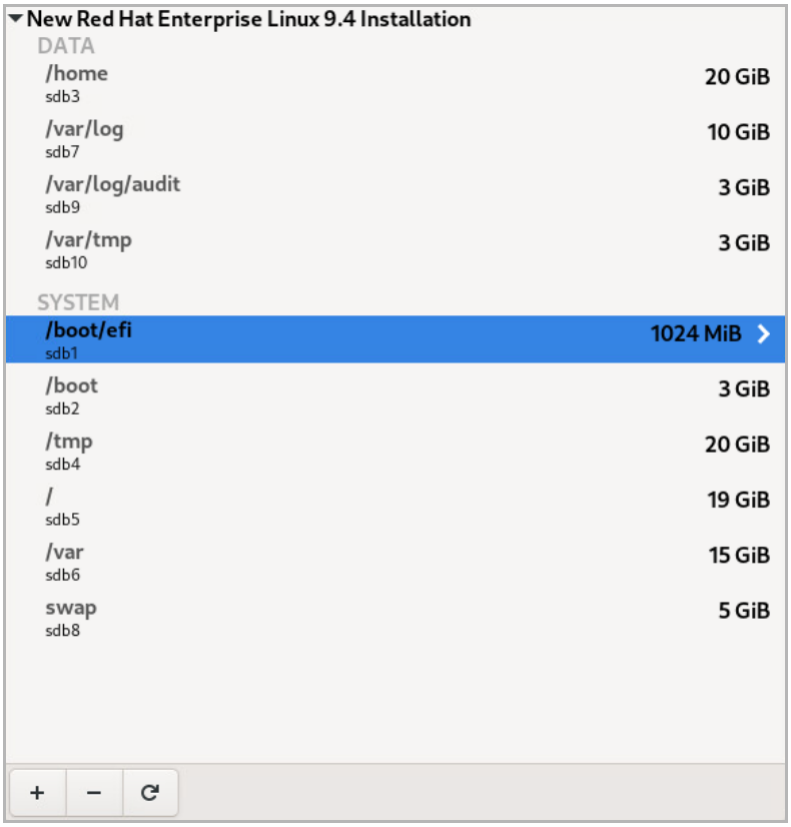


> **NOTE**
>
> If you are using a larger disk, adjust the values proportionally.

10. Add the disk for backup data using the `/mnt/backup` mount point. Leave the **Desired Capacity** field empty to use the whole disk. If you have multiple disks, use your preferred naming convention.
11. Click **Done** and confirm the formatting. Note that everything on the disks will be deleted during the installation process.
12. Go to the **Software Selection** screen and select **Minimal Install**.
13. Go to the **Security Profile** screen, select the DISA STIG profile, and click **Done**.
14. Go to the **User Creation** screen and create a user account with administrator privileges. When setting a password, ensure that it meets the following DISA STIG requirements:
    - 15 characters minimum.
    - 1 upper case character.
    - 1 numeric character.

- 1 special character.
- No more than 3 characters of the same class in a row. For example, more than 3 lowercase or 3 numerical characters in sequence.
- Minimum password lifetime - 24 hours.

If it does not, you will be prompted to change the password after your initial login.

15. Click **Begin Installation**. After the installation finishes, remove the installation media and reboot the system.

## Configuring RHEL Server

To configure system settings, do the following:

1. Connect to the server using the user account you created.
2. For the separate directory that you created for the backup data, allow access only for the user account you created during the installation. Use the following commands:
   - To assign the directory's owner:

     ```
     chown -R owner:group <dir_path>
     ```

     Both `owner` and `group` must be the user account you created during the installation.
   - To allow access to the directory only for its owner and the root account:

     ```
     chmod 700 <dir_path>
     ```
3. Enable automatic security updates:

   ```
   sudo dnf install dnf-automatic -y
   ```

   After you enable automatic updates, change the following parameters in the `/etc/dnf/automatic.conf` file:
   - `upgrade_type` to `security`
   - `apply_updates = yes`
4. Enable download timers for the security updates:

   ```
   sudo systemctl enable dnf-automatic-download.timer
   sudo systemctl start dnf-automatic-download.timer
   sudo systemctl enable dnf-automatic-install.timer
   sudo systemctl start dnf-automatic-install.timer
   ```
5. To reduce the risk of server timing and man-in-the-middle attacks, change the `/etc/sysconfig/chronyd` configuration:

```
# Command-line options for chronyd
OPTIONS="-R -F 2"
```

Restart the service:

```
sudo systemctl restart chronyd
```

6. By default, the user account you created during the installation is a member of the `wheel` group and has sufficient privileges to deploy and install the required Veeam Backup & Replication components. For more information, see Adding Hardened Repositories. After the repository is added, you must remove the user account from the `wheel` group and disable SSH access. To do this, perform the following steps:

   a. Allow the user account to reboot and shutdown the operating system:

   ```
   sudo bash -c "echo 'user1 ALL = (root) NOEXEC: /usr/sbin/reboot' >> /et
   c/sudoers"
   sudo bash -c "echo 'user1 ALL = (root) NOEXEC: /usr/sbin/shutdown' >> /
   etc/sudoers"
   ```

   b. Disable SSH access to the server and remove the user account from the `wheel` group:

   ```
   sudo systemctl disable sshd
   sudo systemctl stop sshd
   sudo gpasswd -d user1 wheel
   sudo reboot
   ```

   Note that this user account will lose sudo permissions. if you need to execute commands as a privileged user, you must boot the operating system into the single user mode.

# Preparing Hardened Repository Using Veeam ISO

This section includes security considerations for installing, configuring and maintaining a hardened repository using the Veeam Hardened Repository ISO. Recommendations are based on Security Technical Implementation Guides (STIGs) created and maintained by the Defense Information Systems Agency (DISA) for Rocky Linux. For more information, see DISA STIGs Document Library.

# In This Section

- Requirements and Limitations
- Installing from Veeam Hardened Repository ISO
- Configuring Server as Hardened Repository
- Repairing Veeam Hardened Repository ISO
- Using Live System

## Requirements and Limitations

> **IMPORTANT**
>
> Veeam offers experimental support for the Veeam Hardened Repository ISO. For more information, see this Knowledge Base article.
>
> To be eligible for support, you must use an unmodified version of the Veeam Hardened Repository ISO on a machine that meets all the system requirements.

In addition to the general backup repository and hardened repository system requirements, the Veeam Hardened Repository ISO has specific requirements and limitations.

# Requirements

The following requirements must be met:

- The Veeam Backup & Replication version must be 12.2 or later.
- All hardware must be on the Red Hat compatibility list or CIQ certified hardware list.
- UEFI secure boot must be enabled.
- Third-party security software must not be installed on the server.
- Any USB devices you intend to use must be connected before you start the installation. After the installation finishes, any new USB devices will be blocked by USBGuard. For more information about other limitations related to USB devices, see Limitations.
- Only hardware RAID controllers must be used. Software RAID, Intel VMD VROC, and FakeRAID controllers are not supported.
- RAID controllers must have write-back cache enabled.
- The server must have at least two storage volumes:
  - A separate volume for the operating system (minimum 100GB).
  - At least one additional volume for data. All additional data volumes must be larger than the operating system volume.
  - All data volumes must have the same sector size.

  > **TIP**
  >
  > It is strongly recommended that you use at least a dual parity RAID configuration.

- Internal or direct attached storage volumes must be used.
- For installations that use a remote console, the network connection must be at least 4Mbit/s.

  > **NOTE**
  >
  > With a 4Mbit/s connection, the graphical installer will take approximately 10 minutes to load.

- In addition to the standard set of ports that must be opened for a backup repository, an outgoing direct or HTTP proxy connection to `repository.veeam.com` on port 443 is also required for security and operating system updates. Without this connection, the GNU Privacy Guard (GPG) keys will eventually expire. Once these keys have expired, no further updates will be possible and a full re-installation of the operating system will be required.
- To prevent unauthorized access or deletion of the hardened repository, the BMC (base management controller) port on your server hardware must be secured using appropriate measures such as firewalls and strong passwords.

# Limitations

The solution has the following limitations:

- Multipathing for storage volumes is not supported.
- A USB storage volume cannot be used either as the operating system volume or a data volume.
- Updates to the server occur automatically on Mondays at 8:00 AM of the configured time zone. However, the server will not automatically reboot. Additionally, there are no notifications about required reboots. To make sure that all updates are installed and applied properly, reboot the hardened repository regularly.
- iSCSI or Fibre Channel LUNs provisioned to the server are not supported.
- Wireless network connections are not supported.
- The **Use Text Mode** option is not supported.
- If the bandwidth between the server and installation media is less than 4 MBit/s, the graphical installer will not start. To load the graphical installer to RAM, add the `rd.live.ram=1` option to GRUB.
- Port 546 cannot be closed on hardened repositories created using the Veeam Hardened Repository ISO. This port is required for connections to DHCPv6 servers.
- Hardened repositories created with the Veeam Hardened Repository ISO have an incoming ICMP ping limit of 5 packets per second. Pings above this limit will fail and appear as packet loss.

## Installing from Veeam Hardened Repository ISO

The Veeam Hardened Repository ISO is delivered as bootable ISO which can be downloaded here. You can use a remote console or create a bootable USB stick from the ISO. Before you start, see Requirements and Limitations.

To create a bootable USB stick, it is recommended that you use Rufus with the default settings. Note that you need to select **Write in DD Image mode** option.

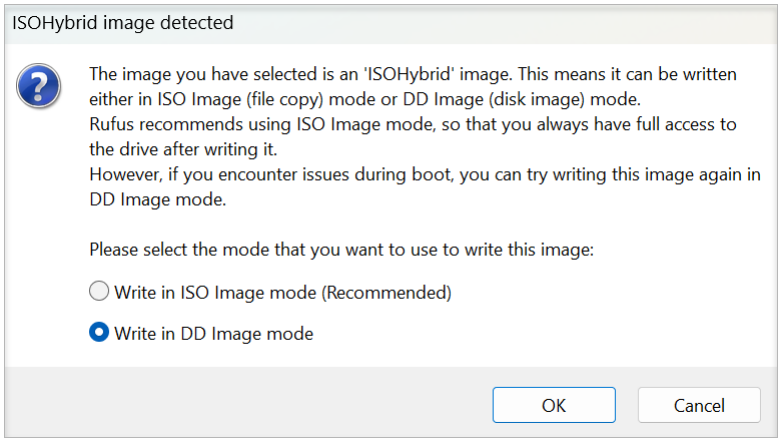

To install a hardened repository, on the boot screen, select **Install Hardened Repository**. Then, do the following:

1. On the **Keyboard** screen, add your preferred keyboard layouts.
2. On the **Time and Date** screen, set your time zone and specify any NTP or NTS servers you want to use.
3. On the **Network & Host Name** screen, specify a hostname or leave it empty. Then, configure your network interface.

   If you use a single network interface, do the following:

   a. Select the network interface and click **Configure**.

   b. In IPv4 or IPv6 settings, specify your IP address and DNS. A static IP address should be used if possible.

   c. Click **Save**.

   If you use multiple network interfaces without failover, do the following:

   a. Select a network interface. Enable it with the toggle if necessary. Click **Configure**.

   b. In IPv4 or IPv6 settings, specify your IP address and DNS. A static IP address should be used if possible.

   c. In the **General** tab, select the **Connect automatically with priority** check box.

   d. Click **Save**.

   Repeat this process for any additional network cards.

   If you use multiple network interfaces with failover, do the following:

   a. Select a network interface. Enable it with the toggle if necessary. Click **Configure**.

   b. In the **General** tab, select the **Connect automatically with priority** check box. Repeat this step for each network interface.

   c. Click the plus icon. In the drop-down menu, select **Bond**.

   d. Specify a name for the connection and the interface.

   e. Click **Add**.

   f. In the drop-down menu, select *Ethernet* as the connection type. Then, select the network interface and click **Save**. Repeat this step for each network interface.

   g. Select the bonding mode: *Round Robin* (EtherChannel without LACP), *802.3ad* (EtherChannel with LACP) or *Active-backup* (for other configurations).

   h. Click **Save**.

   i. Select the bond and click **Configure**.

   j. In IPv4 or IPv6 settings, specify your IP address and DNS. A static IP address should be used if possible.

   k. Click **Save**.
4. Click **Begin Installation**. After the process finishes, remove the media and reboot the system.

## Configuring Server as Hardened Repository

For post-installation, follow these steps:

1. Log in to the server using the `vhradmin` account. By default, the password is `vhradmin`.
2. Select a new password for the account. It must meet the following DISA STIG requirements:
   a. 15 characters minimum.
   b. 1 upper case character.
   c. 1 numeric character.
   d. 1 special character.
   e. No more than 3 characters of the same class in a row. For example, more than 3 lowercase or 3 numerical characters in sequence.
   f. Minimum password lifetime - 24 hours.
3. Accept the license agreement.
4. In the Veeam Hardened Repository Configurator, configure the following settings as required:
   - **Network settings** — To configure IPv4 addresses, DHCP, and DNS for network interfaces, select **Standard configuration**. If DHCP is not enabled, select **Advanced configuration** to configure your IP address using `nmtui`.
   - **Proxy settings** — Specify an HTTP or HTTPS proxy. Note that any proxy you add will be saved in the `/etc/environment` file in the `http://ip_address:port` format.

     > **NOTE**
     >
     > If you use self-signed HTTPS proxies or HTTPS proxies that use certificates signed by an internal CA, HTTPS traffic will be forwarded through an HTTP tunnel.

   - **Time settings** — Add an NTP server. When adding the NTP server, consider the following:
     - `chronyd` is used as the NTP client.
     - NTP servers over DHCP are allowed. This setting cannot be disabled.
     - NTP servers are added with the `iburst` parameter. No additional options can be specified.
5. Select the **Start SSH** option. This will generate credentials for the `veeamsvc` user that will be used to add the hardened repository to Veeam Backup & Replication. To generate new credentials, restart the SSH service.
6. Add the server as a hardened repository. For more information, see Adding Hardened Repositories.
7. After you add the hardened repository, select **Stop SSH** in the configurator.

## Repairing Veeam Hardened Repository ISO

You can repair a hardened repository installed from the Veeam Hardened Repository ISO. Any data stored on the data volumes will be preserved. To do this, a remote console or bootable USB stick is required. For more information, see Installing from Veeam Hardened Repository ISO.

Before you repair a hardened repository, see Requirements and Limitations. Also, consider the following:

- You cannot repair repositories that were not created using the Veeam Hardened Repository ISO.
- The operating system volume must be smaller than the additional data volumes. Otherwise, data volumes will be formatted with data loss.

  If all volumes are the same size, the installation will not start.
- During the repair process, only the operating system volume is formatted.
- When you repair a repository, do not change the server IP address or DNS name.
- The repair process cannot be used for migration purposes.

To repair a hardened repository, on the boot screen, select **Repair Hardened Repository**. Then, do the following:

1. On the **Keyboard** screen, add your preferred keyboard layouts.
2. On the **Time and Date** screen, set your time zone and specify any NTP or NTS servers you want to use.
3. On the **Network & Host Name** screen, specify a hostname or leave it empty. Then, configure your network interface.

   If you use a single network interface, do the following:

   a. Select the network interface and click **Configure**.

   b. In IPv4 or IPv6 settings, specify your IP address and DNS. A static IP address should be used if possible.

   c. Click **Save**.

   If you use multiple network interfaces without failover, do the following:

   a. Select a network interface. Enable it with the toggle if necessary. Click **Configure**.

   b. In IPv4 or IPv6 settings, specify your IP address and DNS. A static IP address should be used if possible.

   c. In the **General** tab, select the **Connect automatically with priority** check box.

   d. Click **Save**.

   Repeat this process for any additional network cards.

   If you use multiple network interfaces with failover, do the following:

   a. Select a network interface. Enable it with the toggle if necessary. Click **Configure**.

   b. In the **General** tab, select the **Connect automatically with priority** check box. Repeat this step for each network interface.

   c. Click the plus icon. In the drop-down menu, select **Bond**.

   d. Specify a name for the connection and the interface.

   e. Click **Add**.

   f. In the drop-down menu, select *Ethernet* as the connection type. Then, select the network interface and click **Save**. Repeat this step for each network interface.

   g. Select the bonding mode: *Round Robin* (EtherChannel without LACP), *802.3ad* (EtherChannel with LACP) or *Active-backup* (for other configurations).

   h. Click **Save**.

i. Select the bond and click **Configure**.

j. In IPv4 or IPv6 settings, specify your IP address and DNS. A static IP address should be used if possible.

k. Click **Save**.

4. Click **Begin Installation**. After the process finishes, remove the media and reboot the system.

## Using Live System

You can use Live System for troubleshooting and performance testing. To perform procedures, it is recommended to have sufficient Linux experience.

# Accessing Live System

To access Live System, do the following:

1. Boot the ISO.
2. Select **Live System**.
3. Log in using the `vhradmin` account with the password `vhradmin`. The first time you access Live System, you will be asked to change your password. There are no complexity or length requirements.

> **TIP**
>
> Live System uses the US English keyboard layout.

If you want to access Live System through an SSH connection, perform the following steps:

1. In the Live System mode, check that the server has an IP address. If you do not have DHCP enabled, configure your IP address using `nmtui`.
2. Run the following command:

```
sudo systemctl start sshd
```

3. Connect to the server using your SSH client.

> **NOTE**
>
> When you reboot or shut down Live System, SSH will be automatically disabled.

# Scripts

Live System includes the following predefined scripts:

- `hardware_collection.sh` — Collects information on your system hardware and writes it to the `/tmp/hardware.txt` file.
- `mount_datavol.sh` — Mounts the data volume to `/mnt/veeam-repository01`.
- `mount_systemvol.sh` — Mounts the operating system volume to `/mnt/`.

# Tools

Live System has the following tools installed:

- `fio` — Used to test the performance of your data volumes.
- `iperf3` — Used to test your network performance.
- `tcpdump` — Used to analyze your network traffic.

# Adding Hardened Repositories

This section describes how to add a hardened repository as a backup repository.

Before you begin, check requirements and limitations and prepare a Linux server. Then use the **New Backup Repository** wizard to add the hardened repository:

1. Launch the New Backup Repository wizard.
2. Specify the hardened repository name and description.
3. Specify a Linux server.
4. Configure hardened repository settings.
5. Specify mount server settings.
6. Review and apply settings.

If you upgrade to Veeam Backup & Replication 12 and have Linux backup repositories with enabled **Make recent backups immutable for** check box, see Upgrading or Switching from Linux Repository to Hardened Repository.

## Step 1. Launch New Backup Repository Wizard

To launch the **New Backup Repository** wizard, perform the following steps:

1. Open the **Backup Infrastructure** view. In the inventory pane, right-click the **Backup Repositories** node and select **Add Backup Repository**. Alternatively, you can click **Add Repository** on the ribbon.
2. In the **Add Backup Repository** window, select **Direct attached storage** > **Linux (Hardened Repository)**.

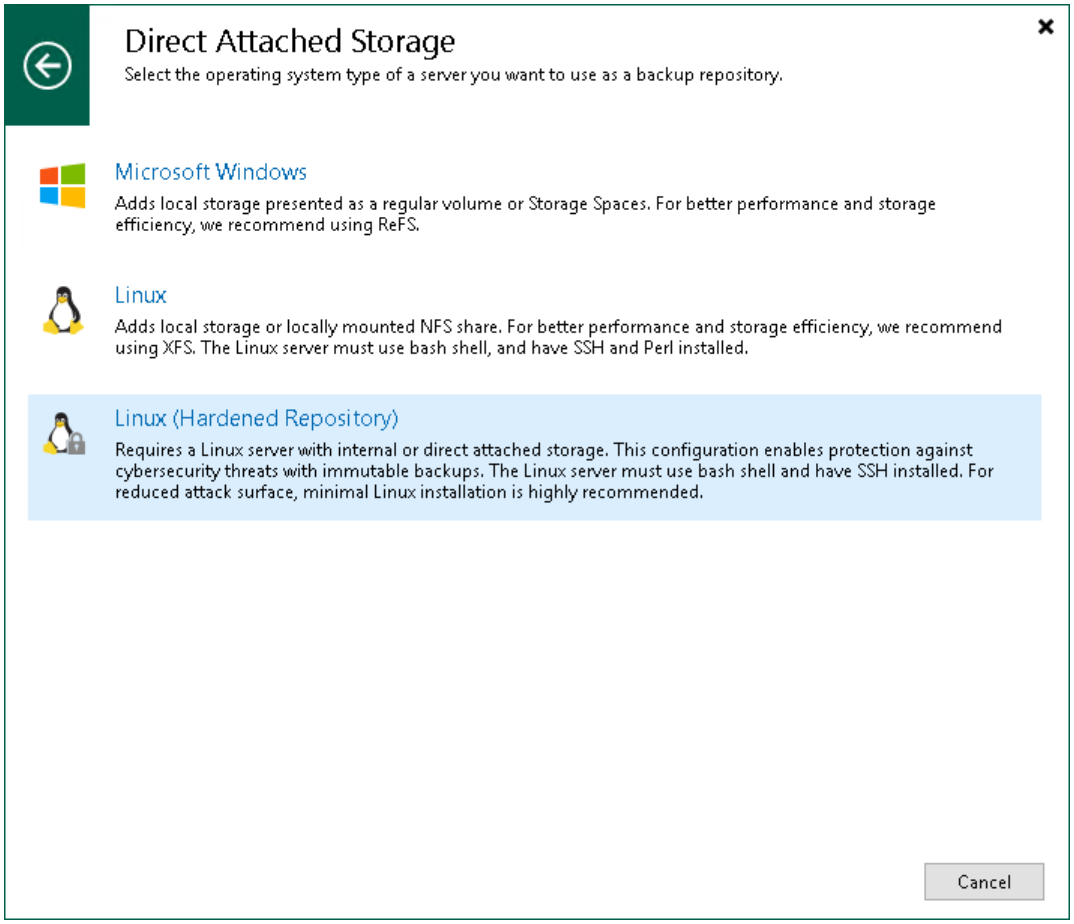

## Step 2. Specify Hardened Repository Name and Description

At the **Name** step of the wizard, specify a name and description for the repository. The default description contains information about the user who added the repository, date and time when the repository was added.

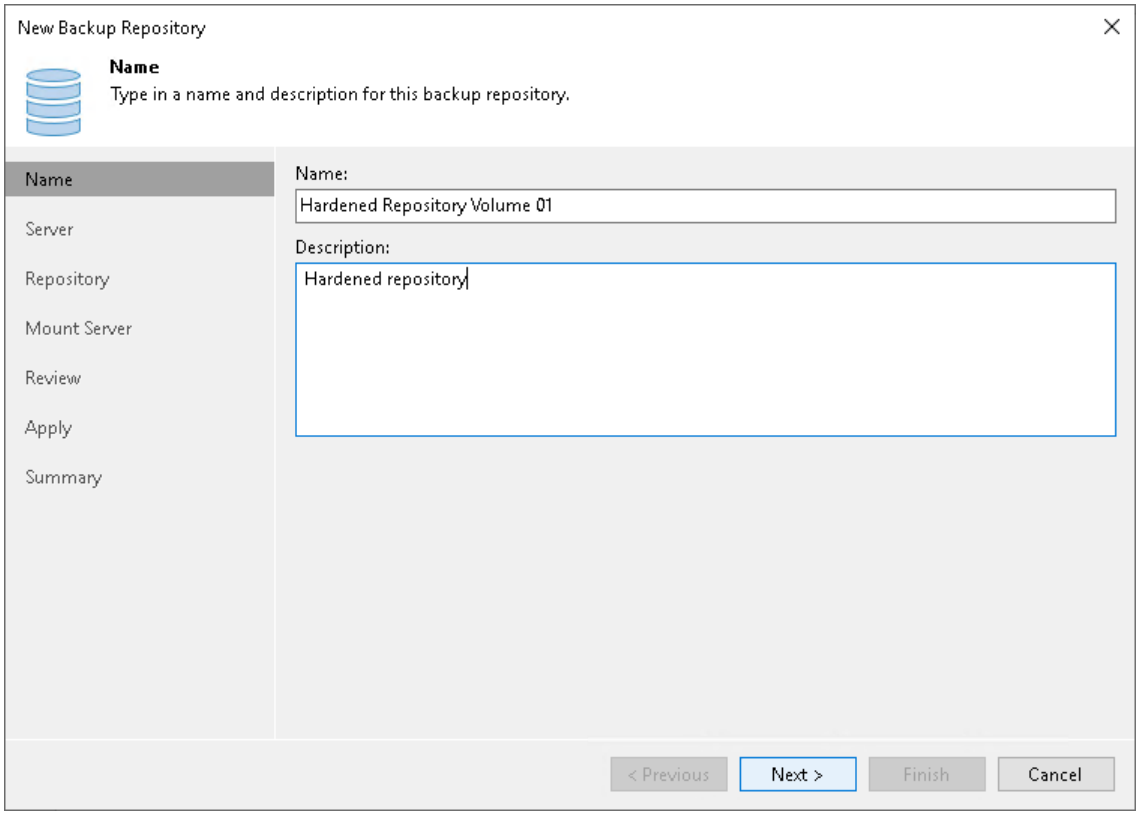

## Step 3. Specify Linux Server

From the **Repository server** list, select a Linux server that you want to use as a hardened repository. Click **Populate** to see a list of disks connected to the server, their capacity and free space.

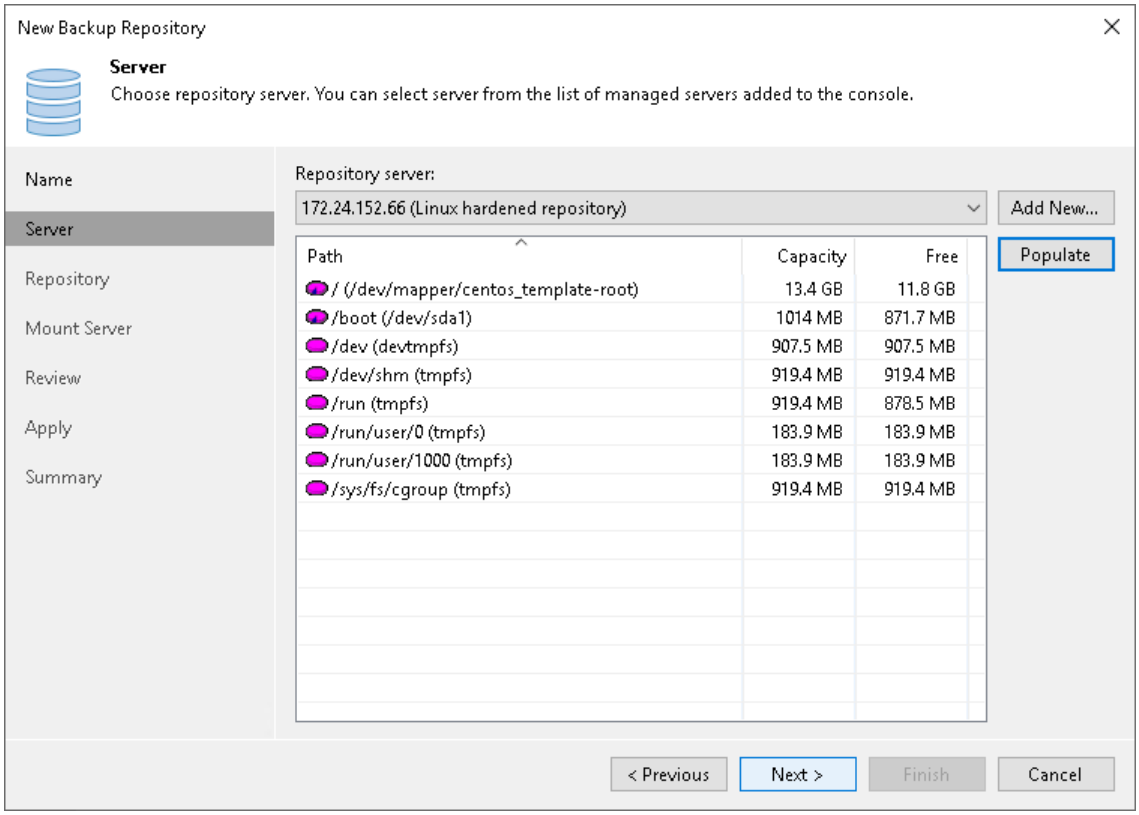


The **Repository server** list contains only those servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, click **Add New** on the right and follow the wizard:

1. At the **Name** step of the wizard, specify a full DNS name or IP address and description of the Linux server. The default description contains information about the user who added the server, date and time when the server was added.
2. At the **SSH Connection** step of the wizard, specify single-use credentials to connect to the Linux server and deploy Veeam Data Mover. Note that Veeam Backup & Replication does not store these credentials in the configuration database.

> **NOTE**
>
> The user account you specified must be a non-root account. Also, it must have the `home` directory created on the Linux server.

The **Elevate account privileges automatically** check box is used by default. If you did not add the user account to the `sudoers` file, select the **Use "su" if "sudo" fails** check box and enter the password for the root account. For more information on these check boxes, see Linux Accounts (User Name and Password).

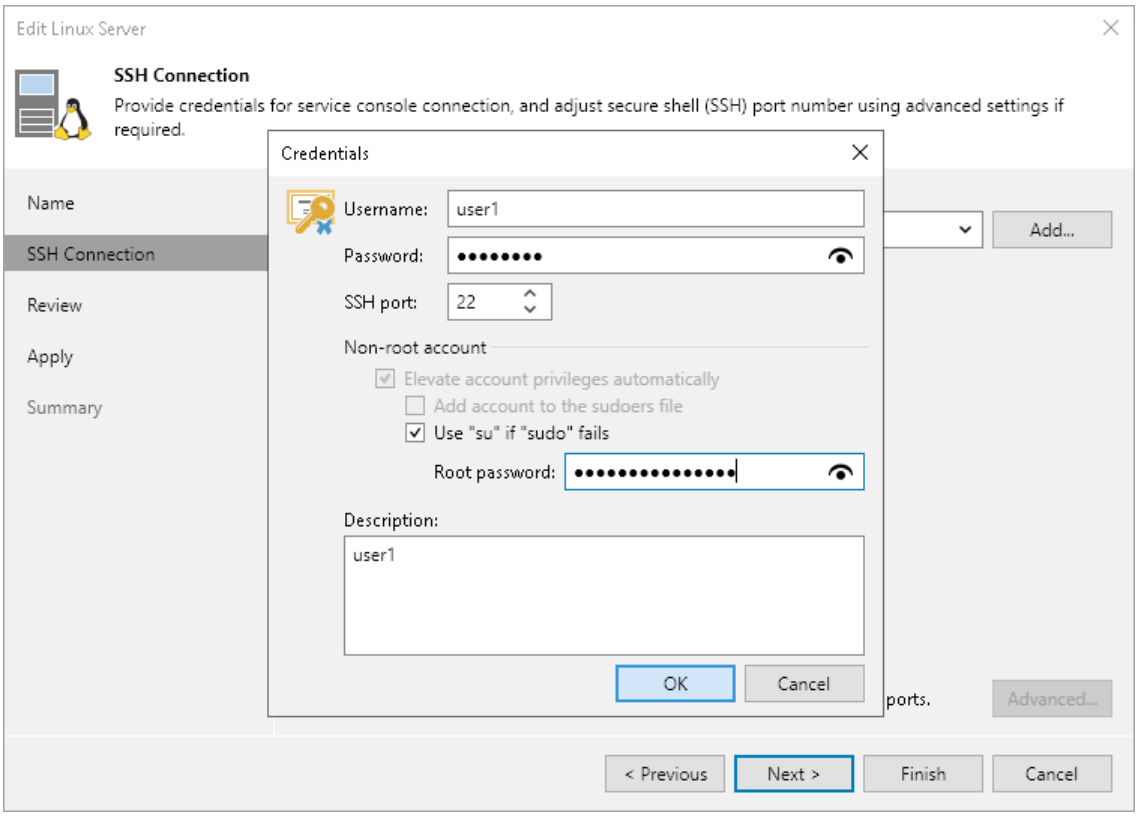


> **IMPORTANT**
>
> If you added the user account to the `sudoers` file, you do not need to select the **Use "su" if "sudo" fails** check box and specify the root password. But after the server is added, you must remove the user account from the file.

If you want to change default SSH settings, click **Advanced**. For more details, see Step 3. Specify Credentials and SSH Settings in the **New Linux Server** wizard.

3. At the **Review** step of the wizard, review what Veeam Backup & Replication components are already installed on the server and what components will be installed. Click **Apply** to add the Linux server to the backup infrastructure.
4. At the **Summary** step of the wizard, review details of the Linux server and click **Finish** to exit the wizard.

## Step 4. Configure Hardened Repository Settings

At the **Repository** step of the wizard, specify path and repository settings:

1. In the **Location** section, specify a path to the directory that you created to store immutable backups when preparing Linux server. Click **Populate** to check capacity and available free space in the selected location.
2. Select the **Use fast cloning on XFS** volumes check box to enable copy-on-write functionality. For more information, see Fast Clone.
3. Specify the immutability period.
4. Specify load control settings to limit the number of concurrent tasks and prevent possible timeouts of storage I/O operations. For more details, see Configure Backup Repository Settings.
5. If you want to configure additional settings for the repository, click **Advanced**. For more details, see Configure Backup Repository Settings.

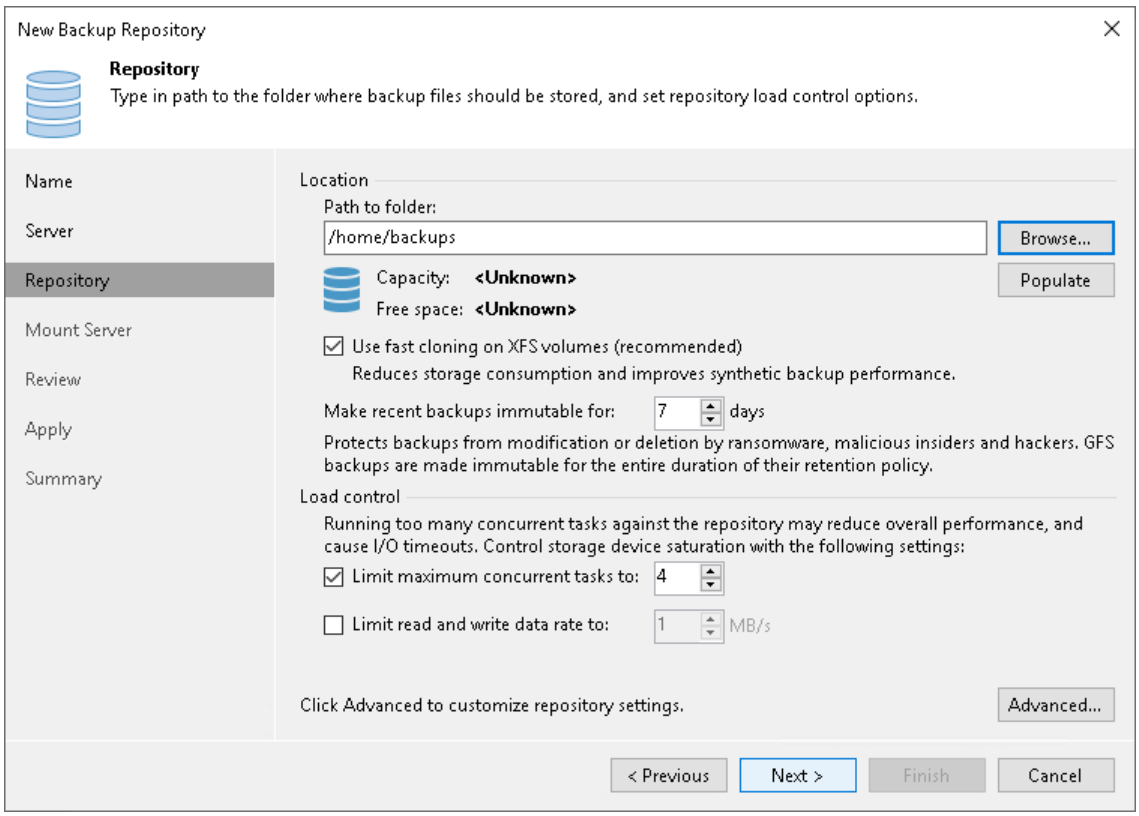

# Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for file-level and application items restore. For more details, see Specify Mount Server Settings.

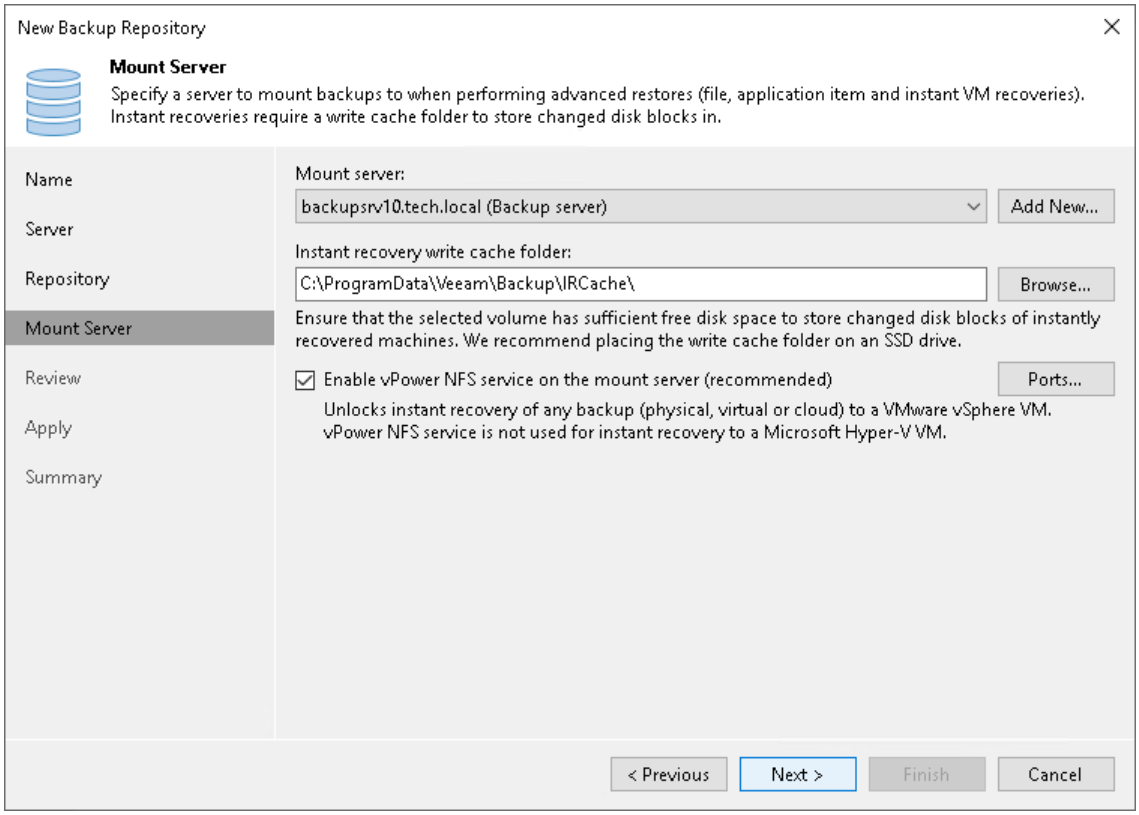

## Step 6. Review and Apply Settings

At the **Review** step of the wizard, review the hardened repository settings and list of components that will be installed.

If the backup repository contains backups that were previously created by Veeam Backup & Replication, select the **Search the repository for existing backups and import them automatically** check box.
Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Disk (Imported)** node.

If the backup repository contains guest file system index files that were previously created by Veeam Backup & Replication, select the **Import guest file system index** check box. Index files will be imported with backup files, and you will be able to search for guest OS files inside imported backups.

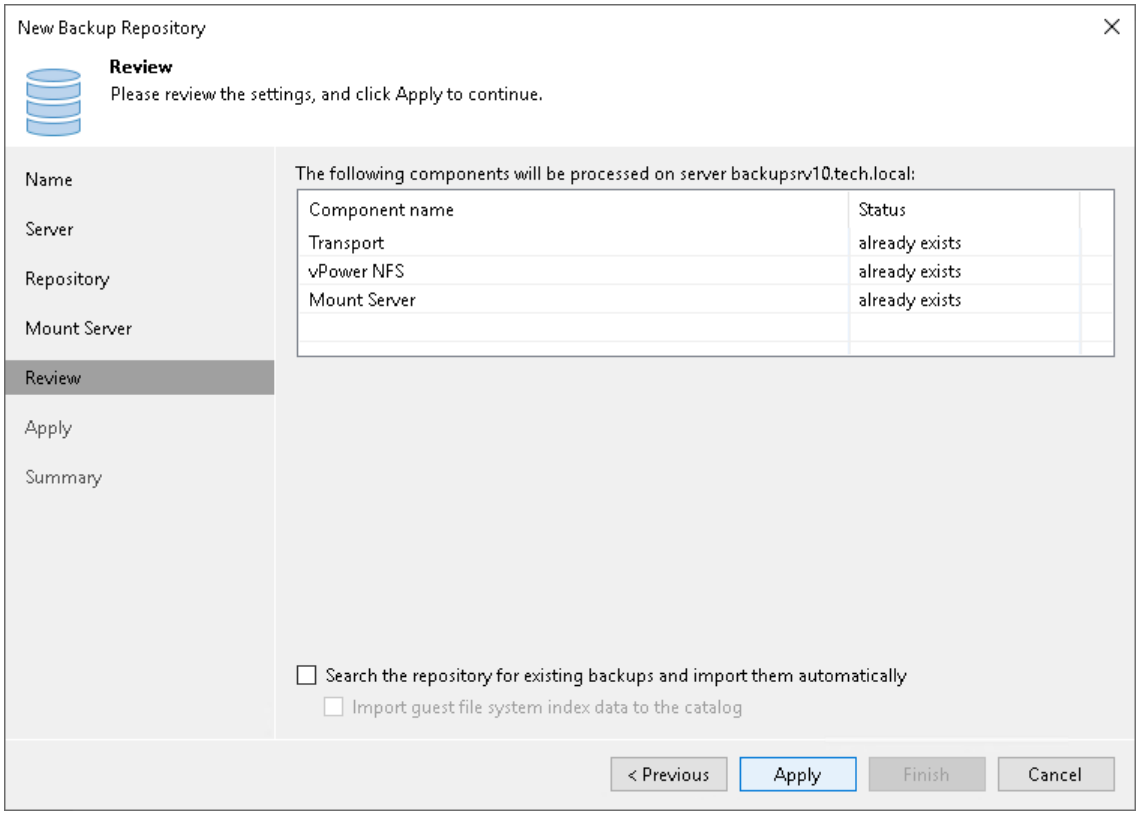


Click **Apply** and wait for Veeam Backup & Replication to install and configure all required components. At the **Summary** step of the wizard, review details of the added hardened repository and click **Finish** to exit the wizard.

To maximize the repository security and protect your data from different attacks, after the deployment of Veeam Data Mover, change file permissions for authentication certificates, so that only the user account you specified to connect to the Linux server and the root account can read the certificate files. Use the following commands:

```
chown owner:group /opt/veeam/transport/certs
chmod 700 /opt/veeam/transport/certs
```

Note that both `owner` and `group` must be the user account you specified to connect to the Linux server. You can also use `chmod 770` to add same permissions to the group.

IMPORTANT

SSH connection is necessary only for the deployment of Veeam Data Mover. For security purposes, after you add the hardened repository, disable SSH connection for the user account you use to connect to the Linux server. If you can work with the server from the console, disable SSH connection for the server itself.

# Managing Hardened Repository

This section describes the process of upgrading a Linux repository to a hardened repository.

# In This Section

- Updating Hardened Repository Components
- Upgrading or Switching from Linux Repository to Hardened Repository
- Upgrading Performance Extent to Hardened Repository

## Updating Hardened Repository Components

If you use Veeam Backup & Replication 12 prior to cumulative patch P20230718 and want to update hardened repository components or remove a hardened repository from the backup infrastructure, consider the following:

- Make sure that the SSH connection is enabled.
- You also need to specify single-use credentials you use to connect to the Linux server and deploy Veeam Data Mover.
- The Linux server must not have Veeam Agent for Linux installed.

IMPORTANT

The SSH connection is also required for updating hardened repository components in the following situations:

- If you upgrade from Veeam Backup & Replication 11a or earlier versions to Veeam Backup & Replication 12 including versions with installed cumulative patches.
- If you upgrade from any previous Veeam Backup & Replication version to Veeam Backup & Replication 12.1 (build 12.1.0.2131).

If you need to specify single-use credentials for multiple hardened repositories, you can use the **Components Update** window or the Set-VBRLinux cmdlet.

If you update Veeam Backup & Replication 12 to any cumulative patch, the SSH connection is not required for updating hardened repository components. For more information about cumulative patches, see this Veeam KB article.

# Upgrading Linux Repository to Hardened Repository

If you upgrade to Veeam Backup & Replication 12, all Linux backup repositories with enabled **Make recent backups immutable for** check box will be automatically converted to hardened repositories. Note that if these repositories were added with persistent credentials, they will be automatically converted to single-use credentials.

> **NOTE**
>
> If a Linux backup repository with enabled **Make recent backups immutable for** check box was added with root credentials, the upgrade process will be interrupted. To continue it, you must specify a user account with non-root permissions for this repository. For more information, see this Veeam KB article.

# Switching from Linux Repository to Hardened Repository

If you have a standalone Linux server added to the backup infrastructure as a backup repository and want to use it as a hardened repository, perform the following steps:

1. On the Linux server, change permissions for the directory where backups are stored. Both owner and group must be the account with non-root permissions you use to connect to the Linux server.

   ```
   chown -R owner:group <dir_path>
   ```

2. Go to the Veeam Backup & Replication console and perform the following steps:

   a. Open the **Backup Infrastructure** view and select **Managed Servers** in the navigation pane. Right-click the Linux server, select **Properties**, go to the **SSH Connection** step of the wizard, and specify the account with non-root permissions you use to connect to the Linux server. Note that these credentials must be single-use.

   b. Disable all backup jobs that use this Linux repository. For more information, see Disabling and Deleting Jobs.

   > **NOTE**
   >
   > If you have Veeam Agent backup jobs managed by Veeam Agent, you need to delete these jobs and configure them once again after you switch to the hardened repository.

   c. Remove backup files stored on the Linux repository from the backup configuration. To do this, go to the **Backups** > **Disk** node and select the backup file. Hold [Ctrl] and right-click the file. Then click **Remove from configuration**.

   > **IMPORTANT**
   >
   > Removing backups from configuration is designed for experienced users only. It is strongly recommended to create encrypted configuration backup before performing this operation.

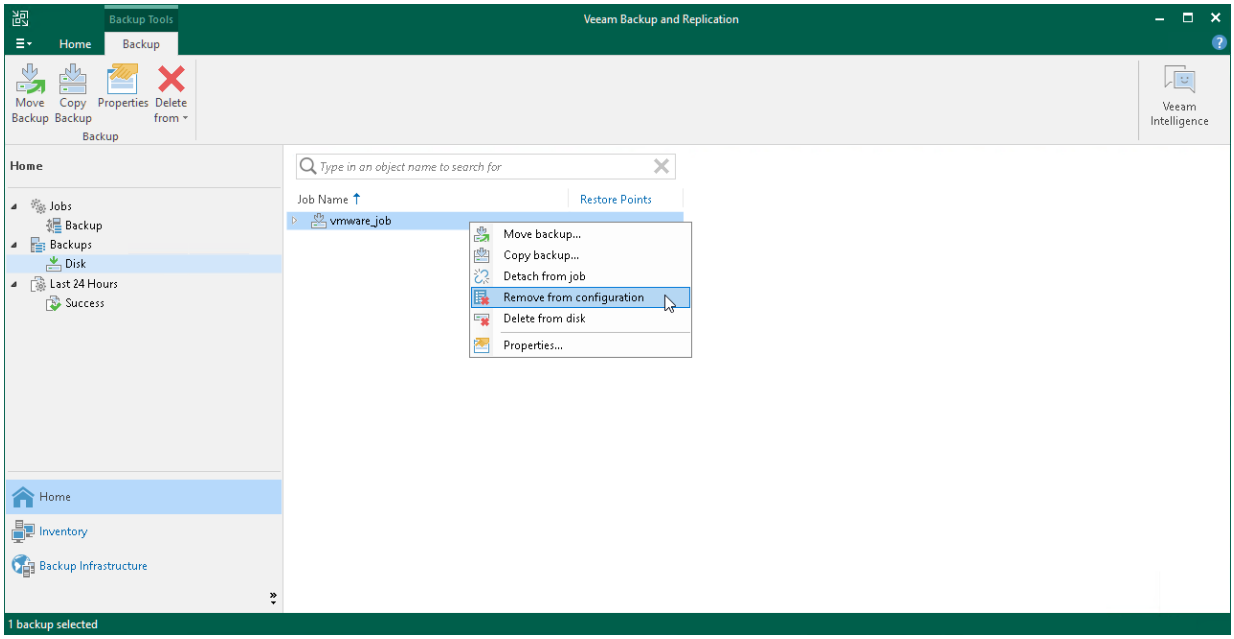


d. Add the same Linux server to the backup infrastructure as a hardened repository using the **New Backup Repository** wizard. At the **Review** step of the wizard, select the **Search the repository for existing backups and import them automatically** check box to detect existing backup files and display them in the Veeam Backup & Replication console under the **Imported** > **Backups** node. For more information, see Adding Hardened Repositories.

> **NOTE**
>
> If you work with encrypted backups, you need to decrypt them after importing. For more information, see How Data Decryption Works.

e. Go to the **Jobs** node and edit the job associated with the Linux repository. At the **Storage** step of the wizard, select the hardened repository from the **Backup repository** list.

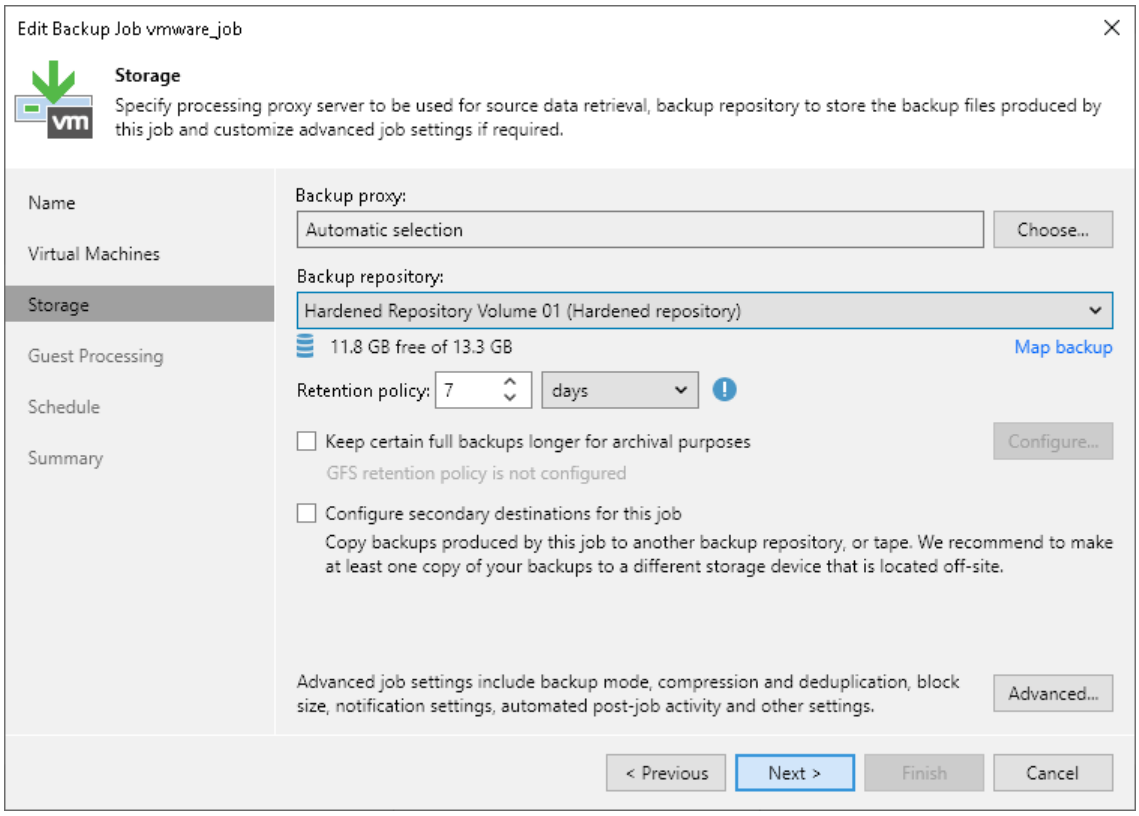


Click **Map backup** and specify imported backups from the previous step.

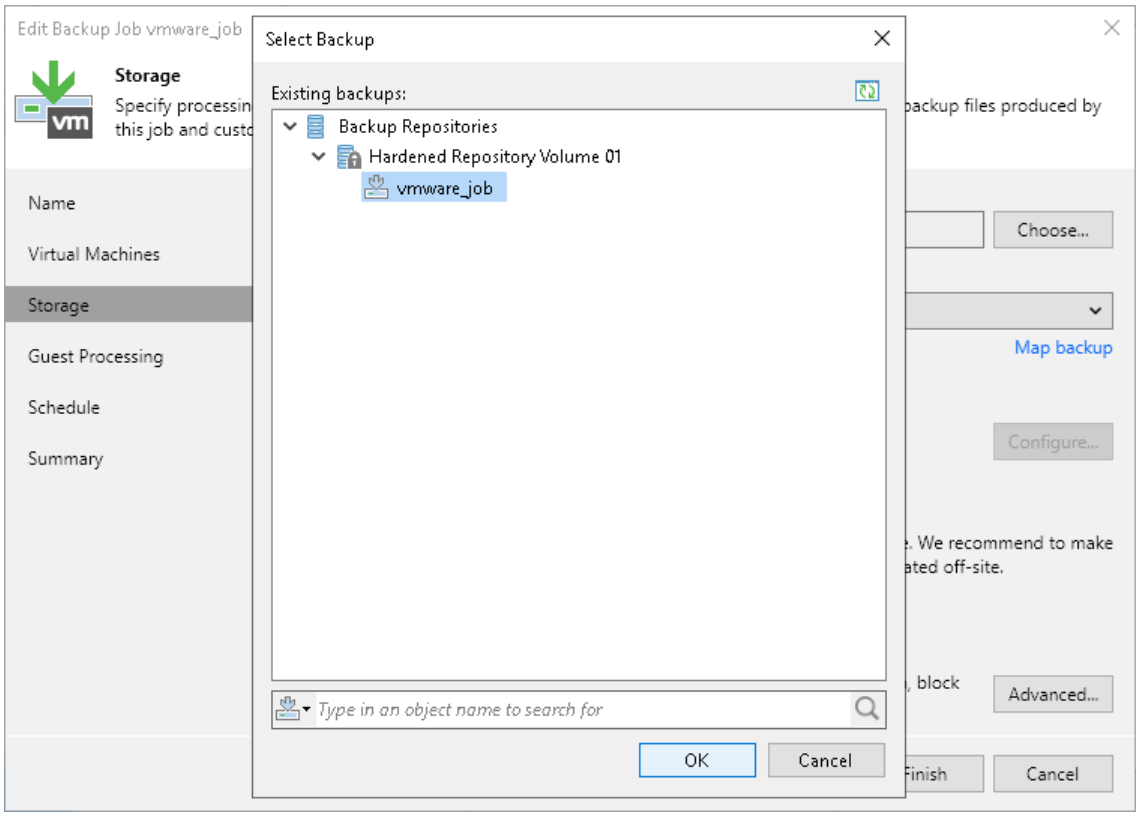


Finish the wizard to apply changes.

f. Open the **Backup Infrastructure** view and remove the Linux backup repository from the backup infrastructure.

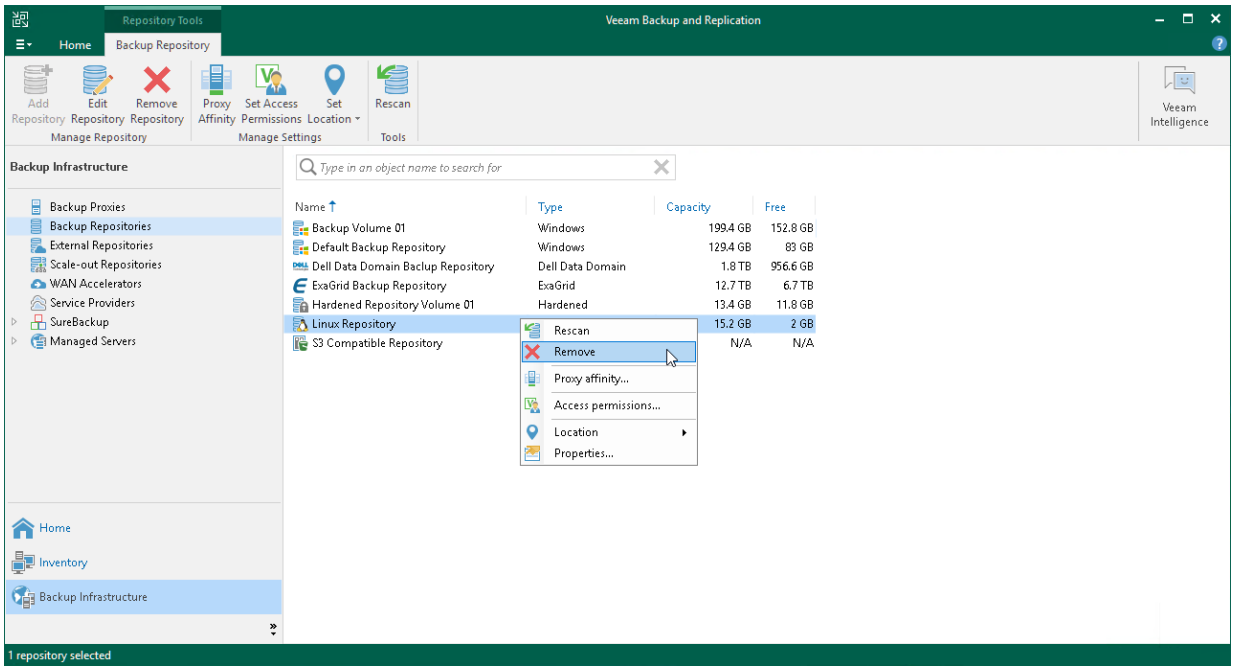


**NOTE**

For more information on enabling immutability for tenant backups, see the Switching from Linux Repository to Hardened Repository section in the Veeam Cloud Connect Guide.

## Upgrading Performance Extent to Hardened Repository

If you have Linux servers added as performance extents to a scale-out backup repository (SOBR), you can upgrade the extents to hardened repositories. To do this, perform the following steps:

1. Disable all backup jobs that use the SOBR. For more information, see Disabling Jobs.

   **NOTE**

   If you have Veeam Agent backup jobs managed by Veeam Agent, you need to delete these jobs and configure them again after you switch to the hardened repository.

2. Switch all Linux extents in the SOBR to maintenance mode. For more information, see Switching to Maintenance Mode.
3. On each Linux extent, change permissions for the directory where the backups are stored. Both owner and group must be the account with non-root permissions you use to connect to the Linux extent.

   ```
   chown -R owner:group <dir_path>
   ```

4. Go to the Veeam Backup & Replication console and perform the following steps:

   a. Open the **Backup Infrastructure** view and select **Managed Servers** in the navigation pane. For each Linux extent, specify the account with non-root permissions you use to connect to the Linux extent. These credentials must be single-use. To do this, right-click the Linux extent, select **Properties** and go the **SSH Connection** step of the wizard.

b. Remove the backup files stored on all Linux extents from the backup configuration. To do this, go to the **Backups** > **Disk** node and select a backup file. Hold [Ctrl] and right-click the file. Then, select **Remove from configuration**.

> **IMPORTANT**
>
> Removing backups from configuration is designed for experienced users only. It is strongly recommended to create encrypted configuration backup before performing this operation.

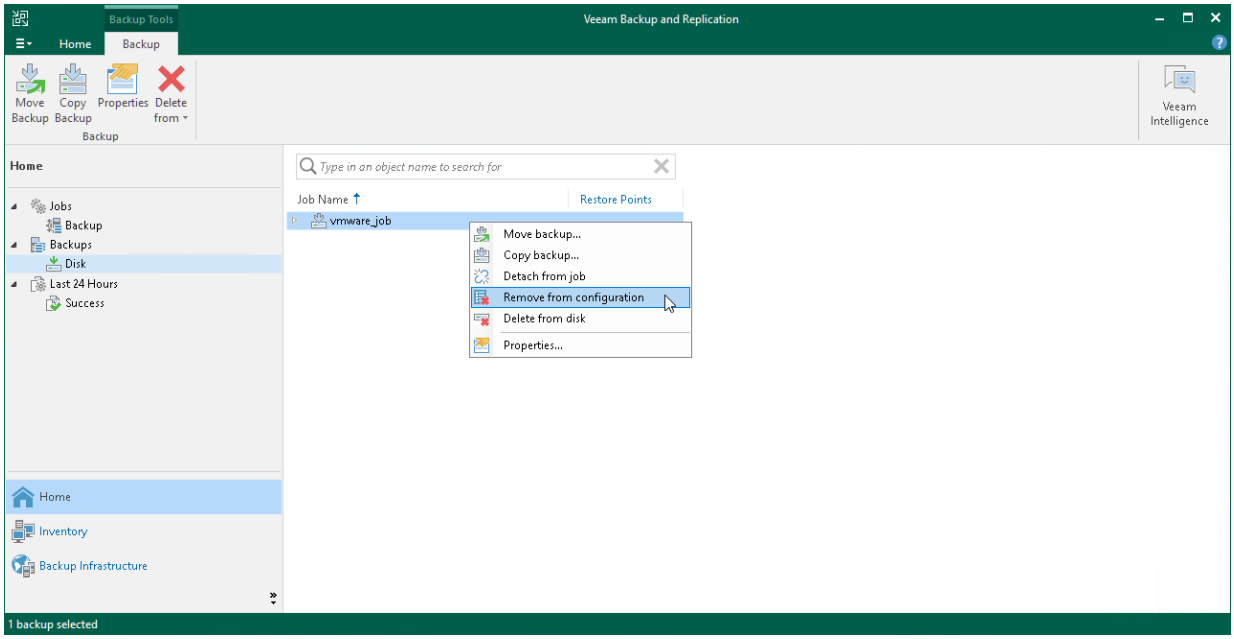


c. Add each Linux extent to the backup infrastructure as a hardened repository using the **New Backup Repository** wizard. For more information, see Adding Hardened Repositories.

5. Go to the **Performance Tier** step of the **Edit Scale-out Repository** wizard and remove all Linux extents. Then, add the newly-created hardened repositories. Go to the last step of the wizard, click **Apply**.

6. Open the **Backup Infrastructure** view and select **Scale-out Repositories** in the navigation pane. Right-click the SOBR and select **Rescan**. Any backup files stored on the added repositories will be automatically imported and displayed in the Veeam Backup & Replication console under the **Imported** > **Backups** node.

> **NOTE**
>
> If you work with encrypted backups, you need to decrypt them after importing. For more information, see How Data Decryption Works.

7. Go to the **Jobs** node and edit the jobs associated with the SOBR. At the **Storage** step of the wizard, select the SOBR from the backup repository list.

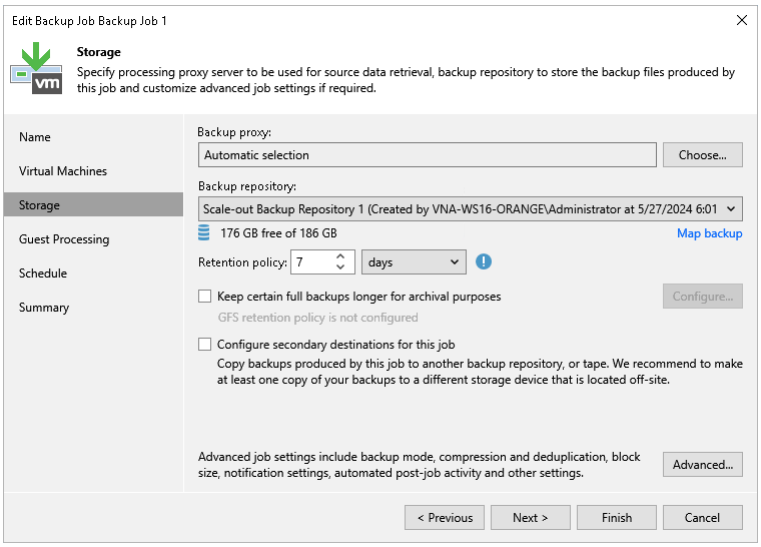


Select **Map Backup** and map the job to the corresponding imported backup. Finish the wizard to apply changes. Repeat this process for each backup job using the SOBR.

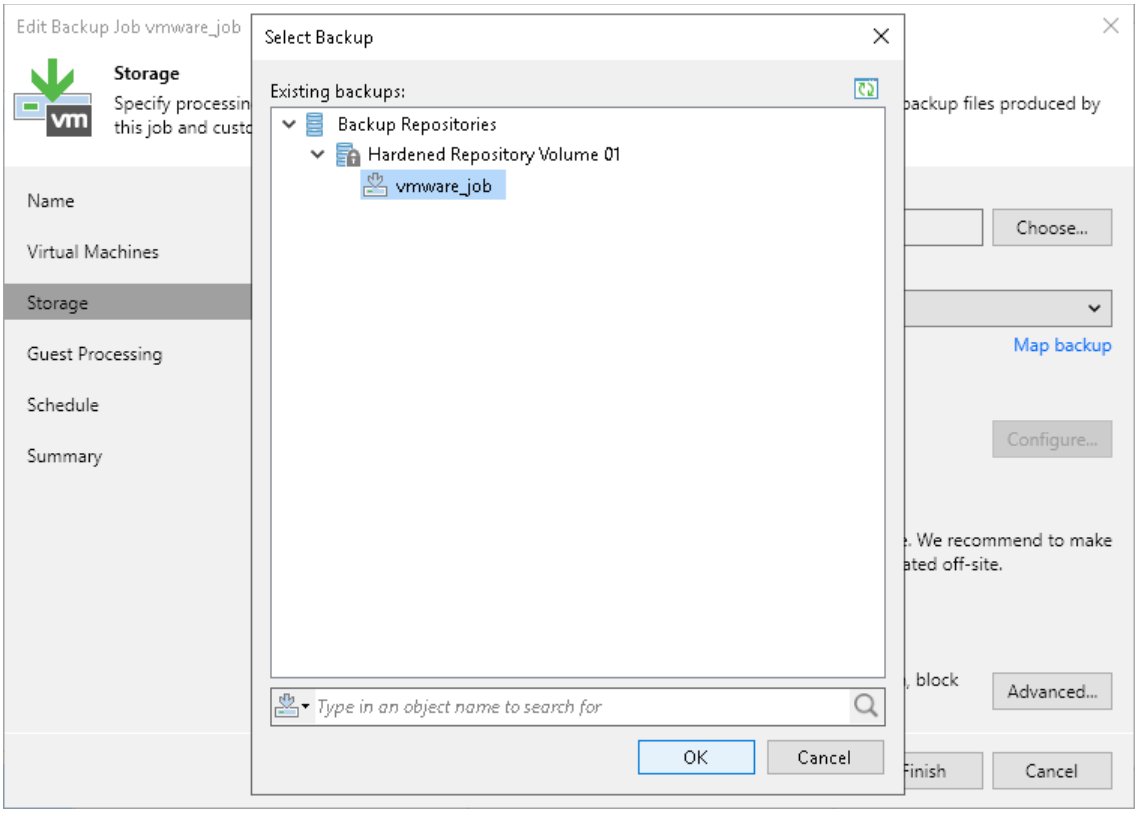


8. Re-enable the jobs that you disabled on the SOBR.
9. Open the **Backup Infrastructure** view and remove the unused Linux repositories from the backup infrastructure.

# SMB (CIFS) Share

You can use SMB (CIFS) shares as backup repositories.

## SMB Backup Repository Deployment

To communicate with an SMB backup repository, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the VMware backup proxy
- Veeam Data Mover on the gateway server

An SMB share cannot host Veeam Data Movers. For this reason, to communicate with the SMB share, you need to deploy a gateway server. Veeam Backup & Replication will automatically deploy a Veeam Data Mover on this gateway server. For more information on gateway server requirements and limitations, see Gateway Servers.

When any job addresses the SMB backup repository, Veeam Data Mover on the gateway server establishes a connection with Veeam Data Mover on the VMware backup proxy, enabling efficient data transfer over LAN or WAN.

If you plan to move VM data to an off-site SMB repository over a WAN link, it is recommended that you deploy an additional gateway server in the remote site, closer to the SMB repository.

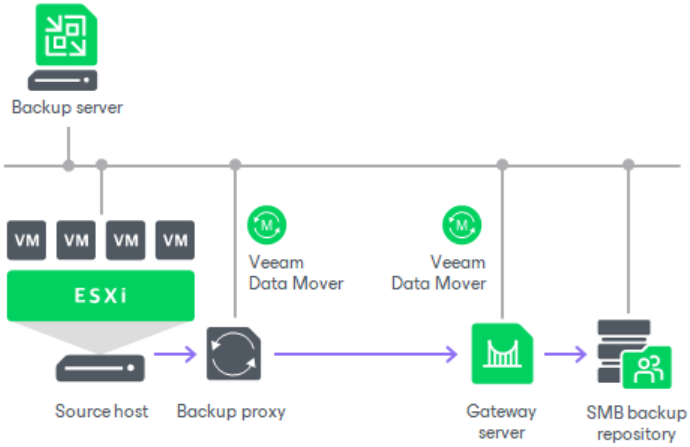


## Requirements for SMB Backup Repositories

The role of an SMB repository can be assigned to a Microsoft Windows machine (physical or virtual). The machine must meet the system requirements. For more information, see System Requirements.

## Adding SMB (CIFS) Repositories

This section describes how to add an SMB (CIFS) share as a backup repository.

To add a backup repository, use the **New Backup Repository** wizard.

## Step 1. Launch New Backup Repository Wizard

To launch the **New Backup Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add Backup Repository**. Alternatively, you can click **Add Repository** on the ribbon.
5. In the **Add Backup Repository** window, select **Network Attached Storage** > **SMB Share**.

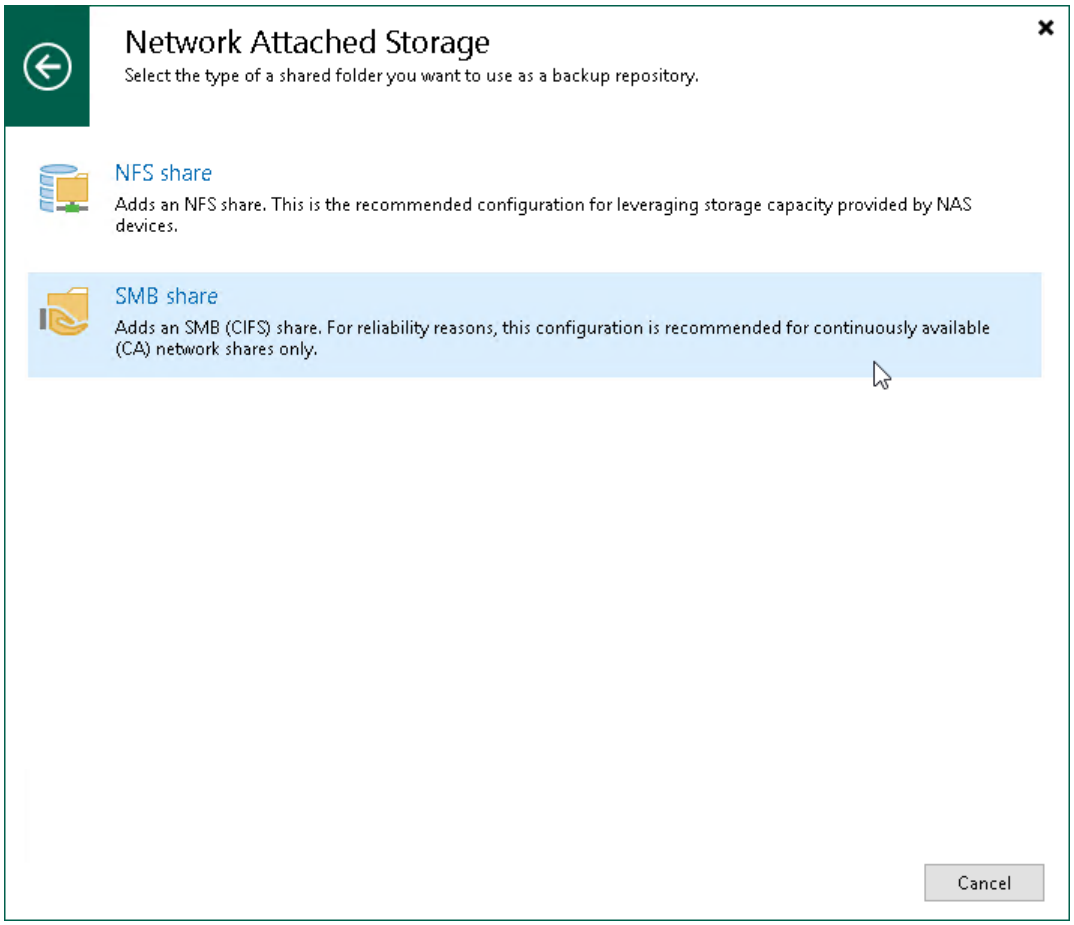

## Step 2. Specify Backup Repository Name and Description

At the **Name** step of the wizard, specify a name and description for the backup repository:

1. In the **Name** field, specify a name for the backup repository.
2. In the **Description** field, provide a description for future reference.

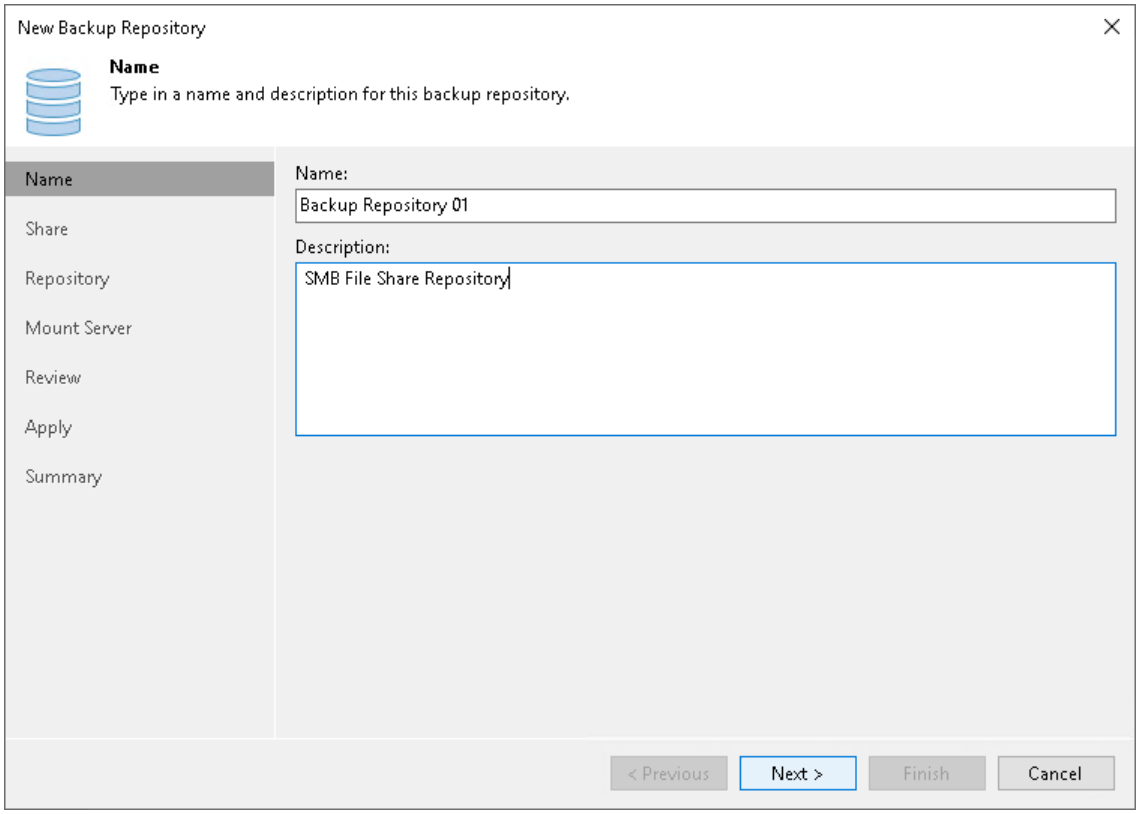

# Step 3. Specify Shared Folder Settings

To configure settings for an SMB share:

1. In the **Shared folder** field, specify a UNC path to the SMB shared folder that you want to use as a backup repository. If you use the IPv6 address, specify the path in a literal format, for example, *||[1080--8-800-200c-417a.ipv6-literal.net]|folder*. For more information, see RFC 2732. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
2. If you must specify user credentials to access the shared folder, select the **This share requires access credentials** check box. From the **Credentials** list, select a credentials record for a user account that has permissions described in section Permissions. Note that the username must be in the down-level logon name format. For example, *DOMAIN|username* or *HOSTNAME|username*.

   If you have not set up credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Credentials Manager.
3. In the **Gateway server** field, specify settings for the gateway server:
   - If you want Veeam Backup & Replication to select a gateway server automatically, leave **Automatic selection**.
   - If you want to select servers that can be used as gateway servers explicitly, click **Choose** next to the **Gateway server** field. In the **Gateway Server** window, click **Use the following gateway servers only** and select servers. The servers must have a direct access to the SMB share and must be located as close to the SMB share as possible. Veeam Backup & Replication will choose the most suitable server.

   For more information on the gateway servers, their requirements and limitations, and how they are selected, see Gateway Servers.

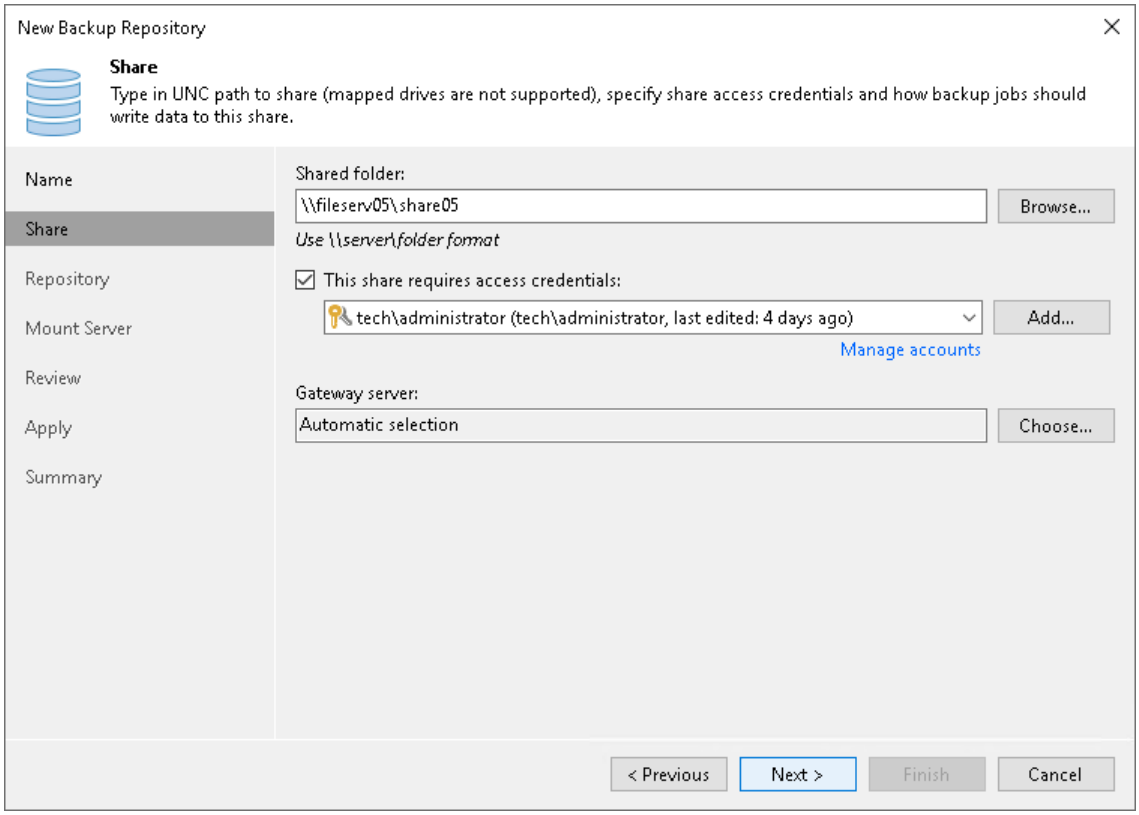

## Step 4. Configure Backup Repository Settings

At the **Repository** step of the wizard, configure general repository settings including path to the repository folder and load control, and also advanced repository settings.

# Configuring General Repository Settings

To configure general repository settings:

1. In the **Location** section, specify a path to the folder where backup files must be stored. Click **Populate** to check capacity and available free space in the selected location.

2. Use the **Load control** section to limit the number of concurrent tasks and data ingestion rate for the backup repository. These settings will help you control the load on the backup repository and prevent possible timeouts of storage I/O operations.

- Select the **Limit maximum concurrent tasks** check box and specify the maximum allowed number of concurrent tasks for the backup repository. If this value is exceeded, Veeam Backup & Replication will not start a new task until one of current tasks finishes. For more information, see Limiting the Number of Concurrent Tasks.

  > **NOTE**
  >
  > Consider the following:
  >
  > - Limitation of concurrent tasks is ignored if the backup repository acts as a target storage for a Veeam Cloud Connect job.
  > - If you use backup repositories with per-machine backup chains, it is recommended to select the **Limit maximum concurrent tasks to N** check box. This option reduces the number of parallel operations performed by synthetic operations (synthetic full backup, backup files merge and transformation). Otherwise, the load on the backup repository may be high.

- Select the **Limit read and write data rate to** check box and specify the maximum rate to restrict the total speed of reading and writing data to the backup repository disk. For more information, see Limitation of Read and Write Data Rates for Backup Repositories.

  > **NOTE**
  >
  > The **Limit read and write data rate to** setting does not apply to health checks performed as part of backup and backup copy jobs. Even if you limit read/write rate for a backup repository, the health check will consume resources of the backup repository regardless of this setting. Consider this limitation when configuring basic and health check schedules for backup and backup copy jobs.

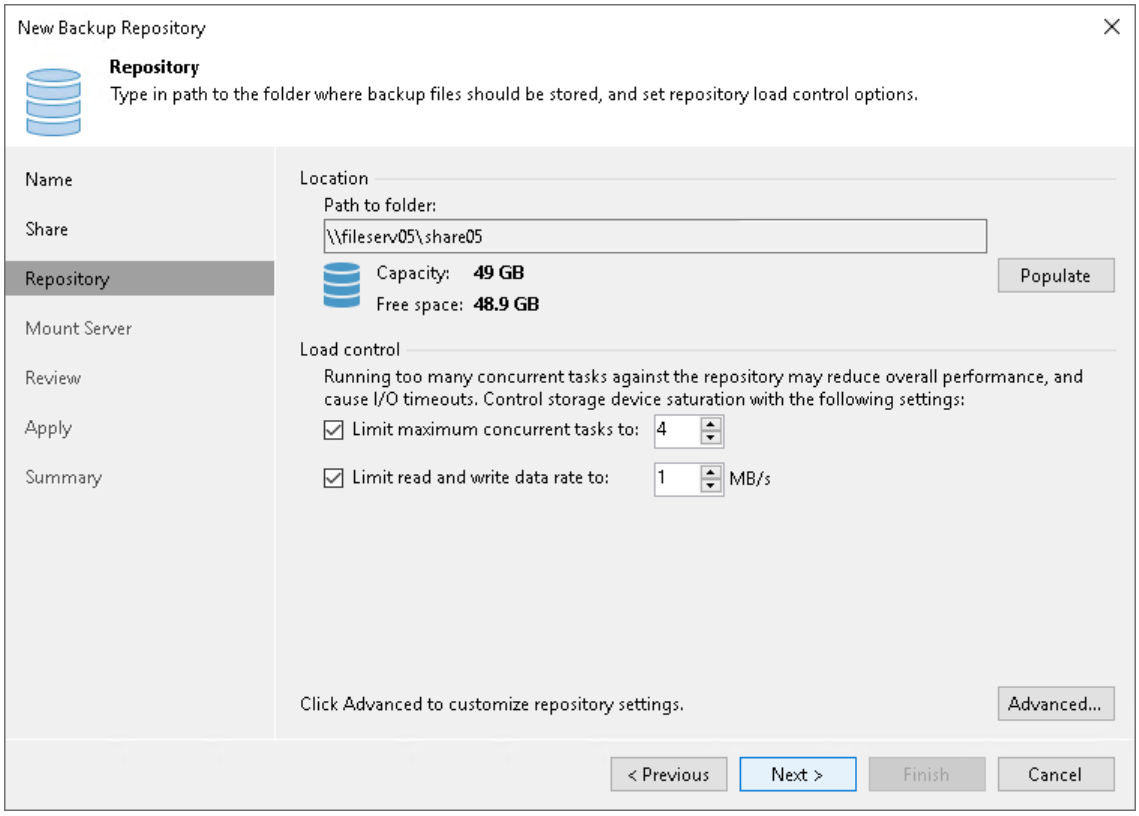


# Configuring Advanced Repository Settings

To configure advanced repository settings:

1. Click **Advanced**.
2. For storage systems using a fixed block size, select the **Align backup file data blocks** check box. Veeam Backup & Replication will align VM data saved to a backup file at a 4 KB block boundary.
3. When you enable compression for a backup job, Veeam Backup & Replication compresses VM data at the source side and then transports it to the target side. Writing compressed data to a deduplicating storage appliance results in poor deduplication ratios as the number of matching blocks decreases. To overcome this situation, select the **Decompress backup data blocks before storing** check box. If data compression is enabled for a job, Veeam Backup & Replication will compress VM data on the source side, transport it to the target side, decompress VM data on the target side and write raw VM data to the storage device to achieve a higher deduplication ratio.

> **NOTE**
>
> Veeam Backup & Replication does not compress VM data if encryption is enabled for a job and the **Decompress backup data blocks before storing** check box is selected in the settings of the target backup repository. Therefore, in the job statistics, you may observe a higher amount of transferred data (the **Transferred** counter) as compared to a job for which encryption is disabled. For more information on job statistics, see Viewing Real-Time Statistics.
>
> In the properties of the backup created with encryption, you may also see that the backup size (the **Backup Size** column) is larger than the original VM size (the **Original Size** column). For more information on backup properties, see Viewing Backup Properties.

4. Select the **This repository is backed by rotated drives** check box if you plan to use a backup repository with rotated drives. For more information on how to configure rotated drives, see Deploying Backup Repositories with Rotated Drives.

5. To create a separate backup file for every machine in the job, make sure that the **Use per-machine backup files** check box is selected. If you clear the check box, Veeam Backup & Replication will create single-file backups. For more information on the backup chain formats and their limitations, see Backup Chain Formats.

> **NOTE**
>
> Changing of the **Use per-machine backup files** setting after the repository was already created does not take any effect. To change backup chain format, follow the instructions provided in Upgrading Backup Chain Formats.

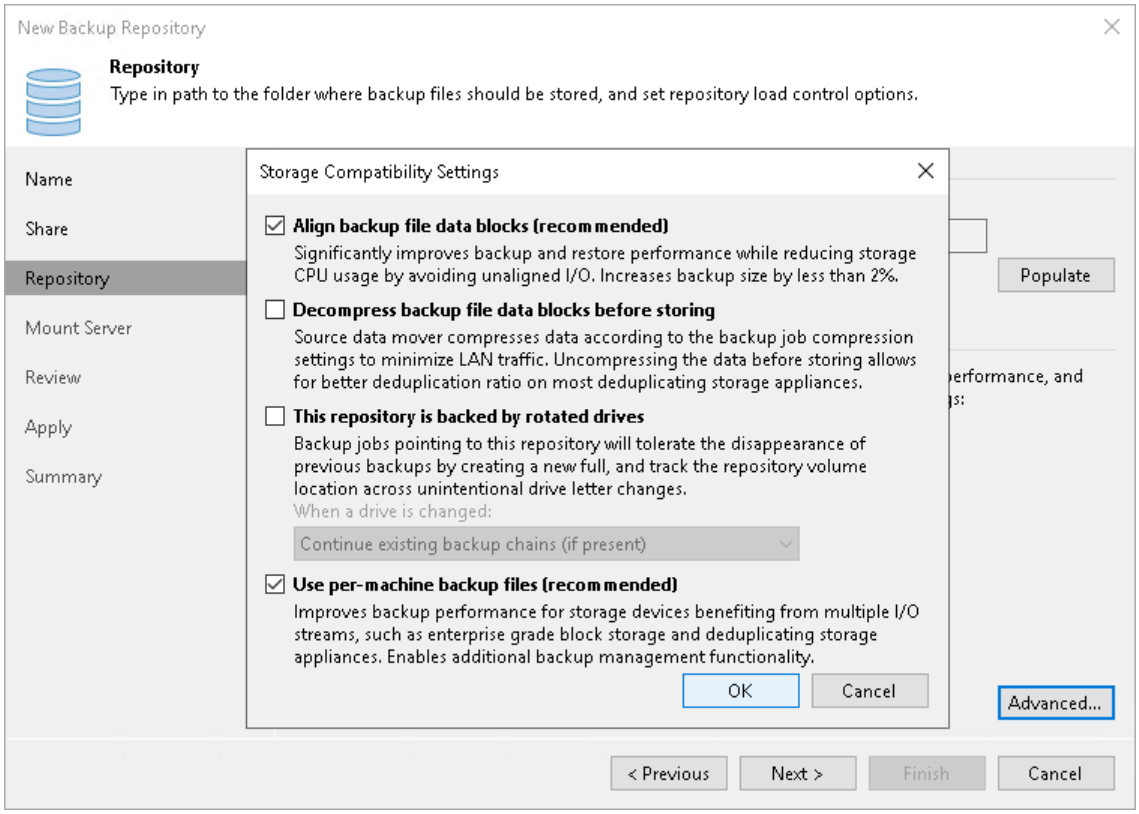

# Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for file-level and application items restore.

1. From the **Mount Server** list, select a server that you want to use as a mount server. The mount server is required for file-level and application items restore. During the restore process, Veeam Backup & Replication mounts the VM disks from the backup file residing in the backup repository to the mount server. As a result, VM data does not have to travel over the network, which reduces the load on the network and speed up the restore process. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers added to the backup infrastructure. If the server is not added to the backup infrastructure, click **Add New** on the right to open the **New Windows Server** wizard. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder that will be used for writing cache during mount operations.

3. To make the backup repository accessible by the Veeam vPower NFS Service, select the **Enable vPower NFS service on the mount server** check box. Veeam Backup & Replication will enable the vPower NFS Service on your selected mount server.

4. To customize network ports used by the vPower NFS Service, click **Ports**. For information on ports used by default, see Ports.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

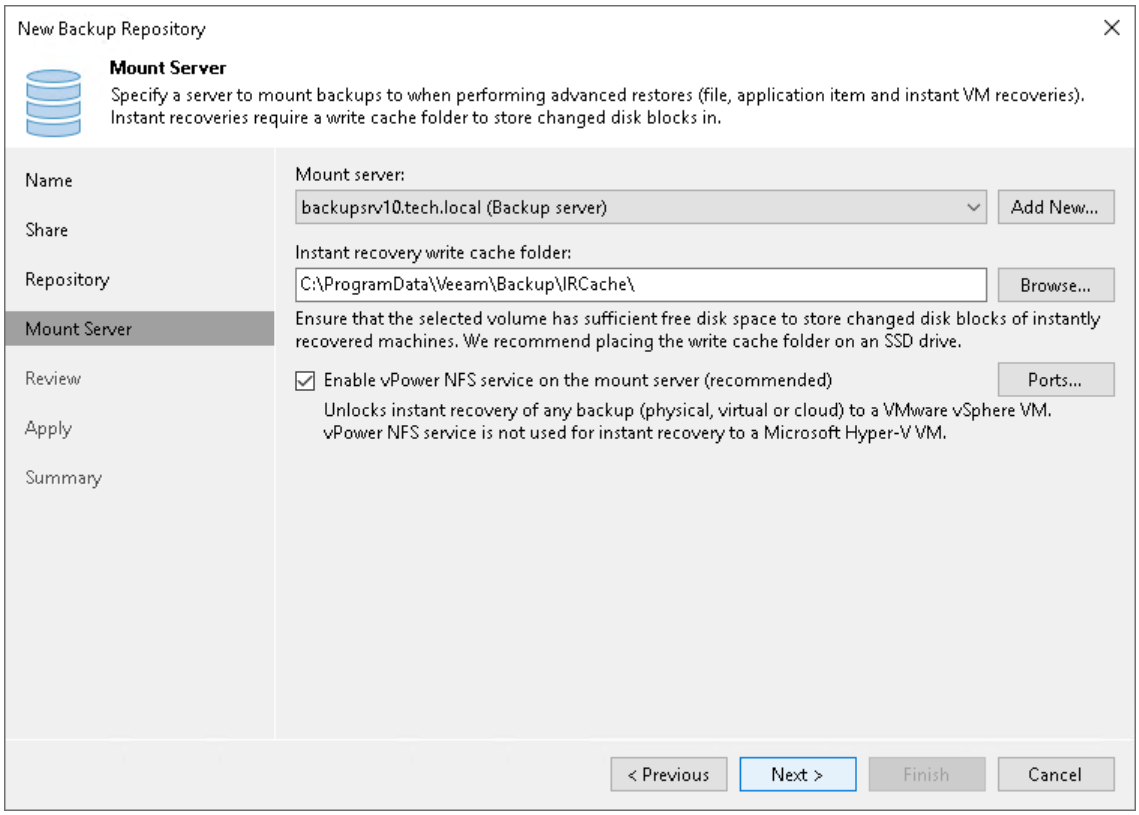

## Step 6. Review Properties and Components

At the **Review** step of the wizard, review details of the backup repository and specify importing settings.

1. Review the backup repository settings and list of components that will be installed on the backup repository server.
2. If the backup repository contains backups that were previously created by Veeam Backup & Replication, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Disk (Imported)** node.
3. If the backup repository contains guest file system index files that were previously created by Veeam Backup & Replication, select the **Import guest file system index** check box. Index files will be imported with backup files, and you will be able to search for guest OS files inside imported backups.

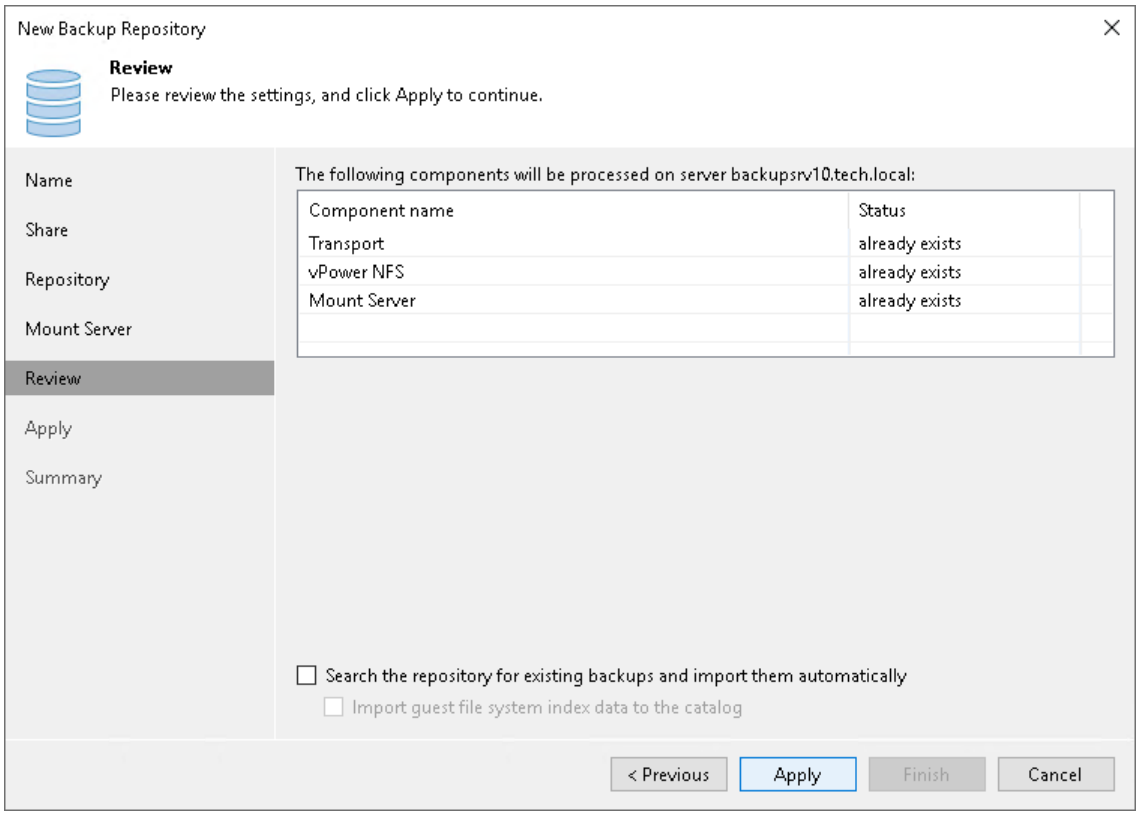

# Step 7. Apply Backup Repository Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to install and configure all required components. Then click **Next** to complete the procedure of adding the backup repository to the backup infrastructure.

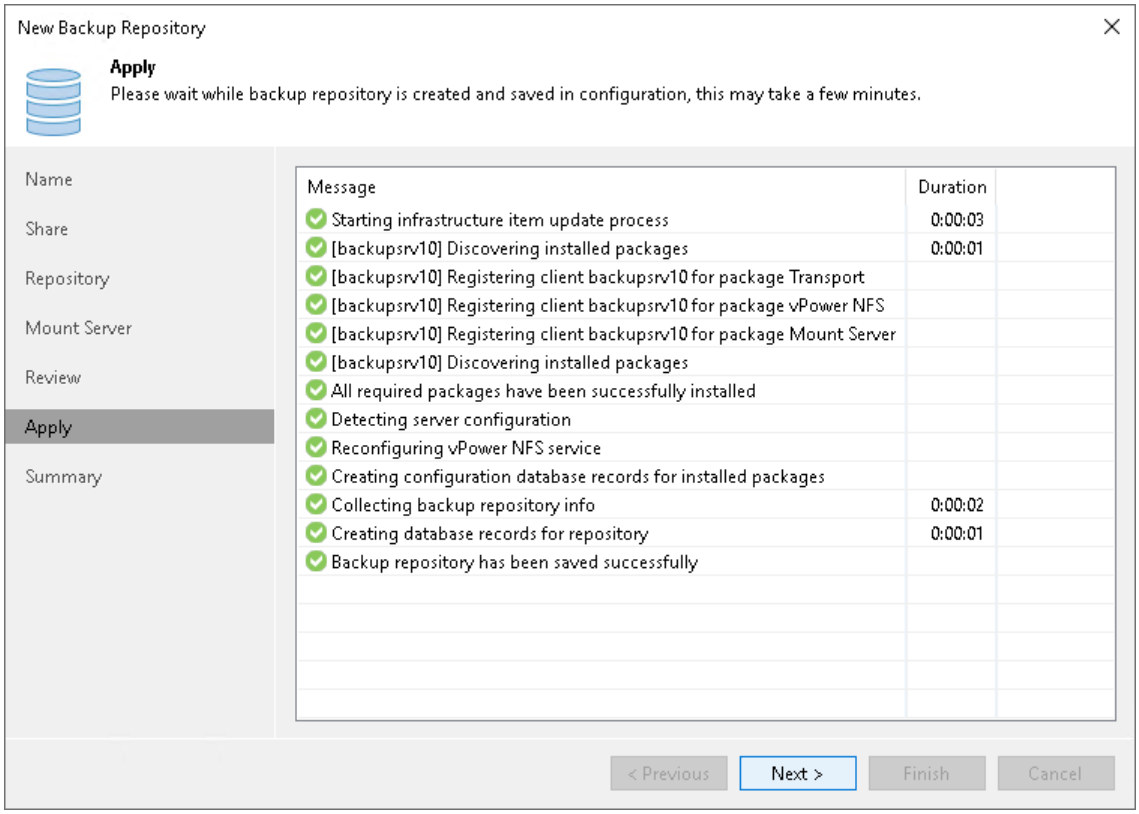

# Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added backup repository. Then click **Finish** to exit the wizard.

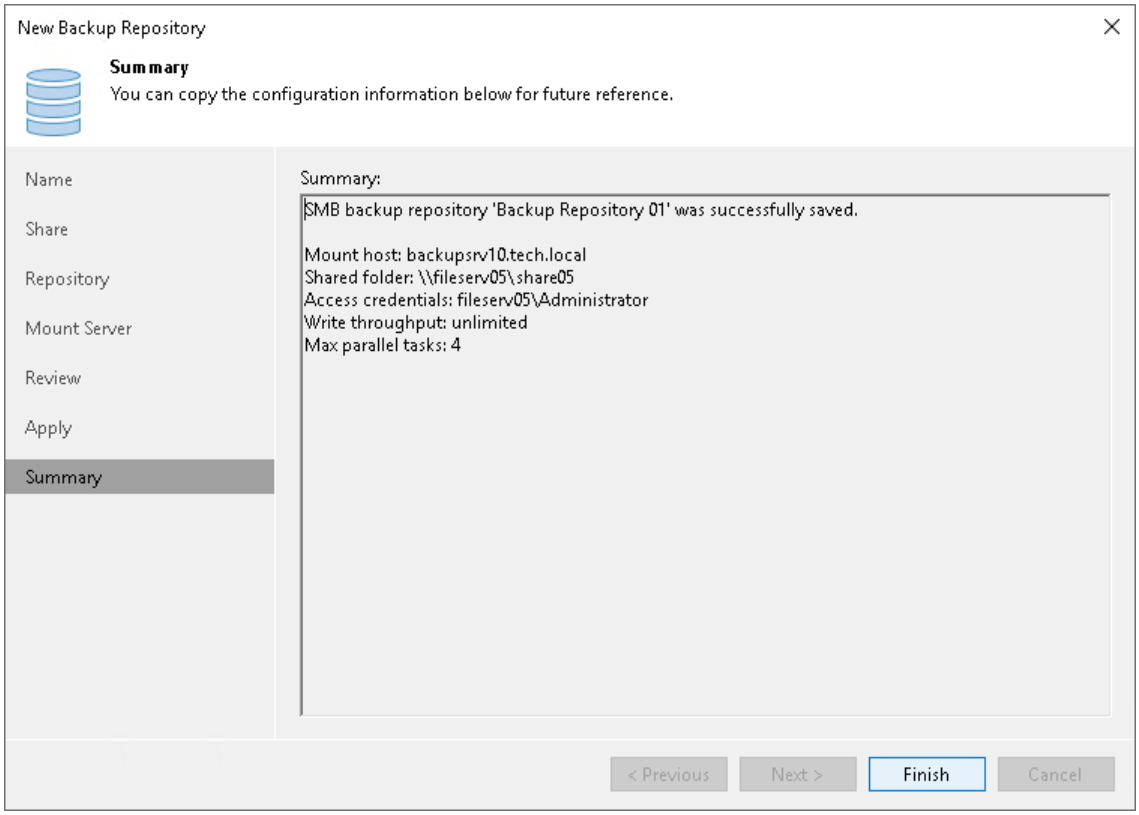

# NFS Share

You can use NFS shares as backup repositories.

## NFS Backup Repository Deployment

To communicate with an NFS backup repository, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the VMware backup proxy
- Veeam Data Mover on the gateway server

> **NOTE**
>
> Windows-based gateway servers cannot be used for NFS shares with **krb5i** and **krb5p** support.

An NFS share cannot host Veeam Data Movers. For this reason, to communicate with the NFS share, you need to deploy a gateway server. Veeam Backup & Replication will automatically deploy a Veeam Data Mover on this gateway server. For more information, see Gateway Servers.

When any job addresses the NFS backup repository, Veeam Data Mover on the gateway server establishes a connection with Veeam Data Mover on the VMware backup proxy, enabling efficient data transfer over LAN or WAN.

If you plan to move VM data to an off-site NFS repository over a WAN link, it is recommended that you deploy an additional gateway server in the remote site, closer to the NFS repository.

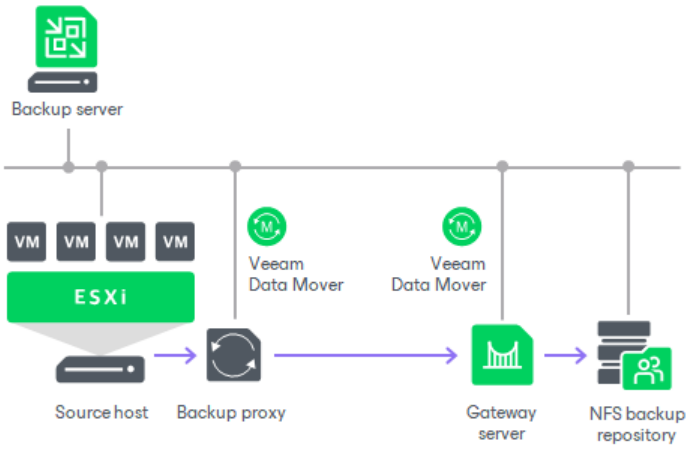


## Requirements and Limitations for NFS Backup Repositories

A machine performing the role of an NFS repository must meet the following requirements:

- The machine must meet the system requirements. For more information, see System Requirements.
- The role of the NFS repository can be assigned to a Microsoft Windows or Linux machine (physical or virtual) or to NAS storage supporting NFS protocol.
- The NFS repository must present read and write access rights to the gateway.
- [For NFS repositories with Microsoft Windows gateway servers] The repository must support NFS protocol versions 3.0 and 4.1. Versions 4.0 and 4.2 and other minor versions are not supported.

- [For NFS repositories with Microsoft Windows gateway servers] For NFS servers using v4.1 protocol, the delegation feature must be disabled.
- [For NFS repositories with Linux gateway servers] The repository supports any NFS protocol version that the Linux gateway server can mount.
- [For NFS servers using v3 protocol] The NFS server must support the READDIRPLUS command.

Note that Veeam Backup & Replication does not support multipathing for NFS repository.

# Requirements for Gateway Server

A machine performing the role of a gateway server for communication with the NFS backup repository must meet the following requirements:

- The role of the gateway server can be assigned to a Microsoft Windows or Linux machine (physical or virtual). The machine must meet the system requirements. For more information, see System Requirements.
- You must add the machine to the Veeam Backup & Replication console as a managed server.
- [For automatic gateway selection] The backup server must have read and write access in the NFS repository.
- [For automatic gateway selection] If you configure automatic gateway selection for NFS repository, Veeam Backup & Replication may use the same machines as gateways for the repository and as proxies for backup jobs. Make sure that the backup proxies meet the following requirements:
  - If you explicitly choose backup proxies for backup jobs, provide read and write access rights to all proxies chosen for backup jobs that are targeted to the NFS repository.
  - If you configure automatic proxy selection for backup jobs, provide read and write access rights to all proxies in the Veeam Backup & Replication infrastructure.
  - If backup jobs that are targeted to the NFS repository use Linux proxies, check that the NFS client package is installed on the Linux proxy server.

## Requirements and Limitations for Linux Gateway Server

In addition to the general requirements, the Linux gateway server must meet the following requirements:

- The Linux gateway server must have NFS client package installed.
- The credentials to authenticate with the Linux gateway server must have root or elevated to root permission.
- Veeam Backup & Replication uses the highest NFS protocol version supported by the gateway and the repository.

  Note that the suffix indicating the NFS version in the NFS share properties may not be displayed correctly, this is a known issue.

> **TIP**
>
> If the NFS protocol version has changed (for example, if you updated the NFS share), click through the **Edit Backup Repository** wizard to update the information in Veeam Backup & Replication.

# Adding NFS Repositories

This section describes how to add an NFS share as a backup repository.

To add a backup repository, use the **New Backup Repository** wizard.

## Step 1. Launch New Backup Repository Wizard

To launch the **New Backup Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add Backup Repository**. Alternatively, you can click **Add Repository** on the ribbon.
5. In the **Add Backup Repository** window, select **Network Attached Storage** > **NFS Share**.

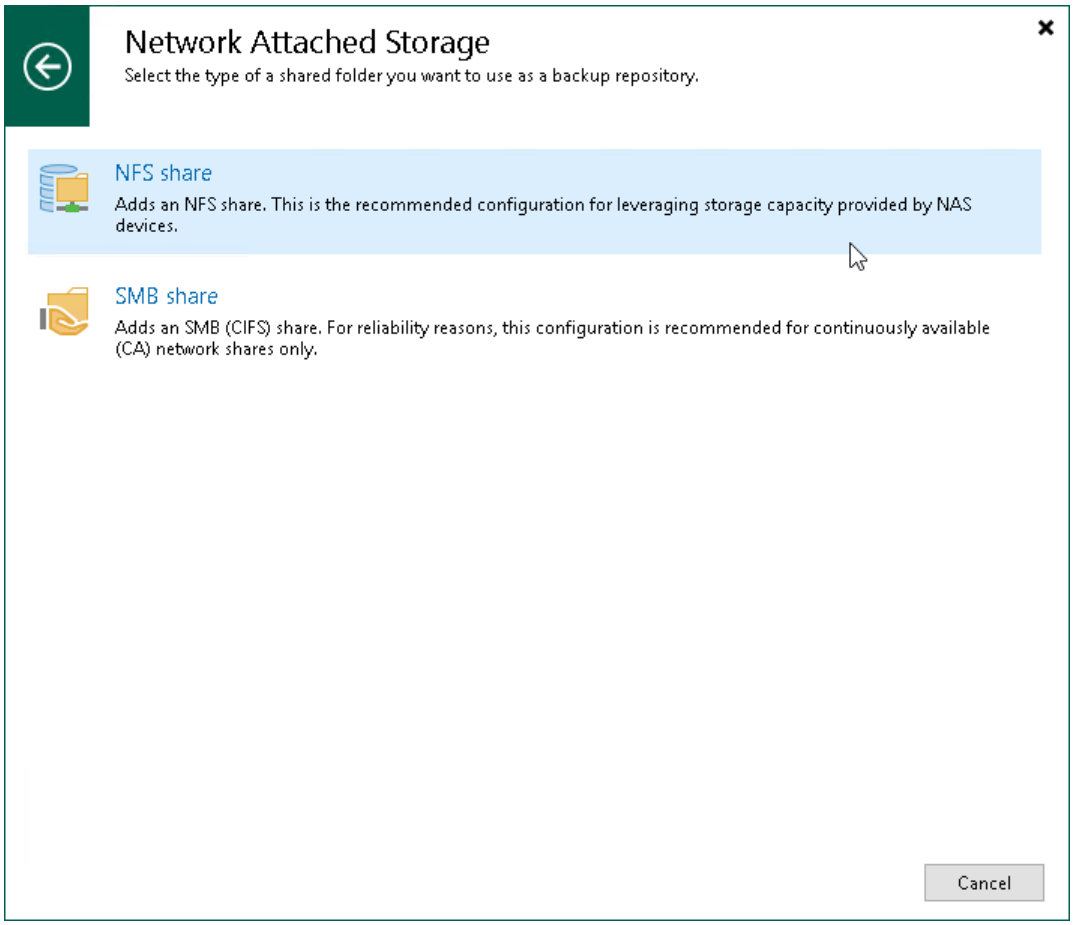

## Step 2. Specify Backup Repository Name and Description

At the **Name** step of the wizard, specify a name and description for the backup repository:

1. In the **Name** field, specify a name for the backup repository.
2. In the **Description** field, provide a description for future reference.

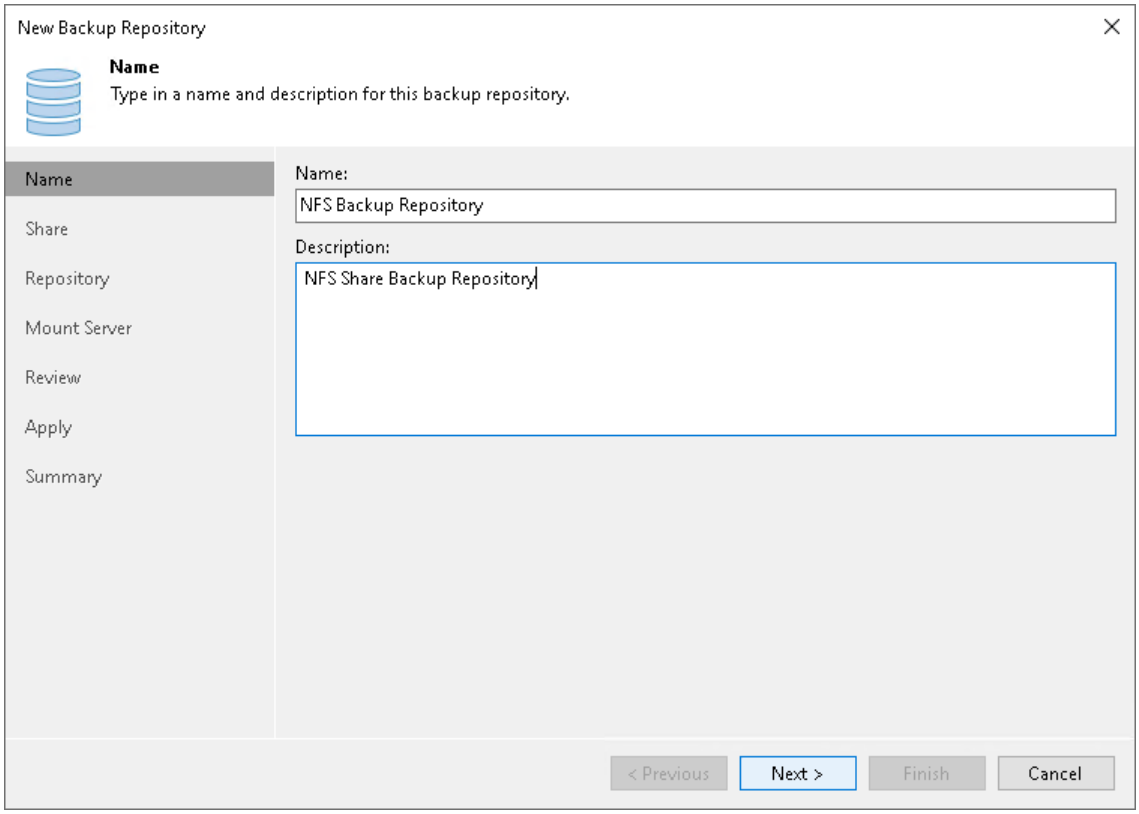

## Step 3. Specify Shared Folder Settings

To configure settings for an NFS share:

1. In the **Shared folder** field, specify a path to the NFS shared folder that you want to use as a backup repository. You can specify the path using an IPv4 or IPv6 address. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
2. In the **Gateway server** field, specify which gateway server you want to use:
   - If you want Veeam Backup & Replication to select a gateway server automatically, leave **Automatic selection**.
   - If you want to select servers that can be used as gateway servers explicitly, click **Choose** next to the **Gateway server** field. In the **Gateway Server** window, click **Use the following gateway servers only** and select servers. The servers must have a direct access to the NFS share and must be located as close to the NFS share as possible. Veeam Backup & Replication will choose the most suitable server.

   For more information on the gateway servers, their requirements and limitations, and how they are selected, see Gateway Servers.

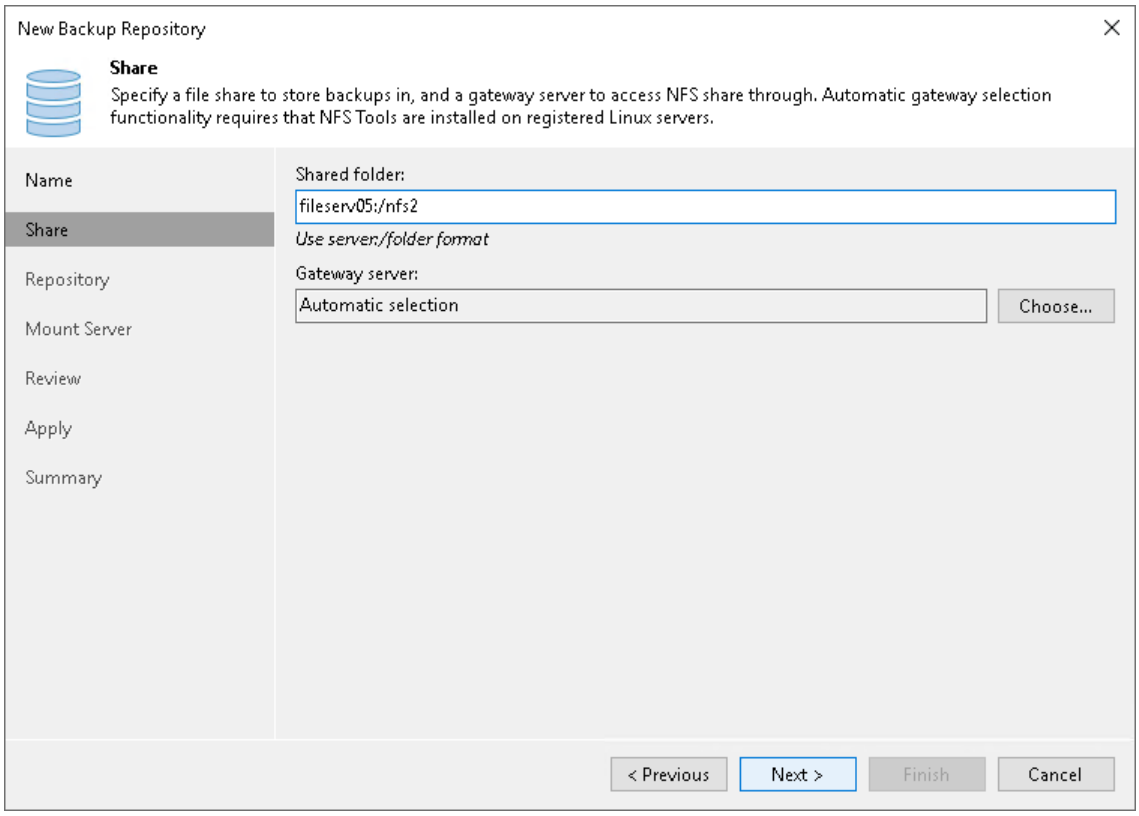

## Step 4. Configure Backup Repository Settings

At the **Repository** step of the wizard, configure general repository settings including path to the repository folder and load control, and also advanced repository settings.

# Configuring General Repository Settings

To configure general repository settings:

1. In the **Location** section, specify a path to the folder where backup files must be stored. Click **Populate** to check capacity and available free space in the selected location.

2. Use the **Load control** section to limit the number of concurrent tasks and data ingestion rate for the backup repository. These settings will help you control the load on the backup repository and prevent possible timeouts of storage I/O operations.

- Select the **Limit maximum concurrent tasks** check box and specify the maximum allowed number of concurrent tasks for the backup repository. If this value is exceeded, Veeam Backup & Replication will not start a new task until one of current tasks finishes. For more information, see Limiting the Number of Concurrent Tasks.

  > **NOTE**
  >
  > Consider the following:
  >
  > - Limitation of concurrent tasks is ignored if the backup repository acts as a target storage for a Veeam Cloud Connect job.
  > - If you use backup repositories with per-machine backup chains, it is recommended to select the **Limit maximum concurrent tasks to N** check box. This option reduces the number of parallel operations performed by synthetic operations (synthetic full backup, backup files merge and transformation). Otherwise, the load on the backup repository may be high.

- Select the **Limit read and write data rate to** check box and specify the maximum rate to restrict the total speed of reading and writing data to the backup repository disk. For more information, see Limitation of Read and Write Data Rates for Backup Repositories.

  > **NOTE**
  >
  > The **Limit read and write data rate to** setting does not apply to health checks performed as part of backup and backup copy jobs. Even if you limit read/write rate for a backup repository, the health check will consume resources of the backup repository regardless of this setting. Consider this limitation when configuring basic and health check schedules for backup and backup copy jobs.

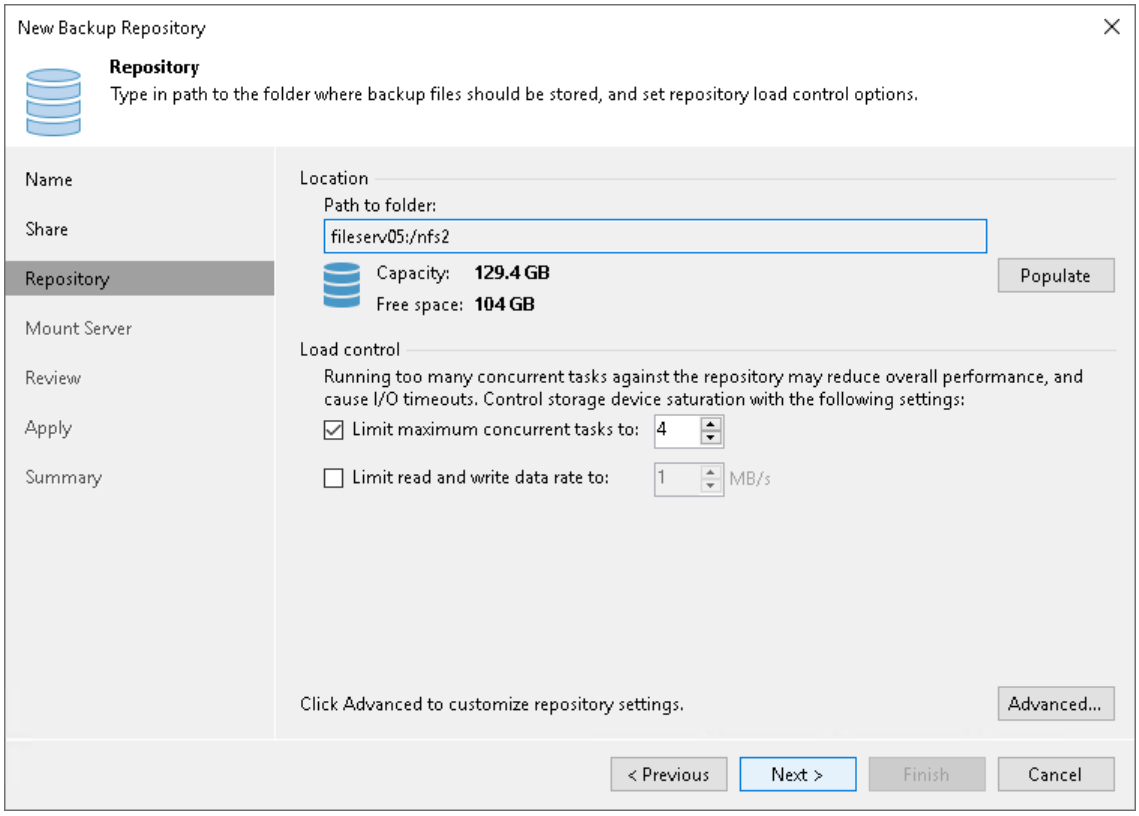


# Configuring Advanced Repository Settings

To configure advanced repository settings:

1. Click **Advanced**.
2. For storage systems using a fixed block size, select the **Align backup file data blocks** check box. Veeam Backup & Replication will align VM data saved to a backup file at a 4 KB block boundary.
3. When you enable compression for a backup job, Veeam Backup & Replication compresses VM data at the source side and then transports it to the target side. Writing compressed data to a deduplicating storage appliance results in poor deduplication ratios as the number of matching blocks decreases. To overcome this situation, select the **Decompress backup data blocks before storing** check box. If data compression is enabled for a job, Veeam Backup & Replication will compress VM data on the source side, transport it to the target side, decompress VM data on the target side and write raw VM data to the storage device to achieve a higher deduplication ratio.

   > **NOTE**
   >
   > Veeam Backup & Replication does not compress VM data if encryption is enabled for a job and the **Decompress backup data blocks before storing** check box is selected in the settings of the target backup repository. Therefore, in the job statistics, you may observe a higher amount of transferred data (the **Transferred** counter) as compared to a job for which encryption is disabled. For more information on job statistics, see Viewing Real-Time Statistics.
   >
   > In the properties of the backup created with encryption, you may also see that the backup size (the **Backup Size** column) is larger than the original VM size (the **Original Size** column). For more information on backup properties, see Viewing Backup Properties.

4. The **This repository is backed by rotated hard drives** option is disabled for NFS repositories. For more information, see Deploying Backup Repositories with Rotated Drives.

5. To create a separate backup file for every machine in the job, make sure that the **Use per-machine backup files** check box is selected. If you clear the check box, Veeam Backup & Replication will create single-file backups. For more information on the backup chain formats and their limitations, see Backup Chain Formats.

> **NOTE**
>
> Changing of the **Use per-machine backup files** setting after the repository was already created does not take any effect. To change backup chain format, follow the instructions provided in Upgrading Backup Chain Formats.

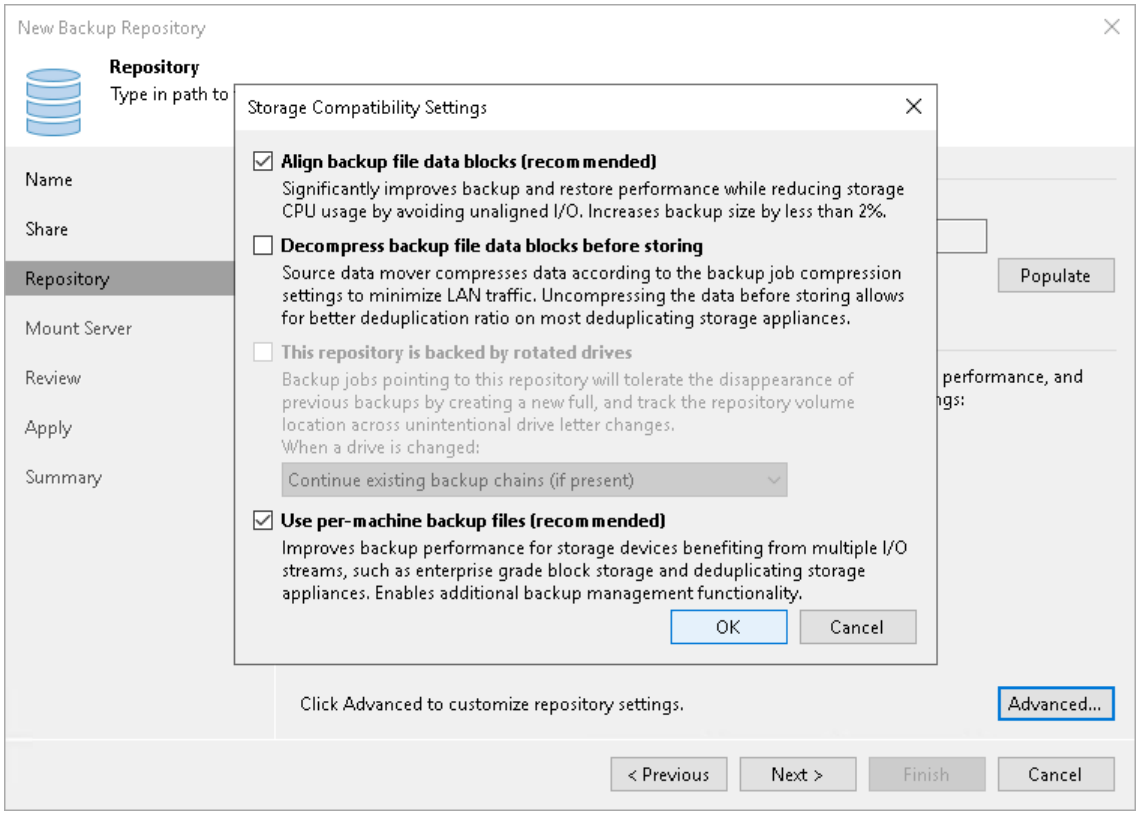

# Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for file-level and application items restore.

1. From the **Mount Server** list, select a server that you want to use as a mount server. The mount server is required for file-level and application items restore. During the restore process, Veeam Backup & Replication mounts the VM disks from the backup file residing in the backup repository to the mount server. As a result, VM data does not have to travel over the network, which reduces the load on the network and speed up the restore process. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers added to the backup infrastructure. If the server is not added to the backup infrastructure, click **Add New** on the right to open the **New Windows Server** wizard. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder that will be used for writing cache during mount operations.

3. To make the backup repository accessible by the Veeam vPower NFS Service, select the **Enable vPower NFS service on the mount server** check box. Veeam Backup & Replication will enable the vPower NFS Service on your selected mount server.

4. To customize network ports used by the vPower NFS Service, click **Ports**. For information on ports used by default, see Ports.

**IMPORTANT**

Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

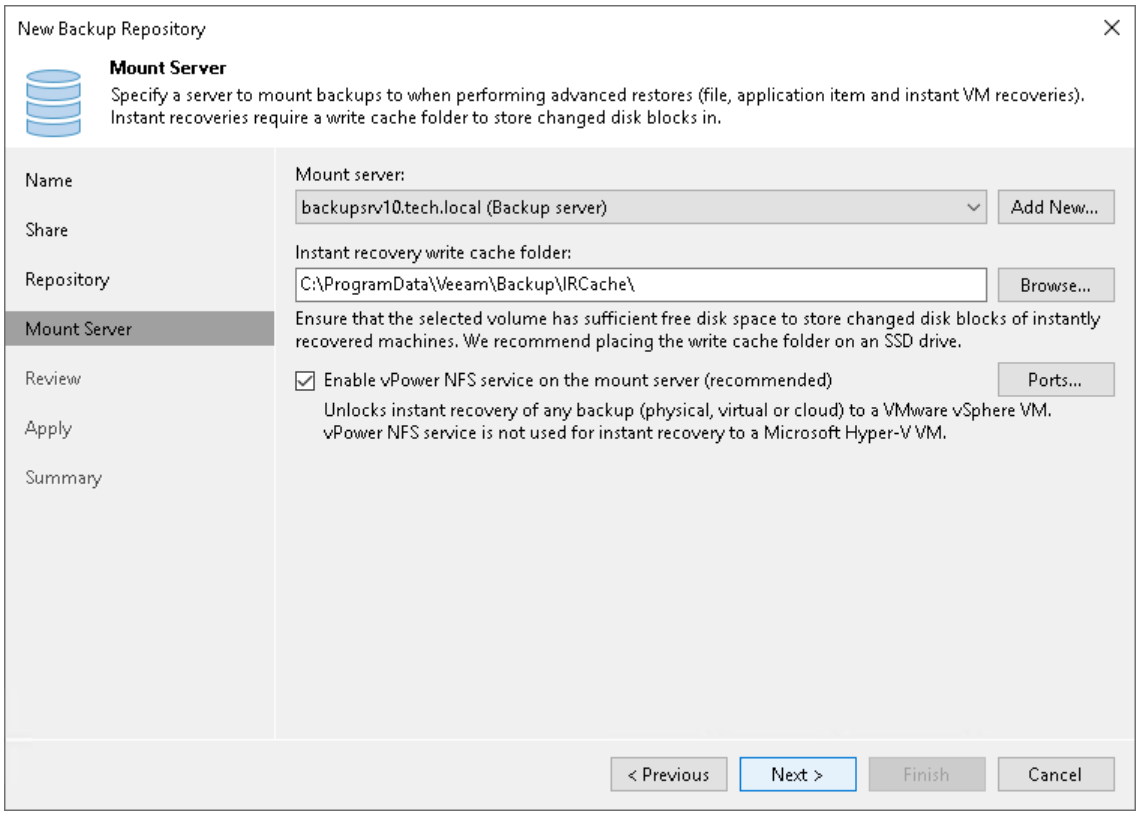

## Step 6. Review Properties and Components

At the **Review** step of the wizard, review details of the backup repository and specify importing settings.

1. Review the backup repository settings and list of components that will be installed on the backup repository server.
2. If the backup repository contains backups that were previously created by Veeam Backup & Replication, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Disk (Imported)** node.
3. If the backup repository contains guest file system index files that were previously created by Veeam Backup & Replication, select the **Import guest file system index** check box. Index files will be imported with backup files, and you will be able to search for guest OS files inside imported backups.

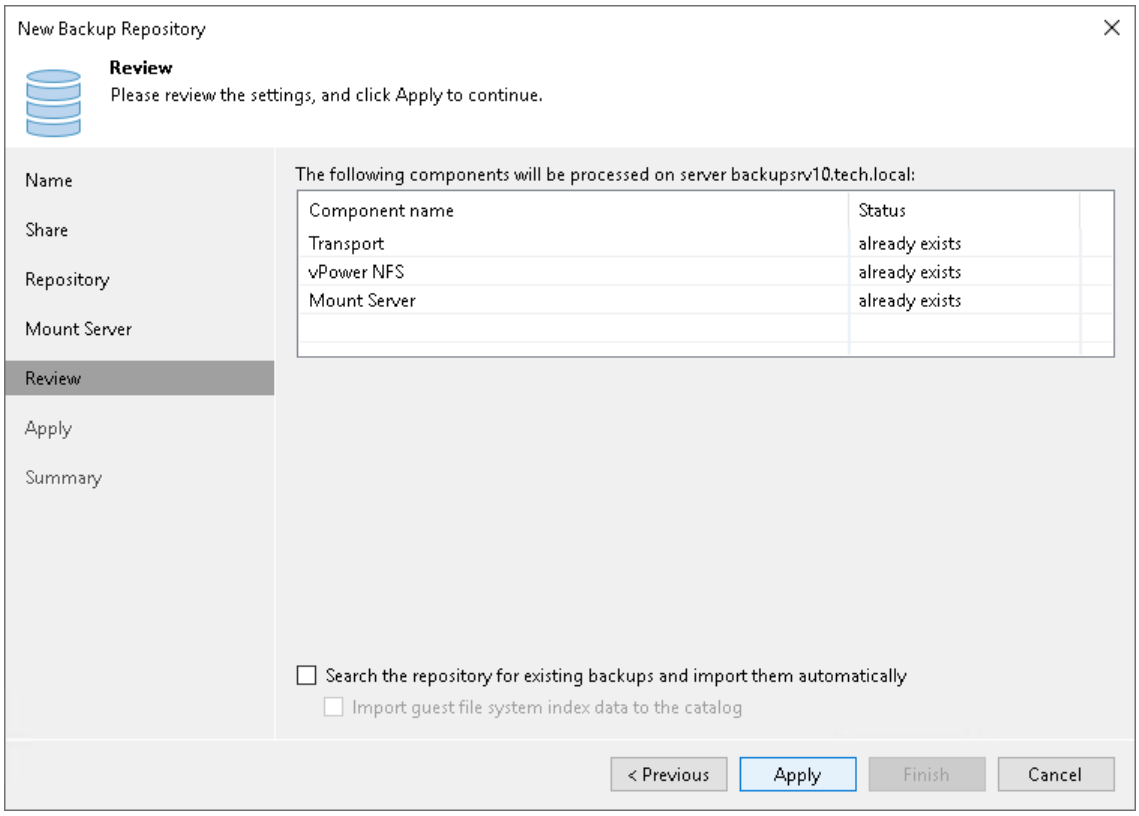

# Step 7. Apply Backup Repository Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to install and configure all required components. Then click **Next** to complete the procedure of adding the backup repository to the backup infrastructure.

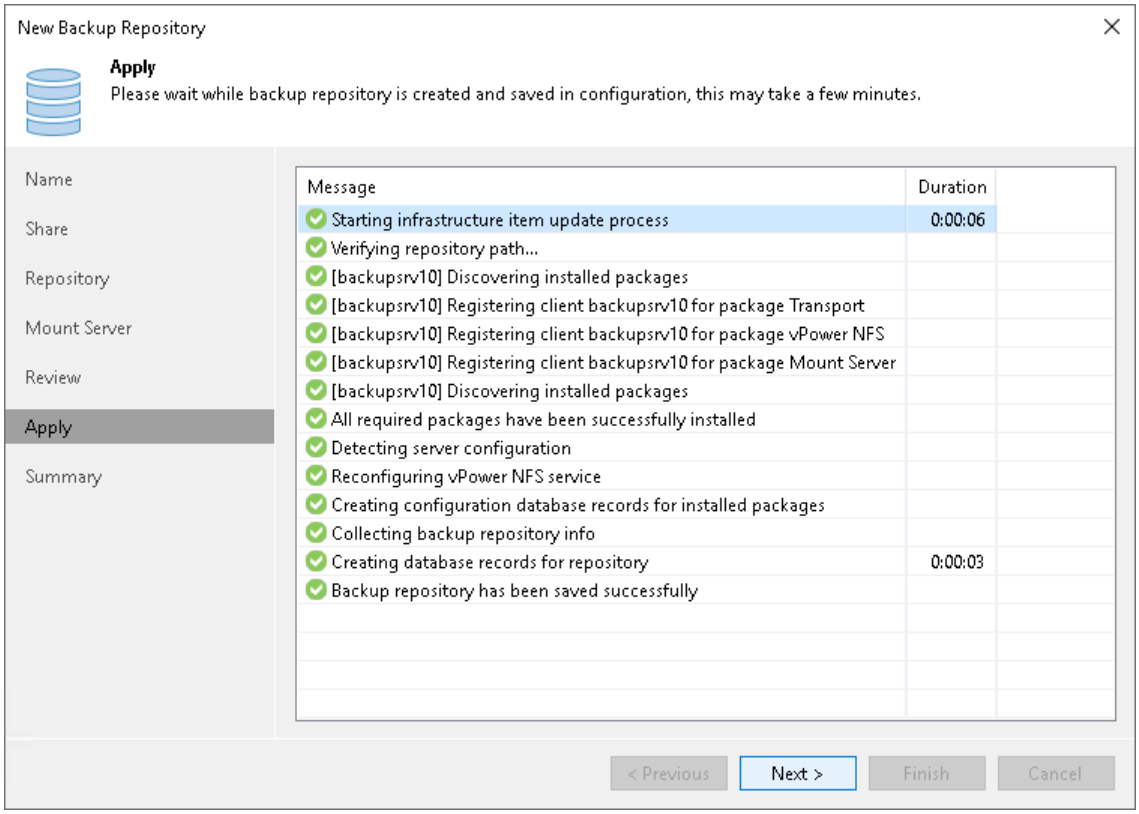

# Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added backup repository. Then click **Finish** to exit the wizard.

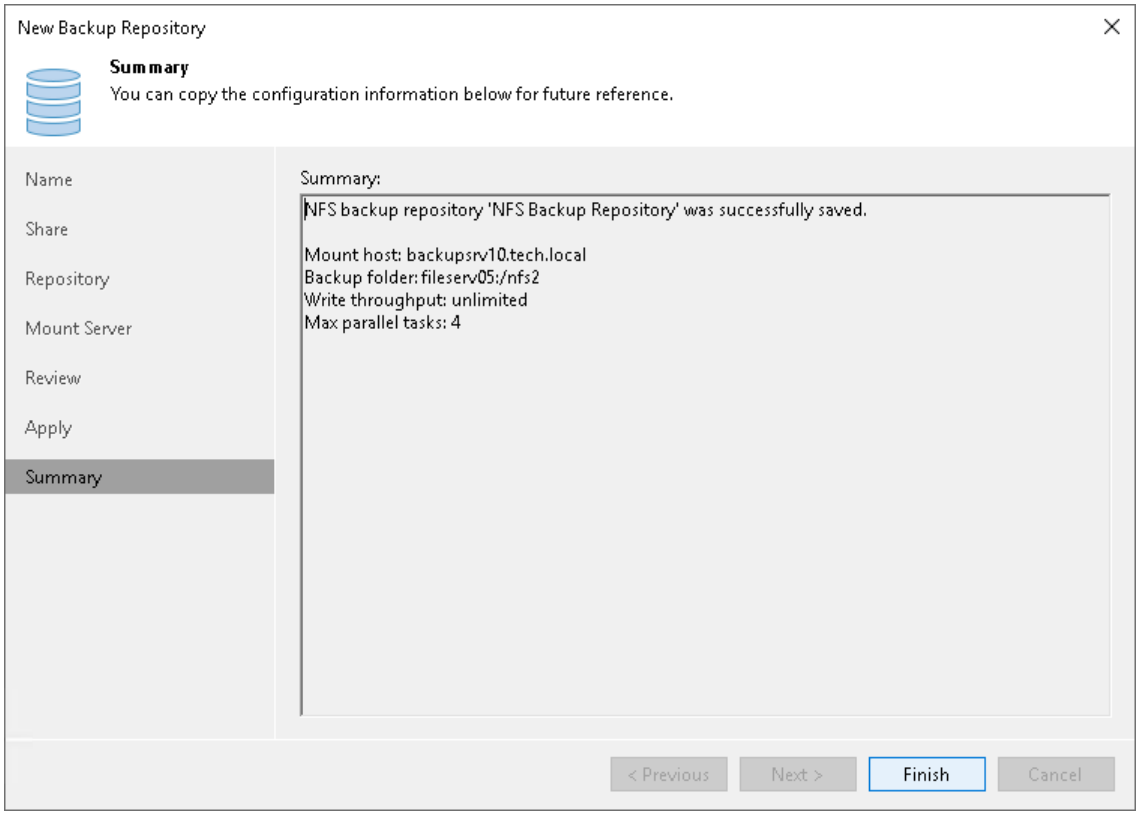

# Deduplicating Storage Appliances

You can add deduplicating storage appliances as backup repositories.

## In This Section


- Dell Data Domain
- ExaGrid
- HPE StoreOnce
- Quantum DXi
- Fujitsu ETERNUS CS800
- Infinidat InfiniGuard
- Adding Deduplicating Storage Appliances


## Dell Data Domain

You can use Dell Data Domain storage systems with Data Domain Boost (DD Boost) as backup repositories.

To support the DD Boost technology, Veeam Backup & Replication leverages the following Dell Data Domain components:

- **DD Boost library**. The DD Boost library is a component of the Dell Data Domain system. The DD Boost library is embedded into the Veeam Data Mover setup. When you add a Microsoft Windows server to the backup infrastructure, the DD Boost Library is automatically installed on the added server together with Veeam Data Mover.
- **DD Boost server**. The DD Boost server is a target-side component. The DD Boost server runs on the OS of the Dell Data Domain storage system.

## Backup Infrastructure

To communicate with Dell Data Domain, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the backup proxy
- Veeam Data Mover on the gateway server

The Dell Data Domain storage cannot host Veeam Data Mover. For this reason, to communicate with the Dell Data Domain storage, you need to deploy a gateway server. Veeam Backup & Replication will automatically deploy Veeam Data Mover on this gateway server. For more information, see Gateway Servers.

When any job addresses the backup repository, Veeam Data Mover on the gateway server establishes a connection with Veeam Data Mover on the backup proxy, enabling efficient data transfer over LAN or WAN.

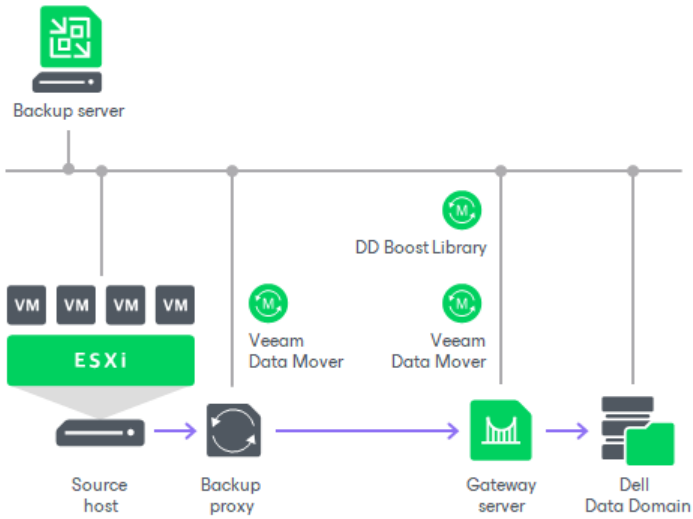


You define what gateway server to use when you assign a backup repository role to Dell Data Domain. You can define the gateway server explicitly or instruct Veeam Backup & Replication to select it automatically.

> **IMPORTANT**
>
> For Dell Data Domain storage systems working over Fibre Channel, you must explicitly define the gateway server that will communicate with Dell Data Domain. As a gateway server, you must use a Microsoft Windows server that is added to the backup infrastructure and has access to Dell Data Domain over Fibre Channel.

# Supported Protocols

Veeam Backup & Replication supports Dell Data Domain storage systems working over the following protocols:

- TCP/IP protocol: Veeam Backup & Replication communicates with the Dell Data Domain server by sending commands over the network.
- Fibre Channel protocol: Veeam Backup & Replication communicates with the Dell Data Domain Fibre Channel server by sending SCSI commands over Fibre Channel.

# Considerations and Limitations for Dell Data Domain

## General

If you plan to use Dell Data Domain as a backup repository, consider the following:

- We strongly recommend that you follow recommendations from this list and also recommendations from this Veeam KB article.
- Use of Dell Data Domain with DD Boost does not guarantee improvement of job performance. It reduces the load on the network and improves the network throughput.
- NFS services must be enabled on Dell Data Domain. Otherwise, Veeam Backup & Replication will not be able to access the storage system.

- Do not enable encryption for the jobs targeted at the deduplication storage appliance. Encryption has a negative effect on the deduplication ratio. For more information, see Job Encryption.

  Encryption can also affect the backup size: the size of a backup can be larger than the size of the original VM. When you enable encryption, set data processing block size to 4 MB, and leave decompression on target and block alignment enabled (enabled by default), the size of the backup increases by 1 MB for each 4 MB data block. This is because Veeam Backup & Replication reads blocks of 4 MB, encrypts them, adds 16 KB metadata to each data block and then aligns data blocks. This results in that each 4 MB data block on the source becomes 5 MB block on the target.

- When you create a backup job targeted at an Dell Data Domain backup repository, Veeam Backup & Replication will offer you to switch to optimized job settings and use the 4 MB size of data block for workload data processing (the **Storage optimization** setting). It is recommended that you use optimized job settings. Large data blocks produce a smaller metadata table that requires less memory and CPU resources to process. For more information on storage optimization, see Data Compression and Deduplication.

- Dell Data Domain does not support the reverse incremental backup method.

- You cannot use Dell Data Domain backup repositories as sources or targets for file copy jobs.

- The length of forward incremental and forever forward incremental backup chains (chains that contain one full backup and a set of subsequent incremental backups) cannot be greater than 120 restore points. To overcome this limitation, schedule full backups (active or synthetic) to split the backup chain into shorter series. For example, to perform backups at 15-minute intervals 24 hours a day, you must schedule synthetic fulls every day. In this scenario, intervals immediately after midnight may be skipped due to duration of synthetic processing. For more information, see How Synthetic Full Backup Works.

- If you connect to an Dell Data Domain backup repository over Fibre Channel, you must explicitly define a gateway server to communicate with Dell Data Domain. As a gateway server, you must use a Microsoft Windows server that is added to the backup infrastructure and has access to the Dell Data Domain backup repository over Fibre Channel.

- During backup repository rescan, Veeam Backup & Replication detects if the hard stream limit is set for a storage unit, and displays this information in backup repository rescan statistics. If the hard stream limit is exceeded when Veeam Backup & Replication runs tasks against the backup repository, Veeam Backup & Replication will fail to create new I/O streams.

## Immutability

If you plan to enable immutability for the Dell Data Domain deduplicating storage appliance, consider the following recommendations and limitations:

- Veeam Backup & Replication supports only the compliance mode of Dell Data Domain retention lock. For more information on retention lock, see Dell Technologies Article.

- Automatic retention lock must be disabled.

- The duration of the immutability period (the Make recent backups immutable setting) as well as the duration of the long-term retention period (the Keep certain backups longer for archival purposes setting) configured in Veeam Backup & Replication must lie in the range between the minimum and maximum retention periods configured in Dell Data Domain. The minimum and maximum values are included into the range.

  If the immutability period configured in Veeam Backup & Replication lies outside the range configured in Dell Data Domain, Veeam Backup & Replication does not allow you to proceed to the next step until the value is changed.

If the long-term retention period configured in Veeam Backup & Replication is greater than the maximum retention period configured in Dell Data Domain, Veeam Backup & Replication keeps backups immutable for the period configured in the Veeam Backup & Replication settings but resets the expiration of the backup immutability flag during each job run. While the actual number of days that the backup must be kept immutable is greater than the maximum, Veeam Backup & Replication sets the expiration period equal to the maximum.

If some operations, for example, evacuate or move, try to set the immutability period less than the minimum period configured in Dell Data Domain, Veeam Backup & Replication sets the expiration period for the immutability flag equal to the minimum retention period.

- [For unstructured data backups] The immutability lock period in Dell Data Domain is set by the backup repository immutability settings. For an example of how this mechanism works, see **Case 8** in the Unstructured Data Backup Retention Scenarios section.
- Immutability is supported only for forward incremental backup chains. Once a backup file becomes immutable, it can be merged or deleted only when the immutability time period expires.
- To use the immutability feature for backup copy jobs, enable the GFS retention policy. For more information, see Long-Term Retention Policy (GFS) for Backup Copy Jobs.

### Dell Data Domain Immutability and Veeam Agents

For more information on how immutability works with Veeam Agents, see the Backup to Deduplicating Storage Appliances section in the Veeam Agent Management Guide.

## Dell Data Domain Supported Features

The DD Boost technology offers a set of features for advanced data processing. Veeam Backup & Replication supports the following features:

- Distributed Segment Processing
- Advanced Load Balancing and Link Failover
- Virtual Synthetics
- In-Flight Data Encryption
- Per Storage Unit Streams
- Retention Lock

> **NOTE**
>
> You cannot configure Managed File Replication using Veeam Backup & Replication. However, you can import and map backups replicated between Data Domain storage systems to backup, backup copy or replication jobs, or perform restore operations from such backups.

# Distributed Segment Processing

Distributed Segment Processing lets Dell Data Domain "distribute" the deduplication process and perform a part of data deduplication operations on the gateway server side.

Without Distributed Segment Processing, Dell Data Domain performs deduplication on the Dell Data Domain storage system. The gateway server sends unfiltered data blocks to Dell Data Domain over the network. Data segmentation, filtering and compression operations are performed on the target side, before data is written to disk.

With Distributed Segment Processing, operations on data segmentation, filtering and compression are performed on the gateway server side. The gateway server sends only unique data blocks to Dell Data Domain. As a result, the load on the network reduces and the network throughput improves.

## Advanced Load Balancing and Link Failover

Advanced Load Balancing and Link Failover allow you to balance data transfer load and route VM data traffic to a working link in case of network outage problems.

Without Advanced Load Balancing, every gateway server connects to Data Domain on a dedicated Ethernet link. Such configuration does not provide an ability to balance the data transfer load across the links. If a network error occurs during the data transfer process, the backup job fails and needs to be restarted.

Advanced Load Balancing allows you to aggregate several Ethernet links into one interface group. As a result, Dell Data Domain automatically balances the traffic load coming from several gateway servers united in one group. If some link in the group goes down, Dell Data Domain automatically performs link failover, and the backup traffic is routed to a working link.

## Virtual Synthetics

Veeam Backup & Replication supports Virtual Synthetic Fulls by Dell Data Domain. Virtual Synthetic Fulls let you synthesize a full backup on the target backup storage without physically copying data from source datastores. To construct a full backup file, Dell Data Domain uses pointers to existing data segments on the target backup storage. Virtual Synthetic Fulls reduce the workload on the network and backup infrastructure components and increase the backup job performance.

## In-Flight Data Encryption

Veeam Backup & Replication supports in-flight encryption introduced in Dell Data Domain Boost 3.0. If necessary, you can enable data encryption at the backup repository level. Veeam Backup & Replication will leverage the Dell Data Domain technology to encrypt data transported between the DD Boost library and Data Domain system.

## Per Storage Unit Streams

Veeam Backup & Replication supports per storage unit streams on Dell Data Domain. The maximum number of parallel tasks that can be targeted at the backup repository (the **Limit maximum concurrent tasks to N** setting) is applied to the storage unit, not the whole Dell Data Domain system.

## Retention Lock

Together with Data Domain Retention Lock, Veeam Backup & Replication allows you to prohibit deletion of data from the Dell Data Domain deduplicating storage appliances by making that data temporarily immutable. This is required for improved security: immutability protects your data from loss as a result of malware activity or unplanned actions.

You can enable the immutability feature and specify the immutable period when adding or editing the Dell Data Domain backup repository. For more information, see Configure Backup Repository Settings. Once imposed, immutability prohibits deletion of data from the deduplicating storage appliance until the immutability expiration date comes. Note that if you enable immutability and Veeam Backup & Replication does not start a new backup chain and still continues the chain, the whole backup chain is marked as immutable. Once you disable immutability, newly created backups are not marked as immutable.

In many respects, Dell Data Domain with enabled immutability works similarly to the hardened repository. For more information on how repositories with immutability operate as a part of a scale-out backup repository, see Hardened Repository as Performance Extent. For more information on how retention periods set in jobs correlate with retention periods set in the deduplicating storage appliance, see Retention Scenarios.

## Accelerated Restore of Entire VM

To speed up entire VM restore on Dell Data Domain, Veeam Backup & Replication uses the mechanism of sequential data reading from backups and parallel VM disks restore.

Dell Data Domain storage systems are optimized for sequential I/O operations. However, data blocks of VM disks in backup files are stored not sequentially, but in the random order. If data blocks of VM disks are read at random, the restore performance from backups on Dell Data Domain degrades.

To accelerate the restore process, Veeam Backup & Replication creates a map of data blocks in backup files. It uses the created map to read data blocks of VM disks from backup files sequentially, as they reside on disk. Veeam Backup & Replication writes data blocks to target in the order in which they come from the target Veeam Data Mover, restoring several VM disks in parallel.

This accelerated restore mechanism is enabled by default, and is used for the entire VM restore scenario.

> **NOTE**
>
> To further accelerate the process of entire VM restore, Veeam Backup & Replication reads VM data from Dell Data Domain in multiple threads.

# How Accelerated Restore Works

Entire VM restore from backups on Dell Data Domain is performed in the following way:

1. Veeam Backup & Replication opens all backup files in the backup chain, reads metadata from these backup files and caches this metadata on the backup proxy that is assigned for the restore task.
2. Veeam Backup & Replication uses the cached metadata to build a map of data blocks. The map contains references to VM data blocks, sorted by VM disks.
3. Every VM disk is processed in a separate task. For each task, Veeam Backup & Replication starts one Veeam Data Mover on the backup proxy. Additionaly, Veeam Backup & Replication starts Data Mover on the gateway server, but one Veeam Data Mover for all tasks.

   Veeam Data Mover on the gateway server reads data blocks of VM disks from the backup repository sequentially, as these blocks reside on disk. Then Veeam Data Mover on the gateway server transfers data blocks to Data Movers on the backup proxy. On the backup proxy, these data blocks are put to the buffer.

4. Data blocks are written to target in the order in which they come from the target Veeam Data Mover.

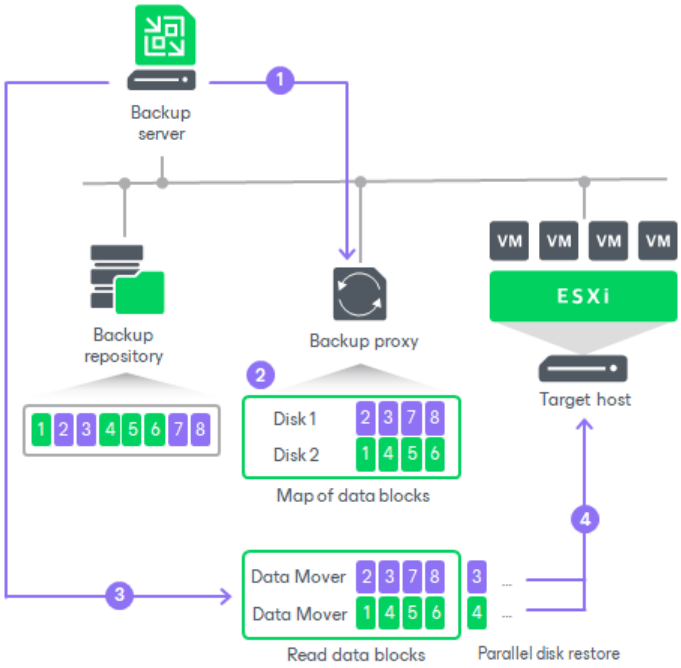


# Backup Proxy for Accelerated Restore

Veeam Backup & Replication restores all disks of a VM through one backup proxy. If you instruct Veeam Backup & Replication to select a backup proxy for the restore task automatically, it picks the least loaded backup proxy in the backup infrastructure. If you assign a backup proxy explicitly, Veeam Backup & Replication uses the selected backup proxy.

For every VM disk, Veeam Backup & Replication starts a separate Veeam Data Mover on the backup proxy. For example, if you restore a VM with 10 disks, Veeam Backup & Replication starts 10 Veeam Data Movers on the backup proxy.

The backup proxy assigned for the entire VM restore task must have enough RAM resources to be able to restore VM disks in parallel. For every VM disk, 200 MB of RAM is required. The total amount of required RAM resources is calculated by the following formula:

```
Total amount of RAM = Number of VM disks * 200 MB
```

Before starting the restore process, Veeam Backup & Replication checks the amount of RAM resources on the backup proxy. If the backup proxy does not have enough RAM resources, Veeam Backup & Replication displays a warning in the job session details and automatically fails over to a regular VM disks processing mode (data of VM disks is read at random and VM disks are restored sequentially).

# Limitations for Accelerated Restore

The accelerated restore of entire VM has the following limitations:

- Accelerated restore works on Dell Data Domain systems with DD Boost.
- If you restore a VM with dynamically expanding disks, the restore process may be slow.
- If you restore a VM using the Network transport mode, the number of VM disks restored in parallel cannot exceed the number of allowed connections to an ESXi host.

- If Dell Data Domain is added as an extent to a scale-out backup repository, you must set the backup file placement policy to Locality. If the backup file placement policy is set to Performance, parallel VM disk restore will be disabled.

# ExaGrid

You can use ExaGrid appliances as backup repositories.

# Adaptive Deduplication

ExaGrid uses adaptive deduplication. Data deduplication is performed on the target storage system. After VM data is written to disk, ExaGrid analyses bytes in the newly transferred data portions. ExaGrid compares versions of data over time and stores only the differences to disk.

ExaGrid deduplicates data at the storage level. Identical data is detected throughout the whole storage system, which increases the deduplication ratio.

# Backup Infrastructure

To communicate with ExaGrid, Veeam Backup & Replication uses Veeam Data Movers. They connect to each other, process data and transfer it over LAN or WAN.

Veeam Backup & Replication uses the following Veeam Data Movers:

- Veeam Data Mover on the ExaGrid appliance
- Veeam Data Mover on the VMware backup proxy

ExaGrid does not host Veeam Data Mover permanently. When any task addresses an ExaGrid storage, Veeam Backup & Replication deploys and starts the Veeam Data Mover on the ExaGrid appliance.

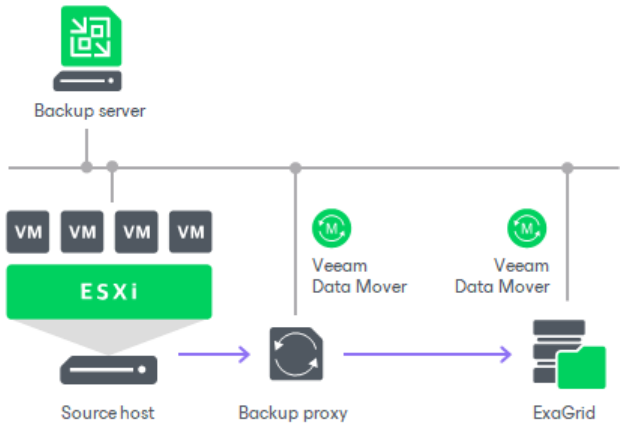


# Requirements and Recommendations for ExaGrid

To perform backup to ExaGrid appliances, it is recommended to configure backup repositories and jobs in the following way:

**Backup repositories**

On the ExaGrid side, do the following:

1. Create a new ExaGrid user account or update an existing one:
   - Use only lower case symbols for the ExaGrid user name.
   - Make sure that the ExaGrid user has the *Backup Only* role or another role with higher permissions.
2. Create a share on the ExaGrid appliance:
   - The share must be located in the home directory of the user you set up in the previous step.
   - Enable the **ExaGrid-Veeam Accelerated Data Mover transport** option for the created share.
   - Leave default compression and deduplication settings for the share.

   For more information, see the ExaGrid documentation.
3. Verify that the user account is added to the Veeam User Access Policy.

On the Veeam Backup & Replication side, do the following:

1. Add credentials for the ExaGrid user account you created in the previous steps:
   - When you add ExaGrid servers to the Veeam backup infrastructure, and you use the UPN format for an Active Directory account user name (for example, john.doe@domain.local), make sure you enter the user name in lowercase letters only.
   - Do not select **Elevate account privileges automatically** check box when setting up credentials for authentication to the Exagrid appliance.

   For more information, see Linux accounts.
2. Add the ExaGrid appliance as a managed server using the created credentials. For more information, see Adding Linux Servers.
3. Configure an ExaGrid backup repository:
   - When adding the repository, make sure the path to the share contains the user home directory.
   - Set the **Limit maximum concurrent tasks to N** option to 10 tasks. This limit can be tuned up or down with assistance from ExaGrid Customer Support.

### Backup Jobs

Configure backup job settings in the following way:

1. Use the forward incremental backup method.
2. Enable synthetic full backups and schedule them to run on a weekly basis.
3. Enable active full backups and schedule them to run on a monthly basis.

NOTE

Consider the following:

- Do not create multiple backup repositories directed at the same folder/path on the same device.
- We recommend against enabling encryption for the jobs targeted at the deduplication storage appliance. Encryption has a negative effect on the deduplication ratio. For more information, see Data Encryption.
- Since Veeam Backup & Replication version 12, deduplicating storage appliances use the TLS connection. You can disable the TLS connection with a registry value for the deduplicating storage appliances that do not support the TLS connection. For more information, see this Veeam KB article.

For more information and recommendations on working with ExaGrid, see this Veeam KB article.

# HPE StoreOnce

You can use HPE StoreOnce storage appliances as backup repositories.

To work with HPE StoreOnce, Veeam Backup & Replication leverages the HPE StoreOnce Catalyst technology and two HPE StoreOnce components:

- **HPE StoreOnce Catalyst agent**. The HPE StoreOnce Catalyst agent is a component of the HPE StoreOnce Catalyst software. The HPE StoreOnce Catalyst agent is embedded into the Veeam Data Mover setup. When you add a Microsoft Windows server to the backup infrastructure, the HPE StoreOnce Catalyst agent is automatically installed on the added server together with Veeam Data Mover.
- **HPE StoreOnce appliance**. The HPE StoreOnce appliance is an HPE StoreOnce storage system on which Catalyst stores are created.

# Backup Infrastructure

To communicate with HPE StoreOnce, Veeam Backup & Replication uses two Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the VMware backup proxy
- Veeam Data Mover on the gateway server

The HPE StoreOnce storage cannot host Veeam Data Mover. For this reason, to communicate with the HPE StoreOnce storage, you need to deploy a gateway server. Veeam Backup & Replication will automatically deploy Veeam Data Mover on this gateway server. For more information, see Gateway Servers. For communicating with the HPE StoreOnce storage appliances, the gateway server must run a 64-bit Microsoft Windows version.

When any job addresses the backup repository, Veeam Data Mover on the gateway server establishes a connection with Veeam Data Mover on the VMware backup proxy, enabling efficient data transfer over LAN or WAN.

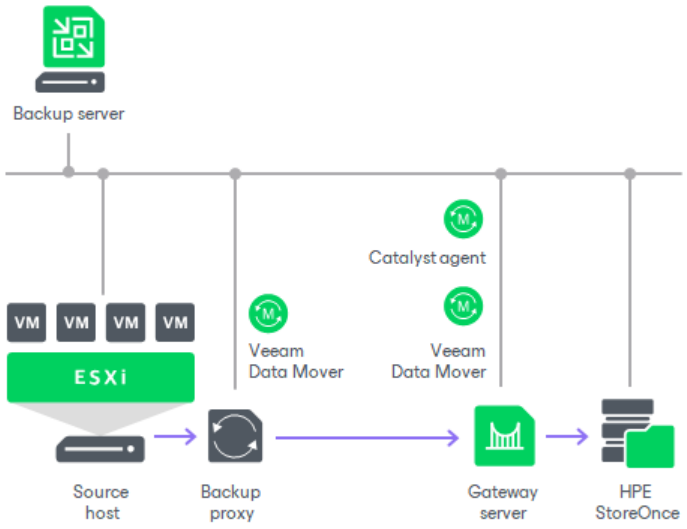


The gateway server is selected when you assign a backup repository role to the HPE StoreOnce appliance. You can define the gateway server explicitly or instruct Veeam Backup & Replication to select it automatically.

> **TIP**
>
> For work with HPE StoreOnce, Veeam Backup & Replication uses the Catalyst agent installed on the gateway server. If you want to reduce the load on the network between the source and target side, assign the gateway server role to a machine on the source side, closer to the VMware backup proxy.

# Supported Protocols

Veeam Backup & Replication supports HPE StoreOnce storage systems working over the following protocols:

- TCP/IP protocol: Veeam Backup & Replication communicates with the HPE StoreOnce appliance by sending commands over the LAN.
- Fibre Channel protocol: Veeam Backup & Replication communicates with the HPE StoreOnce appliance by sending SCSI commands over Fibre Channel.

  Data processing over Fibre Channel (FC) connectivity enables local area network-free backup to HP StoreOnce, eliminates the load from backup activities and increases availability of LAN resources to production workloads.

# Considerations and Limitations for HPE StoreOnce

## General

If you plan to use HPE StoreOnce as a backup repository for jobs other than file backup jobs, object storage backup jobs and Veeam Plug-in for Enterprise Application jobs, consider the following recommendations and limitations. These requirements and limitations apply only if you use HPE StoreOnce in the integration mode, not the shared folder mode.

- Check that HPE StoreOnce that you plan to use is supported. For more information, see Backup Target.
- A Catalyst Store cannot be added to multiple backup servers.

- When you create a job targeted at HPE StoreOnce, Veeam Backup & Replication will offer you to switch to optimized job settings and use the 4 MB size of data block for VM data processing. It is recommended that you use optimized job settings. Large data blocks produce a smaller metadata table that requires less memory and CPU resources to process.
- The HPE StoreOnce backup repository usually works in the **Use per-machine backup files** mode. For more information, see Backup Chain Formats. Note that there are cases when the HPE StoreOnce backup repository can store per-machine backups and non per-machine backups. For example, when backups are evacuated to this repository from other extents in a SOBR. For more information, see Evacuating Backups from Performance Extents. Note that restore from such backups can last longer than from per-machine backups.
- We recommend against enabling encryption for the jobs targeted at the deduplication storage appliance. Encryption has a negative effect on the deduplication ratio. For more information, see Data Encryption.
- HPE StoreOnce does not support the reverse incremental backup method.
- For backup jobs, HPE StoreOnce does not support the forever forward incremental backup method. When creating a backup job, you must enable synthetic or active full backups. For more information on how to enable such backups, see Backup Settings.
- The HPE StoreOnce backup repository does not support the **Defragment and compact full backup file** option.
- You cannot use HPE StoreOnce backup repositories as sources or targets for file copy jobs.
- You cannot copy backup files (VBK, VIB and VRB) manually to the HPE StoreOnce backup repository. To copy such files, use backup copy jobs or evacuate backups if you use a scale-out backup repository.
- You cannot use the HPE StoreOnce backup repository as a cloud repository hosted behind a Cloud Connect Gateway server.
- To optimize data transfer between two HPE StoreOnce repositories, use backup copy jobs for HPE StoreOnce repositories. For more information on how to create jobs and recommendations for them, see Creating Backup Copy Jobs for HPE StoreOnce Repositories.
- You can use HPE Cloud Bank Storage only as a target for backup copy jobs for HPE StoreOnce repositories. Veeam Backup & Replication supports HPE Cloud Bank Storage for HPE StoreOnce software version 4.3.2 or later.
- The HPE StoreOnce Cloud Bank stores do not work with object-versioned buckets. Immutability for these stores is enforced by the HPE StoreOnce solution, not the Cloud Bank object storage. As a result, the object storage administrator can delete data stored in the object storage buckets. To mitigate this risk, see the HPE best practice documentation.
- Veeam Backup & Replication supports fixed block chunking functionality for HPE StoreOnce software version 4.3.2 or later. To be able to use this functionality in Veeam Backup & Replication, check that the Align backup file data blocks option is enabled in the repository settings.
- HPE StoreOnce has a limit on the number of concurrently opened files. For more information, see Storage models and number of restore points.

For more information and recommendations on working with HPE StoreOnce, see this Veeam KB article.

## HPE StoreOnce and Unstructured Data Backup

If you plan to use HPE StoreOnce storage appliances for unstructured data backup, consider the following recommendations for optimal performance:

- A StoreOnce system can have multiple Catalyst stores, and large backup loads (exceeding 1PB) should be spread across more than one Catalyst store on the same StoreOnce system.

- Do not include Catalyst stores in a SOBR intended for unstructured data backups. This will reduce the global deduplication of the StoreOnce system.

## HPE StoreOnce and Veeam Plug-ins for Enterprise Applications

If you plan to use HPE StoreOnce as a backup repository for Veeam Plug-in for Oracle RMAN or Veeam Plug-in for SAP HANA, the total number of stored files (data and metadata) must not exceed 3,000,000 per Catalyst store. If necessary, multiple Catalyst stores may be created on the same StoreOnce system.

## HPE StoreOnce and Immutability

If you plan to enable immutability for the HPE StoreOnce deduplicating storage appliance, consider the following recommendations and limitations:

- HPE StoreOnce software version must be 4.3.2 or later.
- Dual Authorization must be enabled on HPE StoreOnce. For more information on Dual Authorization, see HPE Support Center.
- Check that the **Maximum ISV Controlled Data Retention** setting in HPE StoreOnce is greater than each of the following settings configured in Veeam Backup & Replication:
  - The immutability period (the Make recent backups immutable for setting).
  - The long-term retention period (the Keep certain backups longer for archival purposes setting).

  If a period configured in Veeam Backup & Replication is greater than the **Maximum ISV Controlled Data Retention** setting, Veeam Backup & Replication keeps backups immutable for the period configured in the Veeam Backup & Replication settings but resets the expiration of the backup immutability flag during each job run. While the actual number of days that the backup must be kept immutable is greater than **Maximum ISV Controlled Data Retention**, Veeam Backup & Replication sets the expiration period equal to **Maximum ISV Controlled Data Retention**.
- To use the immutability feature for regular backup copy jobs (without HPE Catalyst Copy), enable the GFS retention policy. For more information, see Long-Term Retention Policy (GFS) for Backup Copy Jobs.
- [For unstructured data backups] The immutability lock period in HPE StoreOnce is set by the backup repository immutability settings. For an example of how this mechanism works, see **Case 8** in the Unstructured Data Backup Retention Scenarios section.
- Immutability applies only if you use HPE StoreOnce in the integration mode, not the shared folder mode.
- Immutability is supported only for forward incremental backup chains. Once a backup file becomes immutable, it can be merged or deleted only when the immutability time period expires.

## HPE StoreOnce Immutability and Veeam Agents

For more information on how immutability works with Veeam Agents, see the Backup to Deduplicating Storage Appliances section in the Veeam Agent Management Guide.

## Storage models and number of restore points

HPE StoreOnce has a limit on the number of concurrently opened files. Due to this limit, the maximum length of backup chains (chains that contain one full backup and a set of subsequent incremental backups) on HPE StoreOnce is also limited and depends on the particular storage model.

For the unstructured data backup, the limitation on the backup chain length cannot be applied. In this case, we recommend that you use a storage appliance that can open at least 64 file sessions.

| Storage model | Maximum number of restore points per backup chain |
|---|---|
| **VSA** | |
| VSA Gen3 | 7 |
| VSA Gen4 | 7 to 14 (varies depending on the amount of available memory) |
| **Proliant Gen7** | |
| 6200 | 14 (per node) |
| **Proliant Gen8** | |
| 2700 | 7 |
| 2900 | 14 |
| 4500 | 14 |
| 4700 | 14 |
| 4900 | 28 |
| 6500 | 28 (per node) |
| **Proliant Gen9** | |
| 3100 | 7 |
| 3500 | 14 |
| 5100 | 21 |
| 5500 | 35 |
| 6600 | 42 (per node) |

| Storage model | Maximum number of restore points per backup chain |
|---|---|
| **Proliant Gen10** | |
| 3620 | 14 |
| 3640 | 14 |
| 5200 | 28 |
| 5250 | 28 |
| 5650 | 42 |
| **Proliant Gen10 Plus** | |
| 3660 | 21 |
| 5260 | 28 |
| 5660 | 42 |

# Several Backup Repositories on HPE StoreOnce

You can configure several backup repositories on one HPE StoreOnce appliance and associate them with different gateway servers.

Consider the following:

- If you configure several backup repositories on HPE StoreOnce and add them as extents to a scale-out backup repository, make sure that all backup files from one backup chain are stored on one extent. If backup files from one backup chain are stored to different extents, the performance of transformation processes will be lower. For more information about transformation performance, see this Veeam blog post.
- HPE StoreOnce has a limit on the number of opened files that applies to the whole appliance. Tasks targeted at different backup repositories on HPE StoreOnce and run in parallel will equally share this limit.
- For HPE StoreOnce working over Fibre Channel, there is a limitation on the number of connections from one host. If you connect several backup repositories to one gateway, backup repositories will compete for connections.
- Deduplication on HPE StoreOnce works within the limits of one object store.
- If your VMs contain similar data, it is recommended to create backup repositories on a single HPE StoreOnce Catalyst Store. This minimizes backup job duration and reduces disk space used for backups. For details, see this Veeam KB article.

## Operational Modes

Depending on the storage configuration and type of the backup target, HPE StoreOnce can work in the following modes:

- Source-side deduplication
- Target-side deduplication
- Shared folder mode

# Source-Side Data Deduplication

HPE StoreOnce performs source-side deduplication if the backup target meets the following requirements:

- You have a Catalyst license installed on HPE StoreOnce.
- You use a Catalyst store as a backup repository.
- The Catalyst store is configured to work in the Low Bandwidth mode (Primary Transfer Policy).
- The HPE StoreOnce Catalyst is added to the backup repository as a deduplicating storage appliance, not as a shared folder.

To deduplicate data on the source side, HPE StoreOnce uses the HPE StoreOnce Catalyst agent. The HPE StoreOnce Catalyst agent is a component of the HPE StoreOnce Catalyst software. It is installed on the gateway server communicating with the HPE StoreOnce appliance.

HPE StoreOnce deduplicates data on the source side, before writing it to target:

1. During the backup job session, HPE StoreOnce analyzes data incoming to the HPE StoreOnce appliance in chunks and computes a hash value for every data chunk. Hash values are stored in an index on disk.
2. The HPE StoreOnce Catalyst agent calculates hash values for data chunks in a new data flow and sends these hash values to target.
3. HPE StoreOnce identifies which data blocks are already saved on disk and communicates this information to the HPE StoreOnce Catalyst agent. The HPE StoreOnce Catalyst agent sends only unique data blocks to target.

As a result, the load on the network reduces, the backup job performance improves, and you can save on disk space.

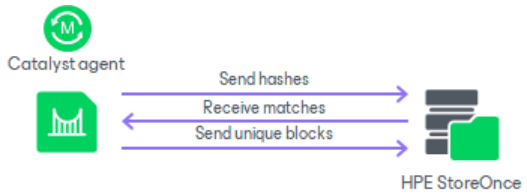


# Target-Side Data Deduplication

HPE StoreOnce performs target-side deduplication if the backup target is configured in the following way:

- For a Catalyst store:
  - The Catalyst store works in the High Bandwidth mode (Primary Transfer Policy is set to High Bandwidth).

- The Catalyst license is installed on the HPE StoreOnce (required).
- The Catalyst store is added to the backup repository as a deduplicating storage appliance, not as a shared folder.

- For an SMB (CIFS) store:
  - The Catalyst license is not required.
  - The SMB (CIFS) store is added as a shared folder backup repository to the backup infrastructure.

  For more information about working with SMB (CIFS) stores, see Shared Folder Mode.

HPE StoreOnce deduplicates data on the target side, after the data is transported to HPE StoreOnce:

1. HPE StoreOnce analyzes data incoming to the HPE StoreOnce appliance in chunks and creates a hash value for every data chunk. Hash values are stored in an index on the target side.
2. HPE StoreOnce analyzes VM data transported to target and replaces identical data chunks with references to data chunks that are already saved on disk.

As a result, only new data chunks are written to disk, which helps save on disk space.

# Shared Folder Mode

If you do not have an HPE StoreOnce Catalyst license, you can add the HPE StoreOnce appliance as a shared folder backup repository. In this mode, HPE StoreOnce will perform target-side deduplication.

If you work with HPE StoreOnce in the shared folder mode, the performance of backup jobs and transformation processes is usually lower (in comparison with the integration mode, when HPE StoreOnce is added as a deduplicating storage appliance).

## HPE StoreOnce Supported Features

The HPE StoreOnce Catalyst technology offers a set of features for advanced data processing. Veeam Backup & Replication supports the following features.

# Synthetic Full Backups

HPE StoreOnce Catalyst improves synthetic full backup file creation and transformation performance. When Veeam Backup & Replication creates or transforms a synthetic full backup, HPE StoreOnce does not physically copy data between the existing backup chain and the target full backup file. Instead, it performs a metadata-only operation — updates pointers to existing data blocks on the storage device. As a result, the operation completes much faster. This mechanism helps improve performance of primary backup jobs and backup copy jobs that are scheduled to create periodic archive full backups (GFS).

# Accelerated vPower-Enabled Operations

If HPE StoreOnce Catalyst is integrated with Veeam Backup & Replication, it helps to benefit from improved performance of vPower-enabled operations — Instant Recovery, SureBackup and On-Demand Sandbox — from backups residing on HPE StoreOnce. To get the maximum vPower performance, HPE StoreOnce must be running firmware version 3.15.1 or later.

# Accelerated Data Recovery

Integration with HPE StoreOnce improves data recovery performance for different restore scenarios: Instant Recovery, file-level recovery and application items recovery with Veeam Explorers.

# WAN-based Catalyst Store Support

Veeam Backup & Replication provides advanced support for WAN-based HPE Catalyst stores. If a WAN connection to HPE StoreOnce is weak, you can instruct Veeam Backup & Replication to compress VM data and calculate checksums for data blocks going from the source side to HPE StoreOnce.

# HPE StoreOnce Replication

HPE StoreOnce replication improves copying data between two HPE StoreOnce backup repositories. For more information on copying, see Creating Backup Copy Jobs for HPE StoreOnce Repositories.

# Independent Software Vendor (ISV) Controlled Data Immutability (ISV-DI)

Together with HPE StoreOnce Independent Software Vendor (ISV) Controlled Data Immutability (ISV-DI), Veeam Backup & Replication allows you to prohibit deletion of data from the HPE StoreOnce deduplicating storage appliances by making that data temporarily immutable. This is required for improved security: immutability protects your data from loss as a result of malware activity or unplanned actions.

You can enable the immutability feature and specify the immutable period when adding or editing the HPE StoreOnce backup repository. For more information, see Configure Backup Repository Settings. Once imposed, immutability prohibits deletion of data from the deduplicating storage appliance until the immutability expiration date comes. Note that if you enable immutability and Veeam Backup & Replication does not start a new backup chain and still continues the chain, the whole backup chain is marked as immutable. Once you disable immutability, newly created backups are not marked as immutable.

In many respects, HPE StoreOnce with enabled immutability works similarly to the hardened repository. For more information on how repositories with immutability operate as a part of a scale-out backup repository, see Hardened Repository as Performance Extent. For more information on how retention periods set in jobs correlate with retention periods set in the deduplicating storage appliance, see How Immutability Works.

# Fixed Block Chunking

Together with fixed block chunking, Veeam Backup & Replication improves performance when creating incremental and synthetic full backup files. Fixed block chunking improves performance up to 4-5 times compared to variable chunking,

Veeam Backup & Replication supports fixed block chunking functionality for HPE StoreOnce software version 4.3.2 or later. To be able to use this functionality in Veeam Backup & Replication, check that the Align backup file data blocks option is enabled in the repository settings.

Fixed block chunking is supported for the following types of backups:

- Backups of VMware vSphere virtual machines created by Veeam Backup & Replication
- Backups of Microsoft Hyper-V virtual machines created by Veeam Backup & Replication
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for Oracle Solaris or Veeam Agent for IBM AIX
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV
- Backups of oVirt VMs created by Veeam Backup for OLVM and RHV

For more information on issue that can occur, see this Veeam KB article.

# Cloud Bank Storage

Veeam Backup & Replication supports using HPE Cloud Bank Storage — an extension of HPE StoreOnce Catalyst that uses external object storage to store backup data. It is possible to use HPE Cloud Bank Storage as a target for backup copy jobs for HPE StoreOnce repositories. This helps reduce long-term retention costs.

Veeam Backup & Replication supports HPE Cloud Bank Storage for HPE StoreOnce software version 4.3.2 or later.

## Accelerated Restore of Entire VM

To speed up entire VM restore on HPE StoreOnce, Veeam Backup & Replication uses the mechanism of sequential data reading from backups and parallel VM disks restore.

HPE StoreOnce storage systems are optimized for sequential I/O operations. However, data blocks of VM disks in backup files are stored not sequentially, but in the random order. If data blocks of VM disks are read at random, the restore performance from backups on HPE StoreOnce degrades.

To accelerate the restore process, Veeam Backup & Replication creates a map of data blocks in backup files. It uses the created map to read data blocks of VM disks from backup files sequentially, as they reside on disk. Veeam Backup & Replication writes data blocks to target in the order in which they come from the target Veeam Data Mover, restoring several VM disks in parallel.

This accelerated restore mechanism is enabled by default, and is used for the entire VM restore scenario.

> **NOTE**
>
> To further accelerate the process of entire VM restore, Veeam Backup & Replication reads VM data from HPE StoreOnce in multiple threads.

# How Accelerated Restore Works

Entire VM restore from backups on HPE StoreOnce is performed in the following way:

1. Veeam Backup & Replication opens all backup files in the backup chain, reads metadata from these backup files and caches this metadata on the backup proxy that is assigned for the restore task.
2. Veeam Backup & Replication uses the cached metadata to build a map of data blocks. The map contains references to VM data blocks, sorted by VM disks.

3. Every VM disk is processed in a separate task. For each task, Veeam Backup & Replication starts one Veeam Data Mover on the backup proxy. Additionaly, Veeam Backup & Replication starts Data Mover on the gateway server, but one Veeam Data Mover for all tasks.

   Veeam Data Mover on the gateway server reads data blocks of VM disks from the backup repository sequentially, as these blocks reside on disk. Then Veeam Data Mover on the gateway server transfers data blocks to Data Movers on the backup proxy. On the backup proxy, these data blocks are put to the buffer.

4. Data blocks are written to target in the order in which they come from the target Veeam Data Mover.

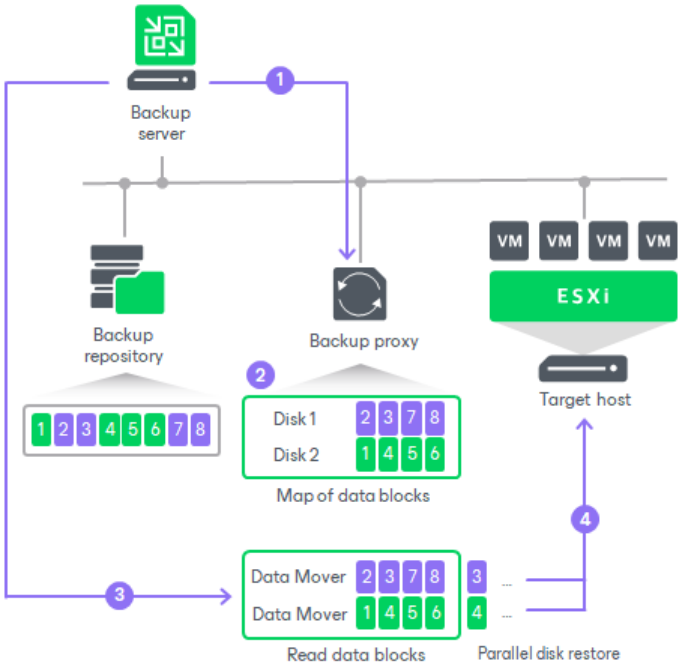


# Backup Proxy for Accelerated Restore

Veeam Backup & Replication restores all disks of a VM through one backup proxy. If you instruct Veeam Backup & Replication to select a backup proxy for the restore task automatically, it picks the least loaded backup proxy in the backup infrastructure. If you assign a backup proxy explicitly, Veeam Backup & Replication uses the selected backup proxy.

For every VM disk, Veeam Backup & Replication starts a separate Veeam Data Mover on the backup proxy. For example, if you restore a VM with 10 disks, Veeam Backup & Replication starts 10 Veeam Data Movers on the backup proxy.

The backup proxy assigned for the entire VM restore task must have enough RAM resources to be able to restore VM disks in parallel. For every VM disk, 200 MB of RAM is required. The total amount of required RAM resources is calculated by the following formula:

```
Total amount of RAM = Number of VM disks * 200 MB
```

Before starting the restore process, Veeam Backup & Replication checks the amount of RAM resources on the backup proxy. If the backup proxy does not have enough RAM resources, Veeam Backup & Replication displays a warning in the job session details and automatically fails over to a regular VM disks processing mode (data of VM disks is read at random and VM disks are restored sequentially).

# Limitations for Accelerated Restore

The accelerated restore of entire VM has the following limitations:

- Accelerated restore works on HPE StoreOnce systems with Catalyst.
- If you restore a VM with dynamically expanding disks, the restore process may be slow.
- If you restore a VM using the Network transport mode, the number of VM disks restored in parallel cannot exceed the number of allowed connections to an ESXi host.
- If HPE StoreOnce is added as an extent to a scale-out backup repository, you must set the backup file placement policy to Locality. If the backup file placement policy is set to Performance, parallel VM disk restore will be disabled.

# Quantum DXi

You can use Quantum DXi appliances as backup repositories.

# Quantum DXi Deduplication

Quantum DXi appliances use Quantum's patented data deduplication technology. During backup data transfer, Quantum analyses blocks in a data stream. Instead of copying redundant data blocks Quantum uses reference pointers to existing blocks on the storage device.

# Backup Infrastructure

To communicate with Quantum DXi, Veeam Backup & Replication uses Veeam Data Movers. They connect to each other, process data and transfer it over LAN or WAN.

Veeam Backup & Replication uses the following Veeam Data Movers:

- Veeam Data Mover on the Quantum DXi appliance
- Veeam Data Mover on the VMware backup proxy

Quantum DXi does not host Veeam Data Mover permanently. When any task addresses a Quantum DXi storage, Veeam Backup & Replication deploys and starts the Veeam Data Mover on the Quantum DXi appliance.

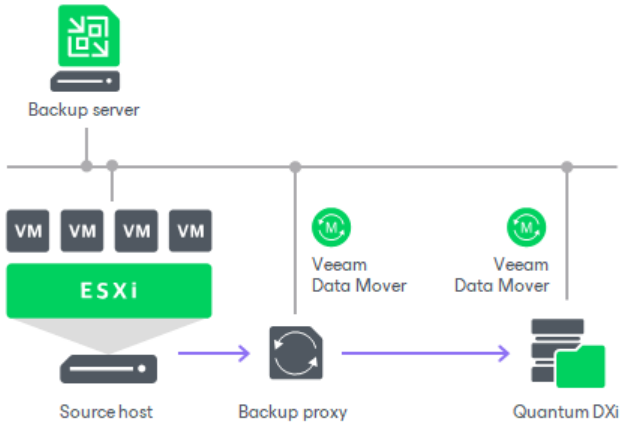

> **NOTE**
>
> Since Veeam Backup & Replication version 12, deduplicating storage appliances use the TLS connection. You can disable the TLS connection with a registry value for the deduplicating storage appliances that do not support the TLS connection. For more information, see this Veeam KB article.

For more information and recommendations on working with Quantum DXi, see this Veeam KB article.

# Backup Repository Configuration

Configure Quantum DXi backup repositories in the following way:

1. Create at least one share on each Quantum DXi appliance. Enable VDMS (Veeam Data Mover Service) for the created share. Leave default compression and deduplication settings for the share. For more information, see Quantum DXi documentation.

   > **NOTE**
   >
   > When adding credentials for Quantum DXi, use only lower case symbols for the user name.

2. Configure Quantum DXi backup repositories and point them at the created shares on each Quantum DXi appliance. Set the **Limit maximum concurrent tasks to N** option to 10 tasks. This limit can be tuned up or down with assistance from Quantum DXi Customer Support.

   When you add Quantum DXi servers to the Veeam backup infrastructure, make sure that you use the UPN format for an Active Directory account user name (for example, john.doe@domain.local), and enter the user name in lowercase letters only.

# Backup Job Configuration

For the backup jobs, use the forward incremental backup method. We recommend that you enable synthetic full and schedule them to run on a weekly basis. Also, enable active full backups and schedule them to run on a monthly basis. For more information, see Backup Settings. We recommend that you do not enable encryption for the jobs targeted at the deduplication storage appliance. Encryption has a negative effect on the deduplication ratio. For more information, see Job Encryption.

# Fujitsu ETERNUS CS800

You can use Fujitsu ETERNUS CS800 deduplicating storage appliances as backup repositories.

# Fujitsu ETERNUS CS800 Deduplication

Fujitsu ETERNUS CS800 appliances use Fujitsu ETERNUS CS800 data deduplication technology. During backup data transfer, Fujitsu analyses blocks in a data stream. Instead of copying redundant data blocks Fujitsu uses reference pointers to existing blocks on the storage device.

# Backup Infrastructure

To communicate with Fujitsu, Veeam Backup & Replication uses Veeam Data Movers. They connect to each other, process data and transfer it over LAN or WAN.

Veeam Backup & Replication uses the following Veeam Data Movers:

- Veeam Data Mover on the Fujitsu appliance
- Veeam Data Mover on the VMware backup proxy

Fujitsu does not host Veeam Data Mover permanently. When any task addresses a Fujitsu storage, Veeam Backup & Replication deploys and starts the Veeam Data Mover on the Fujitsu appliance.

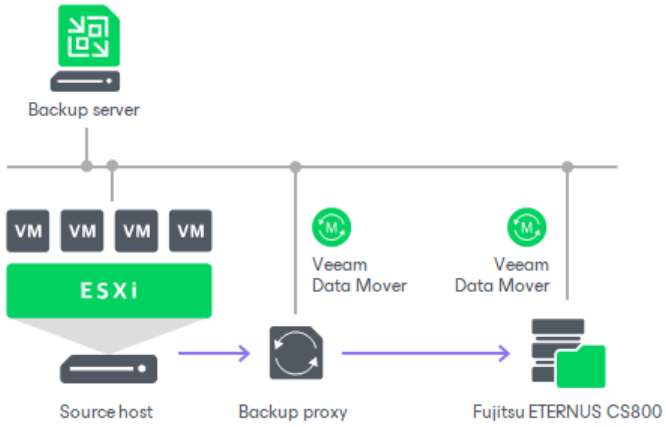


**NOTE**

Since Veeam Backup & Replication version 12, deduplicating storage appliances use the TLS connection. You can disable the TLS connection with a registry value for the deduplicating storage appliances that do not support the TLS connection. For more information, see this Veeam KB article.

# Backup Repository Configuration

Configure Fujitsu backup repositories in the following way:

1. Create at least one share on each Fujitsu appliance. Enable VDMS (Veeam Data Mover Service) for the created share. Leave default compression and deduplication settings for the share.

   **NOTE**

   When adding credentials for Fujitsu, use only lower case symbols for the user name.

2. Configure Fujitsu backup repositories and point them at the created shares on each Fujitsu appliance. Set the **Limit maximum concurrent tasks to N** option to 10 tasks. This limit can be tuned up or down with assistance from Fujitsu Customer Support.

   When you add Fujitsu servers to the Veeam backup infrastructure, make sure that you use the UPN format for an Active Directory account user name (for example, john.doe@domain.local), and enter the user name in lowercase letters only.

# Backup Job Configuration

For the backup jobs, use the forward incremental backup method. We recommend that you enable synthetic full and schedule them to run on a weekly basis. Also, enable active full backups and schedule them to run on a monthly basis. For more information, see Backup Settings. We recommend that you do not enable encryption for the jobs targeted at the deduplication storage appliance. Encryption has a negative effect on the deduplication ratio. For more information, see Job Encryption.

# Infinidat InfiniGuard

You can use Infinidat InfiniGuard appliances as backup repositories.

# Infinidat InfiniGuard Deduplication

Infinidat InfiniGuard appliances use Infinidat InfiniGuard data deduplication technology. During backup data transfer, Infinidat InfiniGuard analyses blocks in a data stream. Instead of copying redundant data blocks Infinidat InfiniGuard uses reference pointers to existing blocks on the storage device.

# Backup Infrastructure

To communicate with Infinidat InfiniGuard, Veeam Backup & Replication uses Veeam Data Movers. They connect to each other, process data and transfer it over LAN or WAN.

Veeam Backup & Replication uses the following Veeam Data Movers:

- Veeam Data Mover on the Infinidat InfiniGuard appliance
- Veeam Data Mover on the VMware backup proxy

Infinidat InfiniGuard does not host Veeam Data Mover permanently. When any task addresses an Infinidat InfiniGuard storage, Veeam Backup & Replication deploys and starts the Veeam Data Mover on the Infinidat InfiniGuard appliance.

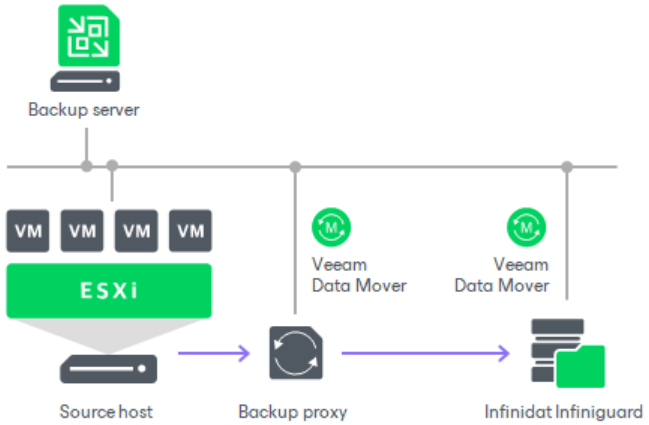


> **NOTE**
>
> Since Veeam Backup & Replication version 12, deduplicating storage appliances use the TLS connection. You can disable the TLS connection with a registry value for the deduplicating storage appliances that do not support the TLS connection. For more information, see this Veeam KB article.

# Backup Repository Configuration

Configure Infinidat InfiniGuard backup repositories in the following way:

1. Create at least one share on each Infinidat InfiniGuard appliance. Enable VDMS (Veeam Data Mover Service) for the created share. Leave default compression and deduplication settings for the share.

   > **NOTE**
   >
   > When adding credentials for Infinidat InfiniGuard, use only lower case symbols for the user name.

2. Configure Infinidat InfiniGuard backup repositories and point them at the created shares on each Infinidat InfiniGuard appliance. Set the **Limit maximum concurrent tasks to N** option to 10 tasks. This limit can be tuned up or down with assistance from Infinidat InfiniGuard Customer Support.

   When you add Infinidat InfiniGuard servers to the Veeam backup infrastructure, make sure that you use the UPN format for an Active Directory account user name (for example, john.doe@domain.local), and enter the user name in lowercase letters only.

# Backup Job Configuration

For the backup jobs, use the forward incremental backup method. We recommend that you enable synthetic full and schedule them to run on a weekly basis. Also, enable active full backups and schedule them to run on a monthly basis. For more information, see Backup Settings. We recommend that you do not enable encryption for the jobs targeted at the deduplication storage appliance. Encryption has a negative effect on the deduplication ratio. For more information, see Job Encryption.

# Adding Deduplicating Storage Appliances

This section describes how to add a deduplicating storage appliance as a backup repository.

Before you add a backup repository, check prerequisites. Then use the **New Backup Repository** wizard to add the backup repository.

## Before You Begin

Before you configure a backup repository, check the following prerequisites.

# Dell Data Domain

- Dell Data Domain must meet software and hardware requirements. For more information, see System Requirements.
- The DD Boost license must be installed on the Dell Data Domain system, DD Boost must be enabled and configured.
- The gateway server that you plan to use for work with Dell Data Domain must be added to the backup infrastructure.

If the Dell Data Domain storage system does not meet these requirements, you can add it as a SMB (CIFS) folder. In this case, Veeam Backup & Replication will not use the DD Boost technology to work with Dell Data Domain. For more information, see Dell Data Domain.

# ExaGrid

- ExaGrid must meet software and hardware requirements. For more information, see System Requirements.
- To use ExaGrid as a backup repository, you must configure an ExaGrid share in ExaGrid Manager. For more information on how to configure the share and repository settings, see ExaGrid.

# HPE StoreOnce

- HPE StoreOnce must meet software and hardware requirements. For more information, see System Requirements.

- The HPE StoreOnce Catalyst license must be installed on the HPE StoreOnce system.
- You must use a Catalyst store as a backup target.
- The gateway server that you plan to use for work with HPE StoreOnce system must be added to the backup infrastructure.
- The client account that you plan to use to connect to HPE StoreOnce must have access permissions on the Catalyst store where backup data will be kept.

If the HPE StoreOnce storage system does not meet these requirements, you can add it as a shared folder. In this case, Veeam Backup & Replication will perform target-side deduplication. For more information, see HPE StoreOnce.

# Quantum DXi

- Quantum DXi must meet software and hardware requirements. For more information, see System Requirements.
- To use Quantum DXi as a backup repository, you must configure a Quantum DXi share. For more information, see Quantum DXi.

# Fujitsu ETERNUS CS800

- Fujitsu ETERNUS CS800 must meet software requirements. For more information, see System Requirements.
- To use Fujitsu ETERNUS CS800 as a backup repository, you must configure a Fujitsu share. For more information, see Fujitsu ETERNUS CS800.

# Infinidat InfiniGuard

- Infinidat InfiniGuard must meet software requirements. For more information, see System Requirements.
- To use Infinidat InfiniGuard as a backup repository, you must configure a Infinidat InfiniGuard share. For more information, see Infinidat InfiniGuard.

# Storage Appliances and Tapes

Storage appliances that are used to store backup data in filer (CIFS/NFS) or block device mode (iSCSI/FC/SAS) are not supported if the backup data is offloaded to tapes and is no longer stored directly on the filer/block device (Hierarchical Storage Management with Tape tier).

To offload data to tapes, make sure that:

- All of the backup data is stored on the appliance altogether (that is, all of the backup chains are stored on the appliance as a whole and not scattered across multiple devices) and only copies are stored on tapes.
- These appliances emulate a tape system (VTL) as an access protocol for.

## Step 1. Launch New Backup Repository Wizard

To launch the **New Backup Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add Backup Repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Deduplicating Storage Appliance** and the type of the backup repository you want to add.

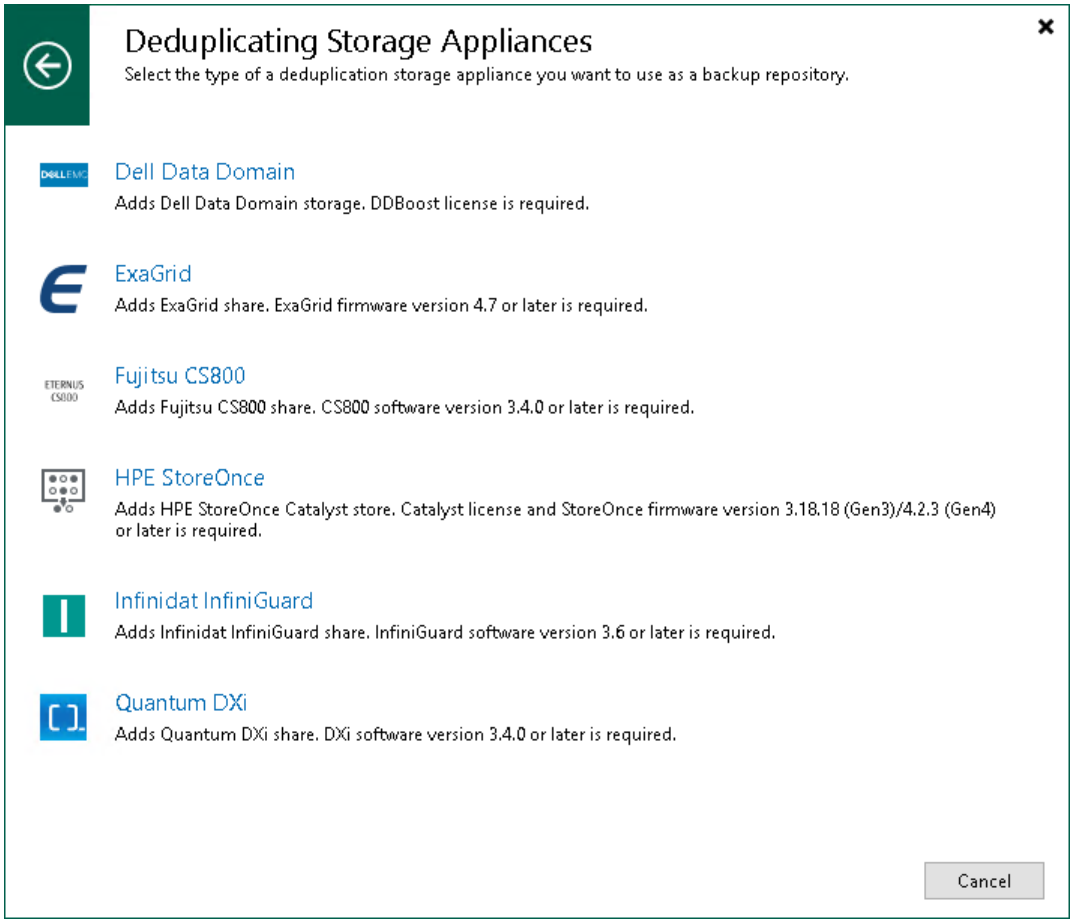

## Step 2. Specify Backup Repository Name and Description

At the **Name** step of the wizard, specify a name and description for the backup repository:

1. In the **Name** field, specify a name for the backup repository.
2. In the **Description** field, provide a description for future reference.

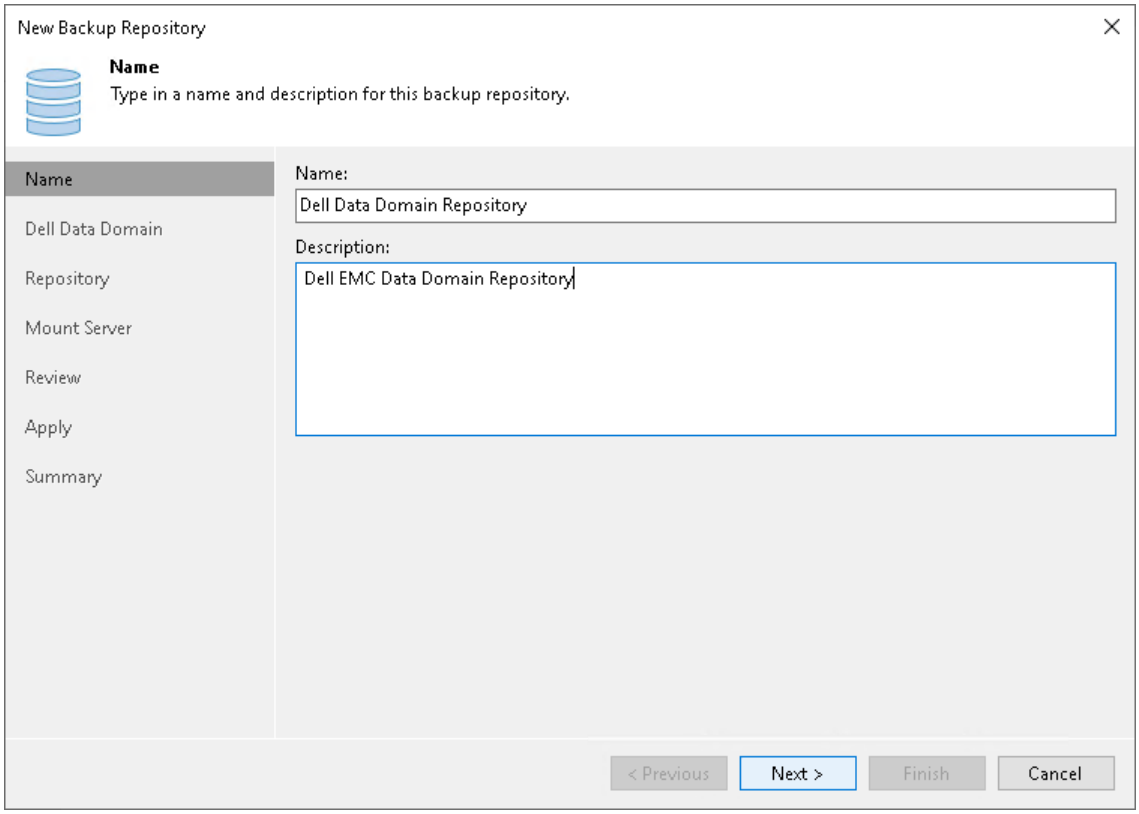

## Step 3. Specify Server Settings

Options that you can specify at the **Server** step of the wizard depend on the type of backup repository you are adding.

### Dell Data Domain

To configure settings for Dell Data Domain:

1. Specify connection settings for Dell Data Domain:
   - If Dell Data Domain works over TCP, in the **Type in Data Domain server name** field enter a full DNS name, or IPv4 or IPv6 address of the Dell Data Domain server. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
   - If Dell Data Domain works over Fibre Channel, select the **Use Fibre Channel (FC) connectivity** check box. In the **Type in Data Domain server name** field, enter a name of the Data Domain Fibre Channel server. To get the Data Domain Fibre Channel server name, in Data Domain System Manager open the **Data Management** > **DD Boost** > **Fibre Channel tab**.
2. In the **Credentials** field, specify credentials of the user account to connect to the Dell Data Domain server or Dell Data Domain Fibre Channel server. The user must have permissions described in section Permissions.

   If you have not set up credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Credentials Manager.
3. To use in-flight encryption between the backup proxy and Dell Data Domain, select the **Enable DDBoost encryption** check box and choose the encryption level — *Medium* or *High*.
4. In the **Gateway server** field, specify which gateway server you want to use:
   - If you want Veeam Backup & Replication to select a gateway server automatically, leave **Automatic selection**.
   - If you want to select servers that can be used as gateway servers explicitly, click **Choose** next to the **Gateway server** field. In the **Gateway Server** window, click **Use the following gateway servers only** and select servers. The servers must have a direct access to the Dell Data Domain appliance and must be located as close to the appliance as possible. Veeam Backup & Replication will choose the most suitable server.

   For more information on the gateway servers, their requirements and limitations, and how they are selected, see Gateway Servers.

> **IMPORTANT**
>
> If you connect to Dell Data Domain over Fibre Channel, you must explicitly define the gateway servers to communicate with Dell Data Domain. The servers you select must be added to the backup infrastructure and must have access to the Dell Data Domain appliance over Fibre Channel.

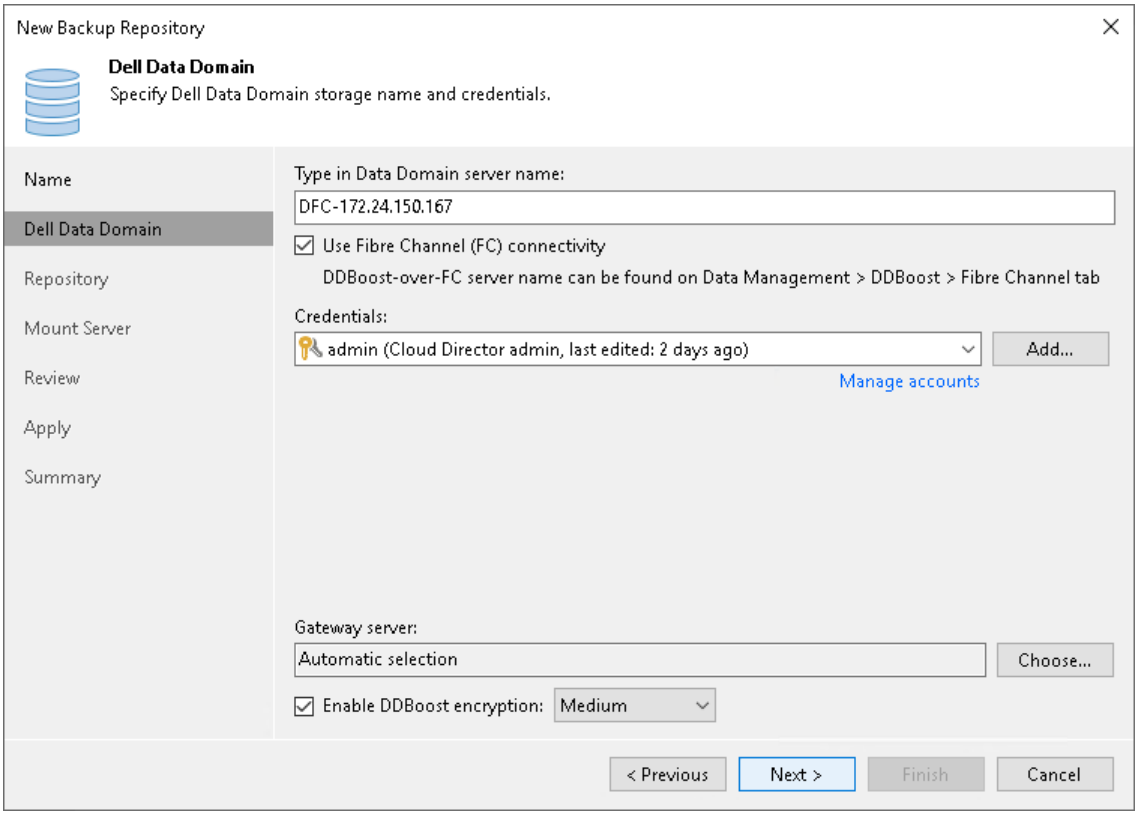


## ExaGrid, Quantum DXi, Fujitsu ETERNUS CS800 and Infinidat InfiniGuard

To configure settings for ExaGrid, Quantum DXi, Fujitsu ETERNUS CS800 or Infinidat InfiniGuard deduplicating appliance:

1. From the **Repository server** list, select an appliance that you want to use as a backup repository. The **Repository server** list contains only those servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, you can click **Add New** to open the **New Linux Server** wizard. For more information, see Virtualization Servers and Hosts.

2. Click **Populate** to see the appliance capacity and available free space.

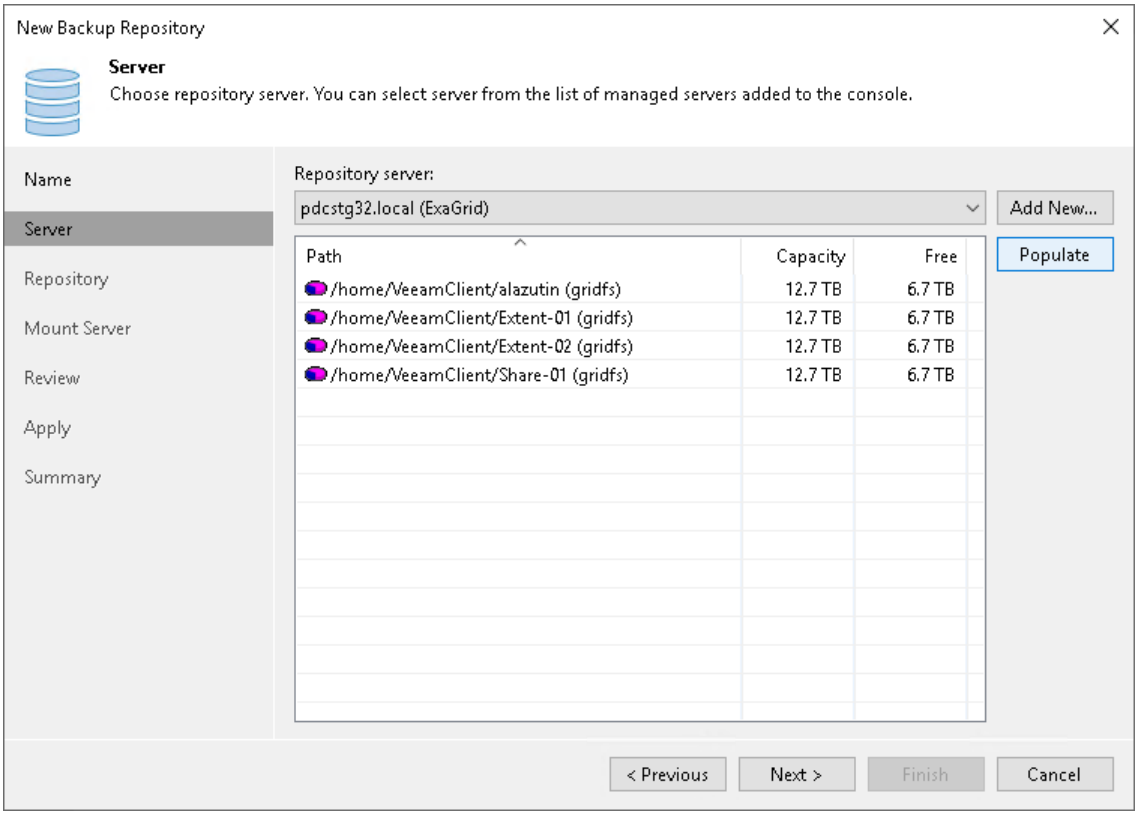


## HPE StoreOnce Deduplicating Appliance

To configure settings for HPE StoreOnce:

1. In the **Type in HPE StoreOnce server name** field, enter a full DNS name, or IPv4 or IPv6 address of the HPE StoreOnce appliance. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
2. If HPE StoreOnce works over Fibre Channel, select the **Use Fibre Channel (FC) connectivity** check box.
3. In the **Credentials** field, specify credentials of the client account to connect to the HPE StoreOnce appliance. The account must have permissions described in section Permissions.

   If you have not set up credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Credentials Manager.
4. In the **Gateway server** field, specify which gateway server you want to use:
   - If you want Veeam Backup & Replication to select a gateway server automatically, leave **Automatic selection**.
   - If you want to select servers that can be used as gateway servers explicitly, click **Choose** next to the **Gateway server** field. In the **Gateway Server** window, click **Use the following gateway servers only** and select servers. The servers must have a direct access to the HPE StoreOnce appliance and must be located as close to the appliance as possible. Veeam Backup & Replication will choose the most suitable server.

   For more information on the gateway servers, their requirements and limitations, and how they are selected, see Gateway Servers.

> **IMPORTANT**
>
> If you connect to HPE StoreOnce over Fibre Channel, you must explicitly define the gateway servers to communicate with HPE StoreOnce appliance. The servers you select must be added to the backup infrastructure and must have access to the HPE StoreOnce appliance over Fibre Channel.

5. If a WAN connection between the gateway server and the HPE StoreOnce appliance is weak, select the **Gateway server and StoreOnce are connected over WAN** check box. Veeam Backup & Replication will compress VM data transported from the gateway server to the HPE StoreOnce appliance, and calculate checksums for data blocks going from the gateway server to the HPE StoreOnce appliance.

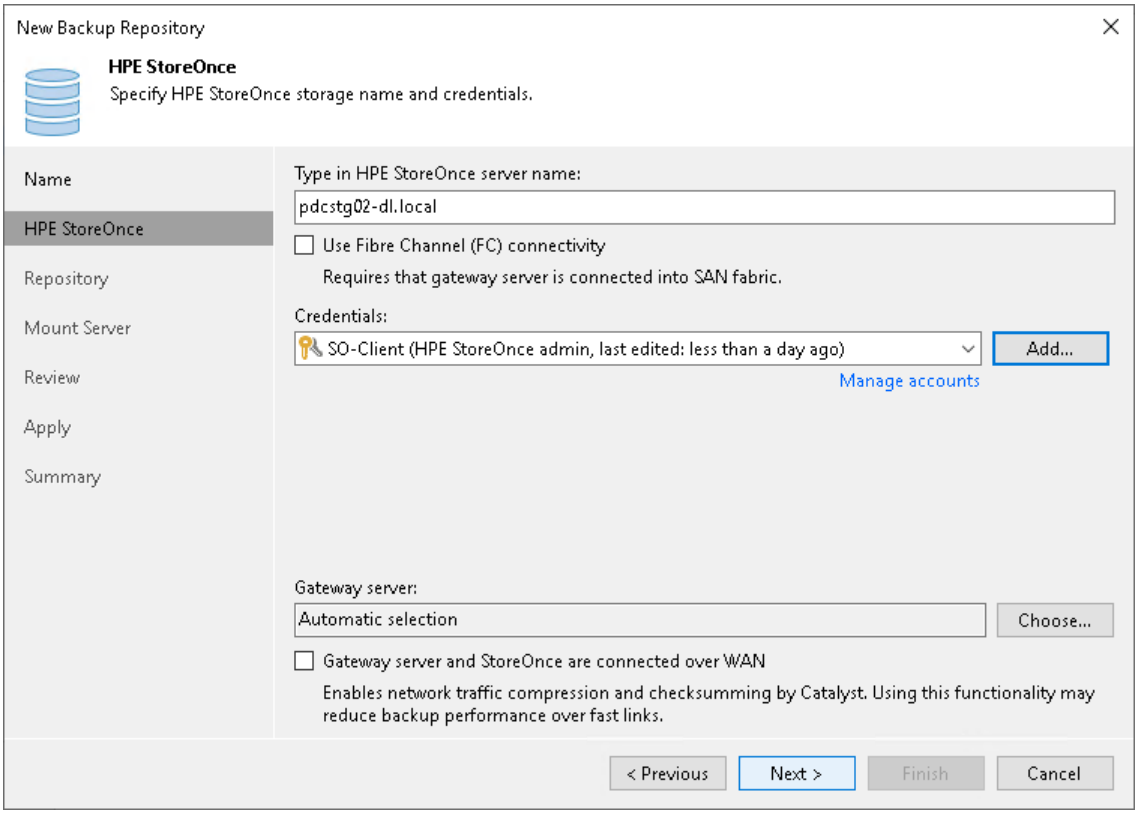

## Step 4. Configure Backup Repository Settings

At the **Repository** step of the wizard, configure general repository settings including path to the repository folder and load control, and also advanced repository settings.

# Configuring General Repository Settings

To configure general repository settings:

1. In the **Location** section, specify a path to the folder where backup files must be stored. Click **Populate** to check capacity and available free space in the selected location.

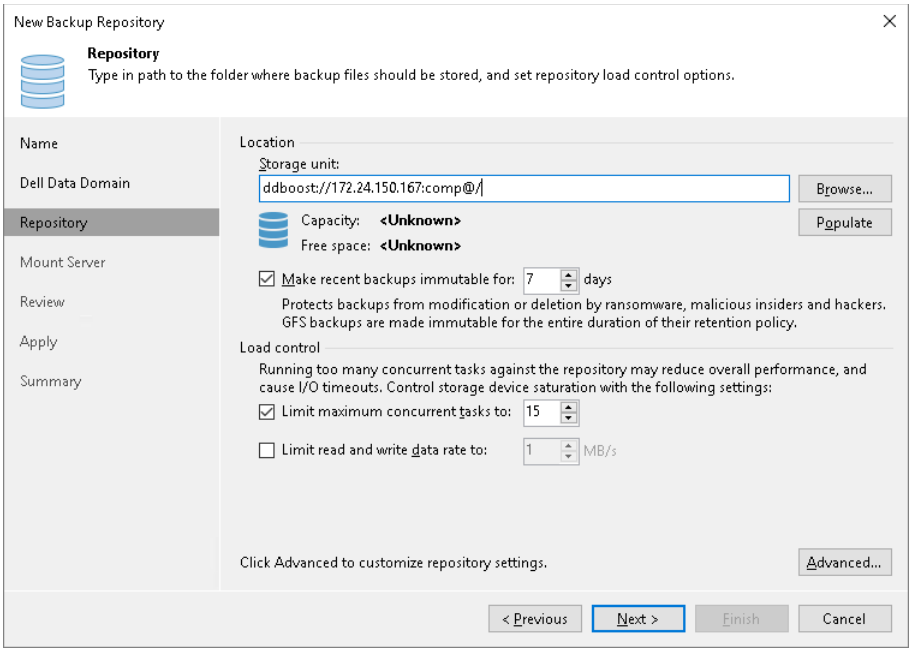

2. [For HPE StoreOnce or Dell Data Domain with immutability] To prohibit deletion of blocks of data from the backup repository, select the **Make recent backups immutable for** check box and specify the immutability period. For more information on limitations and considerations for HPE StoreOnce with immutability, see HPE StoreOnce and Immutability. For more information on limitations and considerations for Dell Data Domain with immutability, see Dell Data Domain.

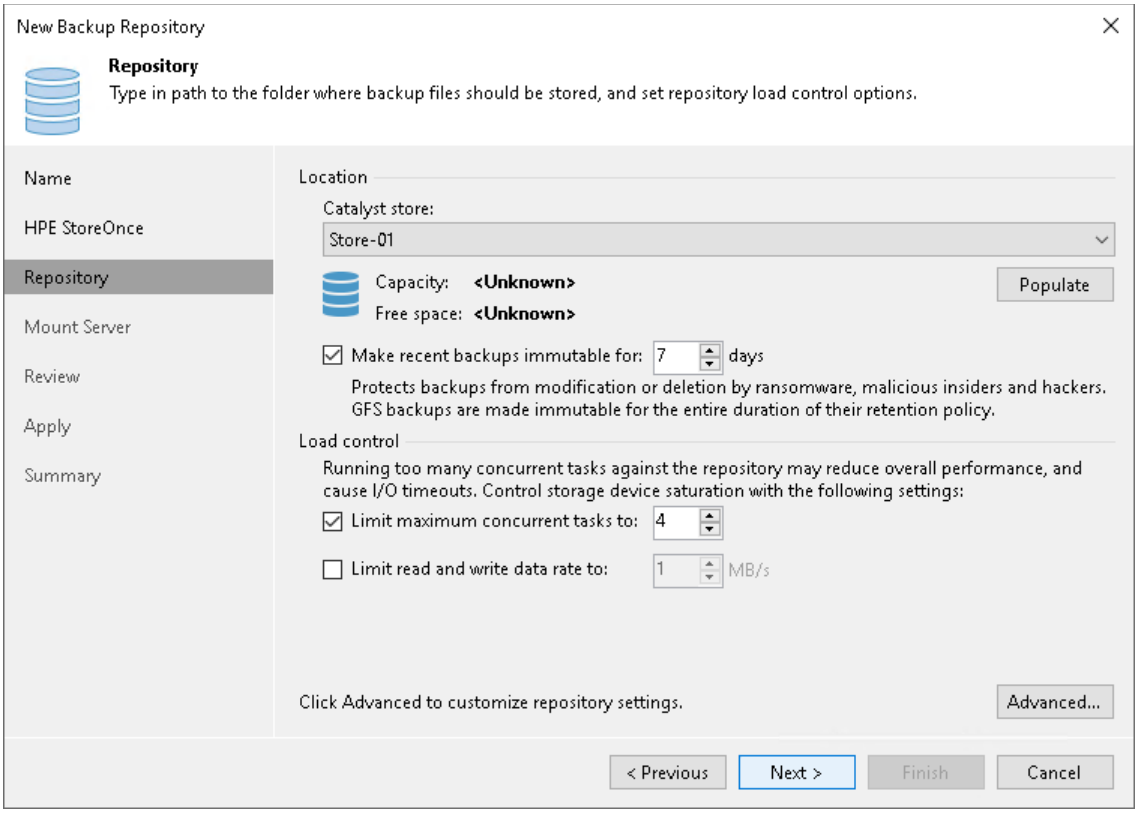


3. [For ExaGrid, Quantum DXi, Fujitsu ETERNUS CS800 and Infinidat InfiniGuard] Select the **Use fast cloning on XFS volumes** check box to enable copy-on-write functionality. In terms of Veeam Backup & Replication, this functionality is known as Fast Clone. For more information, see Fast Clone.

4. Use the **Load control** section to limit the number of concurrent tasks and data ingestion rate for the backup repository. These settings will help you control the load on the backup repository and prevent possible timeouts of storage I/O operations.

- Select the **Limit maximum concurrent tasks** check box and specify the maximum allowed number of concurrent tasks for the backup repository. If this value is exceeded, Veeam Backup & Replication will not start a new task until one of current tasks finishes. For more information, see Limiting the Number of Concurrent Tasks.

  > **NOTE**
  >
  > Consider the following:
  >
  > - Limitation of concurrent tasks is ignored if the backup repository acts as a target storage for a Veeam Cloud Connect job.
  > - If you use backup repositories with per-machine backup chains, it is recommended to select the **Limit maximum concurrent tasks to N** check box. This option reduces the number of parallel operations performed by synthetic operations (synthetic full backup, backup files merge and transformation). Otherwise, the load on the backup repository may be high.

- Select the **Limit read and write data rate to** check box and specify the maximum rate to restrict the total speed of reading and writing data to the backup repository disk. For more information, see Limitation of Read and Write Data Rates for Backup Repositories.

> **NOTE**
>
> The **Limit read and write data rate to** setting does not apply to health checks performed as part of backup and backup copy jobs. Even if you limit read/write rate for a backup repository, the health check will consume resources of the backup repository regardless of this setting. Consider this limitation when configuring basic and health check schedules for backup and backup copy jobs.

# Configuring Advanced Repository Settings

To configure advanced repository settings:

1. Click **Advanced**.
2. Use the **Align backup file data blocks** option to align VM data saved to a backup file at a 4 KB block boundary. The value of the setting depends on the type of backup repository you are adding. For more information, see Preconfigured Advanced Settings.
3. To overcome poor deduplication ratios during writing compressed data, select the **Decompress backup data blocks before storing** check box. In this case, if data compression is enabled for a job, Veeam Backup & Replication compresses data on the source side, transport it to the target side, decompress data on the target side and write raw data to the storage device to achieve a higher deduplication ratio.

   > **NOTE**
   >
   > Veeam Backup & Replication does not compress VM data if encryption is enabled for a job and the **Decompress backup data blocks before storing** check box is selected in the settings of the target backup repository. Therefore, in the job statistics, you may observe a higher amount of transferred data (the **Transferred** counter) as compared to a job for which encryption is disabled. For more information on job statistics, see Viewing Real-Time Statistics.
   >
   > In the properties of the encrypted backup, you may also see that the backup size (the **Backup Size** column) is larger than the original VM size (the **Original Size** column). For more information on backup properties, see Viewing Backup Properties.

4. Deduplicating storage appliances do not support rotated drives. The **This repository is backed by rotated hard drives** check box is disabled for all types of deduplicating storage appliances.
5. To create a separate backup file for every machine in the job, make sure that the **Use per-machine backup files** check box is selected. If you clear the check box, Veeam Backup & Replication will create single-file backups. For more information on the backup chain formats and their limitations, see Backup Chain Formats.

   > **NOTE**
   >
   > If you change the **Use per-machine backup files** setting after the repository is created, the setting does not take any effect. To change backup chain format, follow the instructions provided in Upgrading Backup Chain Formats.

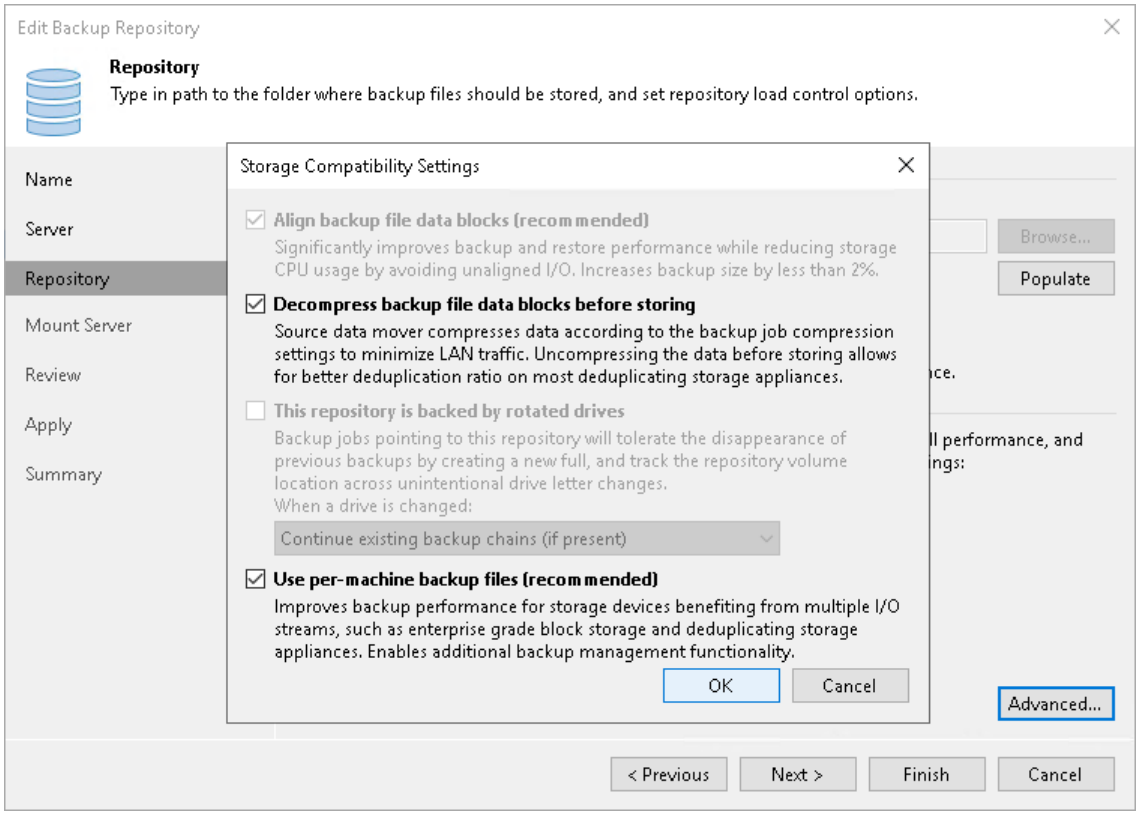


# Preconfigured Advanced Settings

Depending on the type of deduplicating storage appliance you use, Veeam Backup & Replication automatically sets advanced settings to the following ones:

## Dell Data Domain

- The **Align backup file data blocks** option is disabled and cannot be changed.
- The **Decompress backup data blocks before storing** option is enabled.
- The **This repository is backed by rotated hard drives** option is disabled and cannot be changed.
- The **Use per-machine backup files** option is enabled.

## ExaGrid

- When the **Use fast cloning on XFS volumes** option is enabled, the **Align backup file data blocks** option is enabled and cannot be changed.

  When the **Use fast cloning on XFS volumes** option is disabled, we recommend to check that the **Align backup file data blocks** option is disabled.
- The **Decompress backup data blocks before storing** option is enabled.
- The **This repository is backed by rotated hard drives** option is disabled and cannot be changed.
- The **Use per-machine backup files** option is enabled.
- **Limit max concurrent tasks** is equal to 10 (recommended).

## Quantum DXi

- When the **Use fast cloning on XFS volumes** option is enabled, the **Align backup file data blocks** option is enabled and cannot be changed.

  When the **Use fast cloning on XFS volumes** option is disabled, we recommend to check that the **Align backup file data blocks** option is disabled.

- The **Decompress backup data blocks before storing** option is enabled.
- The **This repository is backed by rotated hard drives** option is disabled and cannot be changed.
- The **Use per-machine backup files** option is enabled.

## Fujitsu ETERNUS CS800

- When the **Use fast cloning on XFS volumes** option is enabled, the **Align backup file data blocks** option is enabled and cannot be changed.

  When the **Use fast cloning on XFS volumes** option is disabled, we recommend to check that the **Align backup file data blocks** option is disabled.

- The **Decompress backup data blocks before storing** option is enabled.
- The **This repository is backed by rotated hard drives** option is disabled and cannot be changed.
- The **Use per-machine backup files** option is enabled.

## Infinidat InfiniGuard

- When the **Use fast cloning on XFS volumes** option is enabled, the **Align backup file data blocks** option is enabled and cannot be changed.

  When the **Use fast cloning on XFS volumes** option is disabled, we recommend to check that the **Align backup file data blocks** option is disabled.

- The **Decompress backup data blocks before storing** option is enabled.
- The **This repository is backed by rotated hard drives** option is disabled and cannot be changed.
- The **Use per-machine backup files** option is enabled.

## HPE StoreOnce

- The **Align backup file data blocks** option is enabled and cannot be changed if **Enforce fixed block chunking** is enabled on the HPE StoreOnce. In other cases, the option is disabled and must not be changed.
- The **Decompress backup data blocks before storing** option is enabled.
- The **This repository is backed by rotated hard drives** option is disabled and cannot be changed.
- The **Use per-machine backup files** option is enabled and cannot be changed.

# Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for file-level and application items restore.

1. From the **Mount Server** list, select a server that you want to use as a mount server. The mount server is required for file-level and application items restore. During the restore process, Veeam Backup & Replication mounts the VM disks from the backup file residing in the backup repository to the mount server. As a result, VM data does not have to travel over the network, which reduces the load on the network and speed up the restore process. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers added to the backup infrastructure. If the server is not added to the backup infrastructure, click **Add New** on the right to open the **New Windows Server** wizard. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder that will be used for writing cache during mount operations.

3. To make the backup repository accessible by the Veeam vPower NFS Service, select the **Enable vPower NFS service on the mount server** check box. Veeam Backup & Replication will enable the vPower NFS Service on your selected mount server.

4. To customize network ports used by the vPower NFS Service, click **Ports**. For information on ports used by default, see Ports.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

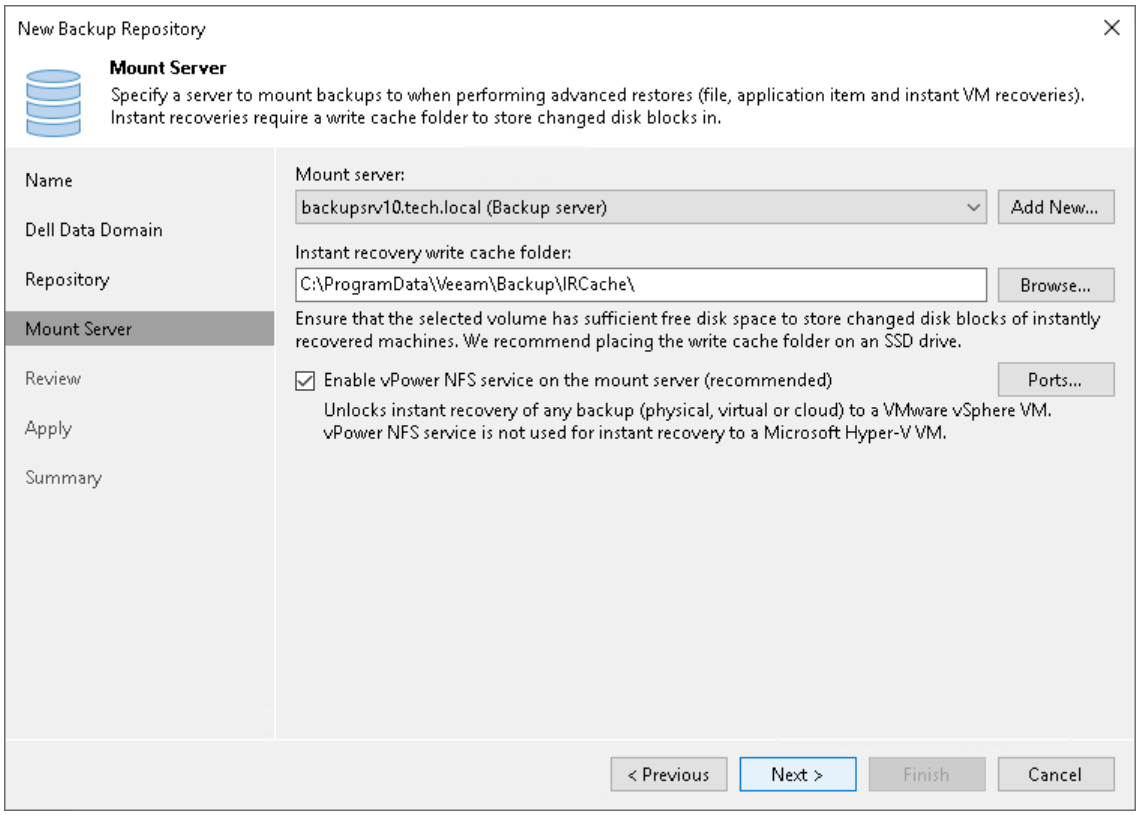

## Step 6. Review Properties and Components

At the **Review** step of the wizard, review details of the backup repository and specify importing settings.

1. Review the backup repository settings and list of components that will be installed on the backup repository server.
2. If the backup repository contains backups that were previously created by Veeam Backup & Replication, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Disk (Imported)** node.
3. If the backup repository contains guest file system index files that were previously created by Veeam Backup & Replication, select the **Import guest file system index** check box. Index files will be imported with backup files, and you will be able to search for guest OS files inside imported backups.

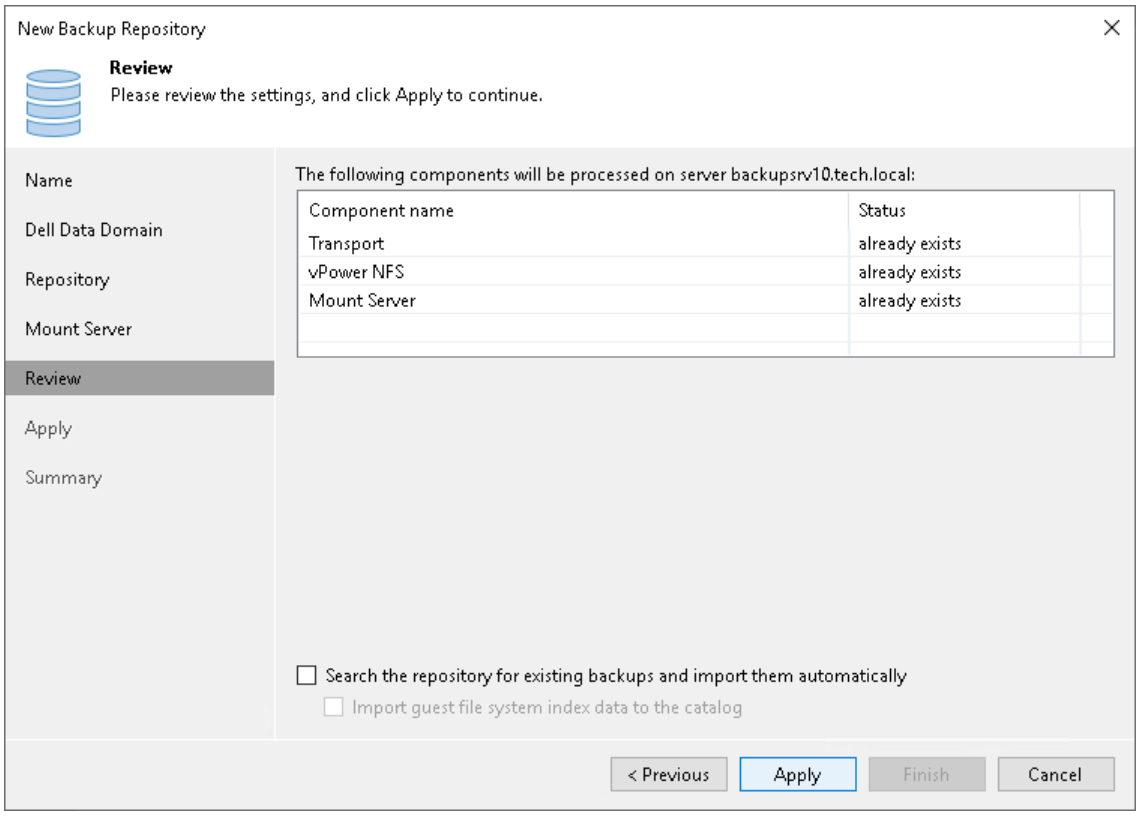

# Step 7. Apply Backup Repository Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to install and configure all required components. Then click **Next** to complete the procedure of adding the backup repository to the backup infrastructure.

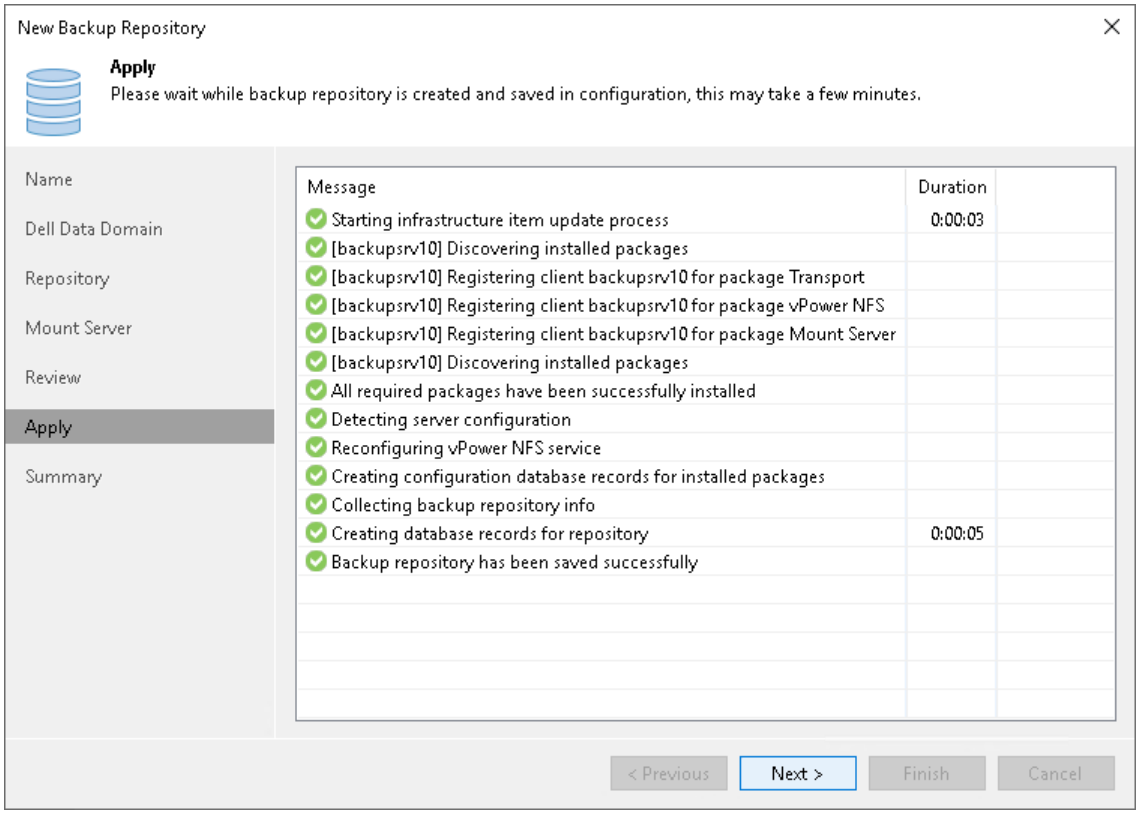

# Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added backup repository. Then click **Finish** to exit the wizard.

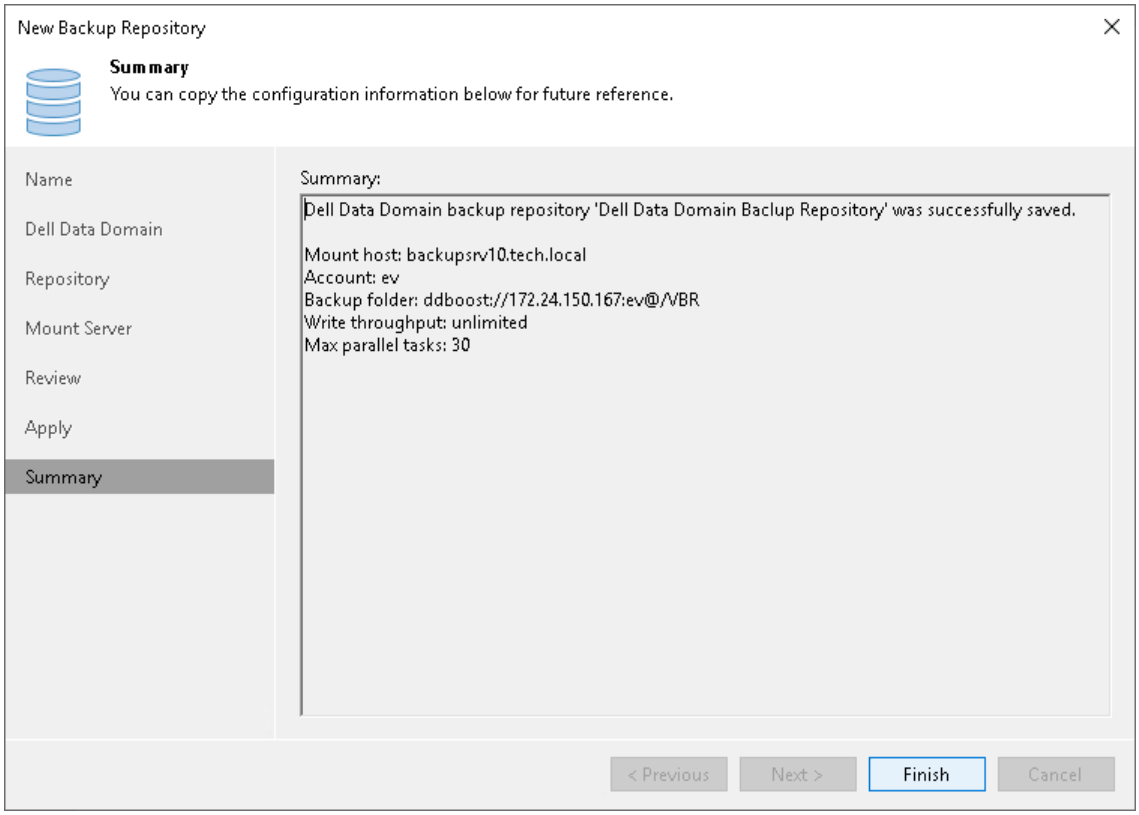

# Backup Repositories with Rotated Drives

A backup repository can use rotated drives. Rotated drives can be detachable USB or eSATA hard drives. This scenario can be helpful if you want to store backups on several external hard drives that you plan to regularly move between different locations.

To use rotated drives, you must enable the **This repository is backed by rotated hard drives** option in the advanced settings of the backup repository. When this option is enabled, Veeam Backup & Replication recognizes the backup target as a backup repository with rotated drives and uses a specific algorithm to make sure that the backup chain created on these drives is not broken.

## Limitations for Backup Repositories with Rotated Drives

Backup repositories with rotated drives have the following limitations:

- On one managed server, you must create only one repository with rotated drives.
- The automount feature must be enabled on the Microsoft Windows server added as a a backup repository with rotated drives. This feature enables Microsoft Windows to assign the partition letter each time you swap the external hard drives. For more information, see Microsoft Docs.
- You cannot store archive full backups (GFS backups) created with backup jobs or backup copy jobs in backup repositories with rotated drives.
- You cannot rescan backup repositories with rotated drives.
- NFS backup repositories do not support rotated drives. If you enable the **This repository is backed by rotated hard drives** setting on the repository, Veeam Backup & Replication will ignore this setting.
- Scale-out backup repositories do not support rotated drives. If you enable the **This repository is backed by rotated hard drives** setting on an extent, Veeam Backup & Replication will ignore this setting and will work with such repository as with a standard extent.
- Repositories with rotated drives are not supported as archive repositories for unstructured data backup.
- Background retention doesn't apply to the backups stored on backup repositories with rotated drives. These backup files will not be deleted by the retention job.
- [For Veeam Cloud Connect users] Veeam Backup & Replication does not support the use of backup repositories with rotated drives as cloud repositories.

## In This Section

- How Repository with Rotated Drives Works
- Deploying Backup Repositories with Rotated Drives

## How Repository with Rotated Drives Works

You can use Microsoft Windows server or Linux server or SMB (CIFS) share as a backup repository with rotated drives.

# Microsoft Windows Backup Repository

Veeam Backup & Replication performs backup jobs and backup copy jobs targeted at a backup repository with rotated drives in different ways.

> **IMPORTANT**
>
> When you specify retention settings for a backup job or a backup copy job targeted at a backup repository with rotated drives, you must define the total number of restore points that you want to retain on one drive. You can specify retention in days or restore points. For example, if you set retention to 10 restore points, the job will keep the total of 10 restore points on each drive in the set.

## Backup Jobs

Backup jobs are performed in the following way:

1. During the first and subsequent runs of the job, Veeam Backup & Replication creates a regular backup chain on the currently attached drive.
2. When the drives are swapped, the behavior depends on how you have configured the repository:
   - If you have chosen to continue an existing backup chain, Veeam Backup & Replication starts the backup chain anew. Veeam Backup & Replication creates a new full backup file on the drive, and this full backup is used as a starting point for subsequent incremental backups until the drives are swapped again.
   - If you have chosen to delete the backups existing on the drive (all backups or backups created by a specific job), Veeam Backup & Replication deletes the existing backup chains. Then Veeam Backup & Replication starts the backup chain anew.
3. [For external drives attached to Microsoft Windows servers] Veeam Backup & Replication checks the retention policy set for the job. If some backup files in the backup chain are outdated, Veeam Backup & Replication removes them from the backup chain.

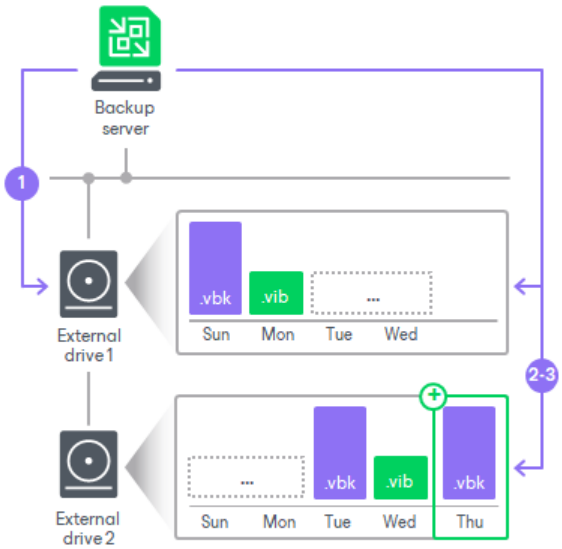

## Backup Copy Jobs

Backup copy jobs are performed in the following way:

1. During the first and subsequent runs of the job, Veeam Backup & Replication creates a regular backup chain on the currently attached drive.
2. When the drives are swapped, the behavior depends on how you have configured the repository:
   - If you have chosen to continue an existing backup chain, the following applies:
     - If the attached drive is empty, Veeam Backup & Replication creates a full backup on it.
     - If there is a full backup or a backup chain on the drive, Veeam Backup & Replication creates a new incremental backup and adds it to the backup chain. The latest incremental backup existing in the backup chain is used as a starting point for the new incremental backup. If the existing backup chain is not consistent, Veeam Backup & Replication starts the backup chain anew. It creates a new full backup file on the drive, and this full backup is used as a starting point for subsequent incremental backups.
   - If you have chosen to delete the backups existing on the drive (all backups or backups created by a specific job), Veeam Backup & Replication deletes the existing backup chains. Then Veeam Backup & Replication starts the backup chain anew.
3. [For external drives attached to Microsoft Windows servers] Veeam Backup & Replication checks the retention policy set for the job. If some backup files in the backup chain are outdated, Veeam Backup & Replication removes them from the backup chain.

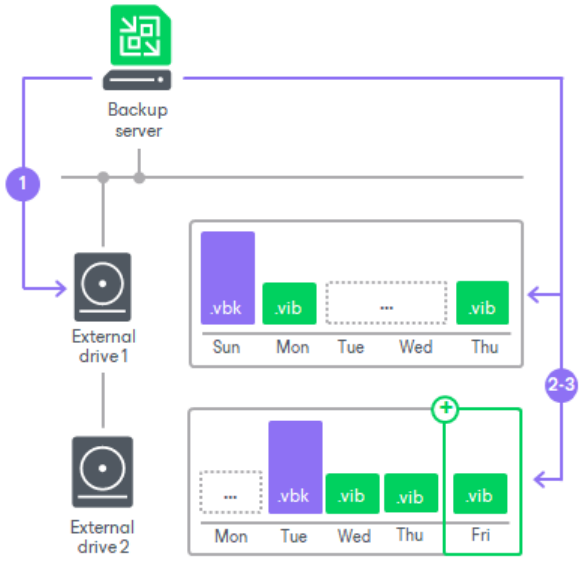


## Drive Detection

Drive letters for external drives may change when you add new volumes or storage hardware such as CD-ROM on the server. In Microsoft Windows backup repositories, Veeam Backup & Replication can keep track of drives and detect them even if the drive letter changes.

To detect a drive correctly, Veeam Backup & Replication must have a record about it in the configuration database. Consider the following requirements:

- When you insert a drive for the first time, the drive is not registered in the configuration database. Such drive must have the same letter as the one specified in the **Path to folder** field in backup repository settings. For more information, see Configuring Path and Load Control Settings.

  If the drive has some other letter, Veeam Backup & Replication will not be able to detect and use it.

- When you insert a drive that has already been used and has some restore points on it, the drive is already registered in the configuration database. Veeam Backup & Replication will be able to detect and use it, even if the drive letter changes.

- If the drive letter of the already used drive changes and a job has not run yet, you cannot change the option that controls the behavior when rotated drives are swapped. To change the behavior, change the drive letter back to the original one and change the option that controls the behavior, then change the drive letter as required.

# Linux and Shared Folder Backup Repository

If you use a Linux server or SMB (CIFS) share as a backup repository with rotated drives, Veeam Backup & Replication employs a "cropped" mechanism of retention with rotated drives. Veeam Backup & Replication keeps information only about the latest backup chain in the configuration database. Information about previous backup chains is removed from the database. For this reason, the retention policy set for the job may not work as expected.

1. During the first and subsequent runs of the job, Veeam Backup & Replication creates a regular backup chain on the currently attached drive.
2. When the drives are swapped, the behavior depends on how you have configured the repository:
   - If you have chosen to continue an existing backup chain, Veeam Backup & Replication starts the backup chain anew. Veeam Backup & Replication creates a new full backup file on the drive, and this full backup is used as a starting point for subsequent incremental backups until the drives are swapped again.

- If you have chosen to delete the backups existing on the drive (all backups or backups created by a specific job), Veeam Backup & Replication deletes the existing backup chains. Then Veeam Backup & Replication starts the backup chain anew.

Veeam Backup & Replication starts a new backup chain on the drive. It removes the information from the configuration database about the restore points from previous backup chains. Backup files related to these previous restore points are not deleted, they remain on disk. This happens because Veeam Backup & Replication applies the retention policy only to the current backup chain, not to previous backup chains.

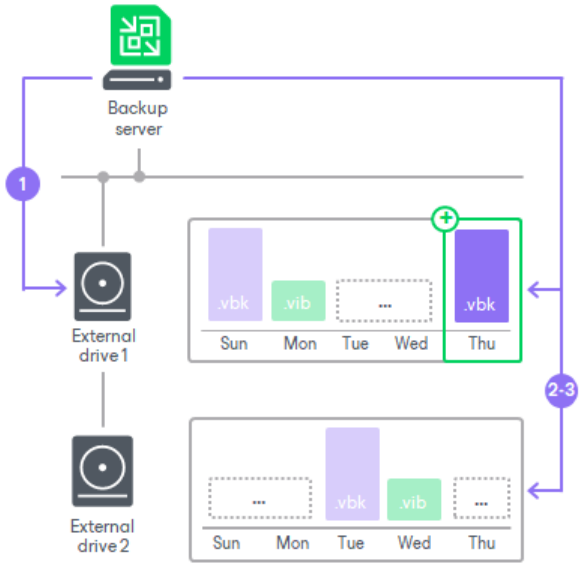


# Deploying Backup Repositories with Rotated Drives

To deploy a backup repository with rotated drives, do the following:

1. Attach one of external drives from the set to a Microsoft Windows or Linux server. The server must be added to the backup infrastructure. For more information on how to add a server, see Virtualization Servers and Hosts.

   You can also attach the external hard drive to the backup server itself. In this case, the VM traffic will path through the backup server, which will produce additional workload on it.

2. Use the **New Backup Repository** wizard to add a direct attached storage. For more information, see Adding Backup Repositories. Pay attention to the following settings:

   a. At the **Server** step of the wizard, select the server to which the drive is attached.

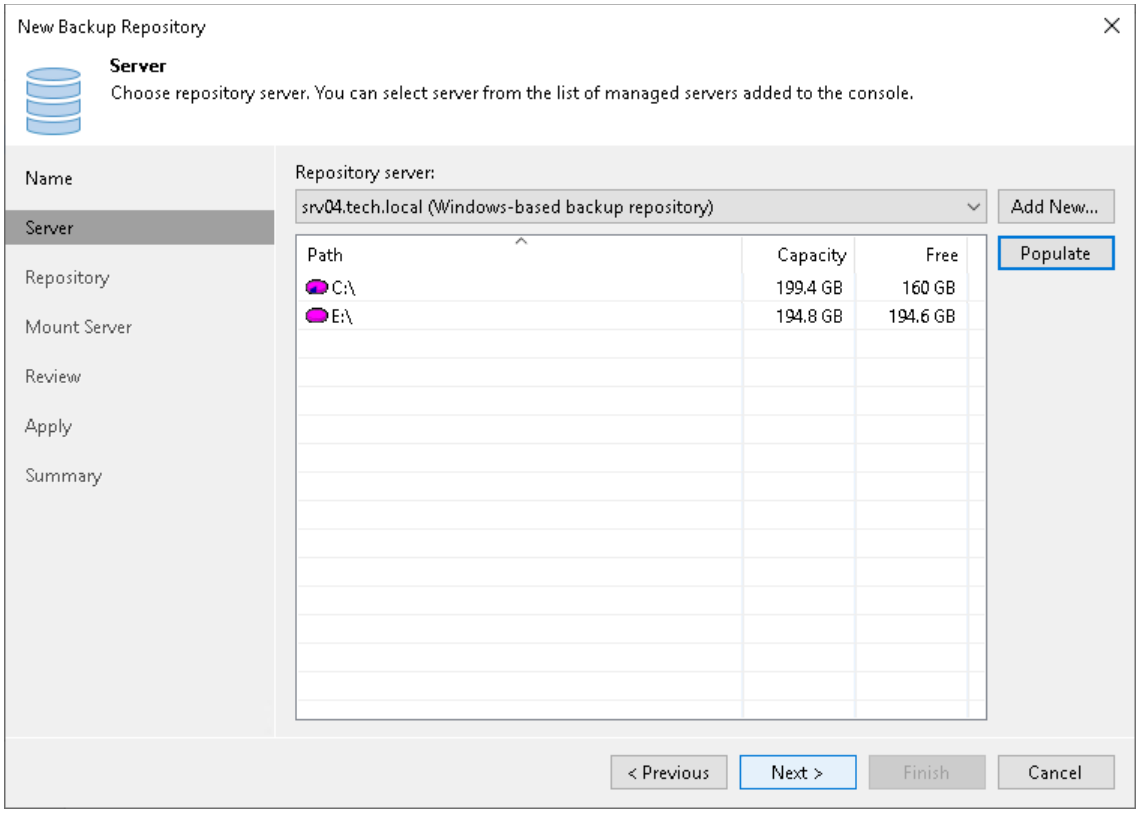


   b. At the **Repository** step of the wizard, in the **Path to folder** field, specify a path to the folder where backup files must be stored.

> **IMPORTANT**
>
> Later, when you attach another external hard drive to the server for the first time, this drive must have the same drive letter as specified in the **Path to folder** field. For more information, see Drive Detection.

   c. Click the **Advanced** button. In the **Storage Compatibility Settings** window, select the **This repository is backed by rotated hard drives** check box and choose how Veeam Backup & Replication will behave when a drive is swapped:

   - *Continue an existing backup chain if present*. Select this option to continue the backup chain found on the drive.
   - *Delete backups belonging to this job*. Select this option to delete all backups created by the currently running job.

- *Delete all backups on the drive*. Select this option to delete all backups created by Veeam Backup & Replication and stored on the drive in the repository folder. Note that Veeam Backup & Replication will delete all backups, even backups created by other jobs.

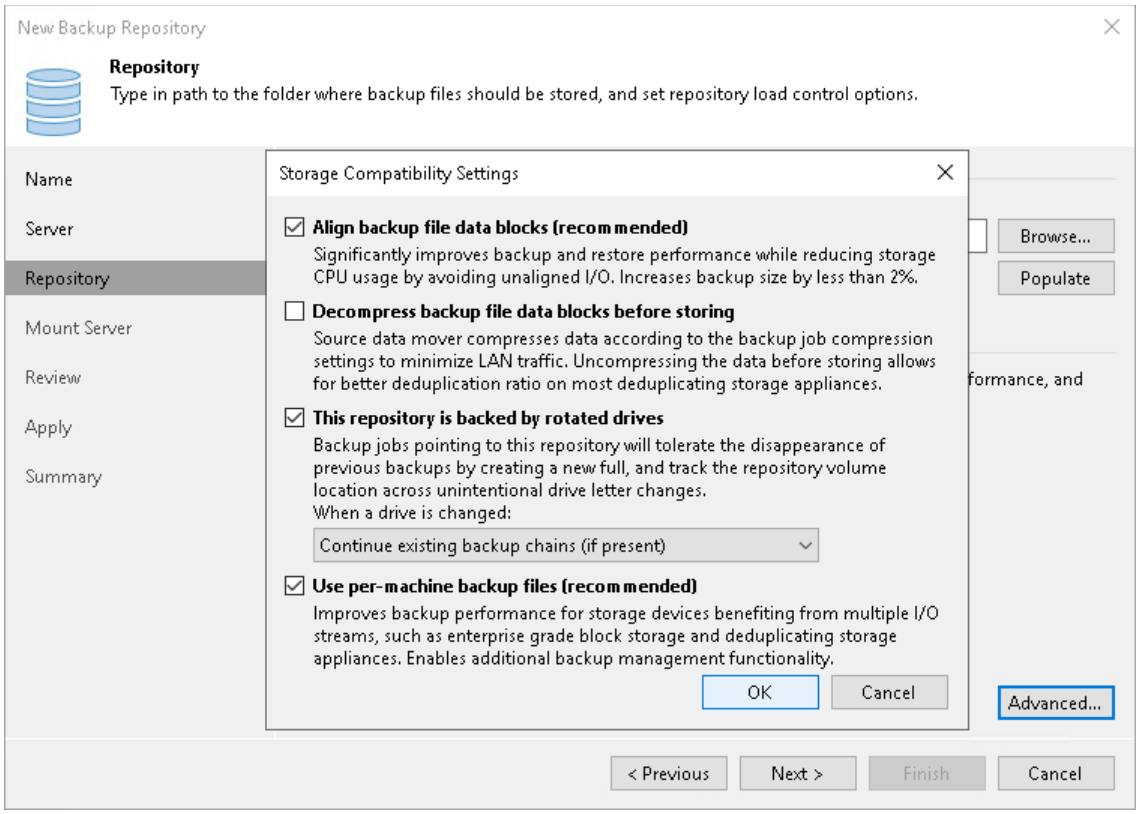


d. Configure other settings of the backup repository as required and finish working with the wizard.

# Object Storage Repositories

An object storage repository is a repository intended for long-term data storage and based on either a cloud solution or S3 compatible on-premises storage solutions.

Veeam Backup & Replication supports the following types of object storage repositories:

- Veeam Data Cloud Vault
- Amazon S3, Amazon S3 Glacier and AWS Snowball Edge
- S3 compatible, S3 compatible with data archiving
- Google Cloud
- IBM Cloud
- Wasabi Cloud Storage
- Microsoft Azure Blob, Azure Archive Storage and Azure Data Box
- 11:11 Cloud Object Storage

You can use object storage repositories as target repositories in the following ways:

- As target repositories for backup jobs and backup copy jobs.

  > **TIP**
  >
  > You can also use object storage repositories as a source location from which backup copy jobs will copy restore points.

- As target repositories for file backup jobs.

  > **IMPORTANT**
  >
  > You cannot keep backups created by file backup jobs in object storage repositories if they are added as performance extents of the scale-out backup repository.

- As target repositories for backups of VMware Cloud Director virtual machines created by Veeam Backup & Replication.
- As target repositories for backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.
- As target repositories for backups of macOS physical machines created by Veeam Agent for Mac.
- As target repositories for backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV.
- As target repositories for backups of oVirt created by Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization.
- As target repositories for backups of applications running on Kubernetes persistent volumes created by Veeam Kasten Plug-in for Veeam Backup & Replication.

> **NOTE**
>
> You can keep backups of Kasten application in object storage repositories only in case they are used as the capacity or archive extents of a scale-out backup repository. You cannot back up Kasten application directly to backup object storage repositories or in case they are used as performance extents of a scale-out backup repository.

- As target repositories for backups of configuration database created by Veeam Backup & Replication.

You can also use object storage repositories as the following parts of the scale-out backup repository:

- As a part of Performance Tier. Performance tier allows quickly accessing the stored data. For more information, see Performance Tier.
- As a part of Capacity Tier. Capacity tier of scale-out backup repository allows offloading the existing backup data directly to cloud-based object storage. For more information, see Capacity Tier.
- As a part of Archive Tier. Archive tier of scale-out backup repository allows transporting the infrequently accessed data from the capacity tier for archive storage. For more information, see Archive Tier.

# Object Storage Repository Deployment

To communicate with an object storage repository, Veeam Backup & Replication uses a VMware backup proxy to transfer data and a mount server to process guest OS applications and perform item recovery.

Depending on the type of a job, a VMware backup proxy connects to the object storage repository using one of the following connection modes:

- **Direct**. In this mode, a VMware backup proxy transfers data directly to the object storage repository. This connection type is used, for example, for backup and restore sessions when your target object storage repository is located in your local infrastructure.
- **Through gateway servers**. In this mode, a VMware backup proxy transfers data to the object storage repository through a gateway server.

> **TIP**
>
> If a backup job backs up multiple VMs, you can use several gateway servers and combine them into a gateway pool to process data. For more information, see Gateway Servers.

# Multiple Buckets for S3 Compatible Object Storage Repositories

To improve performance and provide scalability options, you can deploy S3 compatible object storage in multiple bucket mode.With this option, Veeam Backup & Replication provisions child buckets automatically and distributes backups between these buckets. The number of child buckets depends on the number of per-machine backup chains you want to keep in a single child bucket. Veeam Backup & Replication keeps in a single child bucket only a number of per-machine backup chains that you specify in the repository settings. If the capacity of the existing child buckets is full, Veeam Backup & Replication automatically creates new buckets. To maintain a connection between the backups and child buckets, Veeam Backup & Replication keeps the mapping metadata in the master bucket using the specified prefix.

For every child bucket Veeam Backup & Replication uses the MD5 algorithm to generate a standard name according to the following pattern: *client ID + repository ID + timestamp (f*or example, *dcb740b2c2836cb11f707d63e6ac664f).* In case the creation of a new bucket fails (for example, due to lack of the necessary permissions), Veeam Backup & Replication writes backups to an existing child bucket that is less occupied. At the same time, Veeam Backup & Replication will try to create new child buckets on the next run of a job.

For new S3 compatible object storage repositories this option is enabled by default. To configure it for existing repositories, you must enable the necessary option in the S3 compatible repository wizard. You can use different S3 compatible object storage repositories to store backup mapping metadata in the same master bucket. In case you need to add to the backup infrastructure your S3 compatible object storage repositories with the multiple bucket option enabled, Veeam Backup & Replication will import backups from all child buckets and from the master bucket (in case there are any legacy backups there).

For limitations, see Considerations and Limitations for S3 Compatible Repositories with Multiple Buckets.

For necessary permissions, see Permissions for S3 Compatible Object Storage Repositories.

# Considerations and Limitations

This section lists considerations and known limitations for object storage repositories.

# In This Section

- General Considerations and Limitations
- Direct Backup to Object Storage Considerations and Limitations
- Amazon S3 and S3 Compatible Object Storage Considerations and Limitations
- Microsoft Azure Object Storage Considerations and Limitations
- Google Cloud Object Storage Considerations and Limitations
- Veeam Data Cloud Vault Considerations and Limitations
- Additional Providers Considerations and Limitations

## General Considerations and Limitations

This section contains general limitations for object storage repositories:

- You can add an object storage repository to a second backup server using credentials with the read-only access permissions that allows you to perform data recovery options. If you use credentials with full-access permissions, it will lead to unpredictable behavior and data loss. For more information on permissions, see Permissions.

> **IMPORTANT**
>
> Consider the following:
>
> - This option is not supported for Microsoft Azure Storage and Veeam Data Cloud Vault.
> - This option works for object storage repositories only if they meet the following requirements:
>   - You plan to add these object storage repositories as a performance or capacity extent of a scale-out backup repository.
>   - The object storage repositories do not have data encryption enabled. If encryption is enabled on these repositories, you will not be able to add object storage repositories using credentials with read-only permissions.
> - You can use this option for direct backup object storage repositories added either as a standalone repository or a performance extent of a scale-out backup repository.

- Data in an object storage bucket or container must be managed solely by Veeam Backup & Replication, including retention and data management. If you enable Object Lock and Versioning features on an S3 bucket or version-level WORM on an Azure container, make sure that you do NOT enable the default retention option. Enabling lifecycle rules is not supported, and may result in backup and restore failures.
- You cannot simultaneously store capacity tier and unstructured data backups in the same object storage repository. You must either add this object storage repository as the capacity extent or use it as the backup repository for the unstructured backup job.

  > **IMPORTANT**
  >
  > One object storage repository must not be used across multiple Veeam Backup & Replication servers since it leads to unpredictable system behavior and data loss. For example, you must not use one object storage to store backups of file shares on two different backup servers.

- The backup proxy that processes backup data must meet the following requirements:
  - It must be an on-premises server as close as possible to a backup server.
  - It must have access to the cloud storage that you use as an object storage repository.
- You cannot switch an object storage repository to Sealed Mode and to Maintenance Mode unless it is an extent of a scale-out backup repository.
- Scale-out backup repositories and Veeam Cloud Connect repositories are not supported as a backup destination for cloud machines.

# Ports and Network

Consider the following network limitations:

- Make sure to open required ports to communicate with object storage repositories in advance. Consider that a backup server and a gateway server must have internet access to verify that the certificates installed on object storage repositories are valid. For more details, see Ports.
- If you use default network security configuration for helper appliances, make sure that they are compliant with your internal security policies.

- Object storage gateway appliances that are used to store backup data in filer (SMB (CIFS)/NFS) or block device mode (iSCSI/FC/SAS) are not supported if the backup data is offloaded to object storage and is no longer stored directly on the appliance.

  Such gateway appliances are only supported in the following cases:

  - All of the backup data is stored on the appliance altogether (that is, all of the backup chains are stored on the appliance as a whole and not scattered across multiple devices) and only additional copies of the backup data is transported to object storage.
  - These appliances emulate a tape system (VTL) as an access protocol for Veeam Backup & Replication.

# Security Limitations

Multi-factor authentication (MFA) is not supported for object storage repositories.

# Limitations for Backup Processing

Consider the following limitations for backup processing:

- If a backup chain contains backup files that are marked as corrupted by Health Check, then such corrupted files, as well as all subsequent files that go after the corrupted one are never offloaded. In such a scenario, offload is only possible starting from the full backup file that succeeds the backup chain with corrupted backups.
- For optimal processing, we recommend to set the default block size to 1MB in the storage settings of a backup job. Larger block size can lead to multiple times larger incremental backups, while smaller block sizes will create extra IO pressure on the object storage.
- The periodic compact of a full backup option is not available.
- By default, Veeam Backup & Replication uses the forever forward incremental method to back up directly to object storage repositories. If you want to create a new full backup, enable the long-term retention policy (GFS). In this case, Veeam Backup & Replication will create synthetic full backups and, therefore, will produce a forward incremental backup chain.

  > **NOTE**
  >
  > To produce an independent full backup, you can also run the active full backup manually or specify a periodic schedule for it. Note that this method will significantly increase object storage space consumption.

# Deployment Considerations and Limitations

Consider the following deployment imitations:

- Make sure that a proxy server that you plan to use, meets the following System Requirements.
- You must locate your proxy server as close as possible to the backup source host.
- Veeam Agent transfers data to the object storage repositories without a proxy server. Make sure that you grant Veeam Backup & Replication and Veeam Agent necessary permissions. For more information on how to configure permissions within Veeam Backup & Replication, see Access Permissions. For more information on how to configure permissions for Veeam Agent, see the Permissions section in the Veeam Agent Management Guide. For more information on how Veeam Agent works in direct connection with object storage repositories, see the Access Permissions section in the Veeam Agent Management Guide.

- [For backup copy jobs and file backup copy jobs] Veeam Backup & Replication uses the source backup repository as the gateway server. For more information, see the Automatic Gateway Selection section.
- Veeam Data Cloud Vault does not support immutability for Veeam Agents that use direct connection to transfer data to object storage repositories.
- If you use default network security configuration for helper appliances, make sure that they are compliant with your internal security policies.
- Within a scale-out backup repository, the mount server of a performance extent will act as a gateway server of the capacity extent if all of the following is true:
  a. You use SMB share/NFS share/deduplicating storage appliances as performance extents of your scale-out backup repository.
  b. You have chosen **Automatic selection for the gateway server** at the Specify Shared Folder Settings step of the **New backup repository** wizard.
  c. For the object storage that you use as the capacity extent, you have not selected to connect to object storage using a gateway server at the **Account** step of the **New Object Repository** wizard.

# Limitations for Veeam Solutions

For more information on limitations for Veeam solutions that utilizes object storage repositories functionality, see the following sections of the necessary guide:

- Veeam Agent Management Guide — to check limitations for Veeam Agent management solution.
- Veeam Agent for Microsoft Windows — to check limitations for data protection and disaster recovery solution for physical and virtual machines running Windows-based operating systems.
- Veeam Agent for Linux — to check limitations for data protection and disaster recovery solution for physical endpoints and virtual machines running Linux-based operating systems.
- Veeam Agent for Mac — to check limitations for data protection and disaster recovery solution for physical endpoints and virtual machines running macOS.
- Veeam Cloud Connect Guide — to check limitations for data protection and disaster recovery solution for cloud service providers.

## Direct Backup to Object Storage Considerations and Limitations

This section contains the following limitations for direct backup to object storage:

- The periodic compact of a full backup option is not available.
- By default, Veeam Backup & Replication uses the forever forward incremental method to back up directly to object storage repositories. If you want to create a new full backup, enable the long-term retention policy (GFS). In this case, Veeam Backup & Replication will create synthetic full backups and, therefore, will produce a forward incremental backup chain.

  > **NOTE**
  >
  > To produce an independent full backup, you can also run the active full backup manually or specify a periodic schedule for it. Note that this method will significantly increase object storage space consumption.

- In direct object storage repositories, Veeam Backup & Replication stores backup chain in the per-machine backup with separate metadata files format. For more information, see Backup Chain Formats.

- You cannot back up directly to Amazon S3 Glacier Storage, S3 compatible with Data Archiving and Azure Archive Storage object storage repositories. These types of object storage repositories can only be used as archive extents of the scale-out backup repository. For more information, see Archive Tier.
- You cannot use direct backup to object storage to keep backups of applications running on Kubernetes persistent volumes created by Veeam Kasten Plug-in for Veeam Backup & Replication.
- You cannot use direct backup to object storage to keep backups created with Veeam Plug-ins for Enterprise Applications.
- You cannot use direct backup object storage repositories to keep backups created with Veeam Agent for Oracle Solaris and Veeam Agent for IBM AIX.

  > **NOTE**
  >
  > Note that you can use direct backup object storage repositories to keep backups created with backup copy jobs for Veeam Agent for IBM AIX and Veeam Agent for Oracle Solaris.

## Amazon S3 and S3 Compatible Object Storage Considerations and Limitations

# General Limitations

Consider the following limitations:

- Make sure that you add an S3 compatible object storage device fully compatible with the AWS S3 operations and AWS S3 Signature Version 4 standard.
- Data in an object storage bucket must be managed solely by Veeam Backup & Replication, including retention and data management. If you enable Object Lock and Versioning features on an S3 bucket, make sure that you do NOT enable the default retention option. Enabling lifecycle rules is not supported, and may result in backup and restore failures.
- For some S3 compatible object storage repositories we do not recommended to use more than one bucket for each scale-out backup repository. If you use multiple folders within one bucket for several scale-out backup repositories, it will slow down data processing since the metadata generated by S3 compatible is handled for each bucket. To check requirements for your object storage, contact your S3 compatible vendor.
- [For Amazon S3] Only the Standard, Standard-IA and One Zone-IA storage classes are supported. For more information about Amazon S3 storage classes, see AWS Documentation.
- S3 compatible storage systems with the eventual data consistency model are not supported.
- To be able to work with S3 compatible with Data Archiving object storage, the SOSAPI functionality must be enabled. To enable this, contact your S3 compatible vendor.
- Veeam Backup & Replication does not support native S3 replication across buckets.

# Amazon S3 Immutability Limitations

Consider the following immutability limitations for Amazon S3:

- After you have created an S3 bucket with *Object Lock* enabled, check that the default retention is disabled.

  The default retention may result in an unpredictable system behavior and data loss. However, note that Veeam Backup & Replication will use Compliance object lock mode for each uploaded object. For more information on the retention modes, see AWS documentation. For more information on how to disable retention settings for S3 bucket, see AWS documentation.

- After you have added the buckets to the backup infrastructure, you must NOT enable or disable Versioning and Object Lock as it may lead to unpredictable system behavior and data loss.

- If you plan to use the immutability feature with the existing S3 bucket containing backups created by 9.5 Update 4, keep in mind that both Versioning and Object Lock must be enabled on the bucket simultaneously and immediately before enabling the immutability feature. Any other approach will lead to backup offload failures and inability to correctly interact with backups in the bucket.

# Considerations and Limitations for S3 Compatible Repositories with Multiple Buckets

Consider the following limitations:

- If you add S3 compatible repositories with multiple buckets as capacity extents, Veeam Backup & Replication will ignore this setup and store backups in the master bucket.

- Once you enable the multiple bucket option for S3 compatible object storage repositories, you will not be able to convert these repositories back to a single bucket repository.

- You cannot use S3 compatible repositories with multiple buckets as a target repository for Unstructured Data Backup. This includes backup, archive, or secondary repositories.

- [For Veeam Cloud Connect] You cannot use S3 compatible repositories with multiple buckets as a cloud repository.

- Read-only permissions are not supported for S3 compatible object storage with multiple buckets.

- After you enable the multiple bucket option for existing S3 compatible object storage repositories, the existing backups will still be located in the master bucket. All new backups will be stored in child buckets.

- Child buckets are not displayed in the UI. You can access them using UI of your S3 compatible object storage vendor.

- You cannot use a child bucket to configure another S3 compatible object storage repository.

- Veeam Backup & Replication keeps configuration backups in a separate location and ignores multiple bucket setup.

- Veeam Backup & Replication does not support the multiple-bucket option for S3-compatible repositories that use Veeam Smart Object Storage API (SOSAPI) functionality with load balancing enabled on a smart entity level.

# Microsoft Azure Object Storage Considerations and Limitations

## General Considerations and Limitations

Consider the following limitations:

- Data in an Azure container must be managed solely by Veeam Backup & Replication, including retention and data management. If you enable version-level WORM on the Azure container, make sure that you do NOT enable the default retention option. Enabling lifecycle rules is not supported, and may result in backup and restore failures.
- [For Microsoft Azure Blob storage] Veeam Backup & Replication does not support soft delete for blobs.
- [For Microsoft Azure Archive storage] Microsoft Azure has certain limits (quotas) on maximum amount of resources used. The quotas depend on the type of proxies you have selected. If you exhaust a quota, you will be unable to use Microsoft Azure Archive storage. For more information about Microsoft Azure quotas, see Microsoft Docs.
- Veeam Backup & Replication performs operations only on a blob level. You cannot create Azure containers or storage accounts from the backup infrastructure.
- Veeam Backup & Replication does not support object-level immutability and default immutability policies assigned to Azure storage accounts. You must set immutability for an Azure container where backed-up objects will reside. For more information, see Microsoft Docs.
- Veeam Backup & Replication supports all types of Azure Storage redundancy. For more information, see Microsoft Docs.

  > **NOTE**
  >
  > Note that the support of storage redundancy depends on the type of the Microsoft Azure Storage. For example. Microsoft Azure Archive Storage with the archive access tier does not support Azure accounts with the following redundancy options: zone-redundant storage (ZRS), geo-redundant storage (GZRS) and read-access geo-zone-redundant storage (RA-GZRS). For more information, see Microsoft Docs.

## Azure Storage Access Tier Considerations

Veeam Backup & Replication supports multiple Azure Storage access tiers. The type of the Azure Storage access tier that Veeam Backup & Replication uses depends on the type of the Azure object storage:

- For Azure Blob Storage, Veeam Backup & Replication uses the access tier that you specify at the Container step of the **New Object Storage** wizard.Veeam Backup & Replication supports the following access tiers:
  - *Hot* access tier: to store data that you access frequently.
  - *Cool* access tier: to store data that you access infrequently.
- For Azure Archive Storage, Veeam Backup & Replication uses the access tier that you specify at the Access Tier step of the **New Object Storage** wizard.Veeam Backup & Replication supports the following access tiers:
  - *Cold* access tier: to store data that you access infrequently.
  - *Archive* access tier: to store data that you access rarely.

# Immutability Limitations

Consider the following immutability limitations for Microsoft Azure Storage:

- Veeam Backup & Replication supports the Versioning feature for Microsoft Azure object storage for which immutability is enabled.
- Make sure that the Microsoft Azure Storage settings meet the following requirements when you configure an Azure storage account and an Azure container:
  - When you create a storage account, make sure you enable versioning for blobs.
  - When you create a storage account, do NOT enable version-level immutability. By default, this option is enabled and you must disable it. For more information, see Microsoft Docs.
  - When you create a container, you must enable version-level immutability. For more information, see Microsoft Docs.
  - The default retention is disabled. For more information on how to disable retention settings for Azure container, see Microsoft Docs.

  > **TIP**
  >
  > For instruction on how to configure your Azure storage account with the necessary settings, see this Veeam KB article.

- The default immutability policies are not supported.
- Do NOT enable immutability for already existing containers in the Azure portal. Otherwise, Veeam Backup & Replication will not be able to process these containers properly and it may result in data loss.
- Version-level immutability support for Azure storage accounts is not supported.
- [For backup jobs managed by Veeam Agent] You cannot back up data to Microsoft Azure object storage added in the direct connection mode.
- [For Veeam Cloud Connect] Microsoft Azure object storage cannot be used as a cloud repository in the direct connection mode.

# Google Cloud Object Storage Considerations and Limitations

Consider the following limitations:

- Currently, Veeam Backup & Replication does not support the Object Versioning and Bucket Lock features for Google Cloud object storage.

  > **IMPORTANT**
  >
  > Enabling either any or both of these features on the bucket may result in unpredictable system behavior and data loss, as well as in extra costs for storing objects that have been removed by the retention policy. For more information, see Object Versioning and Bucket Lock.

- Backups stored in Google Cloud object storage must be modified neither manually nor by third-party tools, including native Google Cloud capabilities (for example, the Autoclass feature). Otherwise, Veeam Backup & Replication may fail to restore the backed-up data.

- If you enable the soft delete feature for Google Cloud buckets and the objects in these buckets are deleted, do not restore the objects with the Google console or any third-party capabilities. For more information on the soft delete feature, see the Google Cloud documentation.

## Veeam Data Cloud Vault Considerations and Limitations

Consider the following limitations for Veeam Data Cloud Vault:

- To back up to Veeam Data Cloud Vault from Veeam Backup & Replication and utilize Veeam Data Cloud Vault functionality you must meet the following prerequisites:
  - Obtain Veeam Data Cloud Vault and configure its subscription. For more information, see the Veeam Data Cloud Vault User Guide.
  - Have the License Administrator role. For more information, see this Veeam KB article.

    > **NOTE**
    >
    > Ensure that the email address you use to access the My Account portal or the Veeam Data Cloud Vault portal matches the email address of the License Administrator. Consider that this email address is case-sensitive. If you plan to delegate access to another user, make sure that he also has the License Administrator role.

- Make sure you have WebView2 Runtime installed on your backup server.
- You must enable encryption for every job that you target to Veeam Data Cloud Vault. For more information, see Job Encryption.

  > **NOTE**
  >
  > Job-level encryption is not required if you use Veeam Data Cloud Vault as a capacity extent of a scale-out backup repository.

- By default, immutability is enabled for 30 days for Veeam Data Cloud Vault. You cannot disable immutability and cannot remove data during this period. You can modify the immutability period when you configure Veeam Data Cloud Vault.
- If you back up your configuration, you must enable encryption for it. Otherwise, this option will be disabled when you configure Veeam Data Cloud Vault.
- Veeam Data Cloud Vault does not support China and Government regions.
- Veeam Data Cloud Vault supports the cool access tier only.
- You cannot store backups created by Veeam Plug-ins for Enterprise Applications in Veeam Data Cloud Vault, which you use as a standalone object storage repository. If you want to keep these backups in Veeam Data Cloud Vault, you must add it as a capacity extent of a scale-out backup repository.
- You cannot use Veeam Data Cloud Vault as a source of unstructured data backup.
- You cannot move or copy unencrypted backups to Veeam Data Cloud Vault.
- You must enable the capacity tier encryption for Veeam Data Cloud Vault if you want to use it as a capacity extent.
- You cannot seed data from the Azure Data Box device to Veeam Data Cloud Vault.
- [For backup jobs managed by Veeam Agent] You cannot back up data to Veeam Data Cloud Vault added in the direct connection mode.

- [For Veeam Cloud Connect] Veeam Data Cloud Vault cannot be used as a cloud repository in the direct connection mode.

## Additional Providers Considerations and Limitations

# Limitations for 11:11 Cloud Object Storage

Consider the following limitations:

- You cannot use 11:11 Cloud Object Storage as an archive extent for Archive Tier of a scale-out backup repository.
- You cannot use 11:11 Cloud Object Storage in the same type of a scale-out backup repository tier (either the performance tier or the capacity tier) with Amazon S3 repository.
- If you use 11:11 Cloud Object Storage as a performance and capacity extent, you must use either Amazon S3 Glacier Storage or S3 compatible object storage with data archiving as an archive extent.

# Limitations for IBM Cloud Object Storage

Consider the following limitations:

- For IBM Cloud Object Storage on-premises, Veeam Backup & Replication supports versions starting from 3.15.0.44.
- Veeam Backup & Replication is supported on all IBM Cloud Object Storage (COS) deployment models. This includes on-premises, public cloud, and hybrid models. For the IBM public cloud, the following storage classes are supported: Standard, Vault, Cold Vault and Smart Tier.
- Veeam Backup & Replication does not support Archive and Accelerated Archive storage classes in the IBM public cloud.
- If you use IBM Cloud Object Storage as a performance and capacity extent, you must use S3 compatible object storage with data archiving as an archive extent.

# Limitations for Wasabi Cloud Storage

Consider the following limitations:

- You cannot add as performance extents one Wasabi bucket as an S3 compatible Object Storage and another Wasabi bucket as a Wasabi Cloud Object Storage to the same type of tier (either performance tier or capacity tier) of a scale-out backup repository.
- If you use Wasabi Cloud Object Storage as a performance and capacity extent, you must use S3 compatible object storage with data archiving as an archive extent.

# Health Check for Object Storage Repositories

You can instruct Veeam Backup & Replication to periodically perform a health check for the latest restore point in the backup chain. An automatic health check can help you verify that VM data blocks are present in the object storage repository and check the integrity of these blocks. It helps you avoid a situation where a restore point gets corrupted, making all dependent restore points corrupted, too. The health check helps ensure that the restore point is consistent, and you will be able to restore data from this restore point.

During a health check, Veeam Backup & Replication performs a cyclic redundancy check (CRC) for metadata and a hash check for VM data blocks in the backup file to verify their integrity. If during the health check Veeam Backup & Replication detects corrupted data blocks in the latest restore point in the backup chain, it will start the health check retry. During the health check retry, Veeam Backup & Replication will transport valid data blocks from the source datastore to the object storage repository. After the health check retry completes, the transported data blocks are moved immediately to the latest backup in the backup chain.

To allow Veeam Backup & Replication perform the health check of data blocks located in the object storage repository, you must enable the health check when you configure a backup job and define its schedule. For more information, see Specifying Health Check Settings.By default, the health check is performed on the last Friday of every month. You can change the schedule and run the health check weekly or monthly on specific days.

# Considerations and Limitations

Consider the following limitations for the health check:

- The health check retry will start only if the backups job meets the necessary requirements. Otherwise, Veeam Backup & Replication will add healthy blocks to a new backup file after the next job run.
- We recommend that you configure a helper appliance to perform the health check for data blocks in object storage repositories. If you do not configure the helper appliance, Veeam Backup & Replication will use a mount server (for the direct connection mode) or a gateway server (for the gateway server connection mode) to perform the health check, which may result in additional costs. You can configure the helper appliance at the Mount Server step of the New Object Repository wizard.
- If you perform health check for encrypted backup files, Veeam Backup & Replication will pass encryption keys to the regular backup repository or cloud repository. For more information on encryption, see Data Encryption.
- If you use Veeam Agent to back up machines and store their backups in object storage repositories, you can also perform the health check for Veeam Agent backups. For more information, see the Health Check for Object Storage section of the User Guide for Veeam Agent that you use to back up data.
- The health check runs only after all jobs and operations that process a backup are completed (for example, a backup job, a restore job, synthetic operations with backups and so on). If some of these activities are started, the health check stops.
- If the health check session does not complete until the next scheduled run, this session will stop and a new session will start according to the schedule.
- The health check skips verification of already processed data in case this data has not changed since previous health check session. For example, the source job did not process the backup chain after a health check completes. However, Veeam Backup & Replication verifies the metadata of backup files that did not change between health check sessions.

# How Health Check Works

The health check is performed in the following way:

1. The health check starts according to the schedule.

2. Veeam Backup & Replication performs CRC values check for backup metadata, verifies that blocks of data is present in the object storage repository and checks their integrity.

   During the health check, Veeam Backup & Replication verifies the restore point that represents the latest state of a VM and all data blocks that are required to perform the data recovery (restore point created by the current backup job session — the session during which the health check is performed). In this case, all data blocks that are required for the latest state of the VM in the active part of the backup chain are checked. If this restore point in the backup chain is incomplete, Veeam Backup & Replication checks the last complete restore point preceding the latest restore point.

3. If the health check does not detect data corruption, the backup job session completes in a regular way. If the health check detects corrupted data, Veeam Backup & Replication starts the health check retry process.

   Depending on the revealed data corruption, Veeam Backup & Replication performs the following actions:

   - If the health check has detected corrupted backup metadata, Veeam Backup & Replication marks the whole backup chain as corrupted in the configuration database. In this case, you must detach the corrupted backup from the source job and run a backup job again to create a new backup chain.
   - If the health check detects corrupted blocks of data, Veeam Backup & Replication performs the health check retry and attempts to repair the corrupted data blocks.

# Health Check Retry Mode

If the health check detects corrupted data, the backup job will switch to the *Retry* mode and will start the health check retry process. During the health check retry, Veeam Backup & Replication transports necessary data blocks of the whole VM image from the source datastore and saves transported data blocks to the latest backup file in the object storage repository.

If Veeam Backup & Replication does not perform the health check retry, you must retry the job manually. In this case, Veeam Backup & Replication produces a new backup file with healthy data blocks. You can use this backup file to restore from the latest restore point.

For scheduled jobs, the number of health check retries is equal to the number of job retries specified in the job settings. For jobs started manually, Veeam Backup & Replication performs 1 health check retry.

> **IMPORTANT**
>
> To allow Veeam Backup & Replication add the repaired data blocks to the latest restore point after completing the health check, the backup job must meet the following requirements:
>
> - The backup job must not be disabled.
> - You must schedule the backups job to run automatically. For more information, see the Schedule step of the **New Backup Job** wizard.
> - The target object storage is not set to the Maintenance mode.
> - Backup window settings must allow a backup job to run after the health check completes.
>
> If the backup job does not meet these requirements, Veeam Backup & Replication will add the repaired data blocks to a new restore point, created during a next run of the backup job.

# Object Storage Repository Structure

After backups are moved to an object storage repository, Veeam Backup & Replication creates and maintains the following structure of directories.

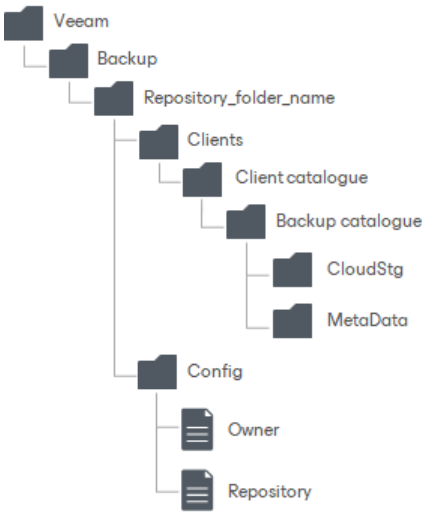


| Directory | Description |
|---|---|
| **Veeam/Backup/** | Standard folders created by Veeam Backup & Replication. |
| **Repository_folder_name** | Contains information on the repository name and the repository ID. |
| **Clients** | Contains backups. |
| **Client catalogue** | Contains information on solutions that create backups to this repository. |
| **Backup catalogue** | Contains backup ID. |
| **CloudStg** | Contains data blocks. |
| **MetaData** | Contains metadata. |
| **Config** | Contains information on the object storage repository infrastructure. |
| **Owner** | Contains information on a repository owner. |
| **Repository** | Contains information on a repository. |

# Repository Structure for Multiple Buckets for S3 Compatible Object Storage Repositories

Veeam Backup & Replication creates and maintains the following structure of directories for S3 Compatible object storage repositories with multiple buckets.

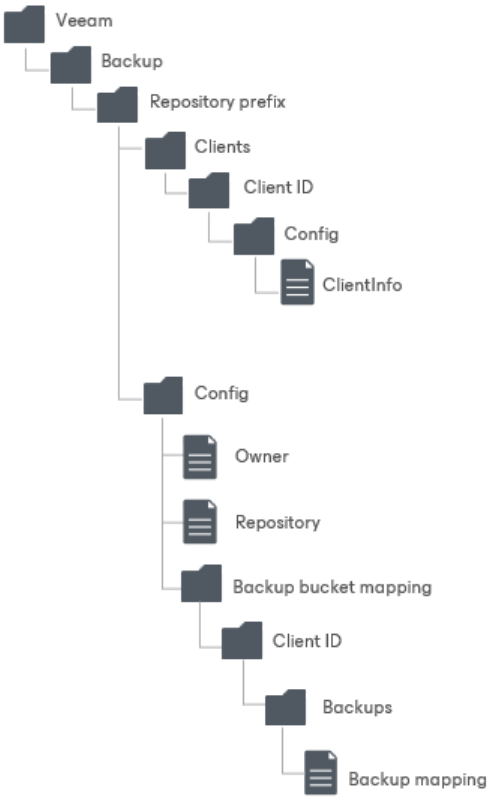


| Directory | Description |
|---|---|
| **Veeam/Backup/** | Standard folders created by Veeam Backup & Replication for object storage repositories. |
| **Repository_prefix** | Contains information on the repository name and the repository ID. |
| **Clients** | Contains information on solutions that create backups to this repository. |
| **Client ID** | Contains information on ID of solutions that create backups to this repository. |
| **Config** | Contains information on the object storage repository infrastructure. |
| **ClientInfo** | Contains information on the object storage immutability settings, tags and metadata. |
| **Config** | Contains information on the object storage repository infrastructure. |

| Directory | Description |
| --- | --- |
| **Owner** | Contains information on a repository owner. |
| **Repository** | Contains information on a repository. |
| **Backup bucket mapping** | Contains information on backup mapping metadata. |
| **Client ID** | Contains information on solutions that create backups to this repository. |
| **Backups** | Contains backups. |
| **Backup mapping** | XML with mapping metadata. |

# Background Checkpoint Removal Job

To improve the performance of backup jobs that target object storage repositories, Veeam Backup & Replication runs a background checkpoint removal job. This job runs continuously as a background activity and removes checkpoints from backups located in object storage repositories.

# How Background Checkpoint Removal Job Works

The algorithm of the background checkpoint removal job depends on the type of backup. For most backups, Veeam Backup & Replication removes outdated checkpoints at the request of the job or when the checkpoints removal job restarts.

## Deleting Checkpoints by Job Request

The checkpoint removal by the job request has the following algorithm:

1. A job produces a backup and starts to process this backup.
2. The job sends a request to a backup service to remove outdated checkpoints and continues to run without waiting for a response from the service. By default, the background checkpoint removal job periodically checks the backup service for requests to remove checkpoints. Once it gets the request, it starts the removal.
3. The background checkpoint removal job gets the backup ID, host ID, and repository ID from the request.
4. The background checkpoint removal job starts an agent to get a list of outdated checkpoints.

> **NOTE**
>
> To reduce the load on the backup server, Veeam Backup & Replication starts agents for checkpoint removal on different hosts depending on the repository connection mode:
>
> - For direct connection mode the agent starts on the mount server.
> - For gateway server connection mode the agent starts on the gateway server configured for the object storage repository.

5. The background checkpoint removal job deletes outdated checkpoints.

### Deleting Checkpoints when Background Checkpoint Removal Job Starts

To delete outdated checkpoints when the background checkpoint removal job restarts, Veeam Backup & Replication applies the backup filtering algorithm. It compares the latest removed checkpoint with the latest checkpoint created by the job. If the latest removed checkpoint does not coincide with the latest created checkpoint, Veeam Backup & Replication identifies the outdated checkpoint and removes it. Otherwise, the checkpoint removal is skipped.

> **NOTE**
>
> The filtering algorithm processes backups based on their type and repository connection mode:
>
> - For backups created by Veeam Agent jobs that operate in managed by agent or standalone mode and directly access object storage repositories, the background checkpoint removal job deletes outdated checkpoints every time it restarts.
> - The background checkpoint removal job on every restart skips from processing the following types of backups:
>   - Backups created by VM jobs.
>   - Backups created by Veeam Agent jobs managed by Veeam Backup & Replication
>   - Backups created by Veeam Agent jobs that operate in managed by agent or standalone mode that use a gateway server to connect to object storage repositories.

## Cloud Connect Checkpoints Removal

Depending on the backup type and repository connection mode, Veeam Backup & Replication deletes checkpoints either on the tenant or on the Service provider side:

- Backups created by any job or Veeam Agent managed by Veeam Backup & Replication jobs that directly access object storage repositories: checkpoints are deleted on the tenant side by the job request.
- Backups created by any job or Veeam Agent jobs managed by Veeam Backup & Replication that use a gateway server to connect to object storage repositories: checkpoints are deleted on the Service provider side by job request.
- Backups created by Veeam Agent jobs managed by agent or in a standalone mode that connect to the object storage directly or using a gateway server: checkpoints are deleted on the Service provider every time the checkpoint removal job restarts.

## Email Reports

Veeam Backup & Replication generates an email report that includes errors from all completed checkpoint removal jobs. The report is sent once a day according to Global Email Notification Settings.

## Immutability for Object Storage Repositories

Veeam Backup & Replication allows you to prohibit deletion of data from the object storage repository by making that data temporarily immutable and to protect data against malware activity by maintaining several versions of a single backup.

The immutability feature can help in the following cases:

- Data on the object storage is corrupted.

- Retention policy is set to keep only one restore point.
- Due to the hacker attack, the retention policy has been modified to a shorter period. For example, instead of keeping data for 5 days, the retention is set to keep it for only 1 day.

Immutability allows you to restore data from the object storage in these or other cases when necessary data is unavailable. To restore data, you need to run Veeam PowerShell. For more information, see Get-VBRObjectStorageRepositorySyncInterval and Sync-VBRObjectStorageRepositoryEntityState cmdlets.

After you enable immutability, you will not be able to perform the following operations with the immutable data stored on object storage repositories:

- Manual data removal, as described in section Deleting Backups from Scale-Out Backup Repositories.
- Removal of data by the retention policy, as described in section Retention Policy.
- Removal of data using any cloud service provider tools, for example an S3 browser.
- Removal of data by the cloud service provider technical support department.
- Removal of data by the **Remove deleted items data after** option, as described in section Maintenance Settings.

You can use immutability for data stored in the following types of object storage repositories:

- Amazon S3
- S3-compatible
- Microsoft Azure Storage
- IBM Cloud Object Storage
- Wasabi Cloud Object Storage
- Veeam Data Cloud Vault
- 11:11 Cloud Object Storage

## Considerations and Limitations

This section lists considerations and known limitations for object storage repositories.

# General Considerations and Limitations

Consider the following immutability limitations:

- The maximum immutability period you can set in the Veeam Backup & Replication UI is 90 days. If you want to set immutability to a longer period, use one of the following Veeam PowerShell cmdlets: Set-VBRAmazonS3CompatibleRepository, Add-VBRAmazonS3CompatibleRepository, Set-VBRAzureBlobRepository.
- Immutable data is preserved as described in Block Generation.
- We recommend that you do not set the immutability period longer than the retention policy of the backup job, otherwise it will result in extra charges.

# Amazon S3 Immutability Limitations

Consider the following immutability limitations for Amazon S3:

- After you have created an S3 bucket with *Object Lock* enabled, check that the default retention is disabled.

  The default retention may result in an unpredictable system behavior and data loss. However, note that Veeam Backup & Replication will use Compliance object lock mode for each uploaded object. For more information on the retention modes, see AWS documentation. For more information on how to disable retention settings for S3 bucket, see AWS documentation.

- After you have added the buckets to the backup infrastructure, you must NOT enable or disable Versioning and Object Lock as it may lead to unpredictable system behavior and data loss.

- If you plan to use the immutability feature with the existing S3 bucket containing backups created by 9.5 Update 4, keep in mind that both Versioning and Object Lock must be enabled on the bucket simultaneously and immediately before enabling the immutability feature. Any other approach will lead to backup offload failures and inability to correctly interact with backups in the bucket.

# Microsoft Azure Storage Immutability Limitations

Consider the following immutability limitations for Microsoft Azure Storage:

- Veeam Backup & Replication supports the Versioning feature for Microsoft Azure object storage for which immutability is enabled.

- Make sure that the Microsoft Azure Storage settings meet the following requirements when you configure an Azure storage account and an Azure container:

  - When you create a storage account, make sure you enable versioning for blobs.
  - When you create a storage account, do NOT enable version-level immutability. By default, this option is enabled and you must disable it. For more information, see Microsoft Docs.
  - When you create a container, you must enable version-level immutability. For more information, see Microsoft Docs.
  - The default retention is disabled. For more information on how to disable retention settings for Azure container, see Microsoft Docs.

  > **TIP**
  >
  > For instruction on how to configure your Azure storage account with the necessary settings, see this Veeam KB article.

- The default immutability policies are not supported.

- Do NOT enable immutability for already existing containers in the Azure portal. Otherwise, Veeam Backup & Replication will not be able to process these containers properly and it may result in data loss.

- Version-level immutability support for Azure storage accounts is not supported.

- [For backup jobs managed by Veeam Agent] You cannot back up data to Microsoft Azure object storage added in the direct connection mode.

- [For Veeam Cloud Connect] Microsoft Azure object storage cannot be used as a cloud repository in the direct connection mode.

# Veeam Data Cloud Vault Limitations

Consider the following immutability limitations for Veeam Data Cloud Vault:

- Immutability is enabled by default for Veeam Data Cloud Vault and you cannot disable it.
- [For backup jobs managed by Veeam Agent] You cannot back up data to Veeam Data Cloud Vault added in the direct connection mode.
- [For Veeam Cloud Connect] Veeam Data Cloud Vault cannot be used as a cloud repository in the direct connection mode.

# Block Generation for GFS Backups

Consider the following immutability limitations for GFS backups:

- Monthly and yearly GFS backups do not have Block Generation unless you change the default values for Block Generation to exceed these backups.
- If you keep backups with multiple GFS flags assigned (weekly, monthly, and yearly) in the same object storage, Veeam Backup & Replication will not assign any Block Generation period for these backups.

## How Immutability Works

Immutability is the state of data that prevents it from being modified or deleted. Veeam Backup & Replication applies immutability to the state of a backup chain within a certain period of time. After you enable immutability, Veeam Backup & Replication prohibits deleting data from object storage repositories until the immutability expiration date comes. Immutability settings configured for object storage repository are applied to the whole backup chain and all its restore points.

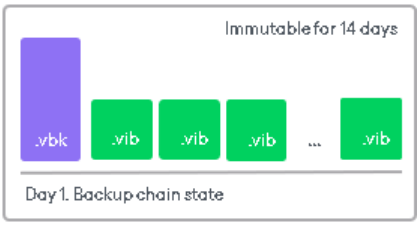


To make data immutable, Veeam Backup & Replication utilizes the technology that prevents data from deletion and allows you to keep several versions of objects. The selected technology depends on the type of object storage:

- Object lock and Versioning — for Amazon S3 Storage, S3 Compatible, IBM Cloud, Wasabi Cloud.
- Version-level WORM and blob versioning — for Azure Storage.

For more information, see Enabling Immutability.

# Immutability Period

Veeam Backup & Replication protects a backup chain and all its restore points during a certain immutability period. The immutability period is the number of days you have to respond to malicious actions. Within this period, you can roll back to the earlier state of your backup chain. To roll back data, you need to run Veeam PowerShell. For more information, see Rolling Back Immutable Data.

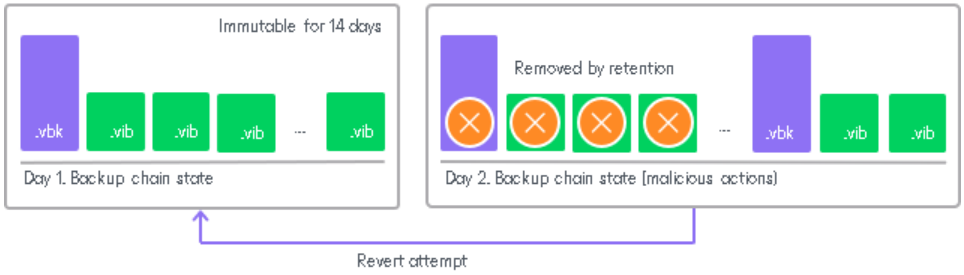


Before you configure immutability settings for your object storage repository, consider how immutability settings can influence the period within which you can restore data in case of malicious actions. The longer the immutability period, the more time you have to perform the necessary actions and roll back to the working state of your backup chain.

In both examples, a user configures a backup job with the same retention period and a backup method but a different immutability period. Thus, in the first scenario, the user has less time to react to malicious actions.

## Example 1. Short Immutability Period

In this scenario, a user configures a backup job with a short immutability period. Then, the user has only 3 days to react to malicious actions. If the user is not aware of these modifications and does not attempt to restore data within 3 days, all data will be lost.

The backup job is created with the following settings:

- The backup method is forever forward incremental.
- The job retention policy is set to 30 days.
- The immutability period is set to 3 days.

On the 2nd day of the month, an intruder changed the backup job settings. This action forced to apply a new job retention policy.

- The job retention policy is set to 1 day.
- An active full and a couple of increments are created.

After that, the previously created backup files are unavailable for restore from the Veeam Backup & Replication UI. Since the immutability period is set to 3 days, the object storage repository keeps all states of a backup chain up to 3 days. On the 5th day, the user attempts to revert back to the state it was before the attack, but since 3 days have already passed, the state of the backup chain is not available anymore.

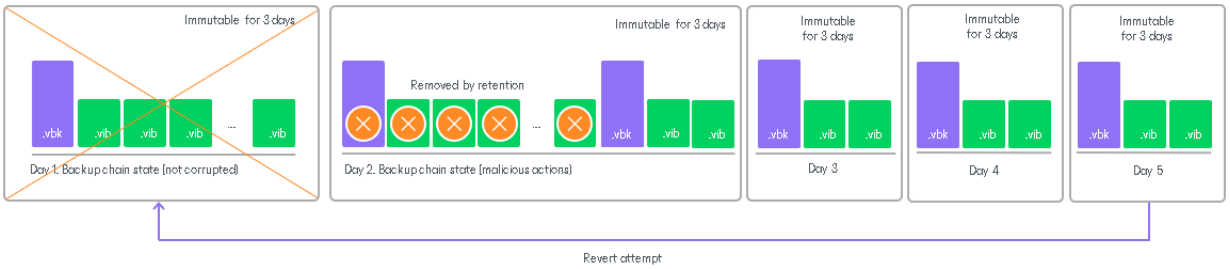

## Example 2. Long Immutability Period

In this scenario, a user has 7 days to react to malicious actions. If the user is not aware of the intruder actions and does not attempt to restore data within 7 days, all data will be lost.

A backup job is created with the following settings:

- The backup method is forever forward incremental.
- The job retention policy is set to 30 days.
- The immutability period is set to 7 days.

On the 2nd day of the month, an intruder changed the backup job settings. This action forced to apply a new job retention policy.

- The job retention policy is set to 1 day.
- An active full and a couple of increments are created.

After that, the previously created backup files are unavailable for restore from the Veeam Backup & Replication UI. The immutability period is set to 7 days, therefore, the object storage repository keeps all restore points for 7 days. On the 6th day, the user finds out the malicious actions took place, however the user still has 2 days to revert to the 1st day, find the necessary state of a backup chain and restore data.

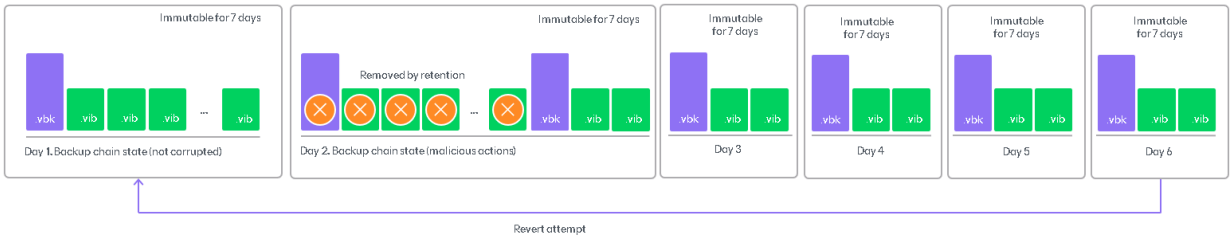


## Object Storage Actual Retention

When you configure an object storage repository, keep in mind that the immutability period affects the total number of restore points stored in the object storage repository.

> **IMPORTANT**
>
> Consider the following:
>
> - Although the immutability period does not depend on the retention of the backup chain, it will preserve more restore points in addition to the restore points stored according to a retention policy. This behavior guarantees consistency of earlier states of the backup chain and the ability to roll back to it.
> - The restore points removed according to the job retention policy are no longer available in the Veeam Backup & Replication UI. To view the immutable data, you need to run Veeam PowerShell. For more information, see Rolling Back Immutable Data.

Apart from the immutability period set for each object storage repository, Veeam Backup & Replication automatically adds several days to the immutability expiration date to reduce I/O operations and associated costs. This period is called Block Generation. You do not have to configure it, the Block Generation setting is applied automatically. For more information, see Block Generation.

Therefore, the actual retention should be calculated according to the following formula:

**Actual retention = job retention policy + immutability period + Block Generation period**

For example, the backup job is created with the following settings:

- The backup method is forever forward incremental.
- The job retention policy is set to 30 days.
- The immutability period is set to 14 days immutability.
- Default Block Generation period is 10 days (for Amazon S3 object storage and IBM Cloud object storage it is 30 days).

According to the backup job retention, an object storage repository will keep backups for 30 days. The immutability period adds extra 14 restore points as a delta that prolongs the job retention policy. Plus, 10 days from the Block Generation are on top. Thus, you need to plan storage that will be able to keep your backups for 30 + 14 + 10 days = 54 days.

> **IMPORTANT**
>
> Consider the following:
>
> - The type of the backup method affects the job retention policy and actual retention. For more information, see Short-Term Retention Policy.
> - For backups with GFS flags, backups created by VeeamZIP jobs, and exported backup files, the immutability period depends on the settings of the object storage repository and the retention settings of these backups:
>   - If immutability settings of the object storage repository exceed the retention period of these backups, VBR applies the immutability retention and ignores the retention of these backups.
>   - If the retention policy for these backups exceeds the maximum immutability period of the object storage repository, VBR applies the retention defined for these backups and ignores the immutability period.
> - If you add an object storage repository as an extent of the performance tier, immutability depends on the scale-out backups repository configuration. For more information, see Immutability for Performance Tier.

# Rolling Back Immutable Data

To restore data removed by the job retention policy, you need to run the Veeam PowerShell cmdlets. It will allow you to get information on a time period within which you can restore the state of a backup chain. This information will help you decide what time period suits you best and roll back the state of a backup chain. For detailed information, see description of the Get-VBRObjectStorageRepositorySyncInterval and Sync-VBRObjectStorageRepositoryEntityState cmdlets in the Veeam PowerShell Reference.

> **IMPORTANT**
>
> Rollback removes from the configuration database the state of the specific backup or an entire object storage backup repository. Then, it imports the state of the necessary backup or object storage backup repository.

For example, the intruder modified your backup job on 6/10/2024 at 5:30:41 PM, and you want to revert to that time and synchronize your object storage repository with this state.

1. Before you start to synchronize the data, you get details on the details on the time period within which you can roll back a backup chain. In this example. you can select any time period within this range: 1/10/2024 4:00:53 PM to 6/11/2024 1:08:39 PM.

```
$repository = Get-VBRObjectStorageRepository
Get-VBRObjectStorageRepositorySyncInterval -Repository $repository
StartDateUtc          EndDateUtc
------------          ----------
1/10/2024 4:00:53 PM 6/11/2024 1:08:39 PM
```

2. Once you get information on the time period, you can roll back the backup chain state before it was corrupted.

```
Sync-VBRObjectStorageRepositoryEntityState -Repository $repository -PointI
nTime "6/10/2024 1:04:17 PM"
 CreationTime : 2/20/2020 6:26:59 AM
 EndTime      : 2/20/2020 6:28:24 AM
 JobId        : 23e57bae-a759-42a2-a47e-06219ac410df
 Result       : Success
 State        : Stopped
 Id           : 10957ec5-adc9-418c-a708-2f24ba11d40e
```

## Block Generation

Block Generation is a time period within which all blocks in backup files (both full and increment backup files) have the same immutability period. Block Generation reduces the number of requests to the object storage repository, thereby saving traffic and reducing costs that can be incurred by your object storage provider. You do not have to configure it, the Block Generation setting is applied automatically.

Depending on the type of the object storage repository, Veeam Backup & Replication will add the following values for the default Block Generation period:

- 30 days — for Amazon S3 object storage and IBM Cloud object storage.
- 10 days — for all other types of object storage repositories.

For example, if you set your immutability period to 30 days for your object storage repository, Veeam Backup & Replication will add 10 days to specific objects to reduce I/O operations with the data blocks over time. Thus, you will have immutability set for 30 days + 10 days of Block Generation set for data blocks in your object storage repositories.

> **TIP**
>
> You can change the default Block Generation period by modifying the registry value on the backup server. For more information, see this Veeam KB article.

# How Block Generation Works

When the first data block (a full backup) arrives, its immutability period by default is set to 30 + 10 = 40 days. The first full backup starts its generation, that will be appended with the incremental backups. All the incremental backups within the generation (that is, within the 10-days period) will have the same immutability expiration date as the full backup. For instance, a data block that was offloaded on day 9 will have the same immutability expiration date as a data block offloaded on day 1. Thus we ensure that the immutability period for all the data blocks within a generation is no less than 30 days.

To maintain the backup consistency, Veeam Backup & Replication can extend immutability expiration for all data blocks in all backup chains (both active and inactive) and assign these blocks to a new generation. For example, within one forward incremental backup chain, a full backup file can not be removed before an incremental backup file. On the other hand, an incremental backup file makes no sense without relevant full backup file. So the immutability period is extended for all data blocks in the backup chain.

> **NOTE**
>
> Consider the following:
>
> - The immutability period differs from the actual retention of the object storage repository. For more information, see Object Storage Actual Retention.
> - For data blocks located in object storage repositories, Veeam Backup & Replication extends the immutability period for every data block of every backup file in the whole backup chain, even in an inactive part.
> - Veeam Backup & Replication will not extended immutability for the data blocks that are not used in any existing backup files.

# Block Generation for GFS Backups

For GFS backups, Veeam Backup & Replication checks the GFS retention period and sets the Block Generation period if it is longer than the minimum GFS retention period. In this case, the immutability period for GFS backups is not less than the GFS retention period and does not exceed the GFS retention period + the Block Generation period.

Considering that the default Block Generation period is either 10 or 30 days, the Block Generation is assigned for weekly GFS backups by default. This behavior is particularly significant for weekly GFS backups, as there is a high probability that full backup files will share a large number of blocks since only one week passes between the first and the second weekly GFS backups. If Veeam Backup & Replication does not assign Block Generation for weekly GFS backups, it will need to extend immutability for the shared blocks each week, resulting in additional I/O operations and costs.

The actual retention of the immutable weekly GFS backups depends on the on the Block Generation period (10 or 30 days). Consider the following examples for weekly GFS backups with different Block Generation periods.

## Example 1. Weekly GFS Backups and Block Generation is 30 Days

The retention period of a weekly GFS backup is 28 days, and the Block Generation period is 30 days, so the actual retention is 58 days. The first weekly GFS backup is created on *1/1/2025* and starts Generation 1, which will be protected until *2/28/2025* and will take 5 weeks to expire. The Block Generation period is reduced every week by 7 days. Generation 2 starts on *2/5/2025* with the 6th weekly GFS backup and will be protected until *4/4/2025*.

### Example 2. Weekly GFS Backup and Block Generation is 10 Days

The retention period of a weekly GFS backup is 28 days, and the Block Generation period is 10 days, so the actual retention is 38 days. The first weekly GFS backup is created on *1/1/2025* and starts Generation 1, which will be protected until *2/8/2025* and will take 2 weeks to expire. Block generation period is reduced every week by 7 days. Generations 2 starts on *1/15/2025* with the 3rd weekly GFS backup and will be protected until *2/22/2025*.

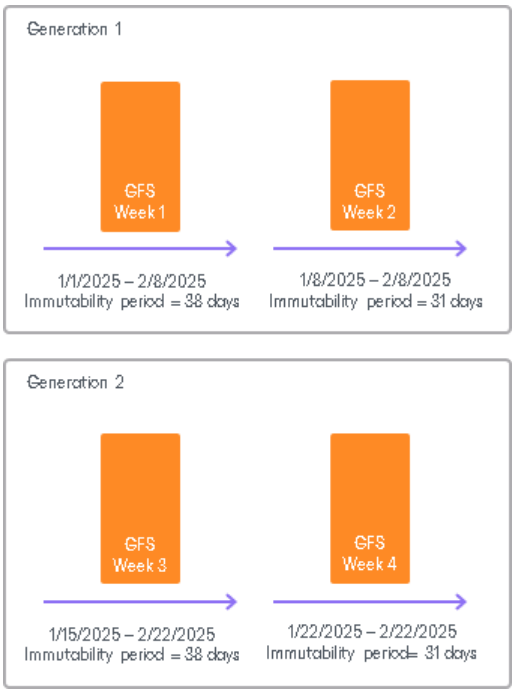


## Enabling Immutability

To enable immutability, you must do the following:

1. Configure the following settings when you create an S3 bucket or Azure container:
   - **Amazon S3 Storage, S3 Compatible, IBM Cloud, Wasabi Cloud** — You must enable the Object Lock and Versioning features on your S3 bucket when you create the bucket.

     > **IMPORTANT**
     >
     > Note that most vendors allow enabling *Object Lock* only at the moment of creating the bucket.

     For more information on enabling the *Object Lock* and *Versioning* features, see these Amazon articles: Creating a bucket, Using S3 Object Lock and Enabling versioning on buckets.
   - **Azure Storage** — You must enable support for version-level WORM on the container and enable blob versioning for your storage account when you create a storage account. For instruction on how to configure your Azure storage account with the necessary settings, see this Veeam KB article.

     For more information on enabling version-level WORM for a container, see Microsoft Docs.

     For more information on blob versioning for a storage account, see Microsoft Docs.

> **IMPORTANT**
>
> When you create the storage account, by default the version-level immutability support option is enabled. You must disable it, otherwise immutability will not be applied for your Azure object storage. For more information, see Microsoft Docs.

2. Enable the immutability option when you add an object storage repository to the backup infrastructure at the **Container** step (for Azure object storage repository) or **Bucket** step (for Amazon S3 or S3 compatible object storage repositories) of the new Object Storage Repository wizard.

# Adding Object Storage Repositories

You can add the following types of object storage repositories:

- Veeam Data Cloud Vault
- S3 Compatible Object Storage, S3 Compatible Object Storage with Data Archiving
- Amazon S3 Object Storage, Amazon S3 Glacier Storage and AWS Snowball Edge
- Google Cloud Object Storage
- Microsoft Azure Object Storage, Microsoft Azure Archive Storage and Data Box
- IBM Cloud Object Storage
- Wasabi Cloud Object Storage
- Veeam Smart Object Storage API (SOSAPI)

## Adding Veeam Data Cloud Vault

This section describes how to add Veeam Data Cloud Vault to the backup infrastructure. Veeam Data Cloud Vault is a secure cloud solution that leverages Azure Blob storage functionality. It provides enhanced security features such as in-built immutability and encryption. Veeam Data Cloud Vault allows you to store and maintain data in a secure and protected environment. Veeam Backup & Replication users can add Veeam Data Cloud Vault storage as an object storage repository, which can be used either for direct backup or as part of a scale-out backup repository.

> **IMPORTANT**
>
> Consider the following:
>
> - To be able to back up to Veeam Data Cloud Vault from Veeam Backup & Replication, you must obtain Veeam Data Cloud Vault and configure its subscription. For more information, see the Veeam Data Cloud Vault User Guide.
> - To be able to use the Veeam Data Cloud Vault functionality, you must have the License Administrator role. For more information, see this Veeam KB article.
> - When you register Veeam Backup & Replication, make sure that the email address that you use is matches the email address of the License Administrator. (case-sensitive, character-perfect)

To add Veeam Data Cloud Vault, use the **New Object Storage Repository** wizard.

## Step 1. Launch New Object Storage Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Veeam Data Cloud Vault**. Alternatively, you can select **Object storage** > **Veeam Data Cloud Vault.**

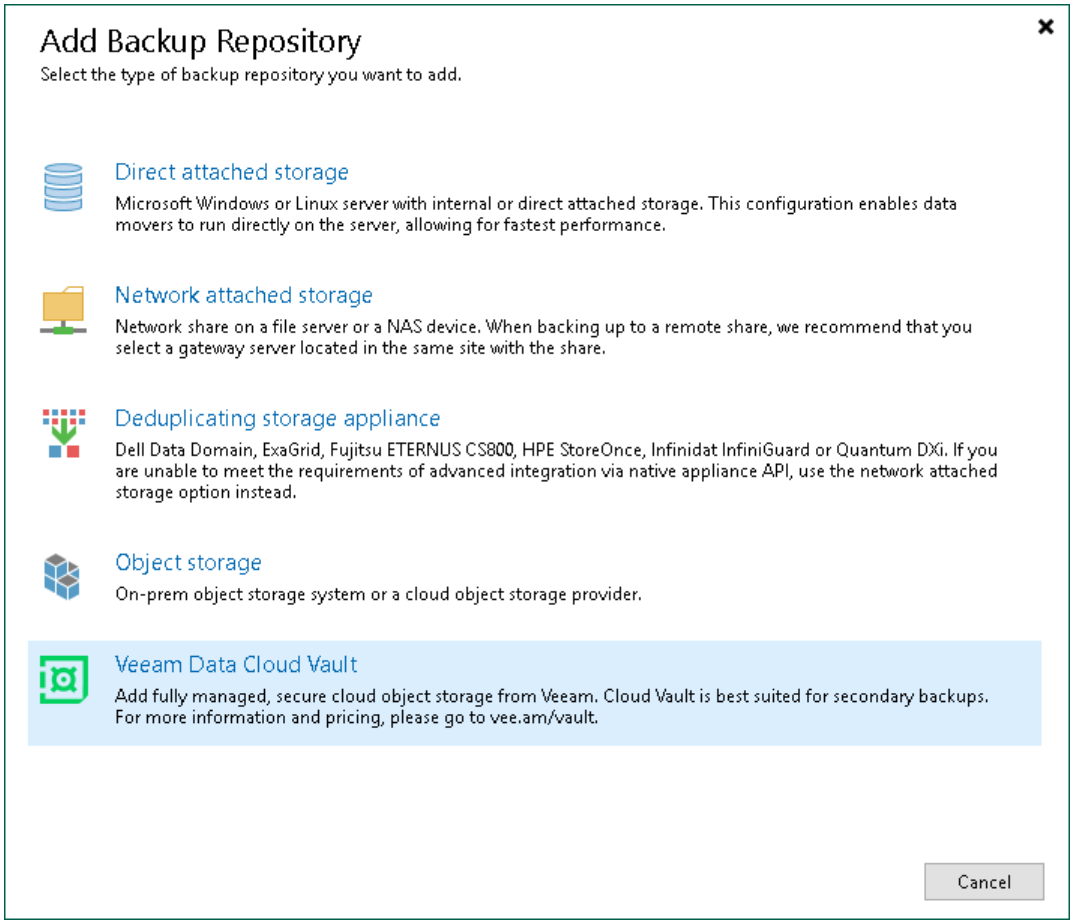

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

If you want to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

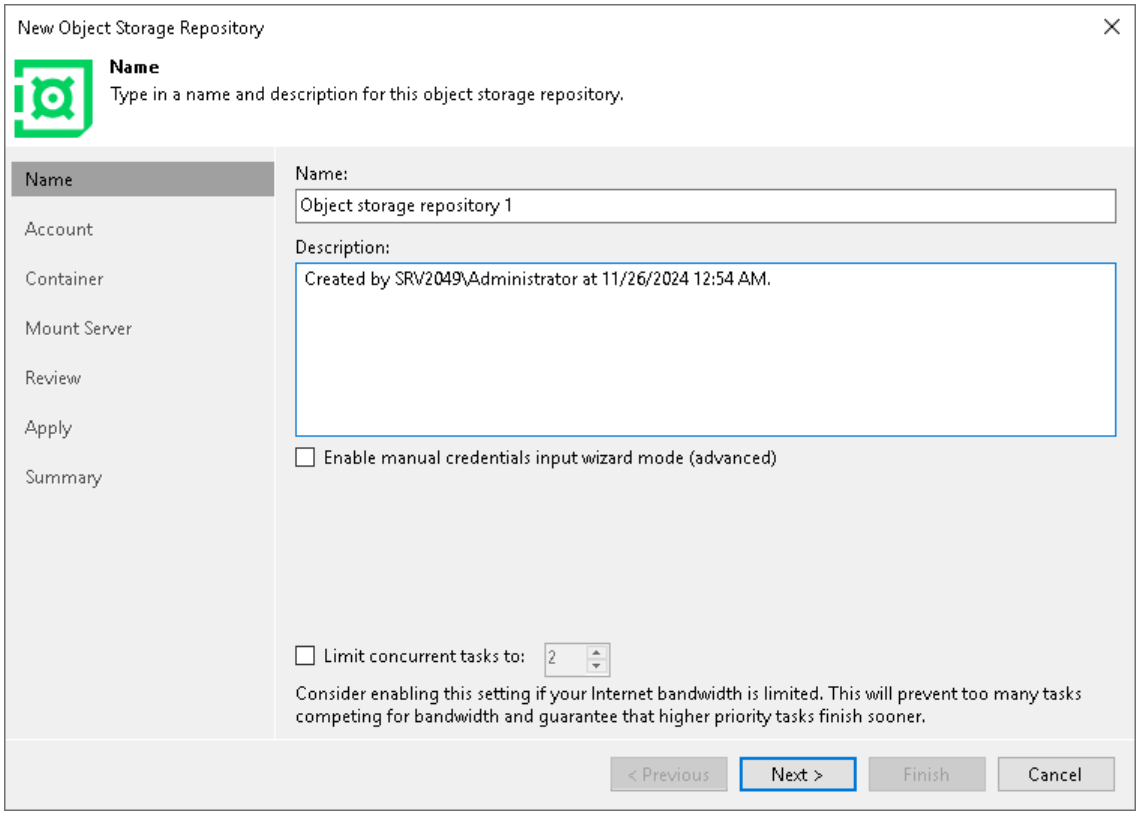

### Step 3. Specify Veeam Account

At the **Account** step of the wizard, specify credentials for the Veeam account and the storage vault.

## Specifying User Credentials

To specify the authorization settings:

1. To authorize using your Veaam credentials, click the **Authorize** link.
2. In the **Veeam authorization** window enter your credentials for the Veeam account. After that, your backup server will be registered in the My Account portal.

   > **NOTE**
   >
   > Ensure that the email address you use matches the email address of the License Administrator. Consider that this email address is case-sensitive. If you plan to delegate access to another user, make sure that he also has the License Administrator role.

3. In the **Vault** drop-down list, specify the storage vault that you want to use. To manage vaults, click the **Manage** link. For more information, see the Managing Subscriptions section in the Veeam Data Cloud Vault User Guide.

   > **NOTE**
   >
   > If you have only one storage vault that has not been associated with any backup server before, this vault will appear in the drop-down list automatically.

   If you do not have any storage vaults associated with your backup server, the **Vault** drop-down list will be empty. For detailed instructions on how to obtain a vault and assign it to your backup server, see the Getting Started with Veeam Data Cloud Vault section in the Veeam Data Cloud Vault User Guide. After you add a storage vault and assign it to the backup repository, click the refresh icon (). The vault will show up in the drop-down list.

   > **IMPORTANT**
   >
   > Keep in mind that it might take several minutes to add a certificate (a public key) on the Microsoft Azure Entra ID side.

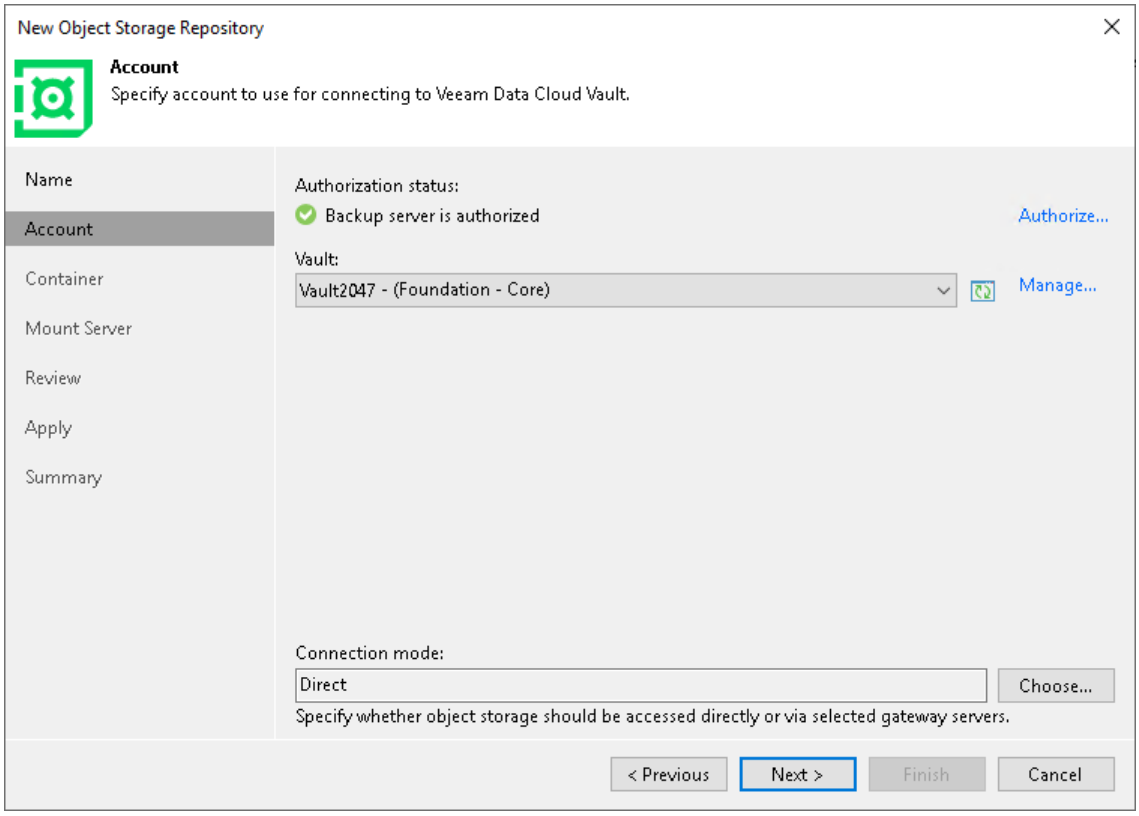


# Specifying Connection Mode

To specify connection settings, next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:

- **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.

- **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the Name list, select gateway servers that you want to use for data transfer operations.

  By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

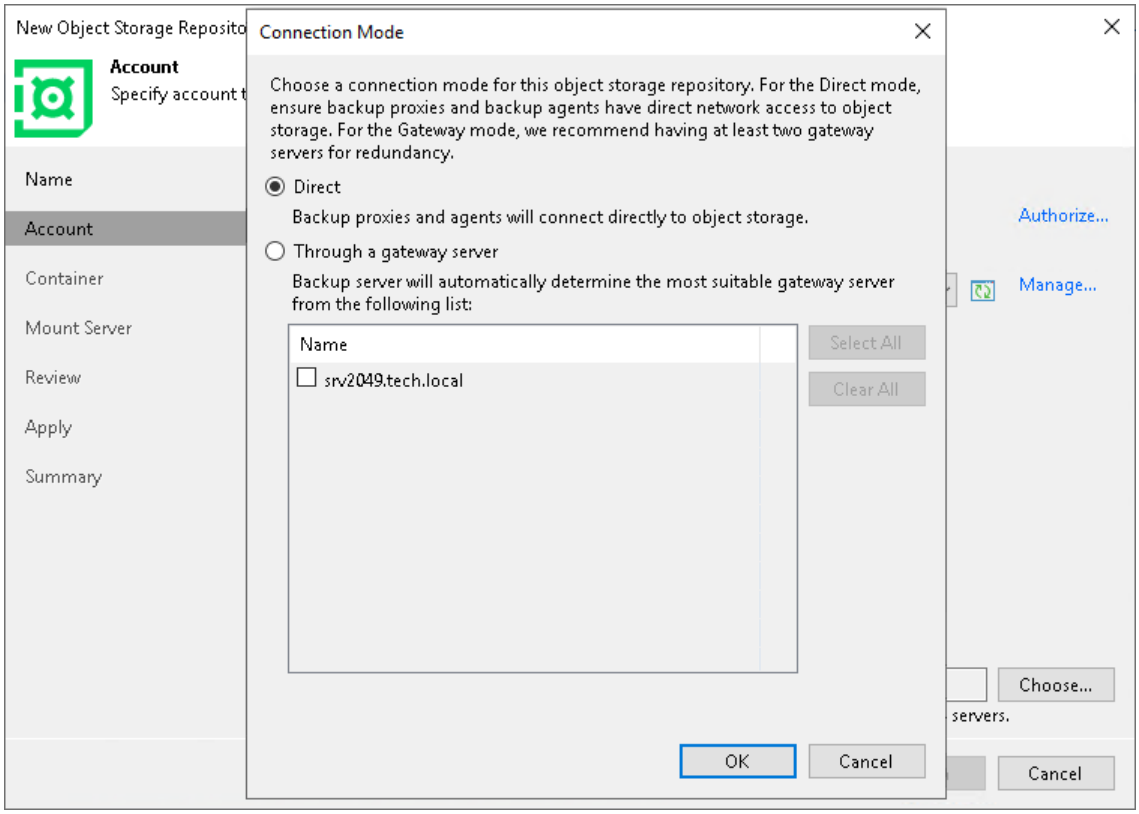

## Step 4. Specify Object Storage Settings

At the **Container** step of the wizard, specify the folder that will be used to store data, the storage consumption and the immutability period.

1. To the right of the **Folder** field, click **Browse** and either select an existing folder or click **New Folder**.
2. Select the **Limit object storage consumption** to check box to define a soft limit for your object storage consumption. If this limit is exceeded during a job run, Veeam Backup & Replication will complete the job. However, a new job will not be able to start unless you remove the extra data that exceeds the limit or change the soft limit settings. Provide the value in TB or PB.
3. Right to the **Make recent backups immutable** for check box, specify the immutability period.

> **NOTE**
>
> Consider the following:
>
> - By default, immutability is enabled for 30 days for Veeam Data Cloud Vault. You cannot disable this option and cannot remove data during this period.
> - You can set the immutability period to different values in the Veeam Backup & Replication UI. The minimum immutability period is 1 day, and the maximum is 90 days. If you want to set immutability to a longer period, use the Set-VBRDataCloudVaultRepository cmdlet.
> - The restore points removed according to the job retention policy are no longer available in the Veeam Backup & Replication UI. To view the immutable data, you need to run Veeam PowerShell. For more information, see Rolling Back Immutable Data. For more information on how immutability works, see Immutability for Object Storage Repositories.

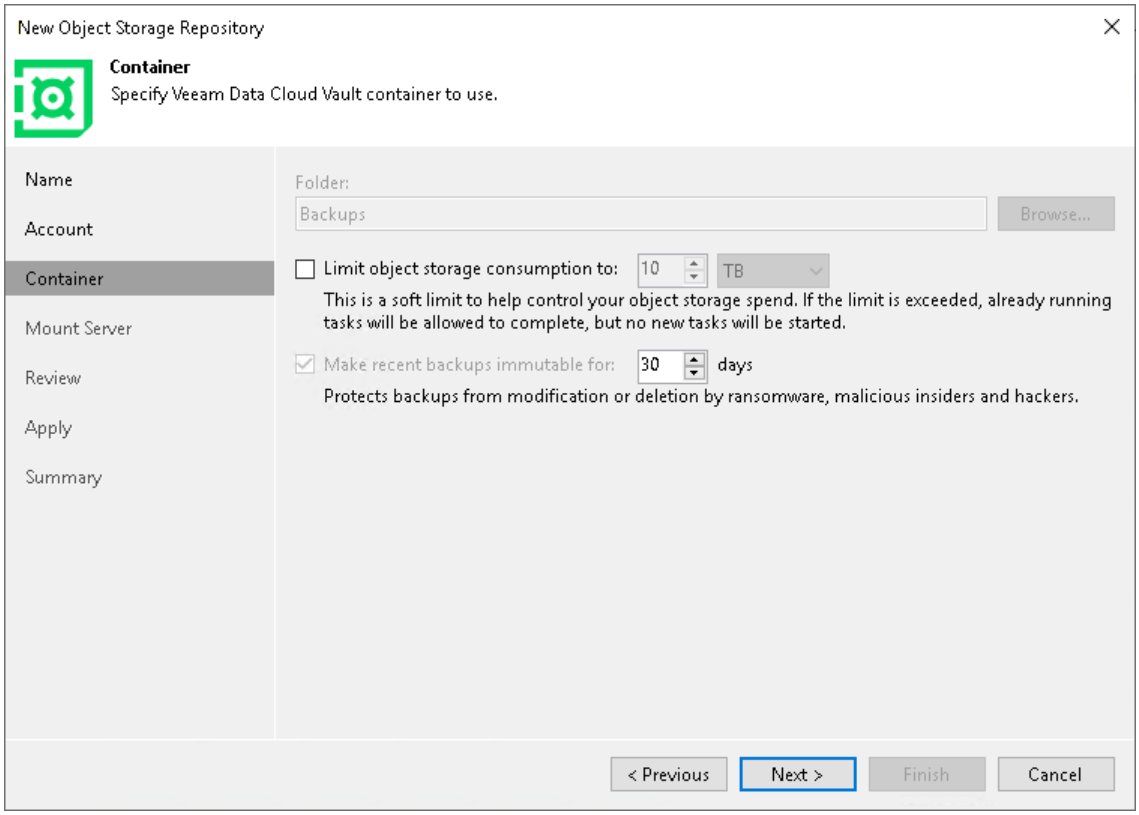

## Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for restore operations.

To specify the mount server settings, do the following:

1. From the **Mount Server** drop-down list, select a server that you want to use as a mount server. Veeam Backup & Replication uses this server during restore operations to mount VM disks directly from objects located in object storage repositories. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, click **Add New** on the right to open the **New Windows Server wizard**. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder to keep cache that is created during mount operations.

3. Select the **Enable vPower NFS service on the mount server** check box to allow the Veeam vPower NFS Service access the object storage repository. Veeam Backup & Replication will enable the Veeam vPower NFS Service on the necessary mount server. For more information, see Veeam vPower NFS Service.

4. Click **Ports** to customize network ports used by the Veeam vPower NFS Service. In the **vPower NFS Port Settings** window, specify the following settings:

   - Next to the **Mount Port** section, specify the port that the Veeam vPower NFS Service will use to mount the vPower NFS datastore to the ESXi host.
   - Next to the **vPower NFS** port section, specify the port that the Veeam vPower NFS Service will use to connect to the target NFS share.

   For information on ports used by default, see Ports.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

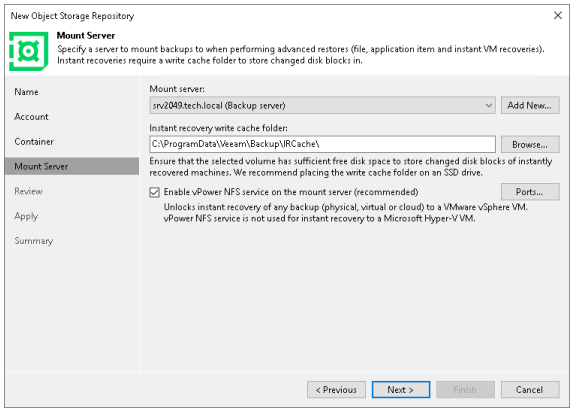

## Step 6. Review Components

At the **Review** step of the wizard, review what components will be processed on the mount server server and their status.

If the backup repository contains backups, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Object Storage (Imported)** node.

If the backup repository contains guest file system index files, select the **Import guest file system index data to the catalog** check box. Veeam Backup & Replication will import index files together with backup files, and you will be able to search for guest OS files inside imported backups. For more information, see the Guest OS File Restore section of the Veeam Backup Enterprise Manager Guide.

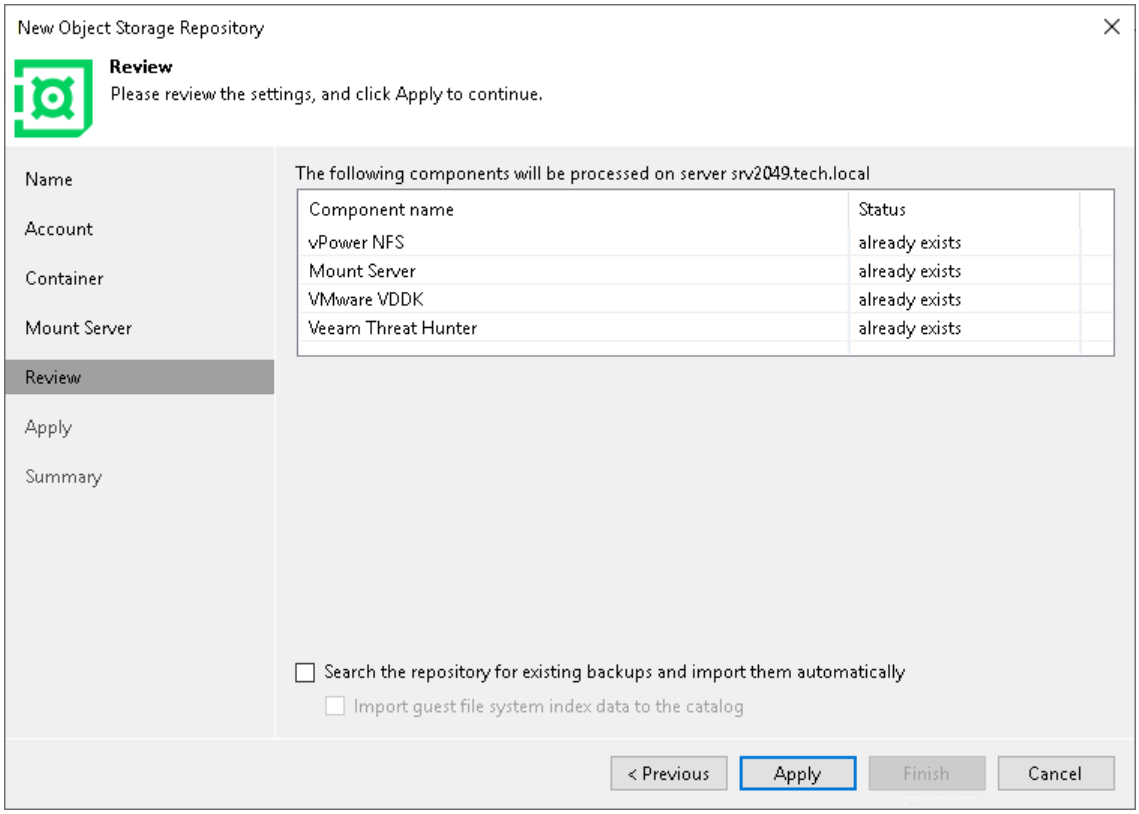

## Step 7. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to complete saving your settings to the configuration database and create backup infrastructure objects.

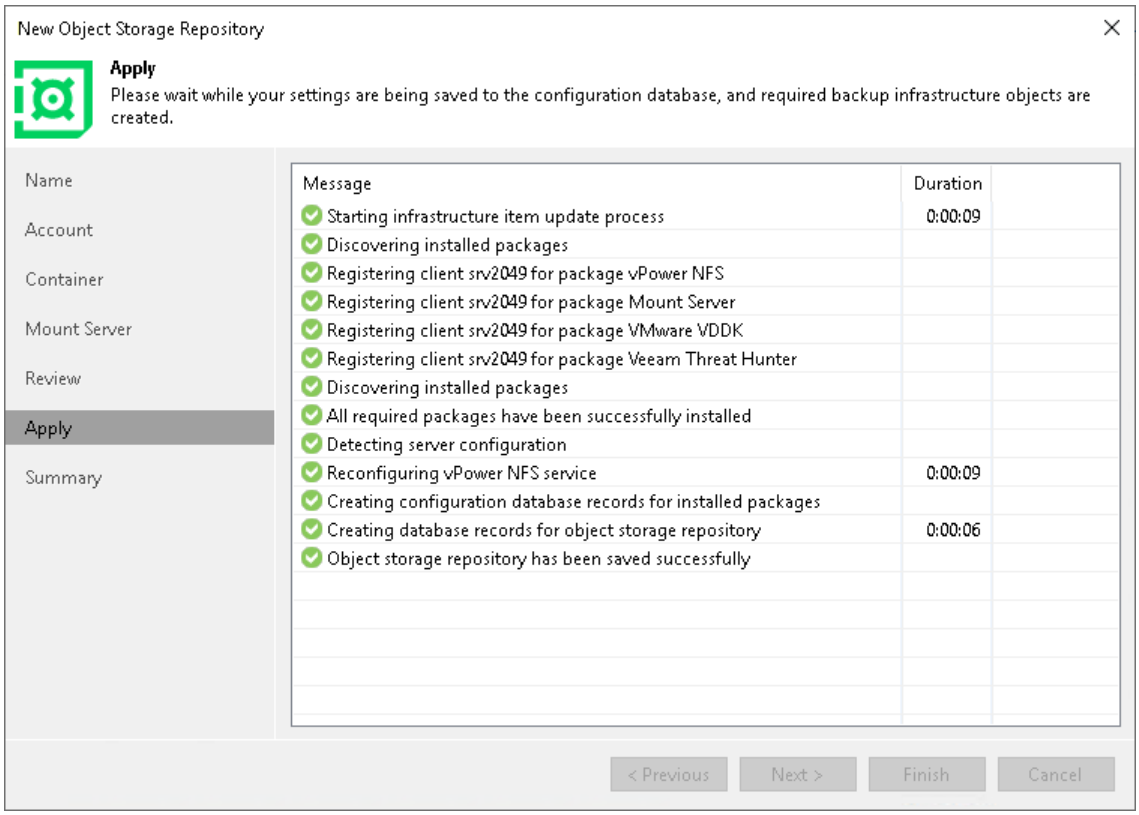

### Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

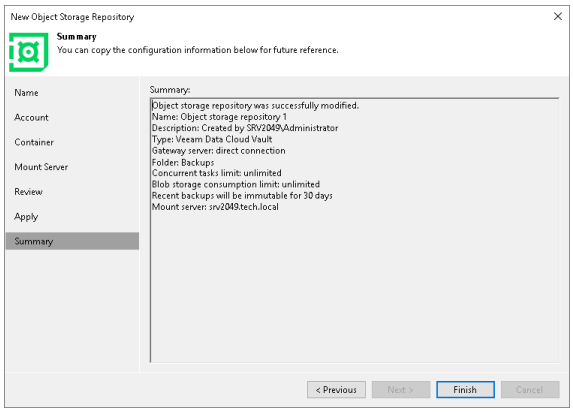


## Adding S3 Compatible Object Storage

Before you add an S3 compatible object storage and S3 compatible object storage with data archiving, check prerequisites and required permissions. After that, use the **New Object Storage Repository** wizard.

### Adding S3 Compatible Object Storage

This section describes how to add S3 compatible object storage to the backup infrastructure.

To add S3 compatible object storage, use the **New Object Storage Repository** wizard.

## Step 1. Launch New Object Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **S3 Compatible** > **S3 Compatible**.

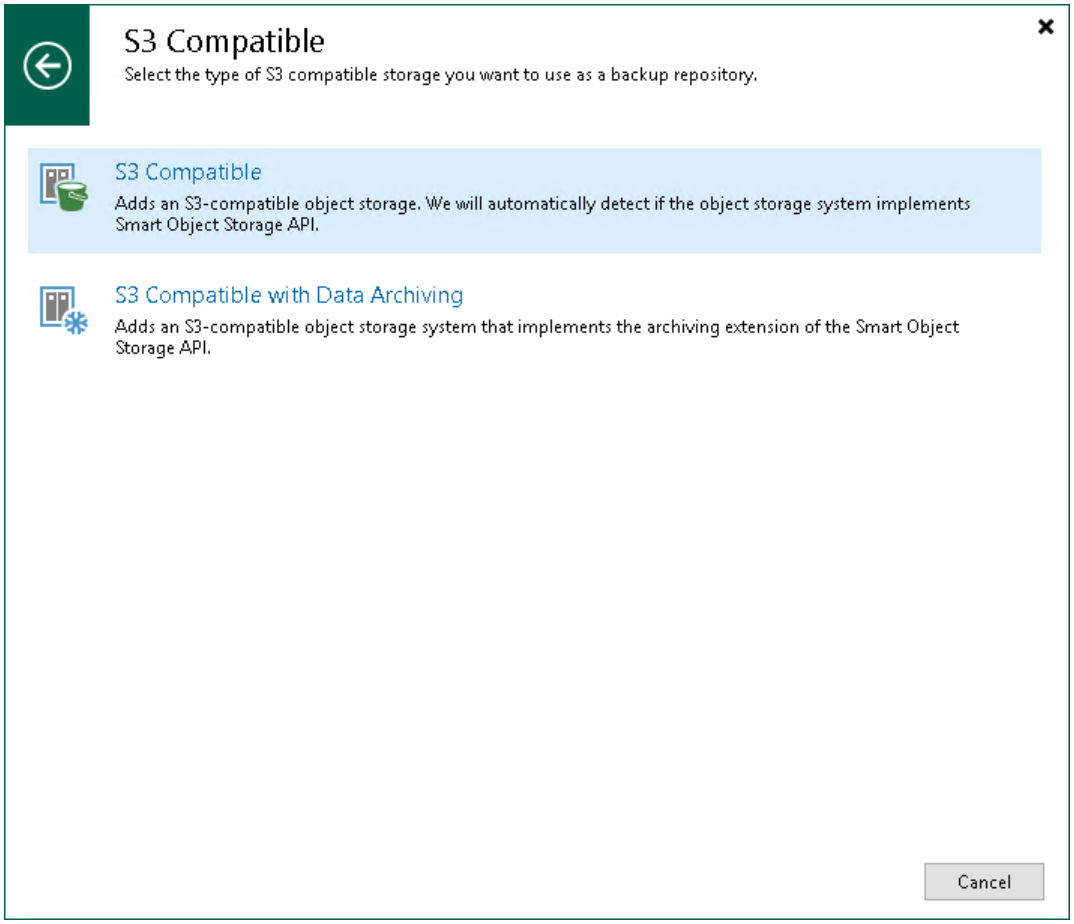

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

If you want to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

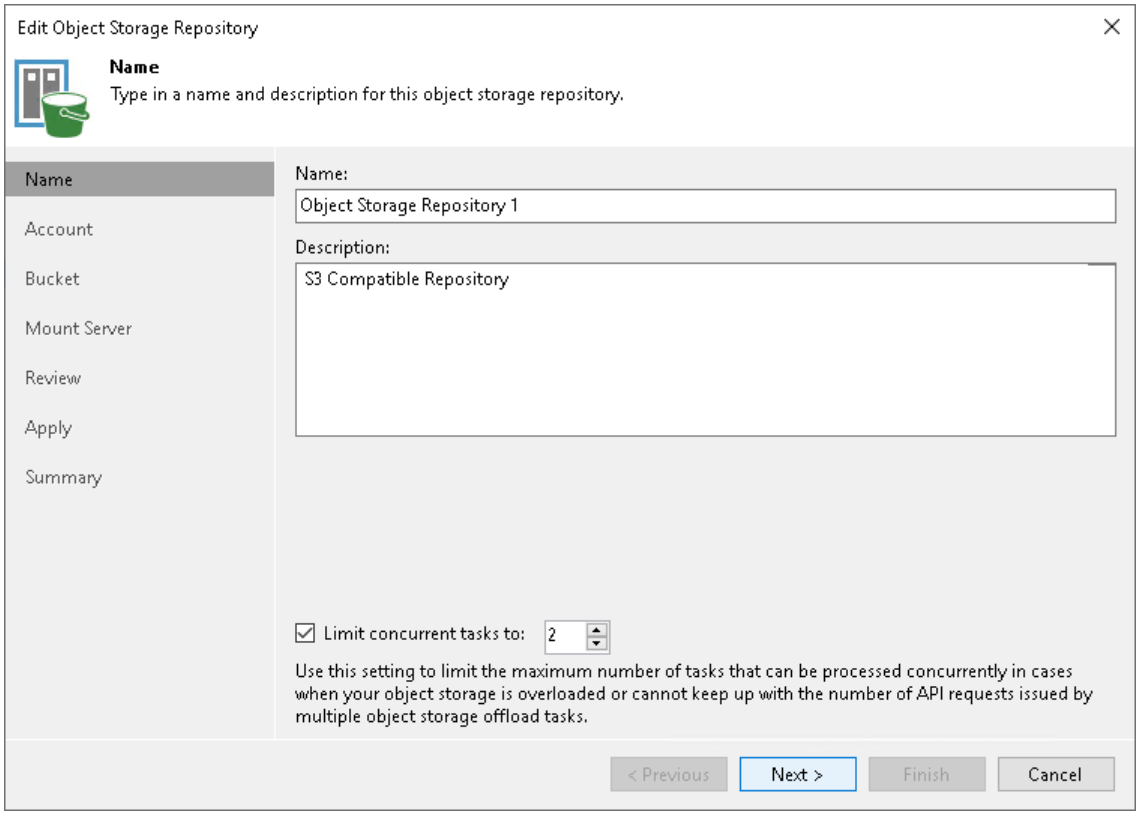

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. In the **Service point** field, specify an endpoint address and a port number of your S3 compatible object storage.
2. In the **Region** field, specify a region.
3. From the **Credentials** drop-down list, select user credentials to access your S3 compatible object storage.

   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Access Keys for AWS Users. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.
4. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:
   - **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.
   - **Through gateway server** – select this option if you want Veeam Backup & Replication to use a gateway server to transfer data from processed VM or file share to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

     By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your Veeam Backup & Replication infrastructure and has internet connection. Note that you must add the server to the Veeam Backup & Replication infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

     > **NOTE**
     >
     > By default, if Veeam Agent stores data in S3 compatible objects storage repositories, it transfers data using a gateway server. If you want Veeam Agent to access repositories directly or using specific credentials, you must specify the Access Permissions settings.

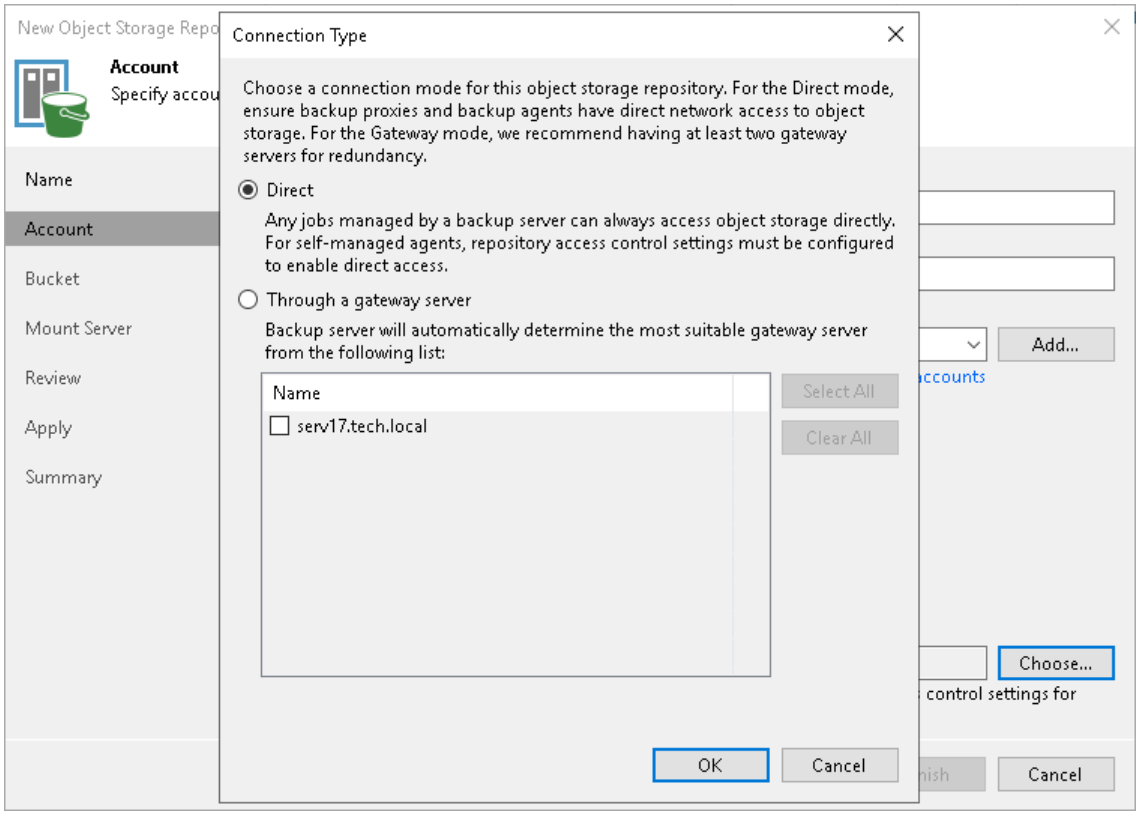
New Object Storage Repo
Account
Specify accou
Name
Account
Bucket
Mount Server
Review
Apply
Summary
Add...
accounts
Choose...
control settings for
Cancel
Connection Type
Choose a connection mode for this object storage repository. For the Direct mode, ensure backup proxies and backup agents have direct network access to object storage. For the Gateway mode, we recommend having at least two gateway servers for redundancy.
Direct
Any jobs managed by a backup server can always access object storage directly. For self-managed agents, repository access control settings must be configured to enable direct access.
Through a gateway server
Backup server will automatically determine the most suitable gateway server from the following list:
Name
serv17.tech.local
Select All
Clear All
OK
Cancel

## Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify the bucket and folder where you will store data, and define storage limits and immutability settings that Veeam Backup & Replication will apply to data in the object storage.

1. In the **Bucket** field, enter a name of the bucket or click **Browse** to get the necessary bucket.

   Note that you must create the bucket where you want to store your backup data beforehand.

2. By default, the multiple child buckets creation is enabled. To manage its settings, do the following:

   a. Click the **Automatic bucket creation disabled** link.

   b. To specify a number of per-machine backups chains that you want to keep in a single child bucket, check the **Create new buckets automatically** check box. Veeam Backup & Replication will create child buckets according to this amount of backup chains.

2. To the right of the **Folder** field, click **Browse** and either select an existing folder or click **New Folder**.

3. Select the **Limit object storage consumption to** check box to define a soft limit for your object storage consumption. If this limit is exceeded during a job run, Veeam Backup & Replication will complete the job. However, a new job will not be able to start unless you remove the extra data that exceeds the limit or change the soft limit settings. Provide the value in TB or PB.

4. Select the **Make recent backups immutable** for check box to prohibit deletion of blocks of data from object storage. Specify the immutability period. For more information, see Immutability for Object Storage Repositories.

   Note that the maximum immutability period you can set in the Veeam Backup & Replication UI is 90 days. If you want to set immutability to a longer period, use the Set-VBRAmazonS3CompatibleRepository cmdlet.

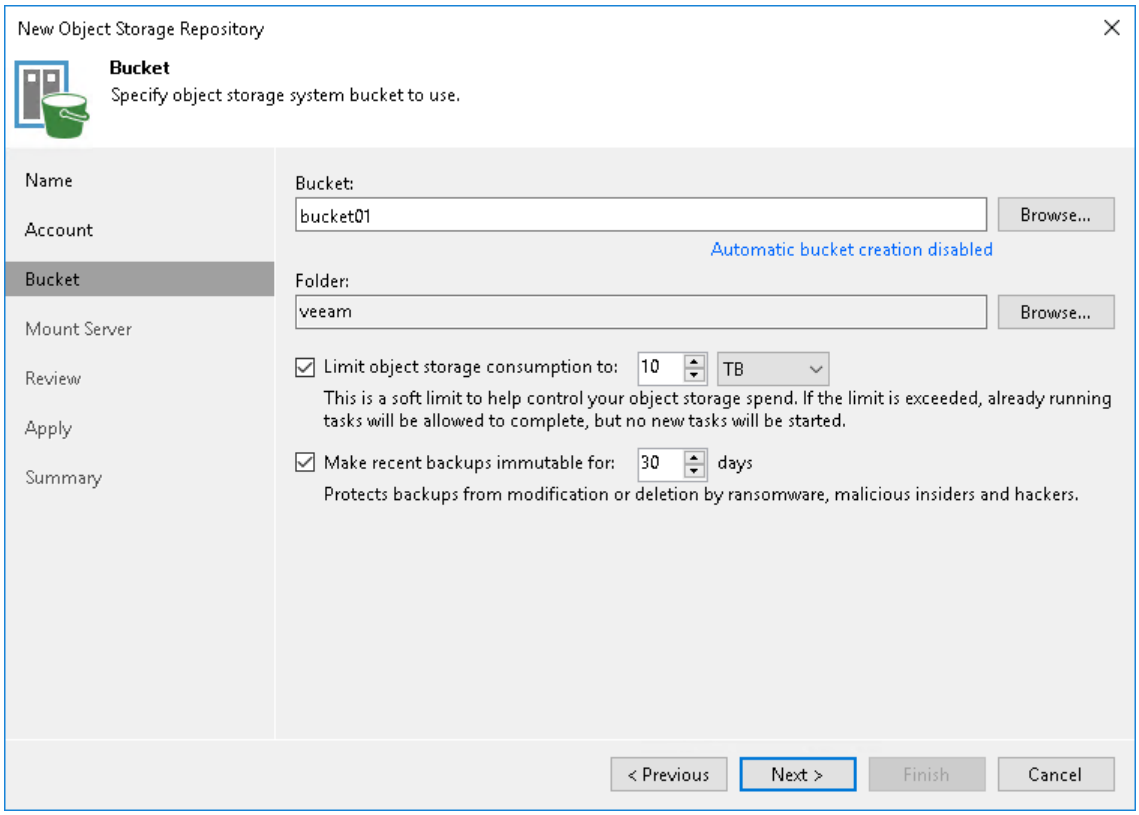

### Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for restore operations and configure a helper appliance. The helper appliance is a Windows-based or Linux-based virtual or physical server, added to the backup infrastructure, that Veeam Backup & Replication uses to perform a health check of backup files and apply retention to unstructured data backup files. For more information, see Health Check for Object Storage Repositories and Helper Appliance in Unstructured Data Backup.

> **IMPORTANT**
>
> Consider the following:
>
> - If you do not configure a helper appliance, Veeam Backup & Replication will use local resources to perform the health check and apply retention to unstructured data backup files. It will consume more cloud resources and can result in additional costs.
> - To perform the health check, you must enable this option when you configure a job. For more information, see Health Check for Backup Files.

# Specifying Mount Server Settings

To specify the mount server settings, do the following:

1. From the **Mount Server** drop-down list, select a server that you want to use as a mount server. Veeam Backup & Replication uses this server during restore operations to mount VM disks directly from objects located in object storage repositories. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, click **Add New** on the right to open the **New Windows Server wizard**. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder to keep cache that is created during mount operations.

3. Select the **Enable vPower NFS service on the mount server** check box to allow the Veeam vPower NFS Service access the object storage repository. Veeam Backup & Replication will enable the Veeam vPower NFS Service on the necessary mount server. For more information, see Veeam vPower NFS Service.

4. Click **Ports** to customize network ports used by the Veeam vPower NFS Service. In the **vPower NFS Port Settings** window, specify the following settings:

   - Next to the **Mount Port** section, specify the port that the Veeam vPower NFS Service will use to mount the vPower NFS datastore to the ESXi host.
   - Next to the **vPower NFS** port section, specify the port that the Veeam vPower NFS Service will use to connect to the target NFS share.

   For information on ports used by default, see Ports.

5. To specify the helper appliance settings, click **Configure**. From the **Managed server** drop-down list, select a server that you want to use as the helper appliance.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

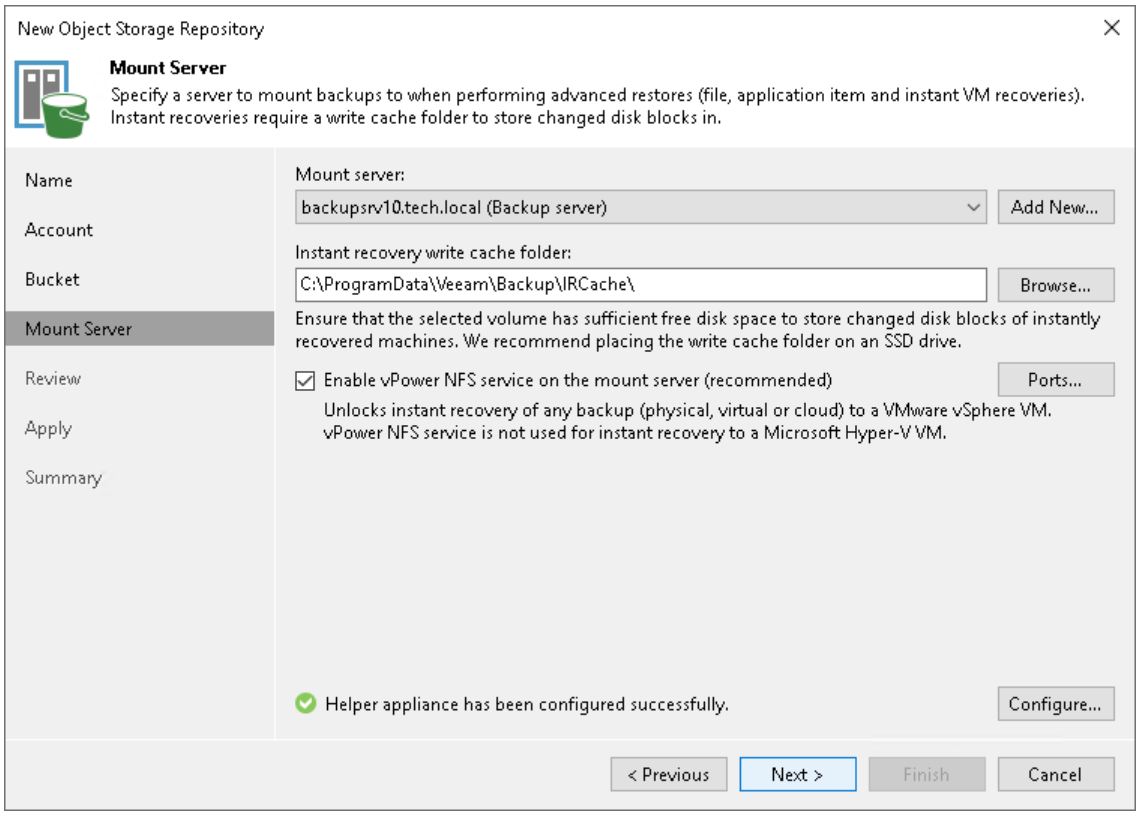
New Object Storage Repository
Mount Server
Specify a server to mount backups to when performing advanced restores (file, application item and instant VM recoveries). Instant recoveries require a write cache folder to store changed disk blocks in.
Name
Account
Bucket
Mount Server
Review
Apply
Summary
Mount server:
backupsrv10.tech.local (Backup server)
Add New...
Instant recovery write cache folder:
C:\ProgramData\Veeam\Backup\IRCache\
Browse...
Ensure that the selected volume has sufficient free disk space to store changed disk blocks of instantly recovered machines. We recommend placing the write cache folder on an SSD drive.
Enable vPower NFS service on the mount server (recommended)
Ports...
Unlocks instant recovery of any backup (physical, virtual or cloud) to a VMware vSphere VM. vPower NFS service is not used for instant recovery to a Microsoft Hyper-V VM.
Helper appliance has been configured successfully.
Configure...
< Previous
Next >
Finish
Cancel

## Step 6. Review Components

At the **Review** step of the wizard, review what components will be processed on the mount server server and their status.

If the backup repository contains backups, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Object Storage (Imported)** node.

If the backup repository contains guest file system index files, select the **Import guest file system index data to the catalog** check box. Veeam Backup & Replication will import index files together with backup files, and you will be able to search for guest OS files inside imported backups. For more information, see the Guest OS File Restore section of the Veeam Backup Enterprise Manager Guide.

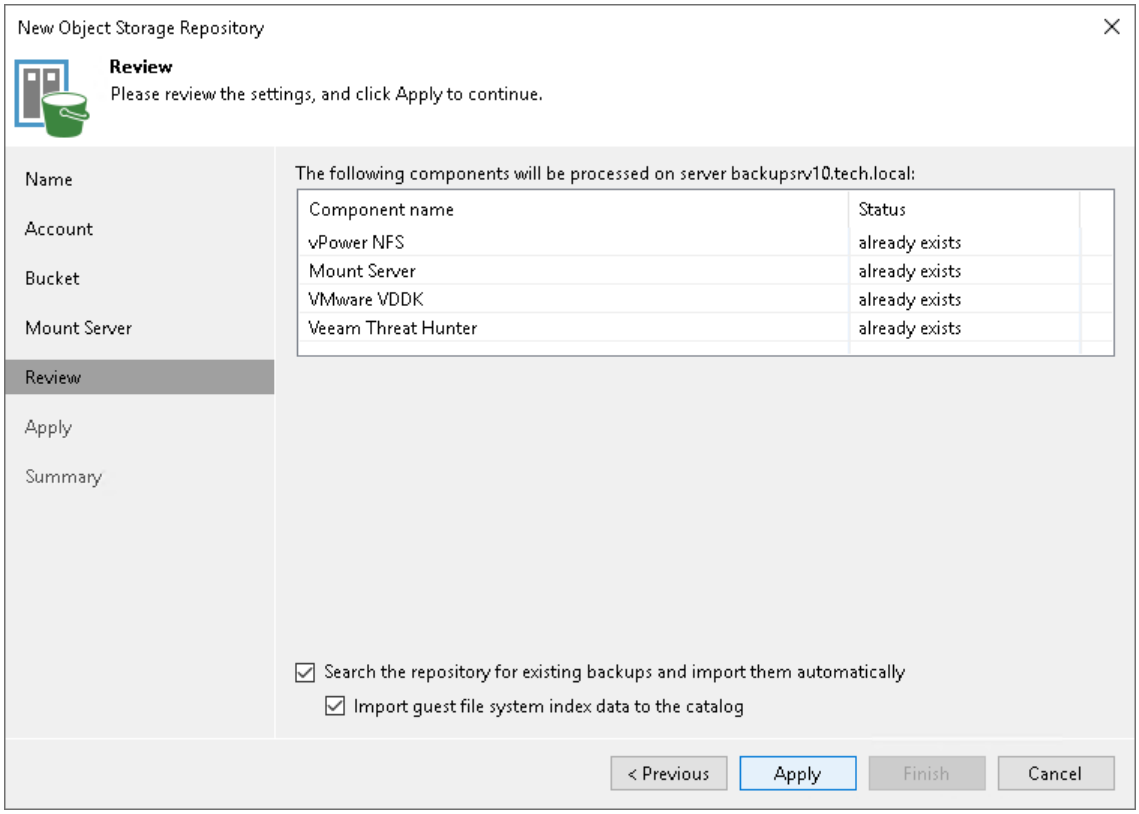

## Step 7. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to complete saving your settings to the configuration database and create backup infrastructure objects.

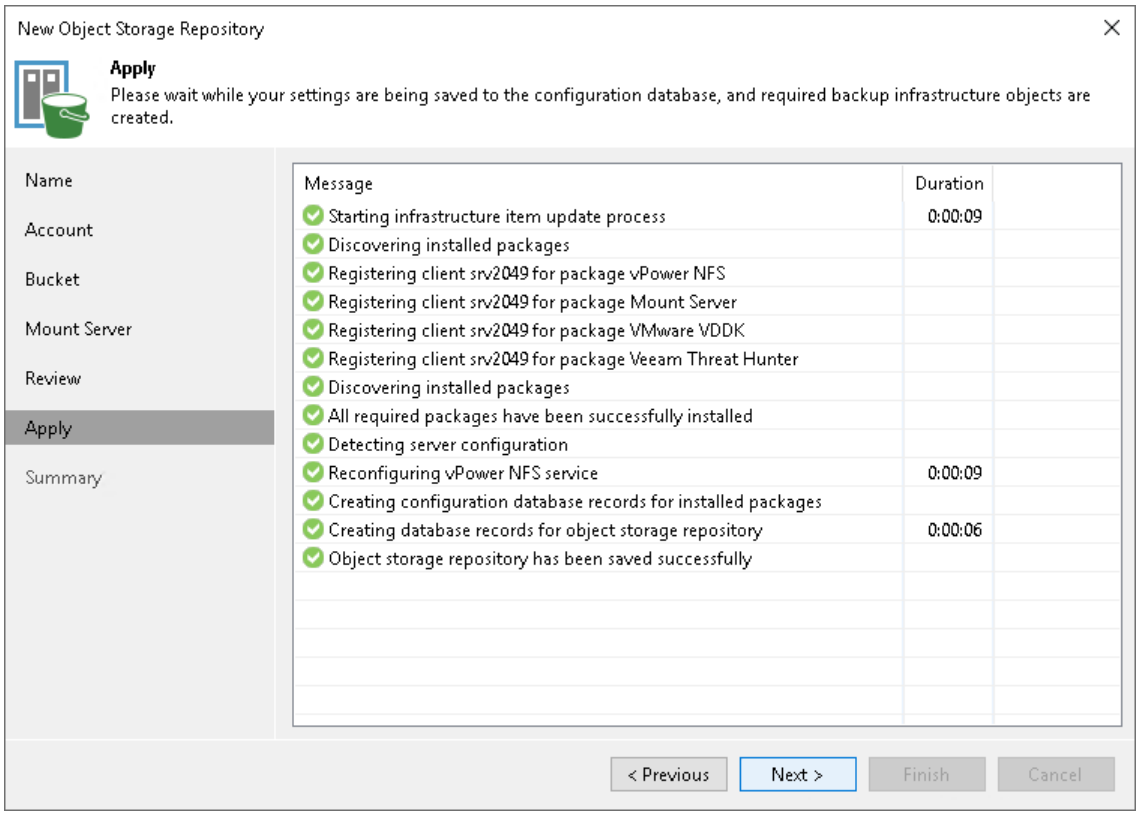

## Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

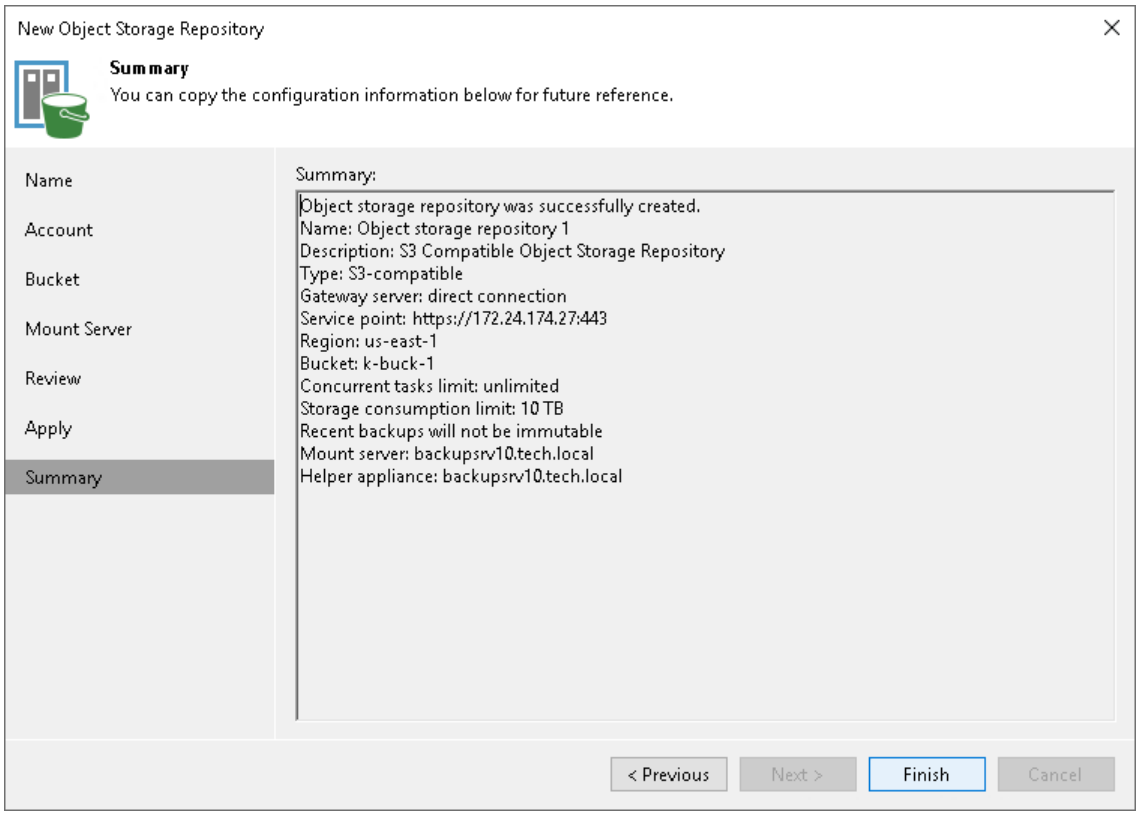


## Adding S3 Compatible with Data Archiving

This section describes how to add S3 compatible object storage with data archiving to the backup infrastructure. You can only use this repository as an archive extent of the scale-out backup repository. For more information, see Archive Tier.

To add S3 compatible object storage with data archiving, use the **New Object Storage Repository** wizard.

## Step 1. Launch New Object Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **S3 Compatible** > **S3 Compatible with Data Archiving**.

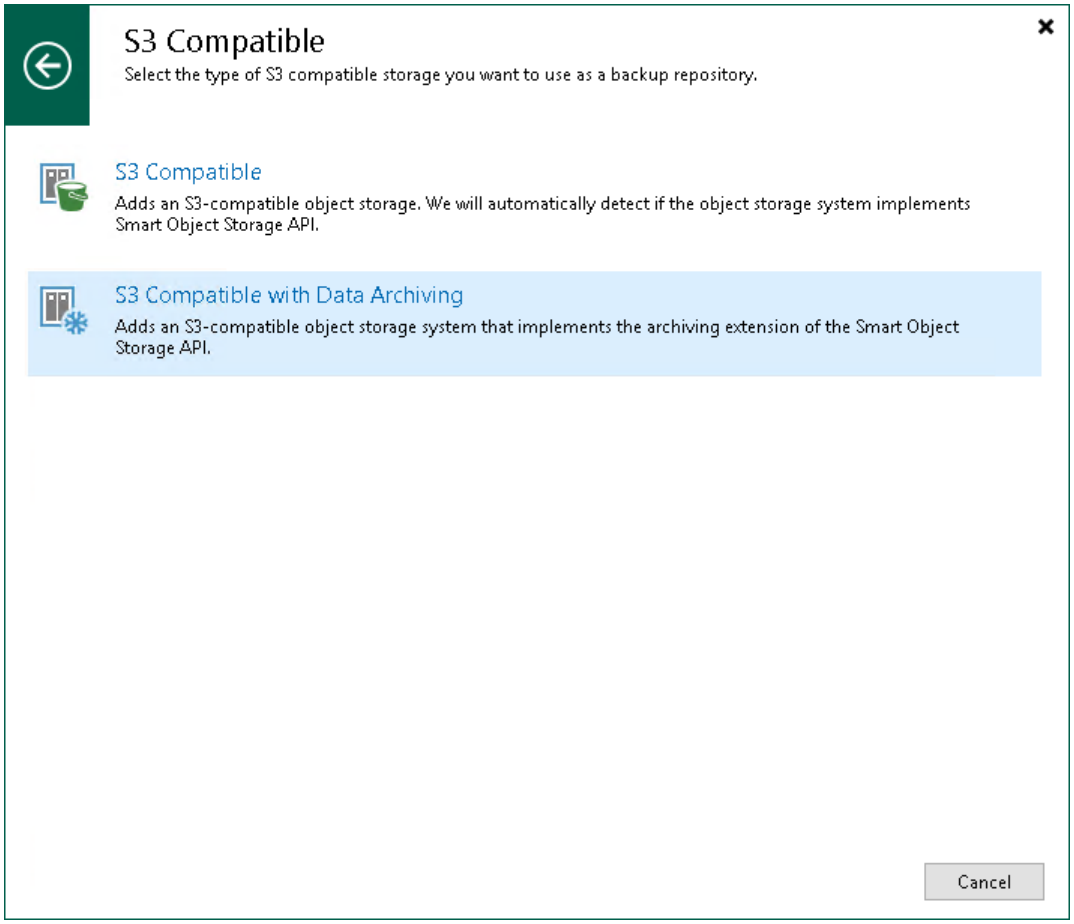

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

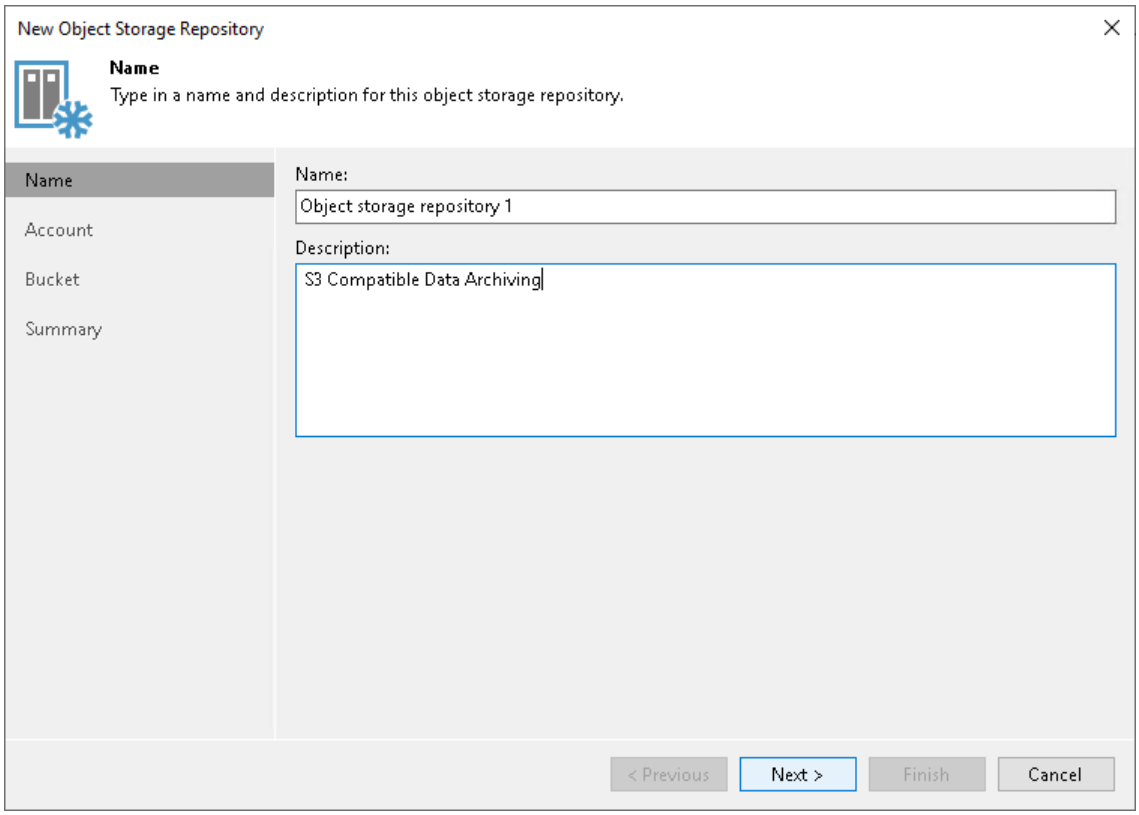

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection and archiver appliance settings. An archiver appliance is an auxiliary machine that transfers data from S3 compatible object storage to S3 compatible object storage with data archive. For more information, see the Archiver Appliances section.

1. In the **Service point** field, specify an endpoint address and a port number of your S3 compatible object storage with data archive.
2. In the **Region** field, specify a region.
3. From the **Credentials** drop-down list, select user credentials to access your S3 compatible object storage with data archive.

   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in Access Keys for AWS Users. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.
4. From the **Archiver appliance** drop-down list, select the archiver appliance. Click **Add New** to add a new archiver appliance to your backup infrastructure. For information on how to add a server, see Adding Microsoft Windows Servers and Adding Linux Servers.

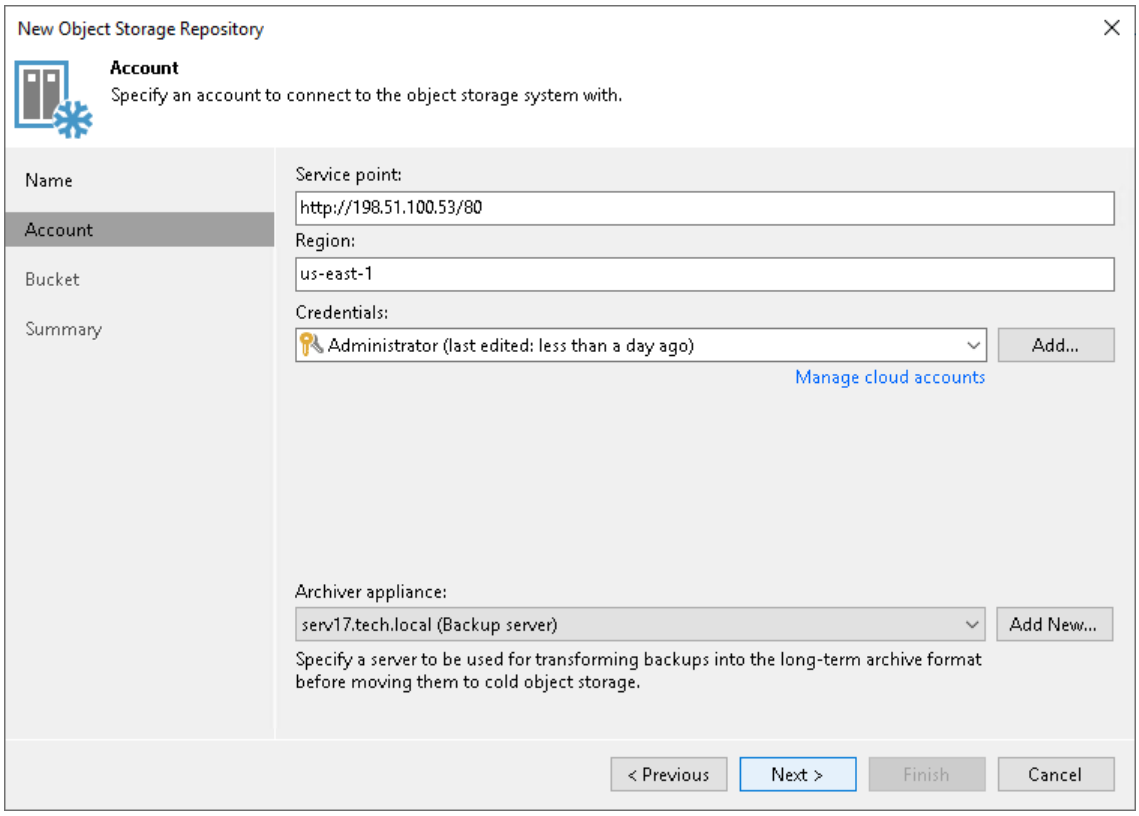

## Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify the bucket and folder where you will store data, and define storage limits and immutability settings that Veeam Backup & Replication will apply to data in the object storage.

1. In the **Bucket** field, enter a name of the bucket or click **Browse** to get the necessary bucket.

   Note that you must create the bucket where you want to store your backup data beforehand.

2. In the **Folder** field, enter a cloud folder name to which you want to map your S3 compatible object storage with data archive. Alternatively, click **Browse** and either select an existing folder or click **New Folder**.

3. To prohibit deletion of blocks of data from S3 compatible object storage with data archive, select the **Make backups immutable for the entire duration of their retention policy** check box. The immutability period will be equal to the retention period (if any) of the data blocks. All the types of files that are eligible for archive storage can be made immutable. For more information on the immutability feature and the retention policy for each file type, see Immutability for Object Storage Repositories.

   Keep in mind that to use immutability, you must enable the *Object Lock* and *Versioning* features on your S3 bucket at the time when you create the bucket. For more information, see Enabling Immutability.

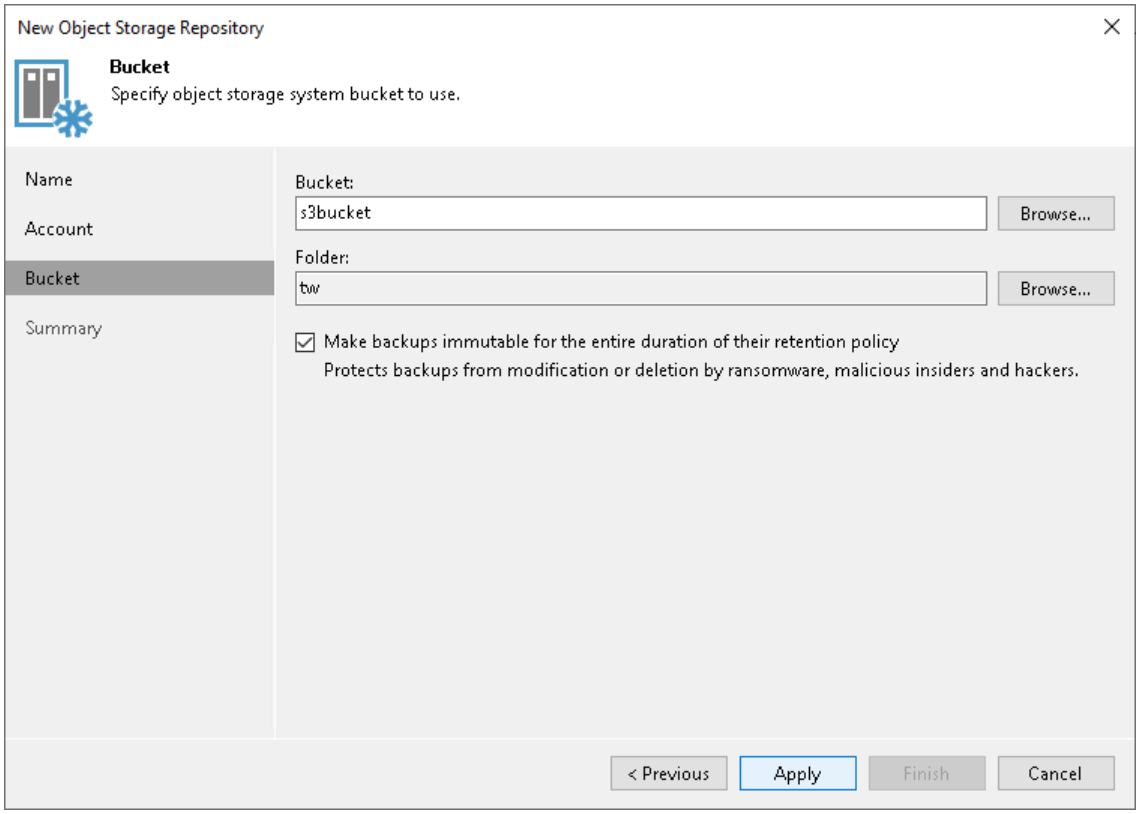

### Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

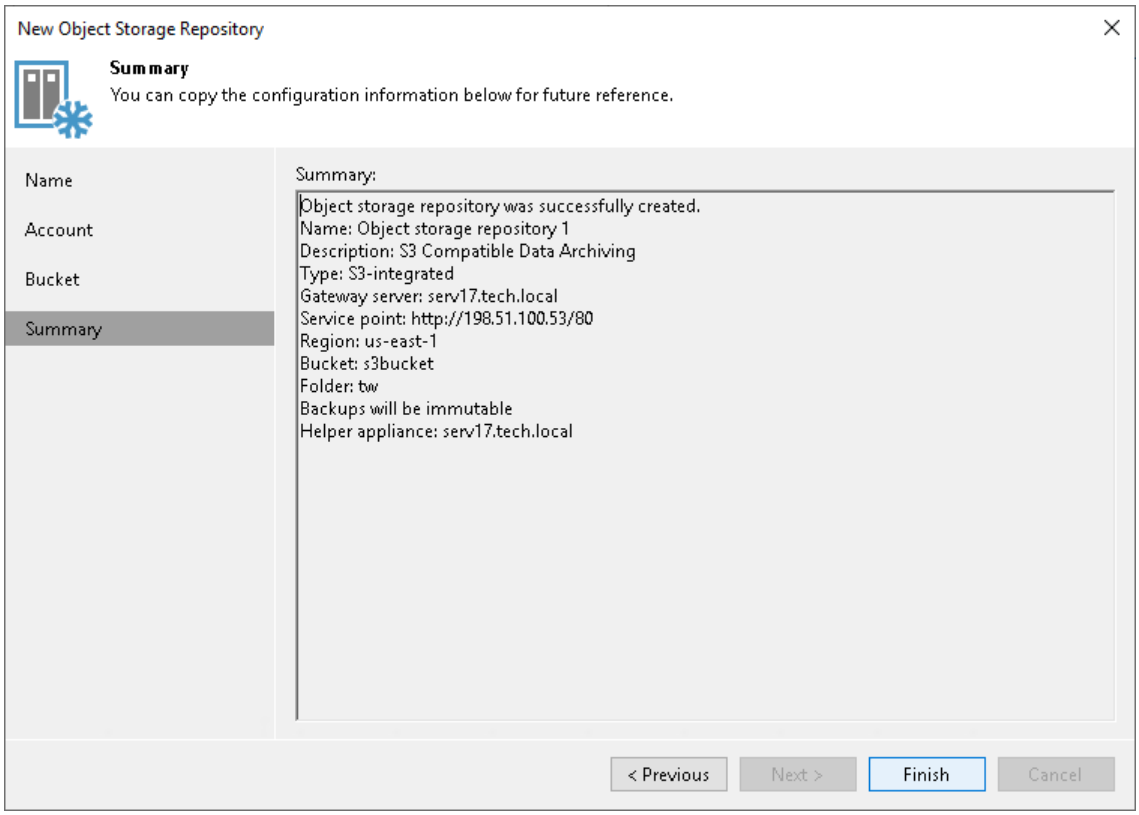


## Adding Amazon S3 Object Storage, Amazon S3 Glacier Storage and AWS Snowball Edge

Before you add an Amazon S3 storage, Amazon Glacier or Amazon Snowball Edge to the backup infrastructure, check prerequisites and required permissions. After that, use the **New Object Storage Repository** wizard.

### Adding Amazon S3 Storage

This section describes how to add Amazon S3 object storage to the backup infrastructure. For information on Amazon S3 object storage, see AWS Documentation.

To add Amazon S3 object storage, use the **New Object Storage Repository** wizard.

## Step 1. Launch New Object Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **Hyperscalers** > **Amazon S3** > **Amazon S3**.

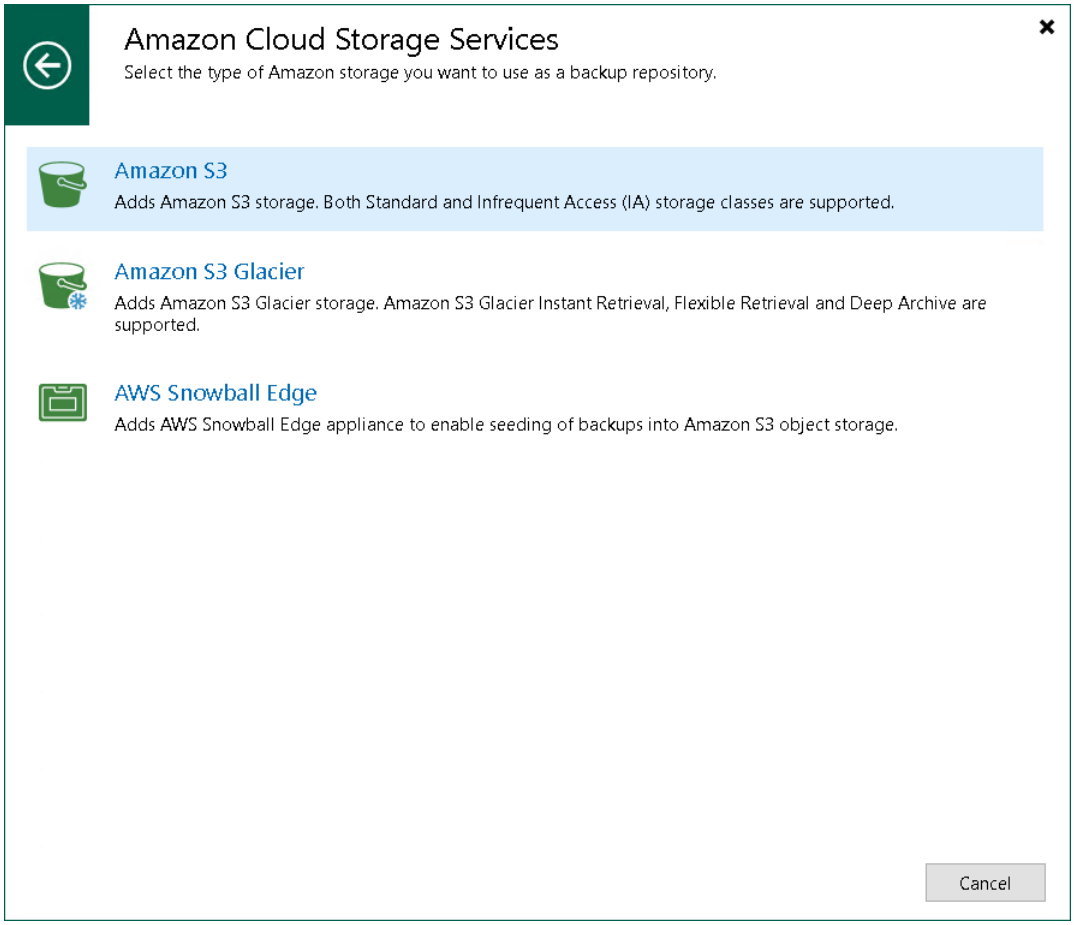

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

If you want to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

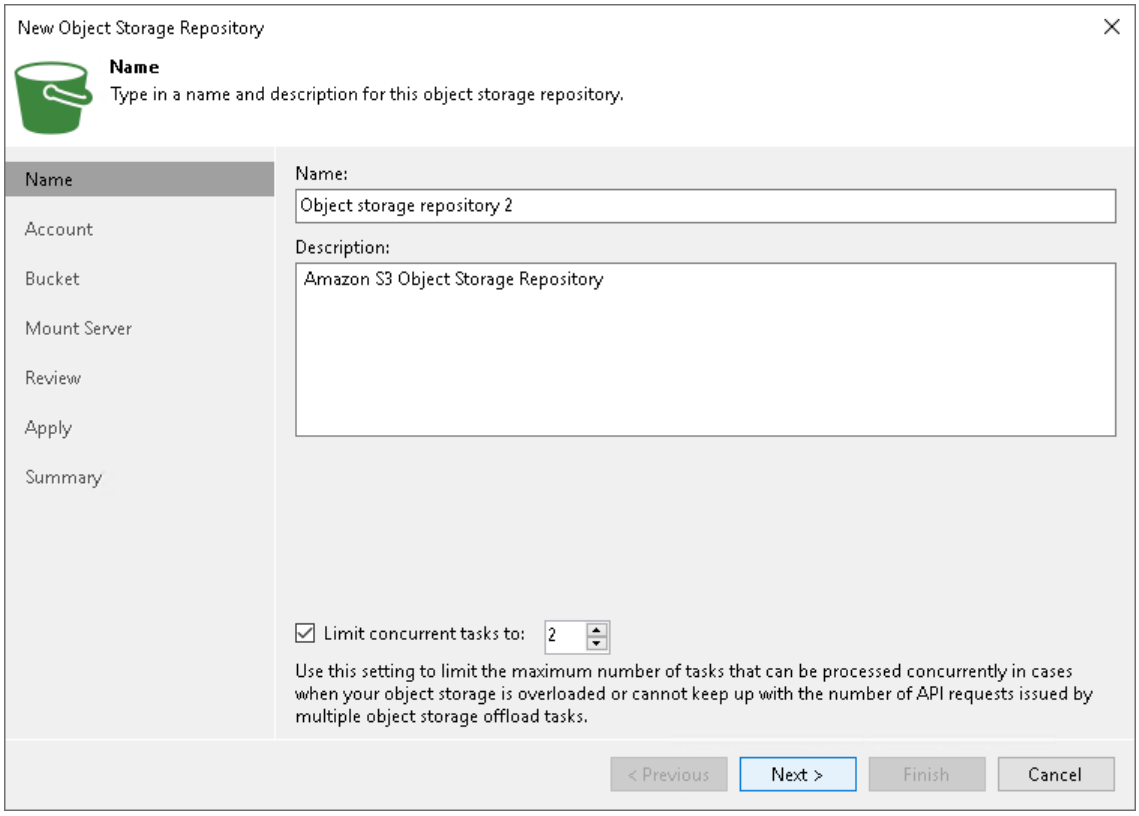

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. From the **Credentials** drop-down list, select user credentials to access your Amazon S3 object storage.

   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Access Keys for AWS Users. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

   The user account must have permissions listed in section Permissions.

2. From the **AWS region** drop-down list, select the AWS region where the Amazon S3 bucket is located.

3. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:

   - **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.
   - **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

     By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

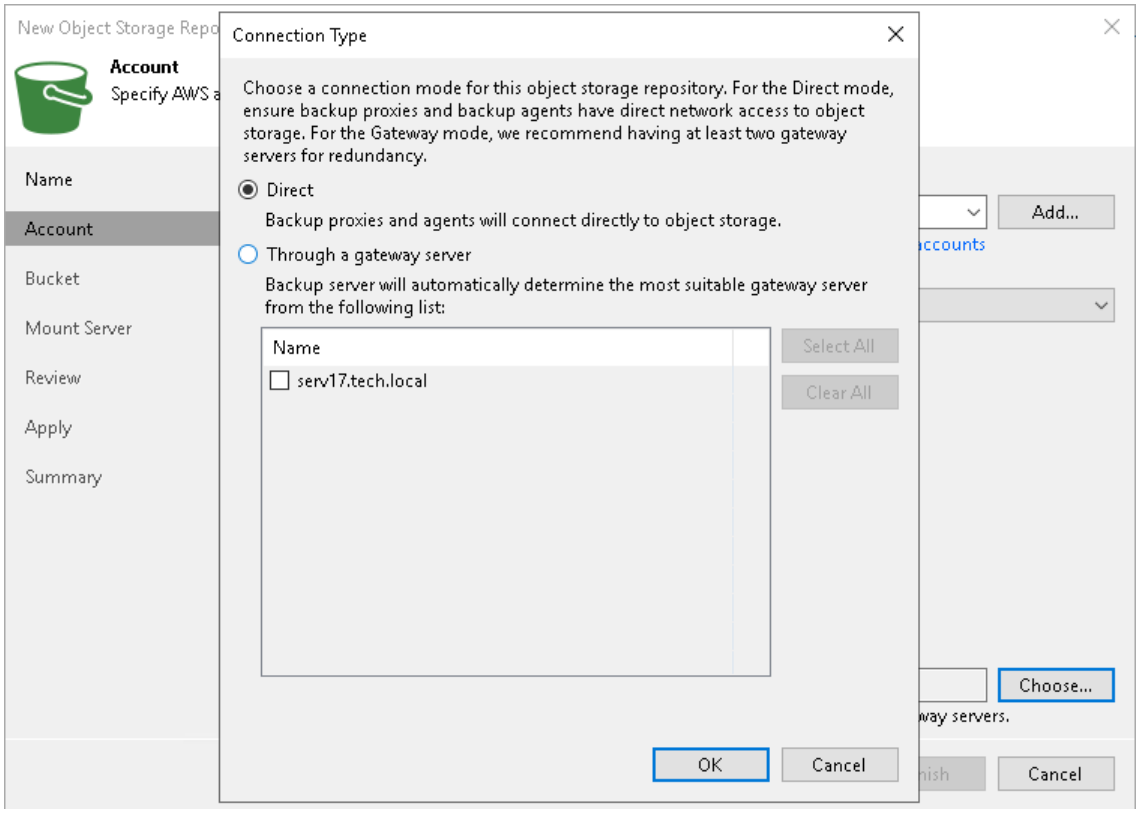

## Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify Amazon S3 bucket and folder that will be used to store data:

1. From the **Data center** drop-down list, select the region where the bucket is located.
2. In the **Bucket** field, enter a name of the bucket or click **Browse** to get the necessary bucket.

   > **IMPORTANT**
   >
   > You must create the bucket where you want to store your backup data beforehand. When you create a bucket, consider Amazon bucket naming rules. It is not recommended that you use dots (.) in the bucket name. For more information on bucket naming rules, see AWS Documentation.

   If the FIPS-compliant operation mode is enabled and the bucket you want to add is non-FIPS compliant, the warning will be displayed. For more information, see FIPS Compliance.
3. To the right of the **Folder** field, click **Browse** and either select an existing folder or click **New Folder**.

   > **IMPORTANT**
   >
   > Veeam Backup & Replication supports specific storage classes. For more information, see Considerations and Limitations.
4. Select the **Limit object storage consumption to** check box to define a soft limit for your object storage consumption. If this limit is exceeded during a job run, Veeam Backup & Replication will complete the job. However, a new job will not be able to start unless you remove the extra data that exceeds the limit or change the soft limit settings. Provide the value in TB or PB.
5. Select the **Make recent backups immutable for** check box to prohibit deletion of blocks of data from object storage. Specify the immutability period. For more information, see Immutability for Object Storage Repositories.

   Note that the maximum immutability period you can set in the Veeam Backup & Replication UI is 90 days. If you want to set immutability to a longer period, use the Set-VBRAmazonS3Repository cmdlet.
6. If you plan to access your backup data in an infrequent manner, select the **Use infrequent access storage class** check box to mark each block as *S3 Standard-IA (S3 Standard Infrequent Access)*.
7. To enable *Amazon S3 One Zone-Infrequent Access*, select the **Store backups in a single availability zone only** check box. For more information, see AWS Documentation.

   > **IMPORTANT**
   >
   > If you enable this option and plan to use this object storage as a performance or capacity tier, do not target to this repository any jobs that constantly send backup data to this storage: scheduled regular backup and backup copy jobs that run without GFS, jobs with transactions logs enabled, jobs created by Veeam Enterprise Plug-ins. Otherwise, it will result in higher costs.

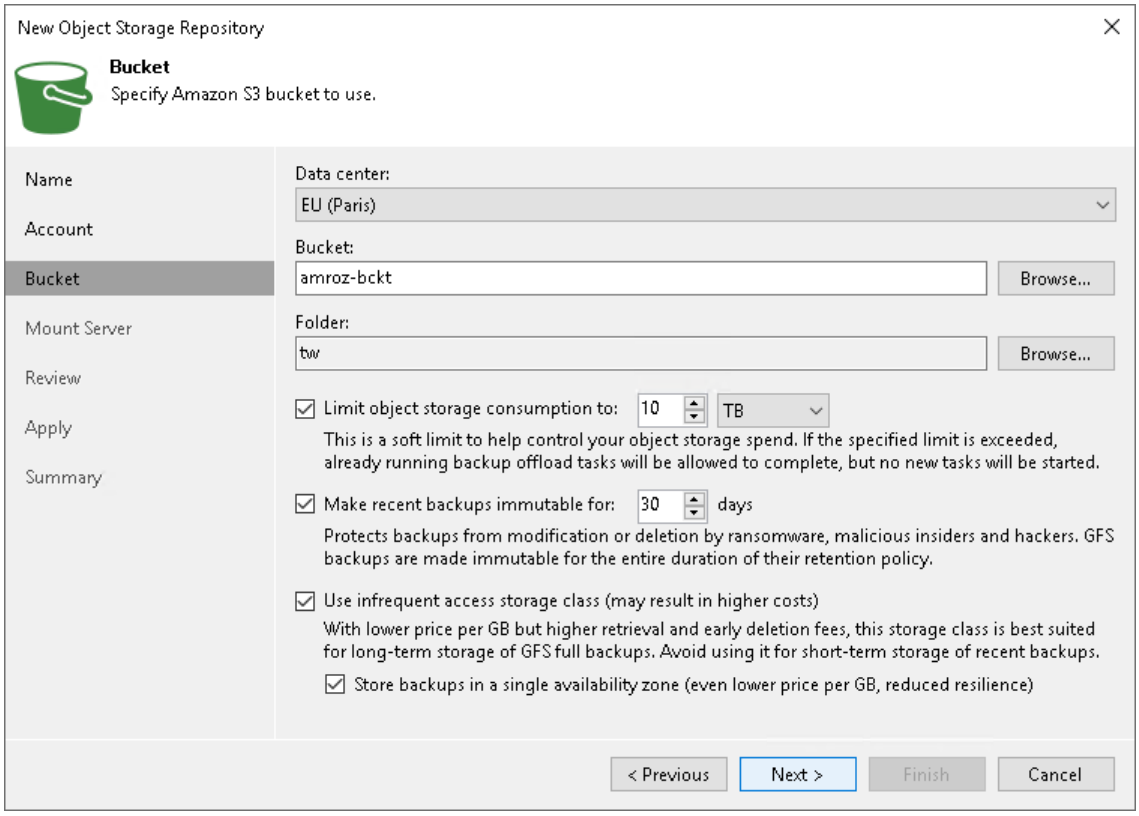
New Object Storage Repository
Bucket
Specify Amazon S3 bucket to use.
Name
Account
Bucket
Mount Server
Review
Apply
Summary
Data center:
EU (Paris)
Bucket:
amroz-bckt
Browse...
Folder:
tw
Browse...
Limit object storage consumption to:
10
TB
This is a soft limit to help control your object storage spend. If the specified limit is exceeded, already running backup offload tasks will be allowed to complete, but no new tasks will be started.
Make recent backups immutable for:
30
days
Protects backups from modification or deletion by ransomware, malicious insiders and hackers. GFS backups are made immutable for the entire duration of their retention policy.
Use infrequent access storage class (may result in higher costs)
With lower price per GB but higher retrieval and early deletion fees, this storage class is best suited for long-term storage of GFS full backups. Avoid using it for short-term storage of recent backups.
Store backups in a single availability zone (even lower price per GB, reduced resilience)
< Previous
Next >
Finish
Cancel

### Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for restore operations, and configure a helper appliance. The helper appliance is a temporary EC2 instance that Veeam Backup & Replication deploys in your Amazon EC2, to perform a health check of backup files and apply retention to unstructured data backup files. For more information, see Health Check for Object Storage Repositories and Helper Appliance in Unstructured Data Backup. After Veeam Backup & Replication completes these operations, it removes the helper appliance from Amazon EC2.

## Specifying Mount Server Settings

To specify the mount server settings, do the following:

1. From the **Mount Server** drop-down list, select a server that you want to use as a mount server. Veeam Backup & Replication uses this server during restore operations to mount VM disks directly from objects located in object storage repositories. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, click **Add New** on the right to open the **New Windows Server wizard**. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder to keep cache that is created during mount operations.

3. Select the **Enable vPower NFS service on the mount server** check box to allow the Veeam vPower NFS Service access the object storage repository. Veeam Backup & Replication will enable the Veeam vPower NFS Service on the necessary mount server. For more information, see Veeam vPower NFS Service.

4. Click **Ports** to customize network ports used by the Veeam vPower NFS Service. In the **vPower NFS Port Settings** window, specify the following settings:

   - Next to the **Mount Port** section, specify the port that the Veeam vPower NFS Service will use to mount the vPower NFS datastore to the ESXi host.
   - Next to the **vPower NFS** port section, specify the port that the Veeam vPower NFS Service will use to connect to the target NFS share.

   For information on ports used by default, see Ports.

5. To specify the helper appliance settings, click **Configure**. From the **Managed server** drop-down list, select a server that you want to use as the helper appliance.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

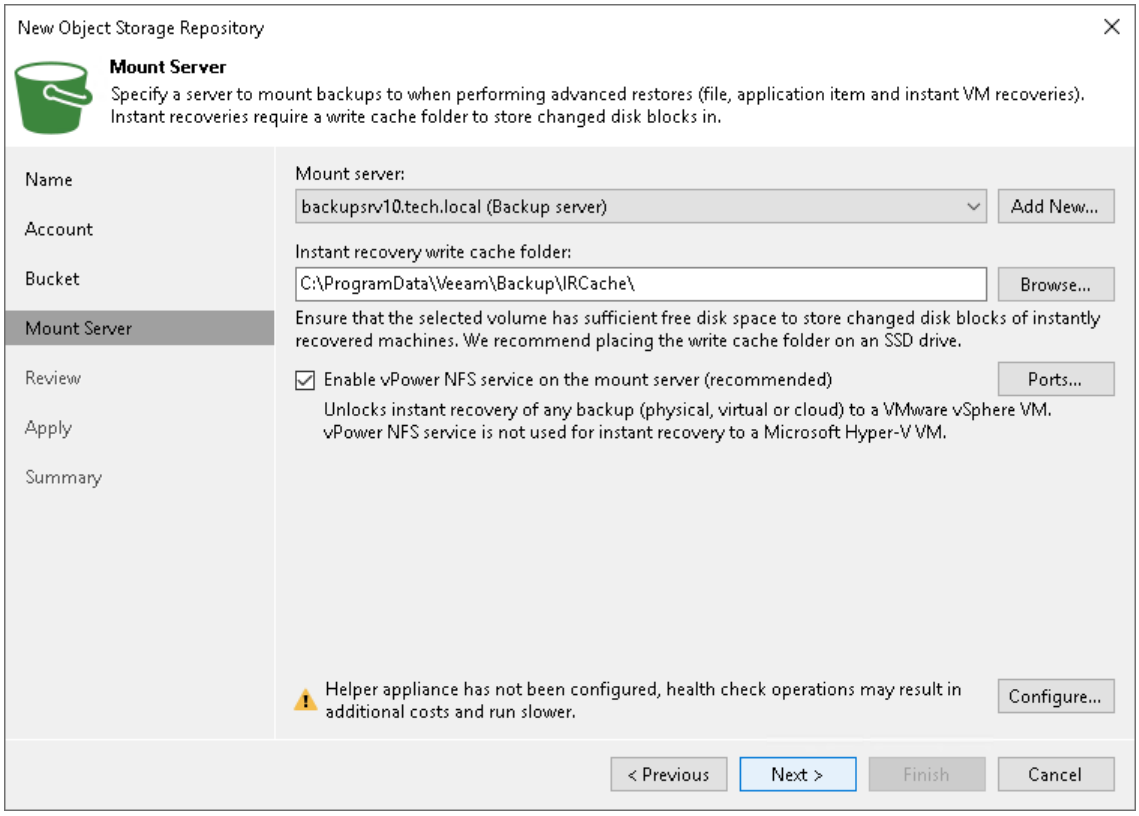


# Configuring Helper Appliance

To configure the helper appliance, at the **Mount Server** step, click **Configure** and in the **Helper Appliance Settings** window, specify the following settings:

1. From the **EC2 instance type** drop-down list, select the EC2 instance type for the helper appliance. The speed and cost of operations that the helper appliance performs depend on the EC2 instance type. For information on EC2 instance types, see AWS Documentation. For details on the EC2 instant types used by Veeam Backup & Replication, see this Veeam KB article.

2. From the **Amazon VPC** drop-down list, select the Amazon VPC where Veeam Backup & Replication will launch the EC2 instance. For information on the Amazon VPC, see AWS Documentation.

   To be able to select the necessary VPC from the drop-down list, you must create it beforehand as described in the AWS Documentation.

3. From the **Subnet** drop-down list, select the subnet where the helper appliance will reside. For information on the subnets for Amazon VPC, see AWS Documentation.

   To be able to select the necessary subnet from the drop-down list, you must create it beforehand as described in the AWS Documentation.

4. From the **Security group** drop-down list, select a security group that will be associated with the helper appliance. For information on security groups for Amazon VPC, see AWS Documentation.

   To be able to select the necessary security group from the drop-down list, you must create it beforehand as described in the AWS Documentation.

> **IMPORTANT**
>
> If you select the **Create new** option, Veeam Backup & Replication will create a new security group with the inbound rules that allow connection using the 443 and 22 ports from everywhere (0.0.0.0/0).

5. In the **Redirector port** field, specify the TCP port that Veeam Backup & Replication will use to route requests between the helper appliance and backup infrastructure components.
6. Click **OK**.

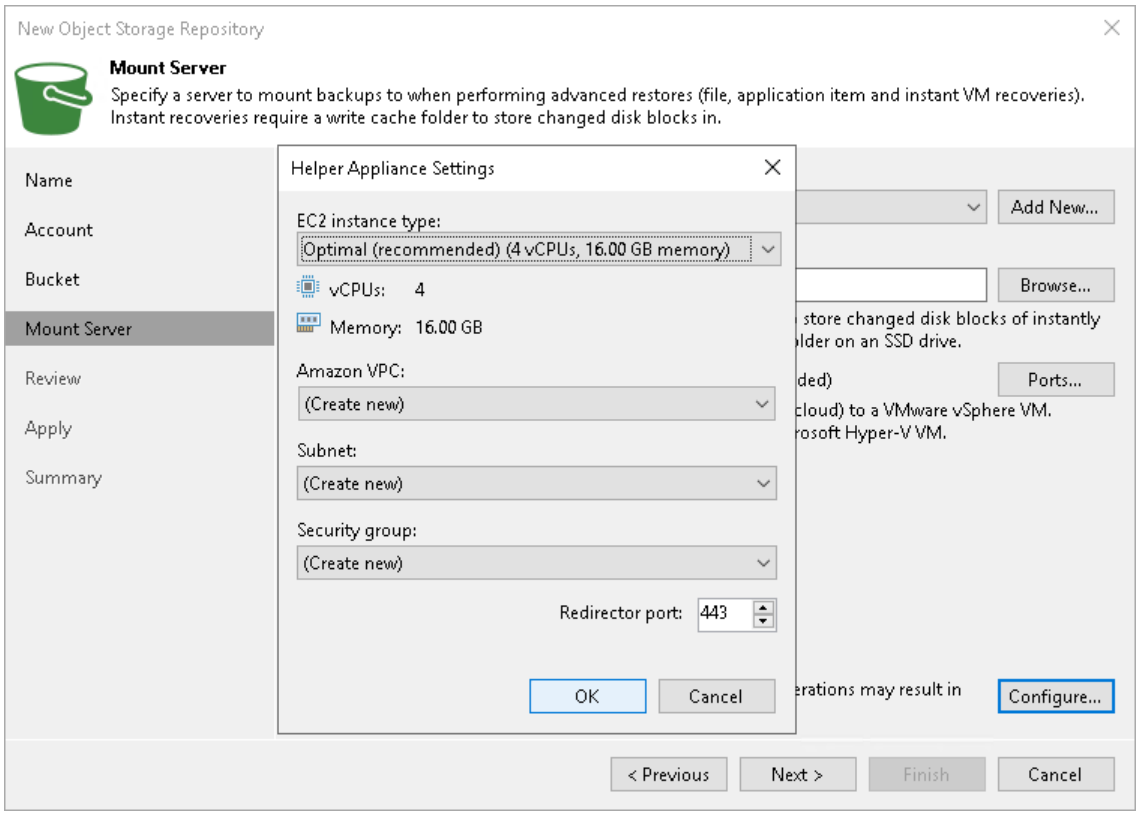

## Step 6. Review Components

At the **Review** step of the wizard, review what components will be processed on the mount server server and their status.

If the backup repository contains backups, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Object Storage (Imported)** node.

If the backup repository contains guest file system index files, select the **Import guest file system index data to the catalog** check box. Veeam Backup & Replication will import index files together with backup files, and you will be able to search for guest OS files inside imported backups. For more information, see the Guest OS File Restore section of the Veeam Backup Enterprise Manager Guide.

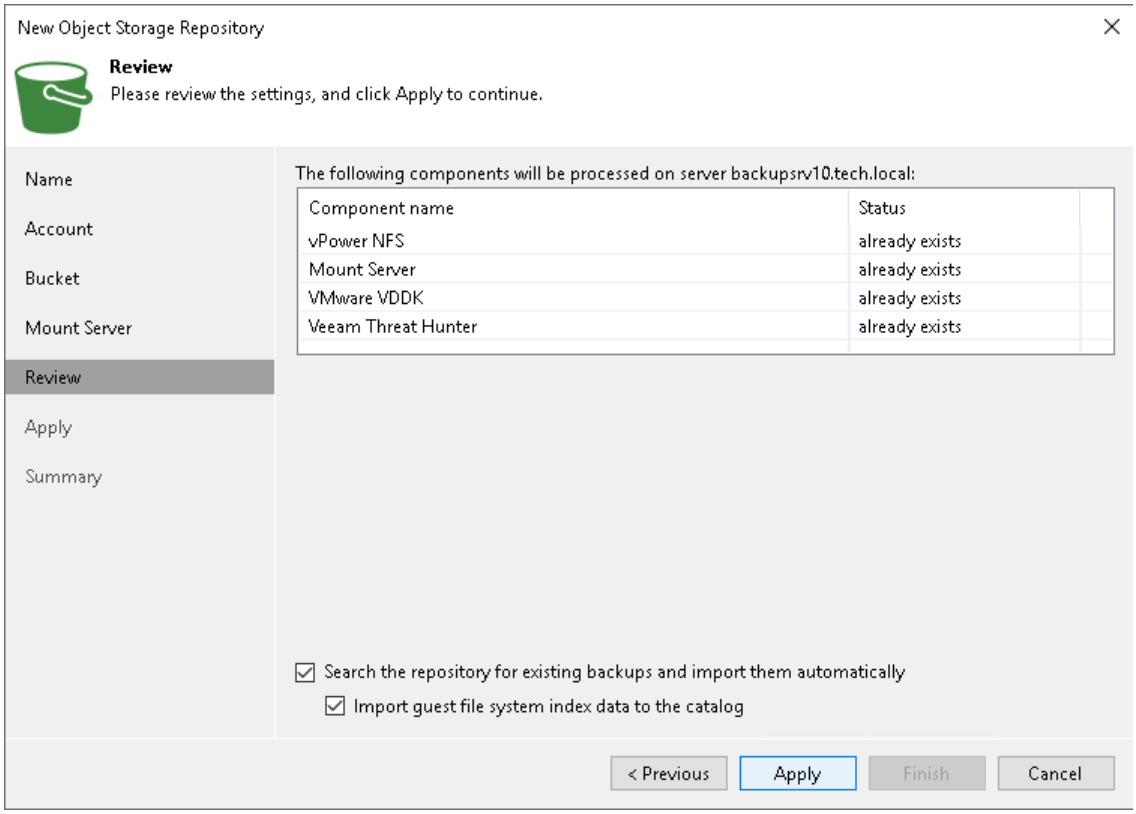

## Step 7. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to complete saving your settings to the configuration database and create backup infrastructure objects.

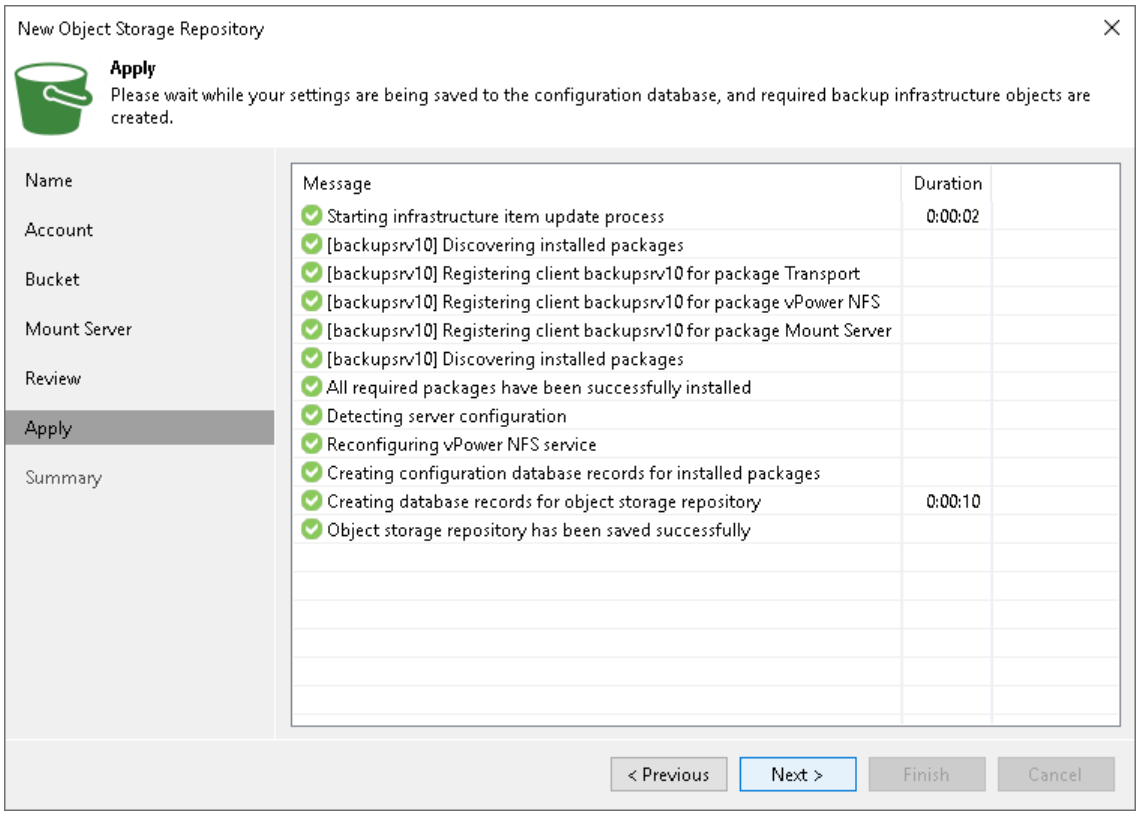

## Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

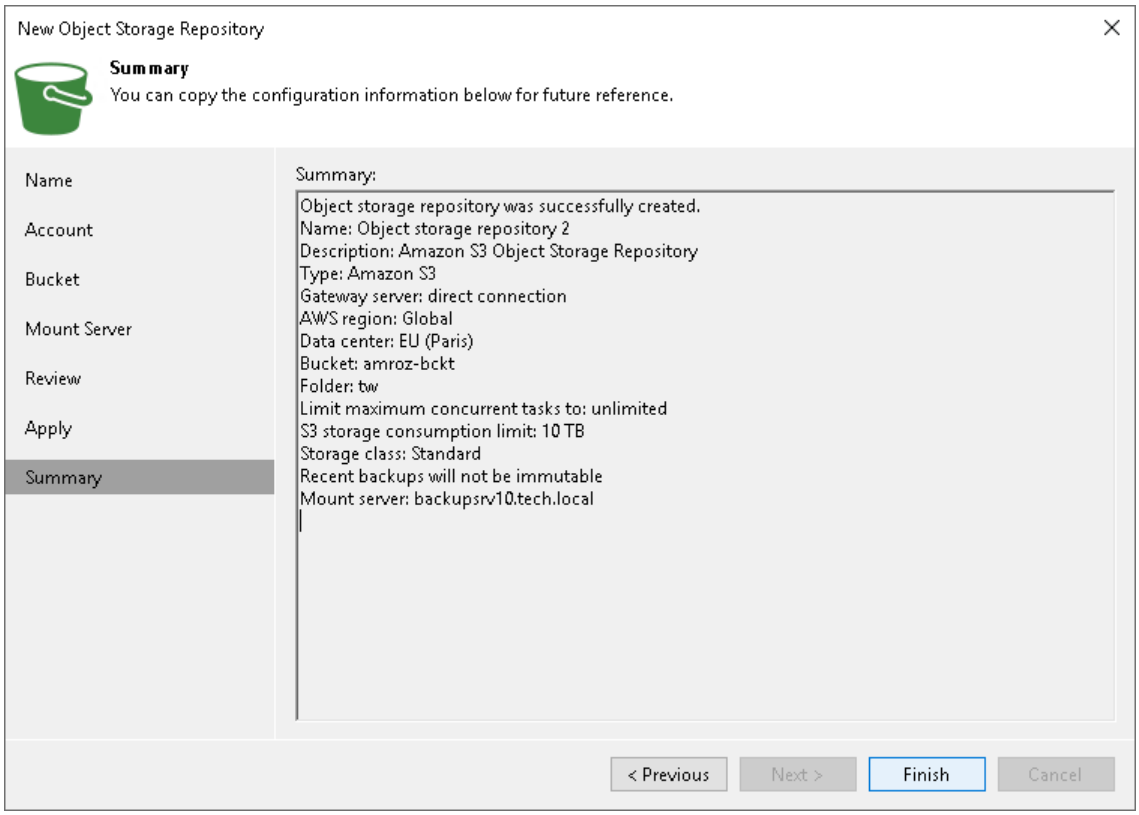


## Adding Amazon S3 Glacier Storage

Veeam Backup & Replication uses Amazon S3 Glacier object storage with Amazon S3 service. For more information about Amazon S3 storage classes that you can use for this object storage repository, see AWS Documentation.

> **NOTE**
>
> Consider the following:
>
> - Veeam Backup & Replication does not create or use any S3 Glacier vaults in your AWS environment. Glacier vaults is an archive storage solution independent from AWS. It uses storage containers named vaults (opposed to S3 buckets) and its own set of APIs to upload and retrieve data.
> - Amazon S3 Glacier storage uses S3 APIs to manage data. It also uses S3 storage as a repository for metadata of the Glacier-stored objects. Veeam Backup & Replication assigns the added storage class to backups stored in the repository. That is why the archived backups remain in Amazon S3 and cannot be accessed directly through the Amazon S3 service.

You can only use this repository as an archive extent of the scale-out backup repository. For more information, see Archive Tier.

To add Amazon S3 Glacier object storage, use the **New Object Storage Repository** wizard.

## Step 1. Launch New Object Storage Repository

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **Hyperscalers** > **Amazon S3** > **Amazon S3 Glacier**.

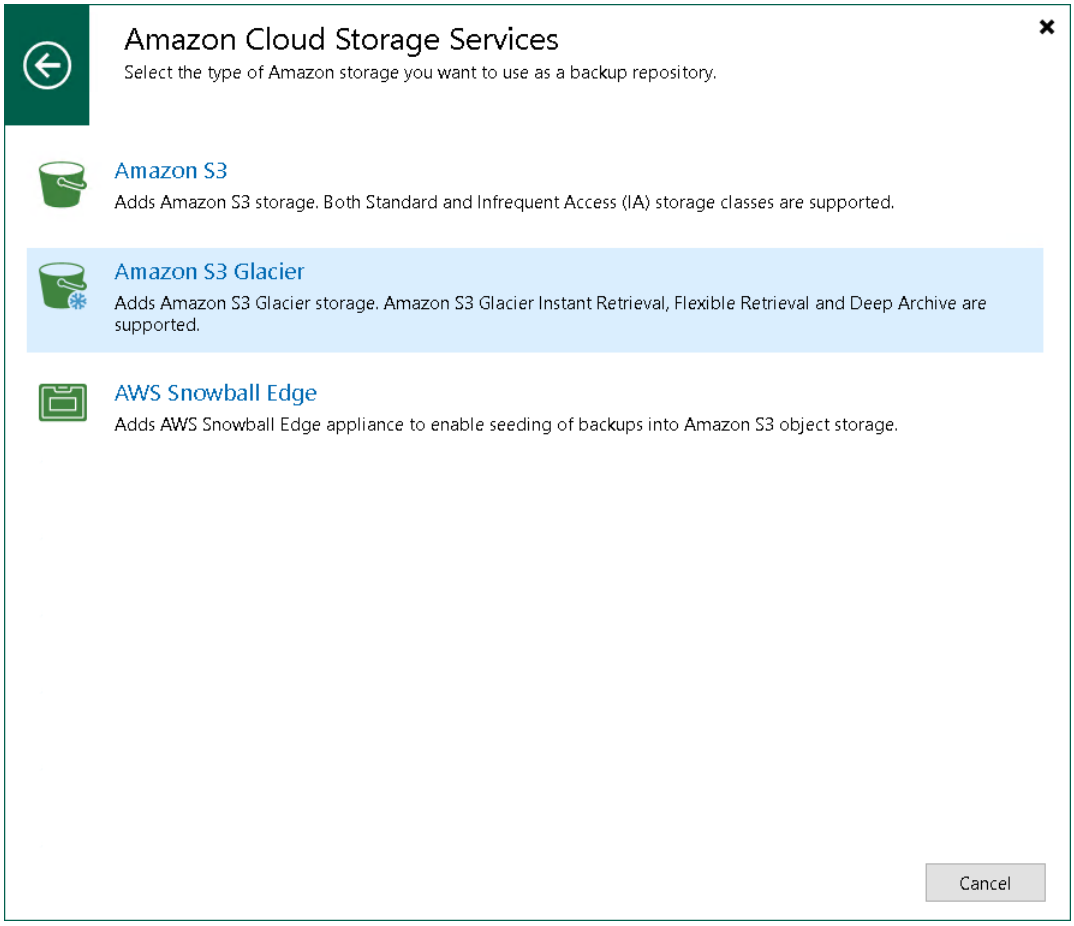

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for the new object storage repository and to provide a description for future reference.

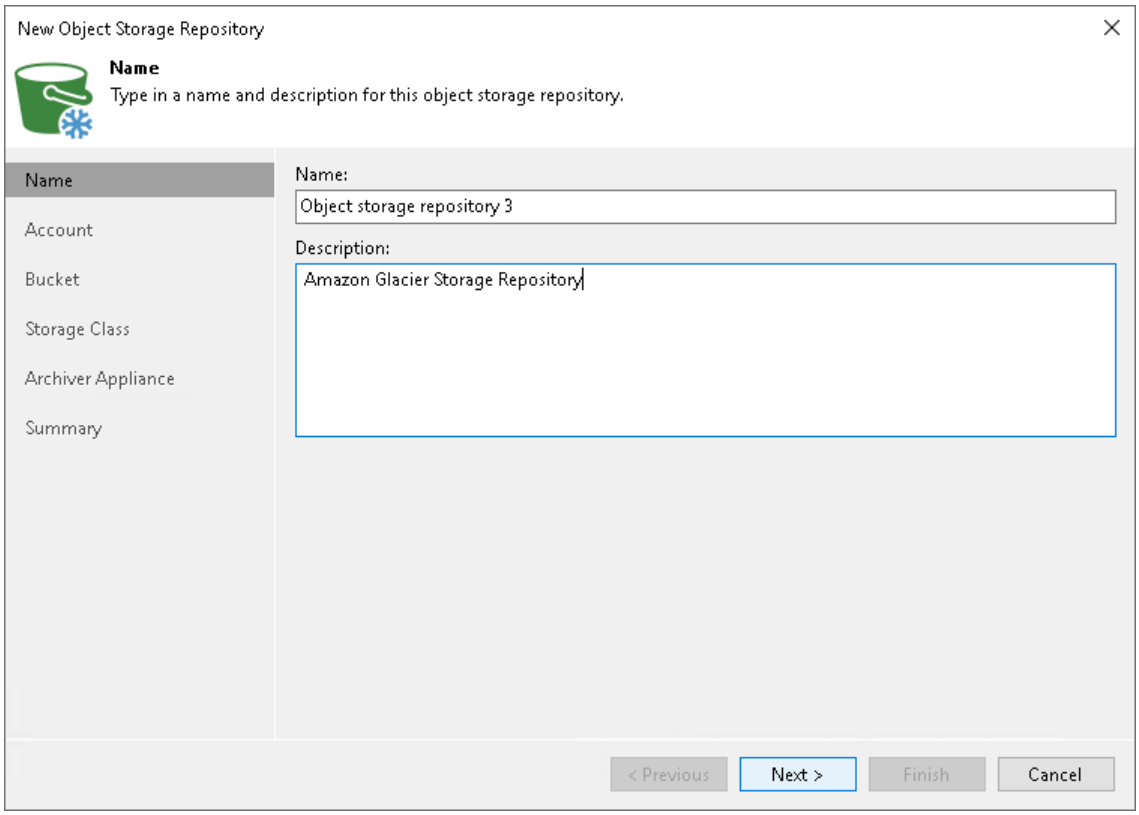

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. From the **Credentials** drop-down list, select user credentials to access your Amazon S3 object storage.

   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Access Keys for AWS Users. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

   The user account must have permissions listed in section Permissions.

2. From the **AWS region** drop-down list, select the geographical location of the Amazon datacenter.

3. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:

   - **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.

   - **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

     > **NOTE**
     >
     > The gateway server should have direct connection to AWS service endpoints. HTTP(S) proxy servers are not supported. For more information, see Ports.

By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

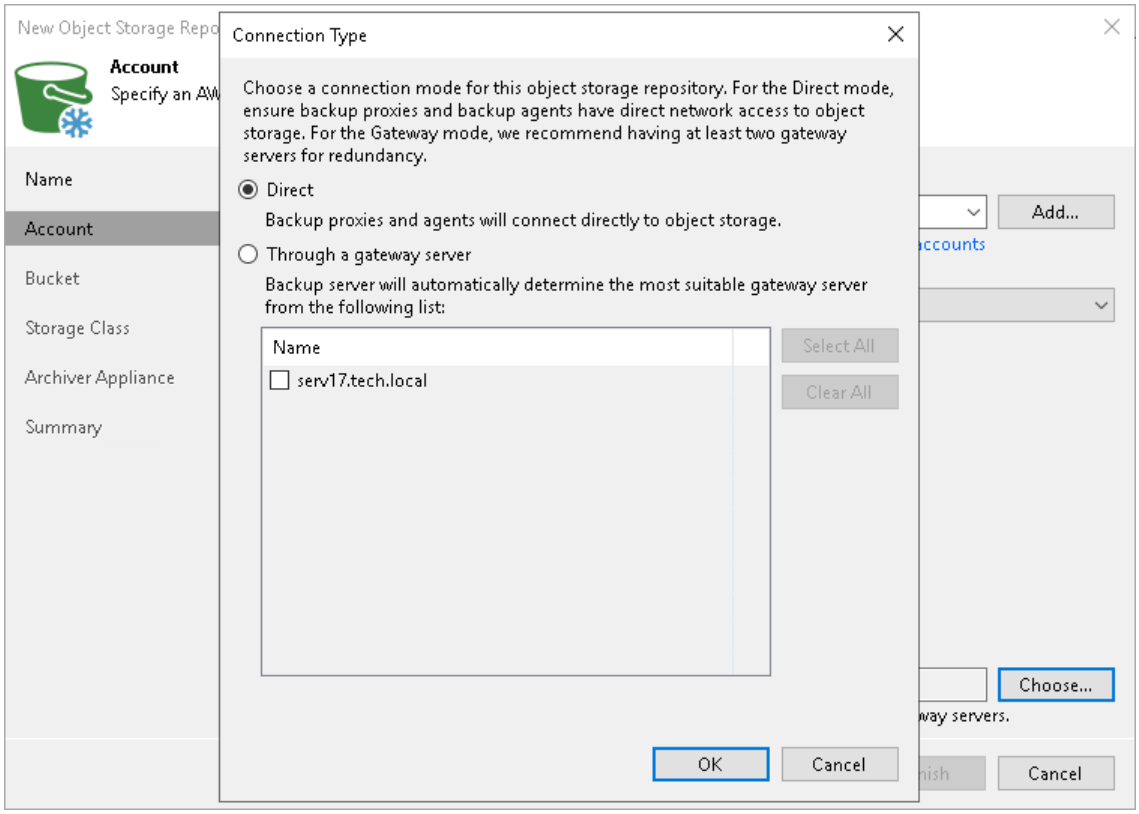

## Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify Amazon S3 bucket and folder that will be used to store data:

1. From the **Data center** drop-down list, select the AWS region where the Amazon S3 bucket is located.
2. In the **Bucket** field, enter a name of the bucket or click **Browse** to get the necessary bucket.

   > **IMPORTANT**
   >
   > You must create the bucket where you want to store your backup data beforehand. When you create a bucket, consider Amazon bucket naming rules. It is not recommended that you use dots (.) in the bucket name. For more information on bucket naming rules, see AWS Documentation.

   If the FIPS-compliant operation mode is enabled and the bucket you want to add is non-FIPS compliant, the warning will be displayed. For more information, see FIPS Compliance.
3. To the right of the **Folder** field, click **Browse** and either select an existing folder or click **New Folder**.
4. To prohibit deletion of blocks of data from object storage, select the **Make backups immutable for the entire duration of their retention policy** check box. The immutability period will be equal to the retention period (if any) of the data blocks. All the types of files that are eligible for archive storage can be made immutable. For more information on the immutability feature and the retention policy for each file type, see Immutability for Object Storage Repositories.

   Keep in mind that to use immutability, you must enable the *Object Lock* and *Versioning* features on your S3 bucket at the time when you create the bucket. For more information, see Enabling Immutability.

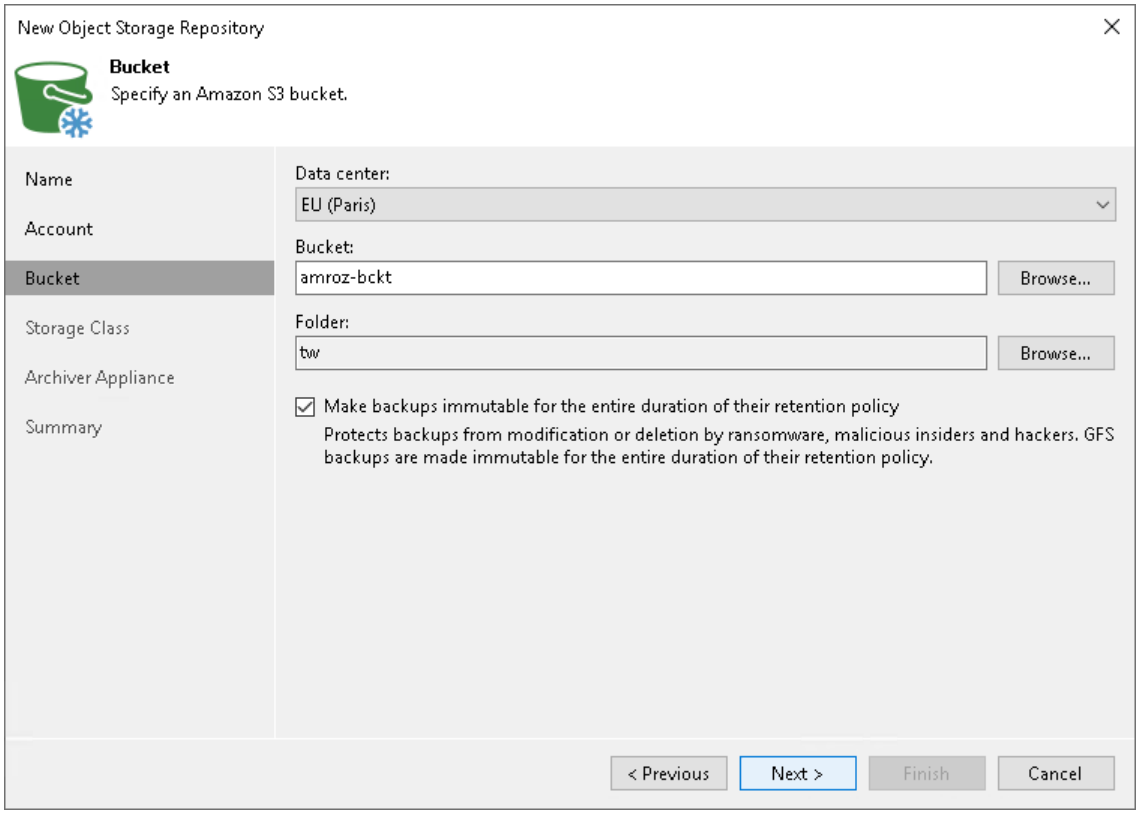

## Step 5. Specify Object Storage Class Settings

At the **Storage Class** step of the wizards, specify a storage class that you want assign to data blocks that you keep in Amazon S3 Glacier object storage. For more information on Amazon S3 archive storage classes, see AWS Documentation.

- Select the **Deep Archive** option to assign the *Amazon S3 Glacier Deep Archive* storage class to data blocks. Use this option if you want to keep data for more than 7 years and do not plan to access it.

  Note that to get data from this type of a backup, you must first retrieve data from Archive Tier. For more information, see Restore from Archive Tier.

- Select the **Flexible retrieval** option to assign the *Amazon S3 Glacier Flexible Retrieval* storage class to data blocks. Use this option if you want to access data in infrequent manner.

  Note that to get data from this type of a backup, you must first retrieve data from Archive Tier. For more information, see Restore from Archive Tier.

- Select the **Instant retrieval** option to assign the *Amazon S3 Glacier Instant Retrieval* storage class to data blocks. Use this option if you plan to access your backup data immediately. In this case, you do not need to retrieve data from Archive Tier and can immediately initiate any type of the data recovery operation.

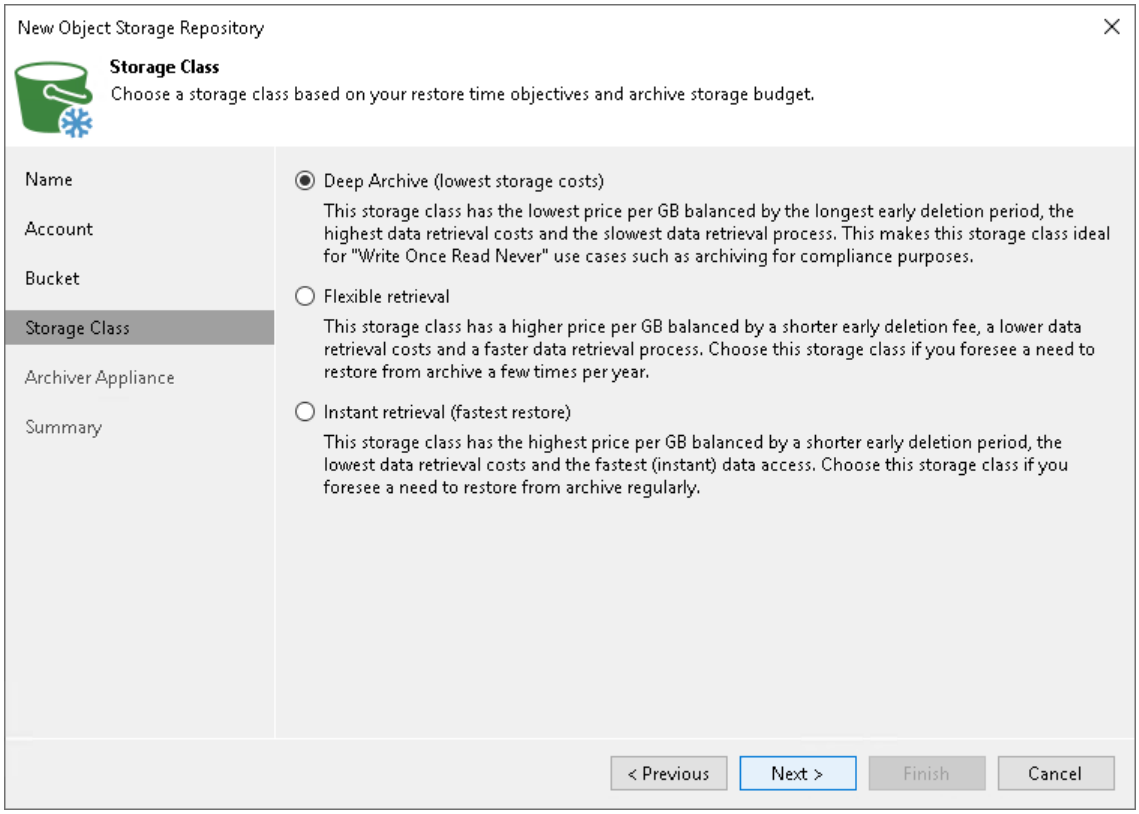

## Step 6. Specify Archiver Appliance

At the **Archiver Appliance** step of the wizard, you can specify the archiver appliance settings. An archiver appliance is an auxiliary instance that is necessary for transferring the data from Amazon S3 to Amazon S3 Glacier. For more information, see the Archiver Appliances section.

> **NOTE**
>
> Veeam Backup & Replication must be able to connect to the machine that you will use as an archiver appliance. Therefore, if your backup server is not located within AWS, you must configure public IP addresses for the subnet in which the appliance resides. For more information on configuring the subnet for Amazon VPC, see AWS Documentation.

You can specify archiver appliance in one of the following ways:

- Use the default settings. In that case, Veeam Backup & Replication will select the necessary settings from those available in your account, or will create for you new settings for the EC2 instance type, Amazon VPC, subnet, security group and redirector port.

  > **IMPORTANT**
  >
  > Veeam Backup & Replication creates a default security group with the inbound rules that allow connection using the 443 and 22 ports from everywhere (0.0.0.0/0).

- Specify the settings manually:

  a. Click **Customize**.

  b. From the **EC2 instance type** drop-down list, select the instance type for the proxy appliance. The EC2 instance type affects the speed and the cost of transferring the backup files to the Archive Tier of a scale-out backup repository. For information on instance types, see AWS Documentation. For details on the EC2 instant types used by Veeam Backup & Replication, see this Veeam KB article.

  c. From the **Amazon VPC** drop-down list, select the Amazon VPC where Veeam Backup & Replication will launch the target instance. For information on the Amazon VPC, see AWS Documentation.

  d. From the **Subnet** drop-down list, select the subnet for the proxy appliance.

  e. From the **Security group** drop-down list, select a security group that will be associated with the proxy appliance. For information on security groups for Amazon VPC, see AWS Documentation.

  f. In the **Redirector port** field, specify the TCP port that Veeam Backup & Replication will use to route requests between the proxy appliance and backup infrastructure components.

g. Click **OK**.

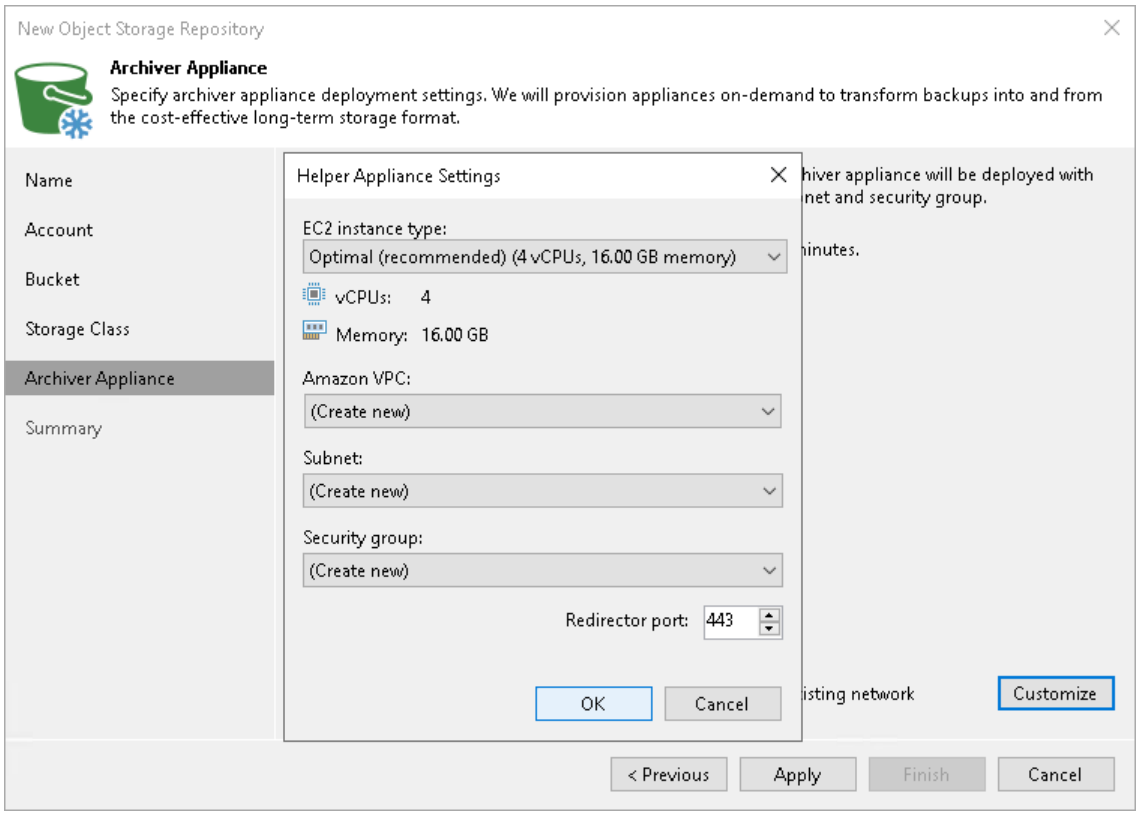
New Object Storage Repository
Archiver Appliance
Specify archiver appliance deployment settings. We will provision appliances on-demand to transform backups into and from the cost-effective long-term storage format.
Name
Account
Bucket
Storage Class
Archiver Appliance
Summary
Helper Appliance Settings
EC2 instance type:
Optimal (recommended) (4 vCPUs, 16.00 GB memory)
vCPUs: 4
Memory: 16.00 GB
Amazon VPC:
(Create new)
Subnet:
(Create new)
Security group:
(Create new)
Redirector port: 443
OK
Cancel
Customize
< Previous
Apply
Finish
Cancel

## Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

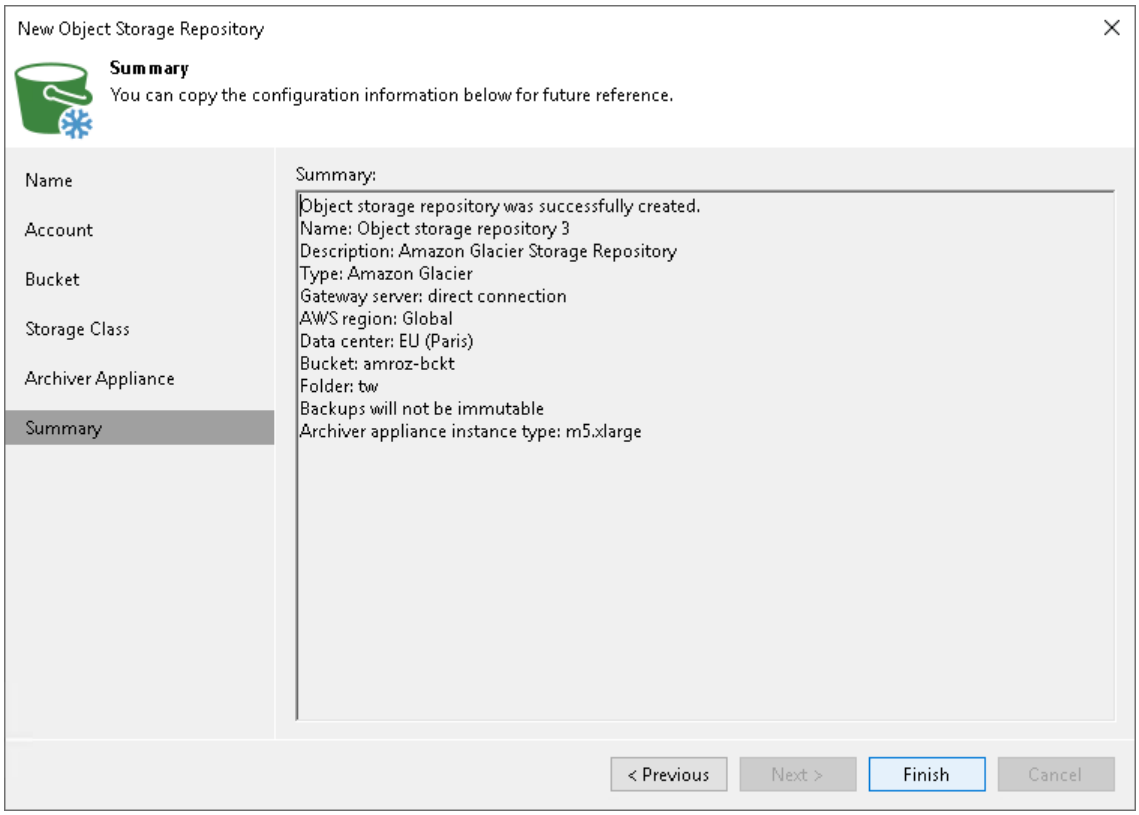


## Adding AWS Snowball Edge Storage

AWS Snowball Edge is a physical device which you can request for a short period of time from Amazon. You can temporarily attach it to the backup infrastructure and use as an object storage repository. For more information about ordering AWS Snowball Edge and preparing to use it, see this Amazon article.

This device may become useful when you need to offload a significant number of backup files occupying storage space on your extents, as offloading data to the AWS Snowball Edge device is much faster than transferring the same amount of data directly to Amazon object storage. Once you have offloaded backups to AWS Snowball Edge, you need to ship the device back to Amazon for further data synchronization with your storage account, as described in section Seeding Backups to Amazon S3 Storage.

To add AWS Snowball Edge storage, do the following:

1. Check prerequisites.
2. Launch the New Object Storage Repository wizard.
3. Specify object storage name.
4. Specify object storage account.
5. Specify object storage settings.
6. Finish working with the wizard.

## Before You Begin

Before you add AWS Snowball Edge storage to the backup infrastructure, check the following prerequisites:

- To maintain the device performance, be sure to select 1 MB or 4 MB as the storage optimization option when you configure a backup job. For more information on the optimization, see Storage Optimization.
- When you configure a scale-out backup repository with AWS Snowball Edge storage or capacity tier in a scale-out backup repository, it is recommended to only use the copy policy. This way you keep copy of the data on your local storage which helps you reduce the risk of data loss if the device is damaged during shipping. It will also ensure that the backup data is available for restore operations while AWS Snowball Edge storage is in shipment.
- For information on FIPS status of AWS Snowball Edge storage, see Amazon official updates.
- You can add only one AWS Snowball Edge storage to a scale-out backup repository.
- Veeam Backup & Replication does not support Versioning and Object Lock for S3 buckets associated with AWS Snowball Edge storage.
- Veeam Cloud Connect service providers can use AWS Snowball Edge only as a capacity extent of a scale-out backup repository.
- You cannot back up data using Veeam Agent backup job or policy to AWS Snowball Edge devices.

## Step 1. Launch New Object Storage Repository

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **Hyperscalers** > **Amazon S3** > **AWS Snowball Edge**.

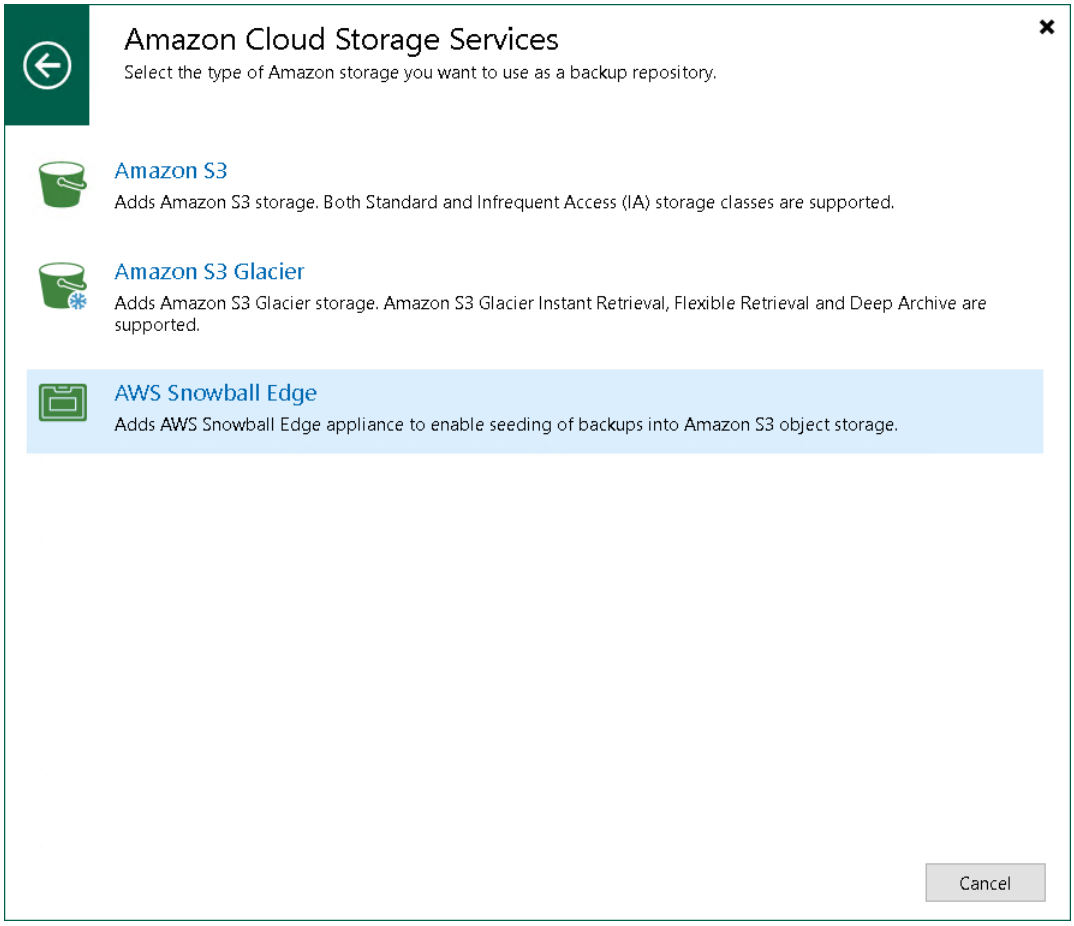

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

If you want to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

> **TIP**
>
> Set the maximum number of concurrent task to a reasonable number to avoid overloading if you plan to upload significant amount of backup chains to the device.

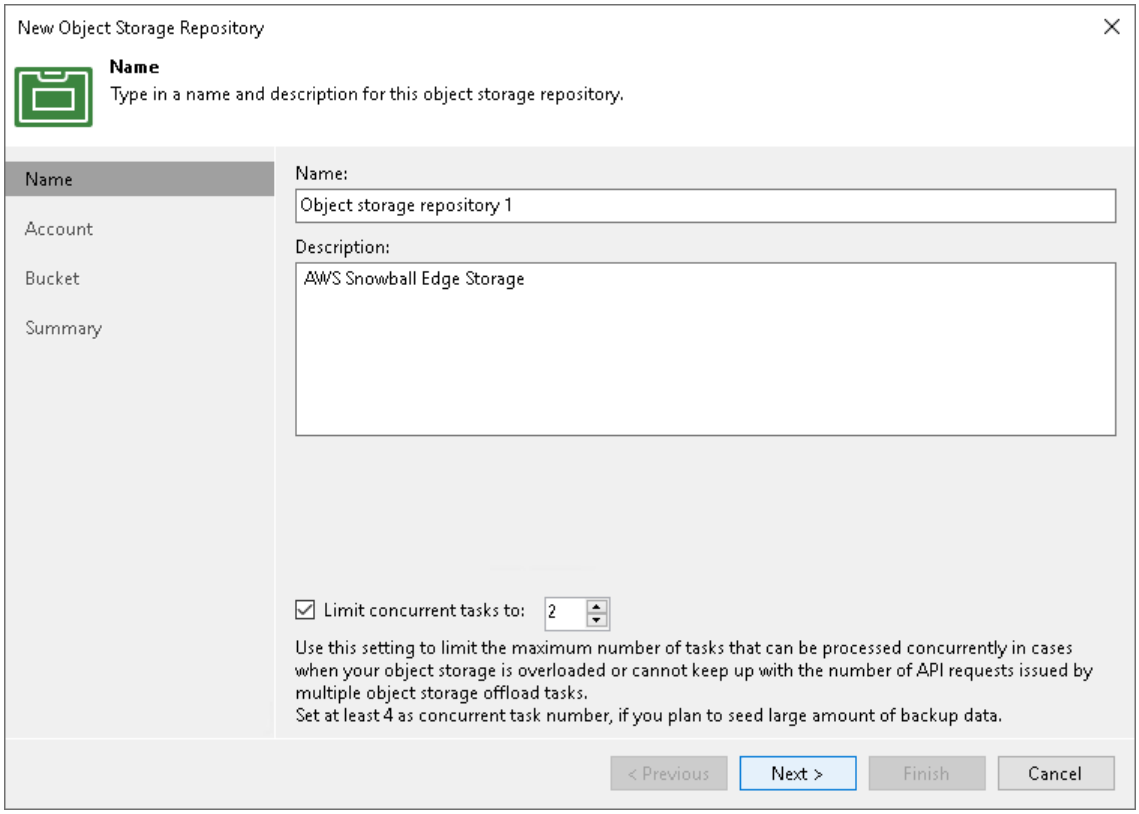

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. In the **Service point** field, specify a service point address of your AWS Snowball Edge device and define a port.

   To get information on the IP address that must be used as a service point, see the AWS Snowball Edge Developer Guide.

   To get information on the port you must specify, see the AWS Snowball Edge Developer Guide. By default, the 8443 port is used.

2. In the **Region** field, type "snow".

3. From the **Credentials** drop-down list, select user credentials to access your AWS Snowball Edge storage.

   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Access Keys for AWS Users. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

4. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:

   - **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.
   - **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

     > **NOTE**
     >
     > The gateway server should have direct connection to AWS service endpoints. HTTP(S) proxy servers are not supported. For more information, see Ports.

By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

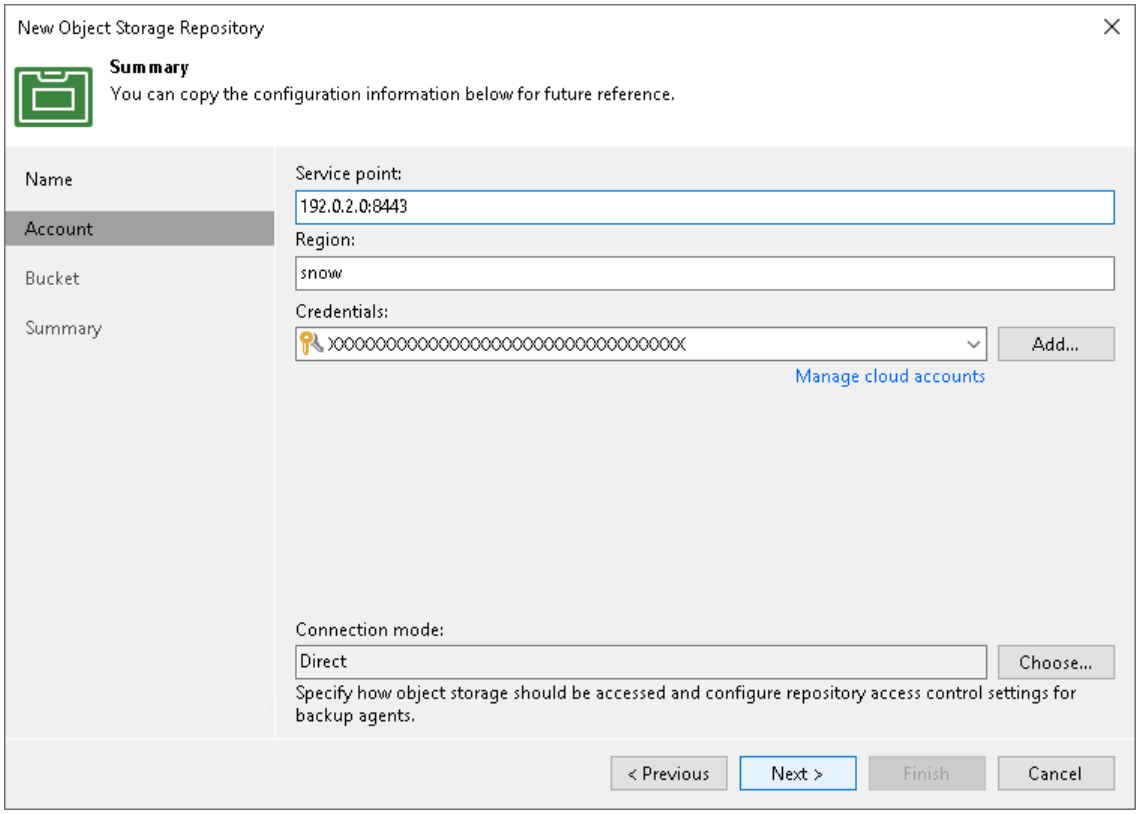

## Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify Amazon S3 bucket and folder that will be used to store data:

1. From the **Bucket** drop-down list, select a bucket.

   Make sure that the bucket where you want to store your backup data was created in advance.

2. In the **Folder** field, select a cloud folder to which you want to map your object storage repository.

   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

3. Select the **Limit object storage consumption to** check box to define a soft limit that can be exceeded temporarily for your object storage consumption. Provide the value in TB or PB.

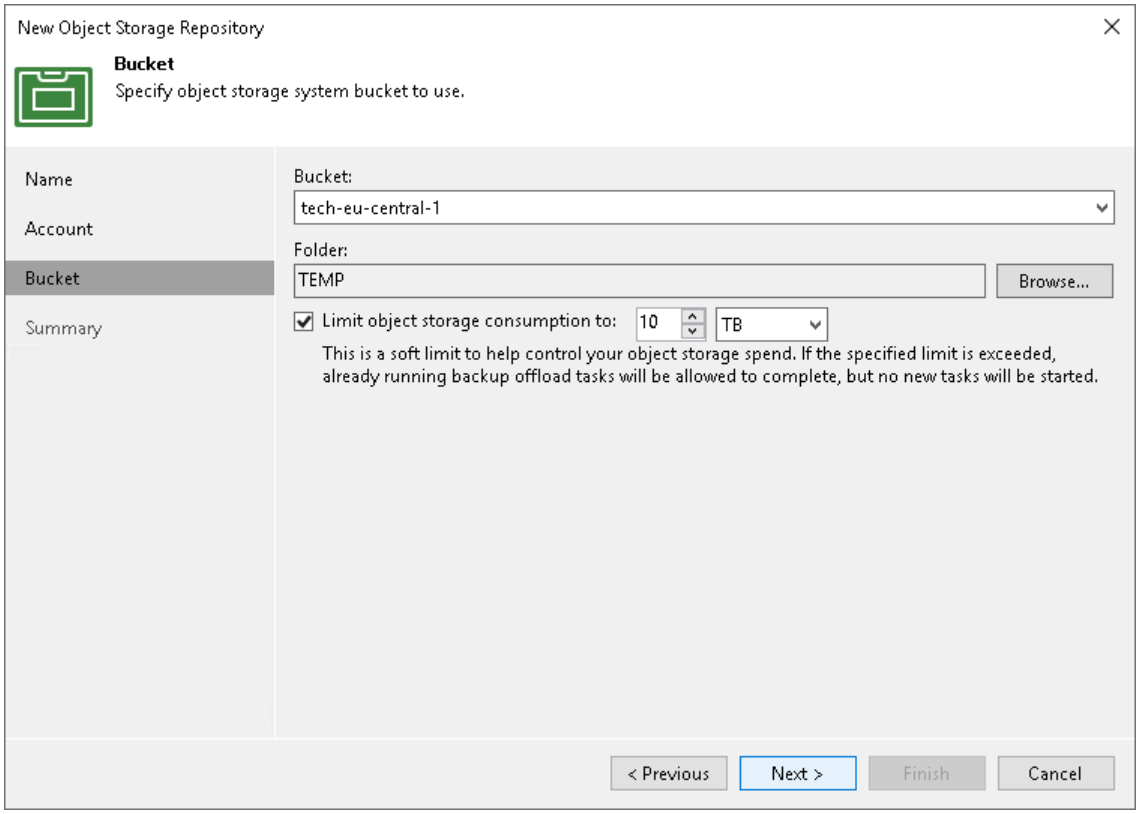

### Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

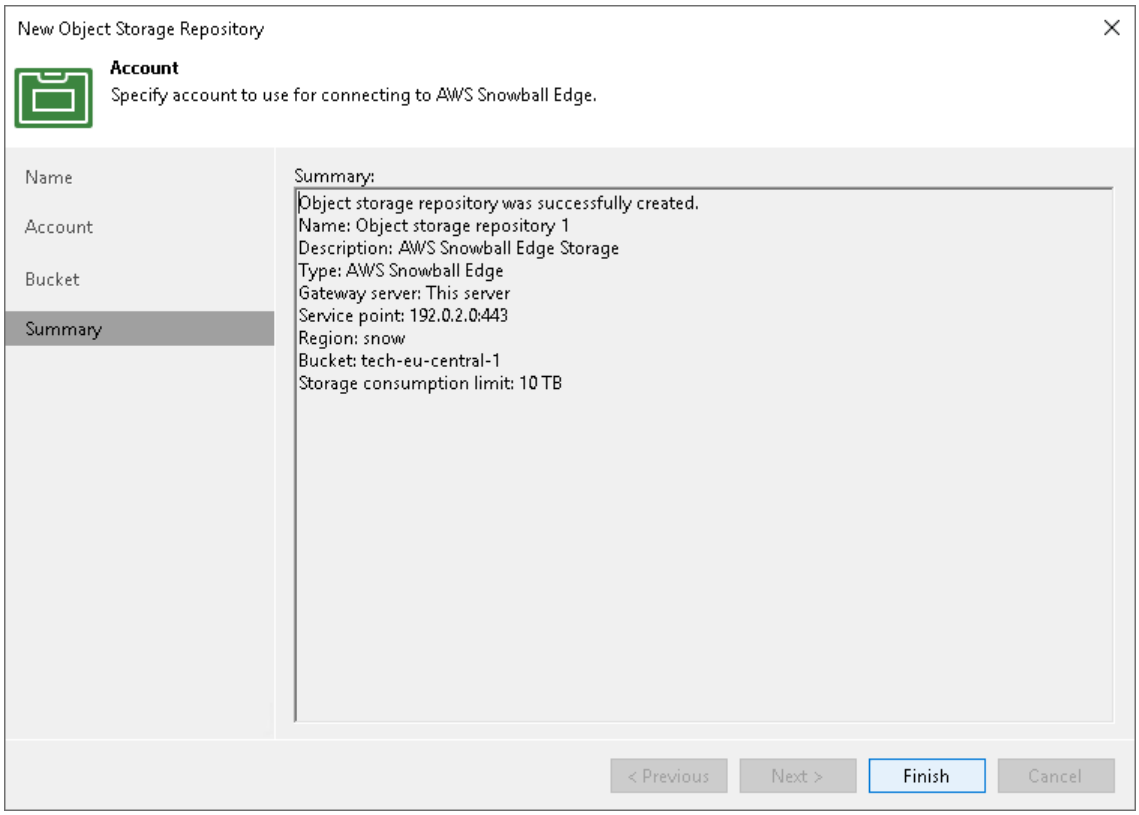


## Adding Google Cloud Object Storage

For more information about Google Cloud object storage, see this Google article.

Before you add a Google Cloud object storage to the backup infrastructure, check prerequisites and required permissions. After that, use the **New Object Repository** wizard.

## Step 1. Launch New Object Storage Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **Hyperscalers** > **Google Cloud Storage**.

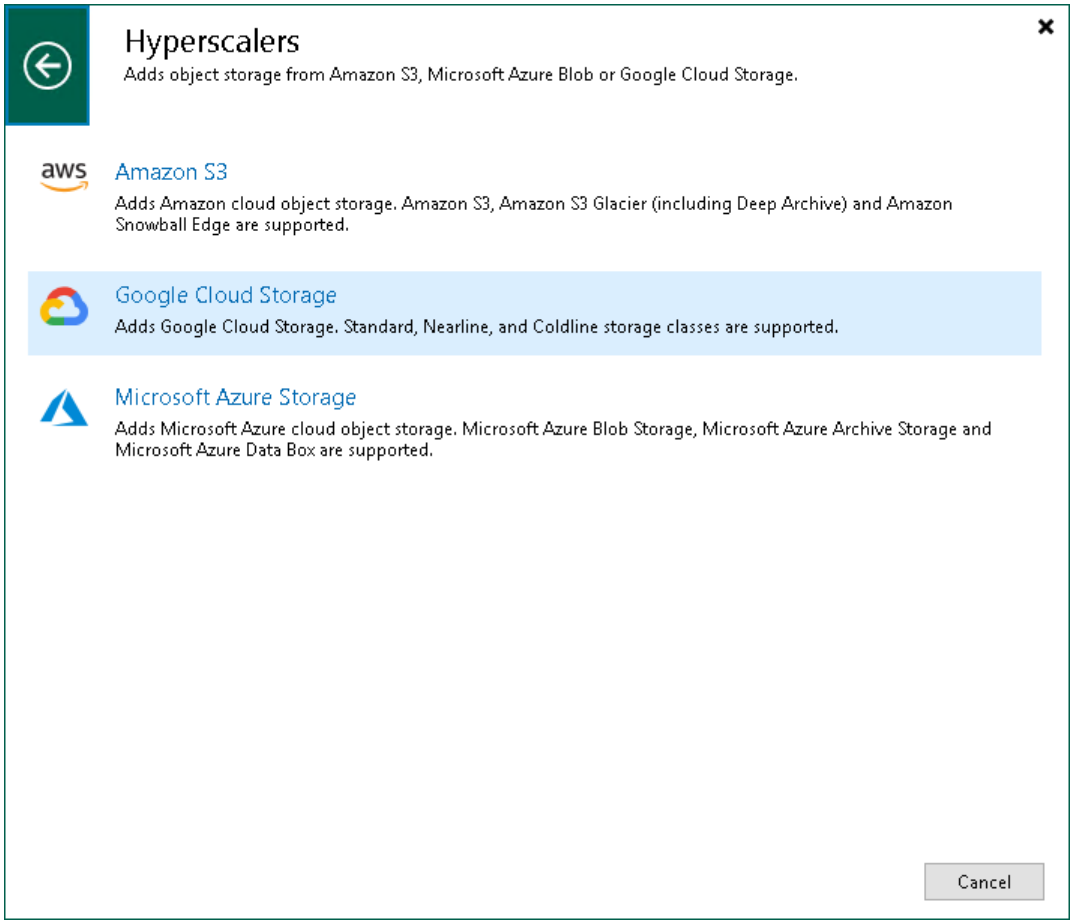

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

If you want to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

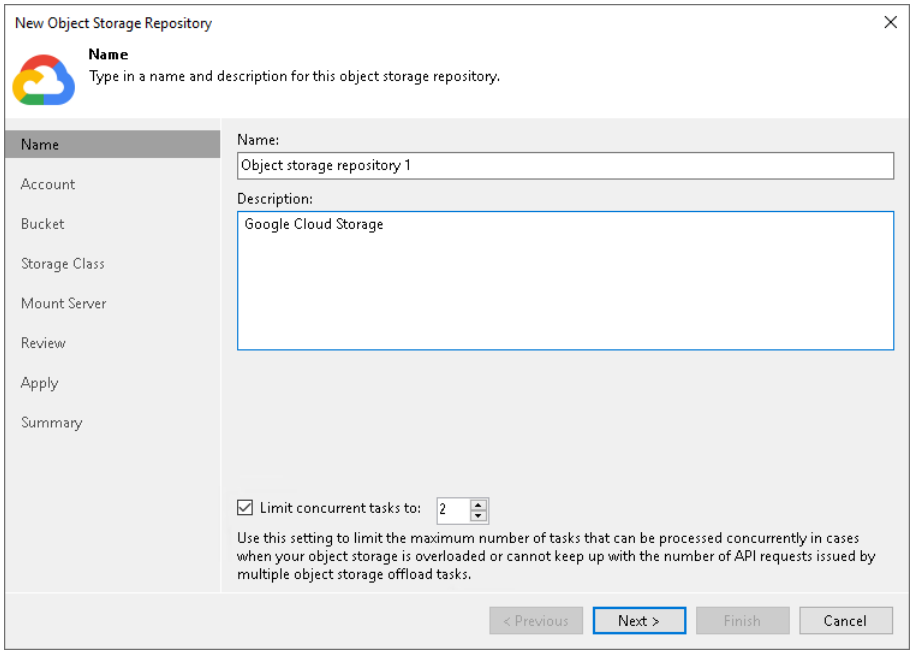

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify Google Cloud credentials and connection settings that Veeam Backup & Replication will use to transfer data to the object storage repository.

1. From the **Credentials** drop-down list, select user credentials to access your Google Cloud object storage.

   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click Add and provide your access and secret keys, as described in section Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

   The user account must have permissions listed in section Permissions.

2. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will access the object storage repository:

   - **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.
   - **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

     By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

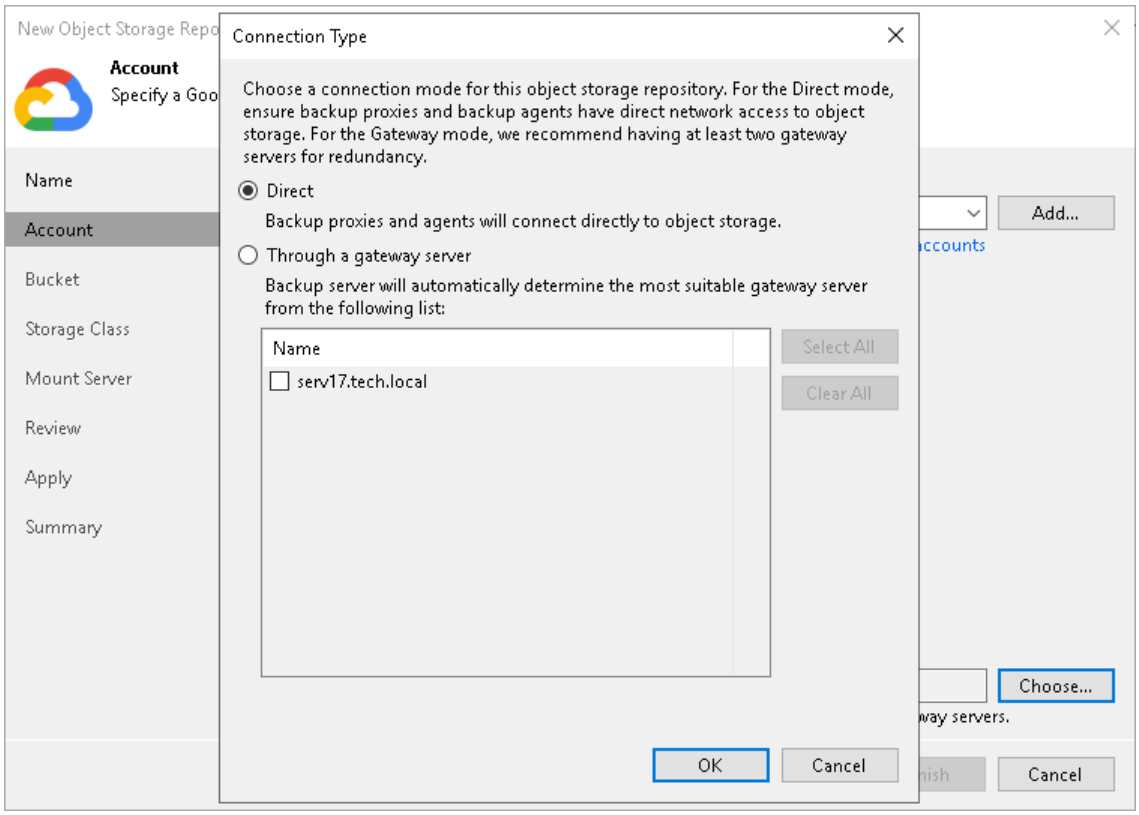

## Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify the bucket and folder that will be used to store data:

1. From the **Data center region** drop-down list, select a region.
2. In the **Bucket** field, enter a name of the bucket or click **Browse** to get the necessary bucket.

   Note that you must create the bucket where you want to store your backup data beforehand.
3. In the **Folder** field, enter a cloud folder name to which you want to map your object storage repository. Alternatively, click **Browse** and either select an existing folder or click **New Folder**.
4. Select the **Limit object storage consumption to** check box to define a soft limit for your object storage consumption. If this limit is exceeded during a job run, Veeam Backup & Replication will complete the job. However, a new job will not be able to start unless you remove the extra data that exceeds the limit or change the soft limit settings. Provide the value in TB or PB.

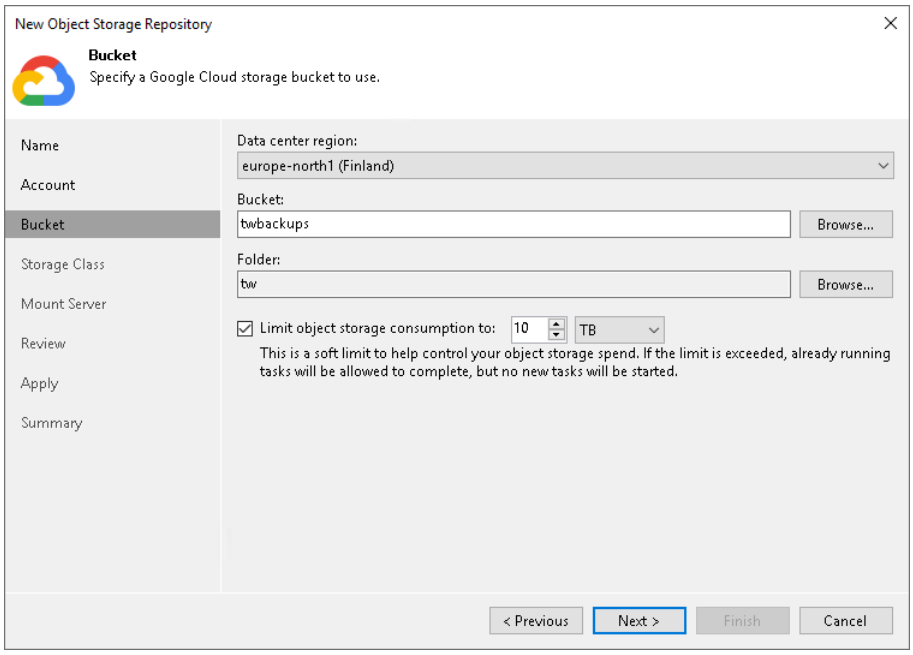

## Step 5. Specify Object Storage Class Settings

At the **Storage Class** step of the wizards, specify a storage class that you want assign to data blocks that you keep in Google Cloud object storage. For more information on Google Cloud object storage classes, see Google Documentation.

- Select the **Standard** option to assign the *Standard storage* class to data blocks. Use this option if you plan to access your data frequently and store it for a short period of time.
- Select the **Nearline** option to assign the *Nearline storage* class to data blocks. Use this option if you want to access data rarely (for example, once in a month or less) and plan to store data at least for 30 days.
- Select the **Coldline** option to assign the *Coldline storage* class to data blocks. Use this option if you plan to access data rarely (for example, once a quarter) and plan to store data minimum 90 days.

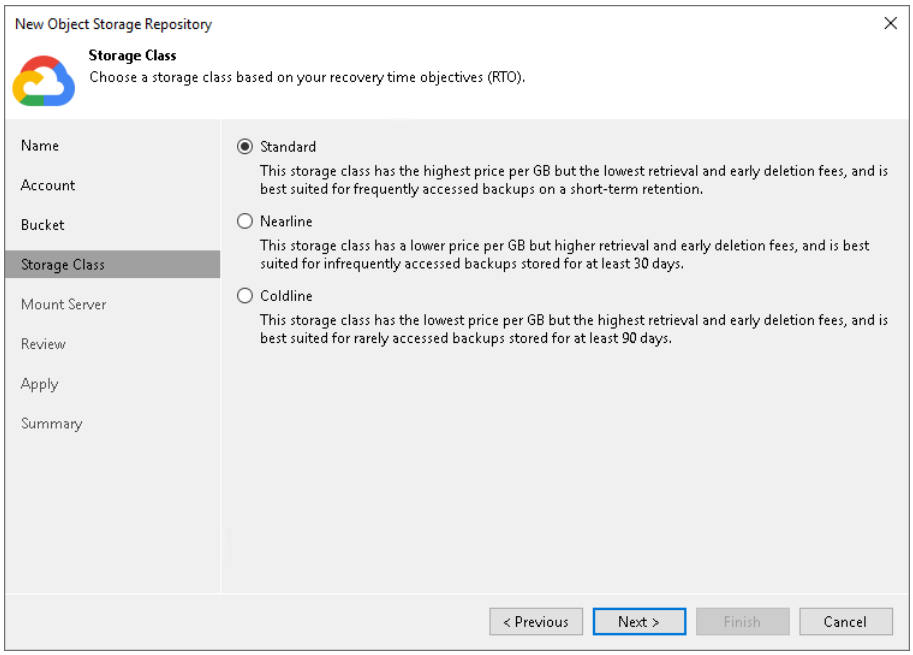

### Step 6. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for restore operations, and configure a helper appliance. The helper appliance is a temporary VM instance that Veeam Backup & Replication deploys in Google Compute Engine to perform a health check of backup files and apply retention to unstructured data backup files. For more information, see Health Check for Object Storage Repositories and Helper Appliance in Unstructured Data Backup. After Veeam Backup & Replication completes these operations, it removes the helper appliance from Google Compute Engine.

## Specifying Mount Server Settings

To specify the mount server settings, do the following:

1. From the **Mount Server** drop-down list, select a server that you want to use as a mount server. Veeam Backup & Replication uses this server during restore operations to mount VM disks directly from objects located in object storage repositories. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, click **Add New** on the right to open the **New Windows Server wizard**. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder to keep cache that is created during mount operations.

3. Select the **Enable vPower NFS service on the mount server** check box to allow the Veeam vPower NFS Service access the object storage repository. Veeam Backup & Replication will enable the Veeam vPower NFS Service on the necessary mount server. For more information, see Veeam vPower NFS Service.

4. Click **Ports** to customize network ports used by the Veeam vPower NFS Service. In the **vPower NFS Port Settings** window, specify the following settings:

   - Next to the **Mount Port** section, specify the port that the Veeam vPower NFS Service will use to mount the vPower NFS datastore to the ESXi host.
   - Next to the **vPower NFS** port section, specify the port that the Veeam vPower NFS Service will use to connect to the target NFS share.

   For information on ports used by default, see Ports.

5. To specify the helper appliance settings, click **Configure**. From the **Managed server** drop-down list, select a server that you want to use as the helper appliance.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

# Configuring Helper Appliance

To configure the helper appliance, at the **Mount Server** step, click **Configure** and in the **Advanced Settings** window, specify the following settings:

1. From the **Account** drop-down list, select a credentials record to access your Google Cloud object storage. Veeam Backup & Replication will use this credentials record to connect to Google Compute Engine within Google Cloud and create the helper appliance.

   If you have not added the credentials record beforehand, click **Manage cloud accounts** or **Add** to add the necessary service account. For more information, see Google Cloud Service Accounts.

2. Next to the **Helper appliance** field, click **Configure**. In the **Appliance Settings** window, specify the following settings:

   a. From the **Machine type** drop-down list, select the machine type for the helper appliance. The speed and cost of operations that the helper appliance performs depend on the machine type. For information on instance types, see Google Cloud documentation.

   b. From the **VPC** drop-down list, select a VPC network where a helper appliance will be launched. For more information on VPC networks, see Google Cloud documentation.

      To be able to select the necessary VPC from the drop-down list, you must create it beforehand as described in the Google Cloud documentation..

   c. From the **Subnet** drop-down list, select a subnet where a helper appliance will reside. For more information on subnets, see Google Cloud documentation.

      To be able to select the subnet from the drop-down list, you must create it beforehand as described in the Google Cloud documentation.

   d. In the **Redirector port** field, specify the TCP port that Veeam Backup & Replication will use to route requests between the helper appliance and backup infrastructure components.

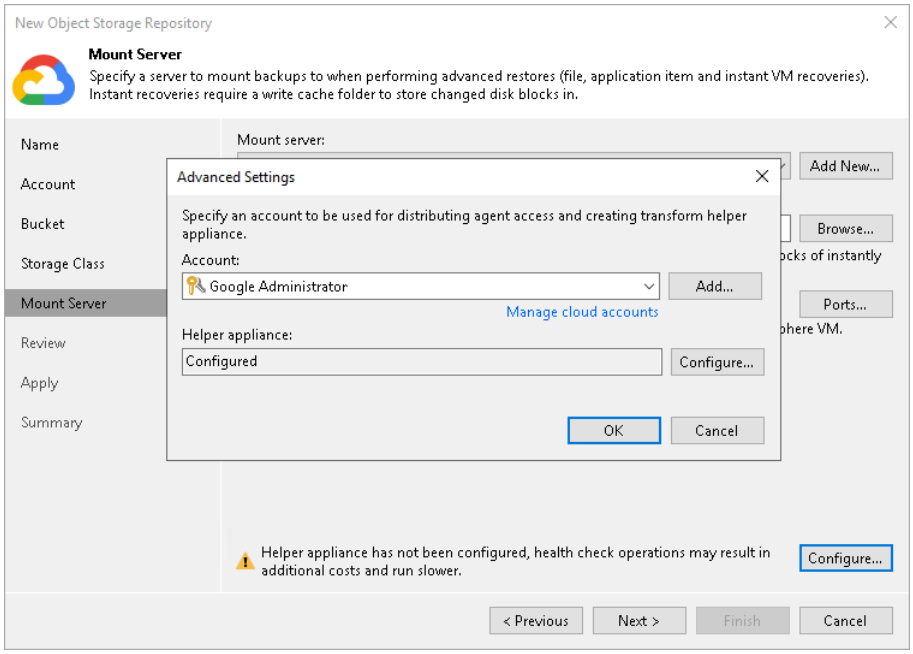

## Step 7. Review Components

At the **Review** step of the wizard, review what components will be processed on the mount server server and their status.

If the backup repository contains backups, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Object Storage (Imported)** node.

If the backup repository contains guest file system index files, select the **Import guest file system index data to the catalog** check box. Veeam Backup & Replication will import index files together with backup files, and you will be able to search for guest OS files inside imported backups. For more information, see the Guest OS File Restore section of the Veeam Backup Enterprise Manager Guide.

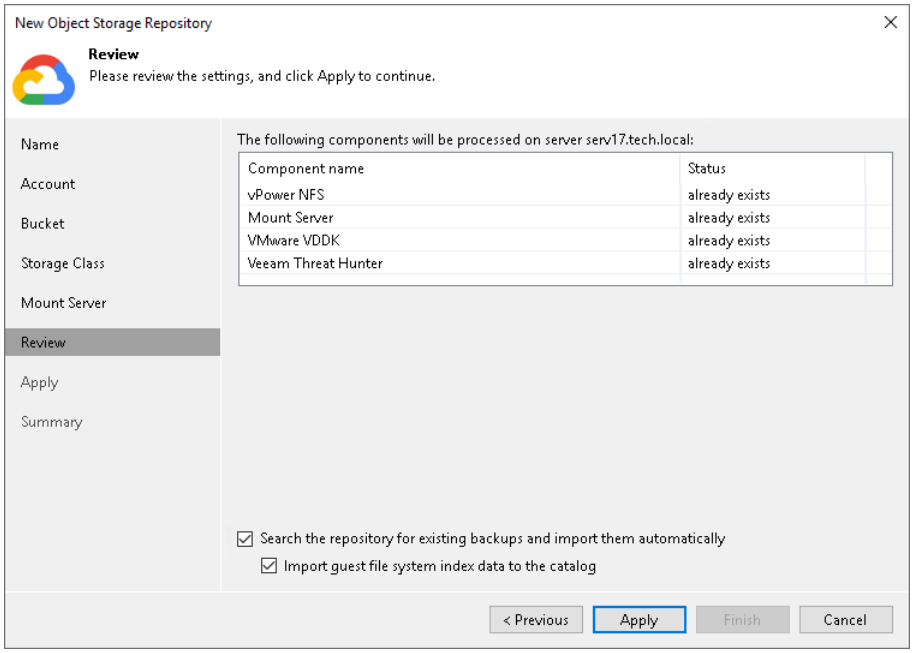

## Step 8. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to complete saving your settings to the configuration database and create backup infrastructure objects.

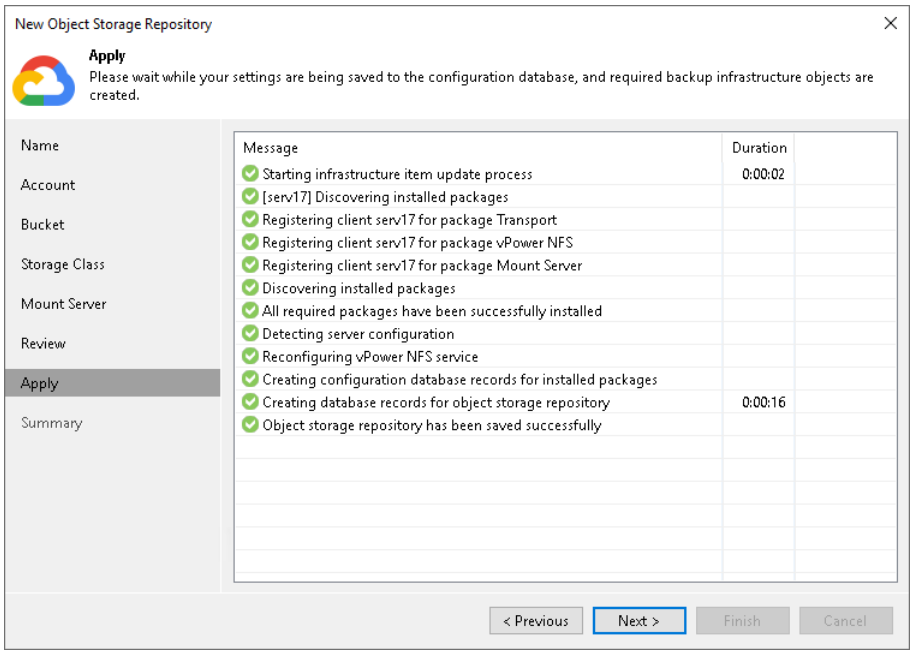

### Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

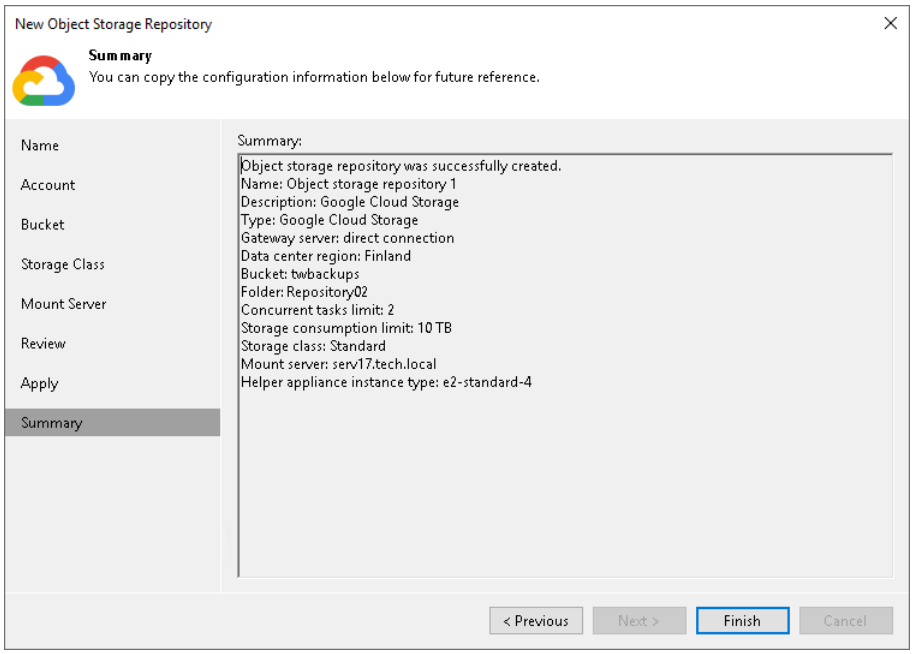


## Adding Microsoft Azure Blob Storage, Microsoft Azure Archive Storage and Microsoft Azure Data Box

Before you add a Microsoft Azure Blob storage, Azure Archive storage or Azure Data Box to the backup infrastructure, check prerequisites and required permissions. After that, use the **New Object Storage Repository** wizard.

### Adding Azure Blob Storage

This section describes how to add Microsoft Azure Blob storage to the backup infrastructure. For more information about Microsoft Azure Blob storage, see Microsoft Docs.

To add Microsoft Azure Blob storage, use the **New Object Storage Repository** wizard.

## Step 1. Launch New Object Storage Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **Hyperscalers** > **Microsoft Azure Storage** > **Azure Blob Storage**.

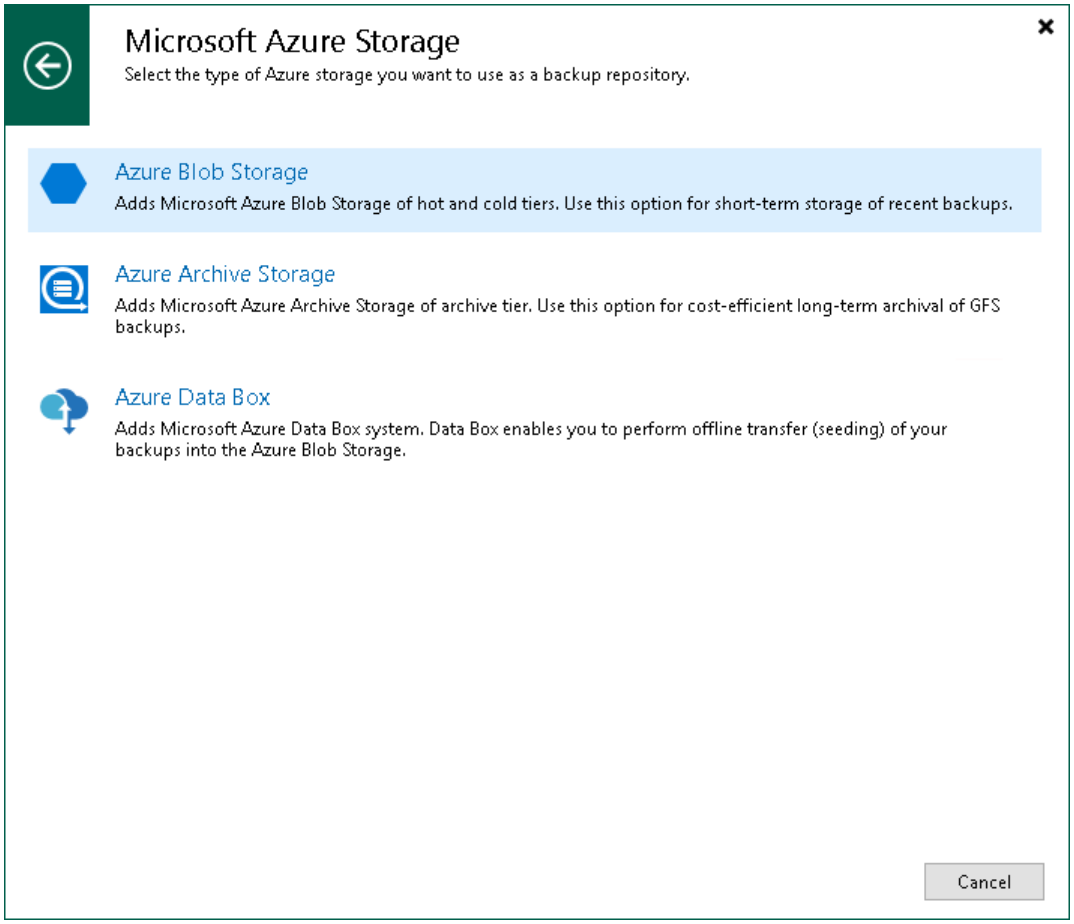

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

If you want to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

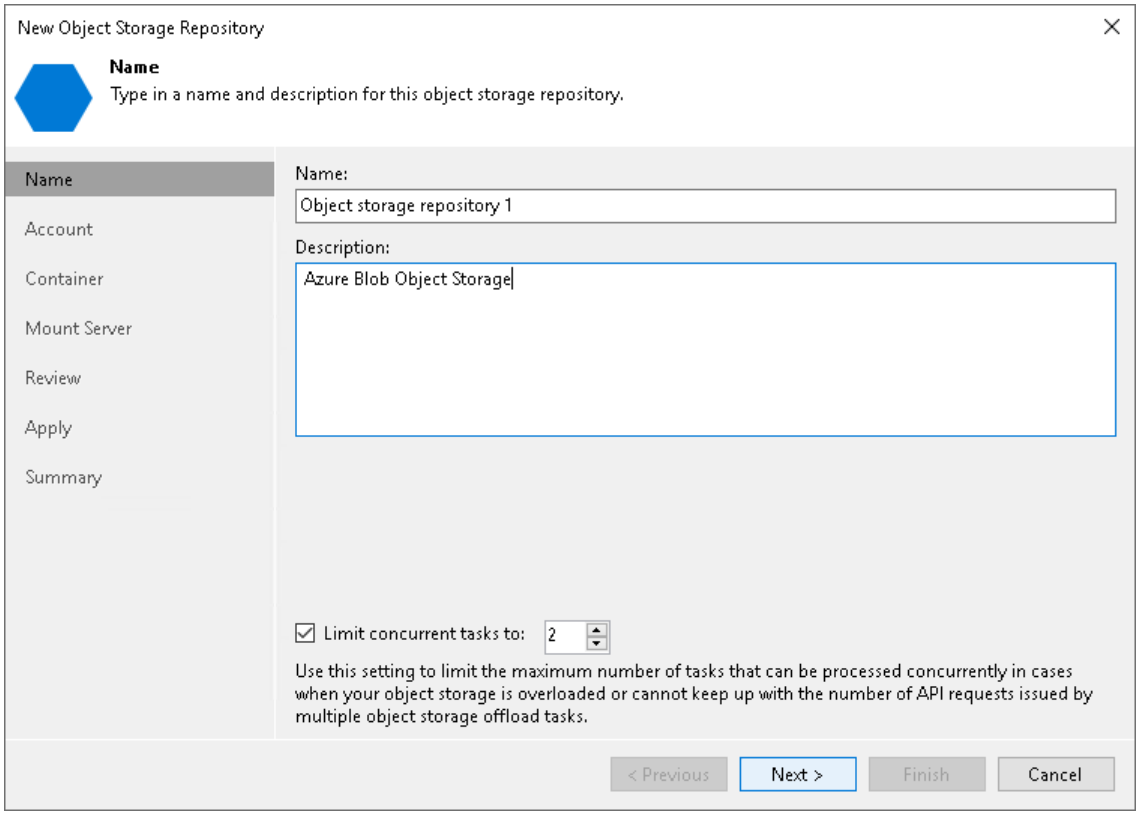

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. From the **Credentials** drop-down list, select user credentials to access your Azure Blob storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, add the credentials record using either your account name and a shared key or specify Microsoft Azure Entra ID storage account. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

   > **IMPORTANT**
   >
   > Keep in mind that in case you add a new Microsoft Azure Entra ID storage account, it takes from 30 to 60 seconds to propagate it on the Microsoft Azure side.

2. From the **Region** drop-down list, select an Azure region.
3. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:

   - **Direct** — select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.

- **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

  By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

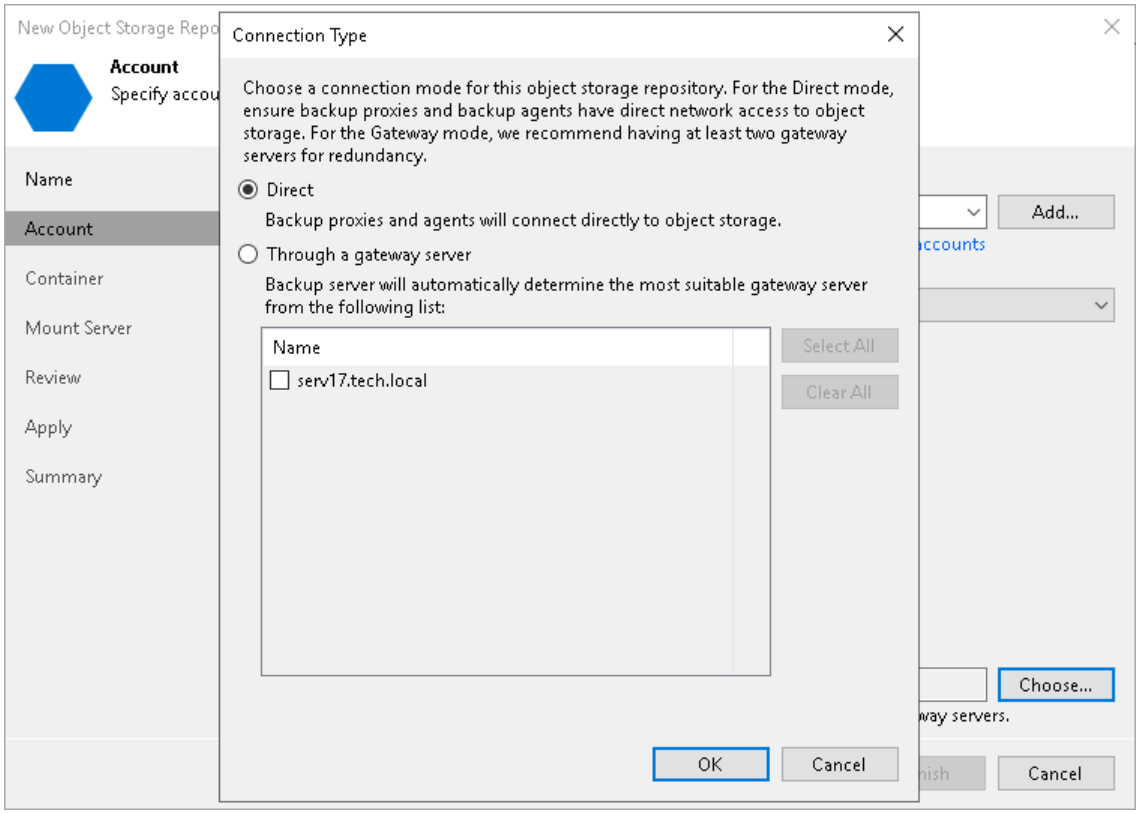

## Step 4. Specify Object Storage Settings

At the **Container** step of the wizard, specify the container, the folder that will be used to store data and select an access tier in which Veeam Backup & Replication will keep blobs. For more information on access tiers, see Microsoft Docs.

1. From the **Container** drop-down list, select a container.

   Make sure that the container where you want to store your backup data was created in advance.

   > **NOTE**
   >
   > The default *Root* container is not supported. For more information about this container, see Microsoft Docs.

2. To the right of the **Folder** field, click **Browse** and either select an existing folder or click **New Folder**.
3. Select the **Limit object storage consumption** to check box to define a soft limit for your object storage consumption. If this limit is exceeded during a job run, Veeam Backup & Replication will complete the job. However, a new job will not be able to start unless you remove the extra data that exceeds the limit or change the soft limit settings. Provide the value in TB or PB.
4. Select the **Make recent backups immutable for** check box to prohibit deletion of blocks of data from object storage. Specify the immutability period. For more information, see Immutability for Object Storage Repositories.

   Note that the maximum immutability period you can set in the Veeam Backup & Replication UI is 90 days. If you want to set immutability to a longer period, use the Set-VBRAzureBlobRepository cmdlet.
5. If you plan to access your backup data infrequently, select the **Use cool blob storage tier (may result in higher cost)** check box. Veeam Backup & Replication will keep blobs in the *cool* access tier.

   > **IMPORTANT**
   >
   > Consider the following:
   >
   > - If you do not enable the *cool* access tier option, Veeam Backup & Replication will use the *hot* access tier.
   > - If you specify the *cold* tier in the Azure storage account settings, Veeam Backup & Replication will ignore this setting and will use the *hot* access tier.
   > - If you enable this option and plan to use this object storage as a performance or capacity tier, do not target to this repository any jobs that constantly send backup data to this storage: scheduled regular backup and backup copy jobs that run without GFS, jobs with transactions logs enabled, jobs created by Veeam Enterprise Plug-ins. Otherwise, it will result in higher costs.

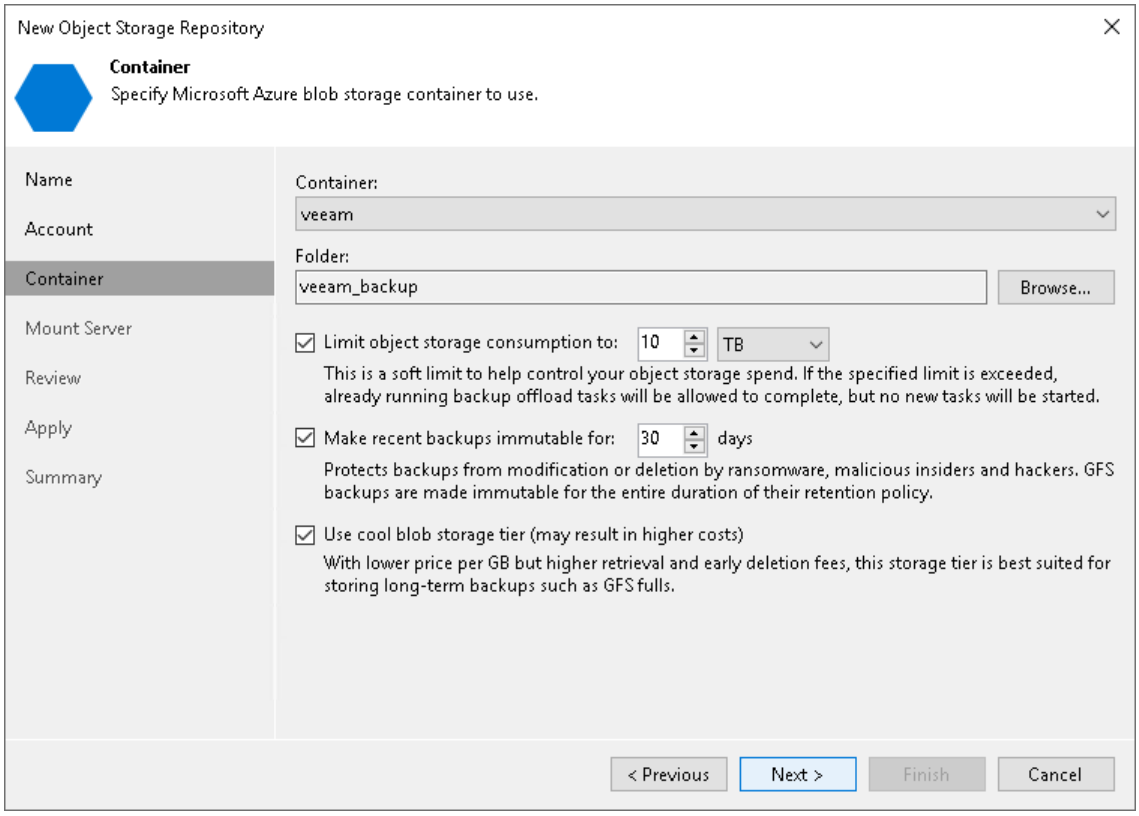
New Object Storage Repository
Container
Specify Microsoft Azure blob storage container to use.
Name
Account
Container
Mount Server
Review
Apply
Summary
Container:
veeam
Folder:
veeam_backup
Browse...
Limit object storage consumption to:
10
TB
This is a soft limit to help control your object storage spend. If the specified limit is exceeded, already running backup offload tasks will be allowed to complete, but no new tasks will be started.
Make recent backups immutable for:
30
days
Protects backups from modification or deletion by ransomware, malicious insiders and hackers. GFS backups are made immutable for the entire duration of their retention policy.
Use cool blob storage tier (may result in higher costs)
With lower price per GB but higher retrieval and early deletion fees, this storage tier is best suited for storing long-term backups such as GFS fulls.
< Previous
Next >
Finish
Cancel

### Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for restore operations, and configure a helper appliance. The helper appliance is a temporary host that Veeam Backup & Replication deploys on your Microsoft Azure Blob storage to perform a health check of backup files and apply retention to unstructured data backup files. For more information, see Health Check for Object Storage Repositories and Helper Appliance in Unstructured Data Backup. After Veeam Backup & Replication completes these operations, it removes the helper appliance from the Microsoft Azure Blob storage.

## Specifying Mount Server Settings

To specify the mount server settings, do the following:

1. From the **Mount Server** drop-down list, select a server that you want to use as a mount server. Veeam Backup & Replication uses this server during restore operations to mount VM disks directly from objects located in object storage repositories. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, click **Add New** on the right to open the **New Windows Server wizard**. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder to keep cache that is created during mount operations.

3. Select the **Enable vPower NFS service on the mount server** check box to allow the Veeam vPower NFS Service access the object storage repository. Veeam Backup & Replication will enable the Veeam vPower NFS Service on the necessary mount server. For more information, see Veeam vPower NFS Service.

4. Click **Ports** to customize network ports used by the Veeam vPower NFS Service. In the **vPower NFS Port Settings** window, specify the following settings:

   - Next to the **Mount Port** section, specify the port that the Veeam vPower NFS Service will use to mount the vPower NFS datastore to the ESXi host.
   - Next to the **vPower NFS** port section, specify the port that the Veeam vPower NFS Service will use to connect to the target NFS share.

   For information on ports used by default, see Ports.

5. To specify the helper appliance settings, click **Configure**. From the **Managed server** drop-down list, select a server that you want to use as the helper appliance.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

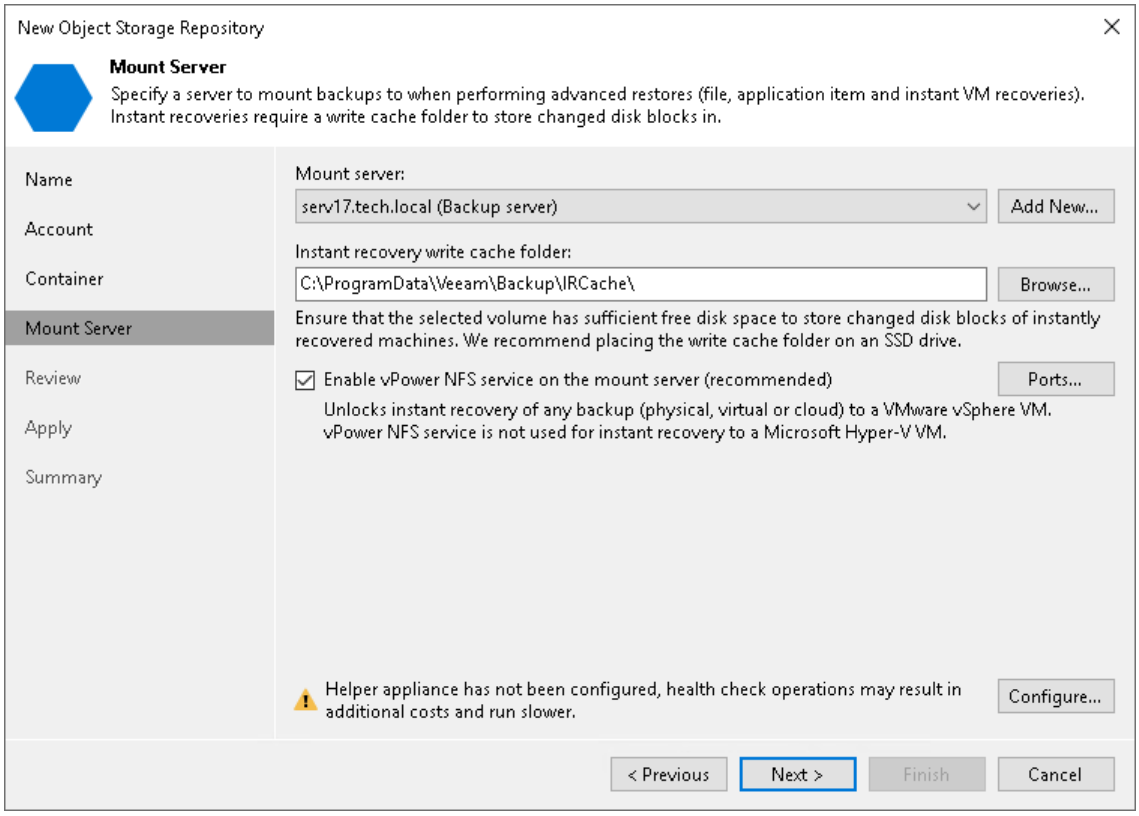


# Configuring Helper Appliance

To configure the helper appliance, at the **Mount Server** step, click **Configure** and in the **Helper Appliance Settings** window, specify the following settings:

1. From the **Subscription** drop-down list, select your Microsoft Azure subscription credentials.

   If you have not set up credentials beforehand, click **Add**. You will be prompted to the Specify Microsoft Azure Compute Account Name wizard. Follow the wizard to add your account. Before adding your Microsoft Azure account, check the prerequisites.

2. From the **Size** drop-down list, select the size of the helper appliance.

3. From the **Resource group** drop-down list, select a resource group that will be associated with the helper appliance.

   To be able to select the necessary resource group from the drop-down list, you must create it beforehand as described in the Microsoft Docs.

4. From the **Virtual network** drop-down list, select a network to which the helper appliance must be connected.

   To be able to select the necessary network from the drop-down list, you must create it beforehand as described in the Microsoft Docs.

   > **IMPORTANT**
   >
   > Veeam Backup & Replication creates a default network security group within a virtual network with the inbound rules that allow connection using the 443 and 22 ports from everywhere (0.0.0.0/0).

5. From the **Subnet** drop-down list, select a subnet for the helper appliance.

   To be able to select the necessary subnet from the drop-down list, you must create it beforehand as described in the Microsoft Docs.

6. In the **Redirector port** field, specify the port that Veeam Backup & Replication will use to route requests between the helper appliance and backup infrastructure components.

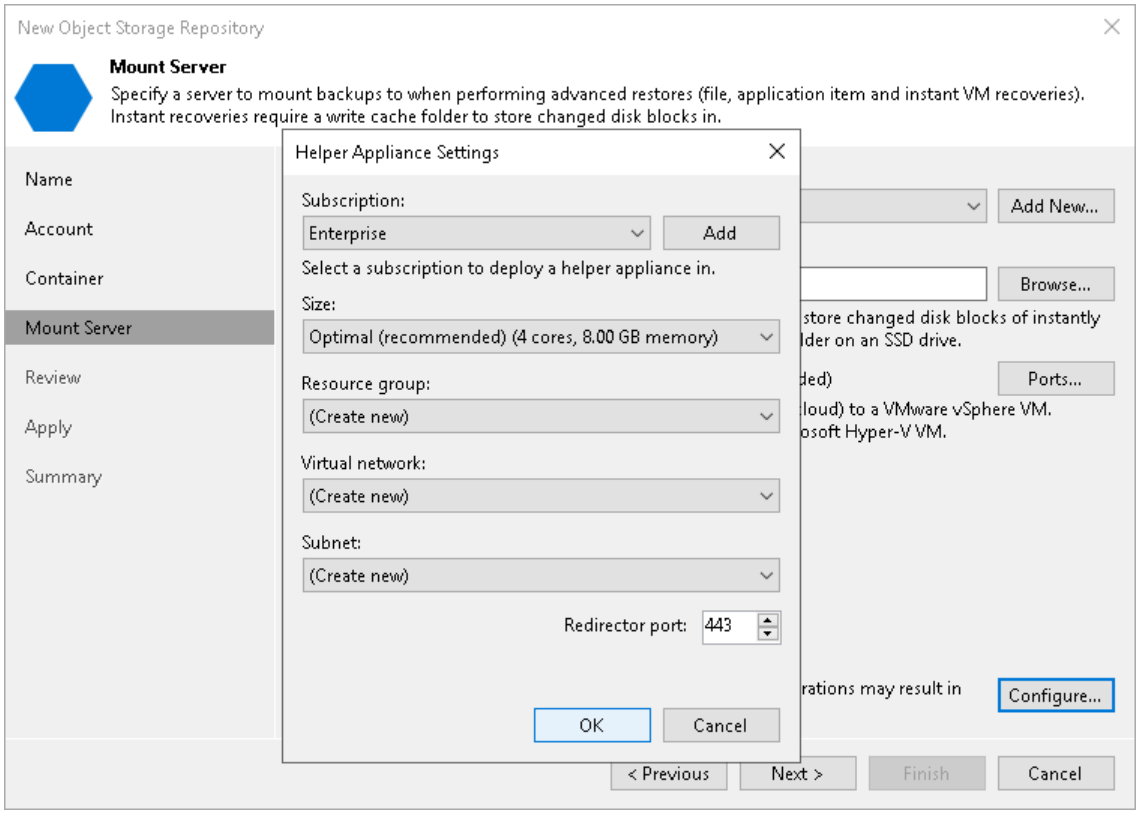

## Step 6. Review Components

At the **Review** step of the wizard, review what components will be processed on the mount server server and their status.

If the backup repository contains backups, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Object Storage (Imported)** node.

If the backup repository contains guest file system index files, select the **Import guest file system index data to the catalog** check box. Veeam Backup & Replication will import index files together with backup files, and you will be able to search for guest OS files inside imported backups. For more information, see the Guest OS File Restore section of the Veeam Backup Enterprise Manager Guide.

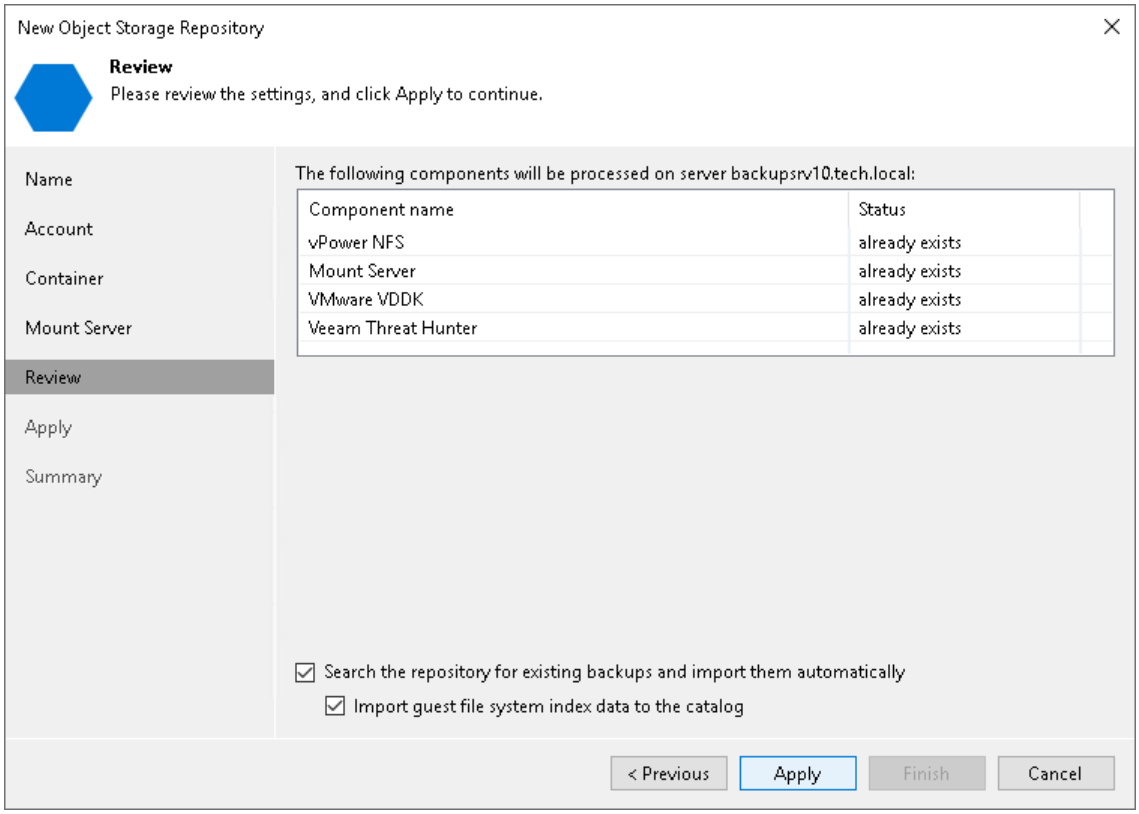

## Step 7. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to complete saving your settings to the configuration database and create backup infrastructure objects.

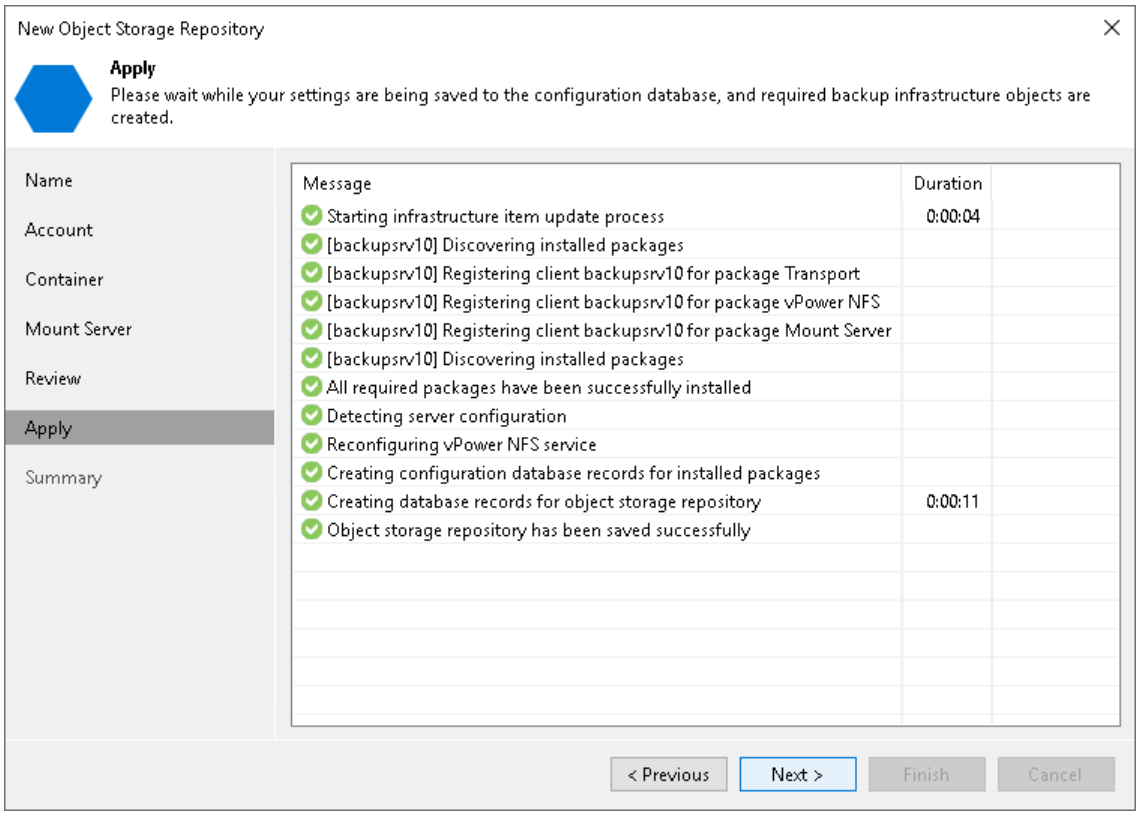

## Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

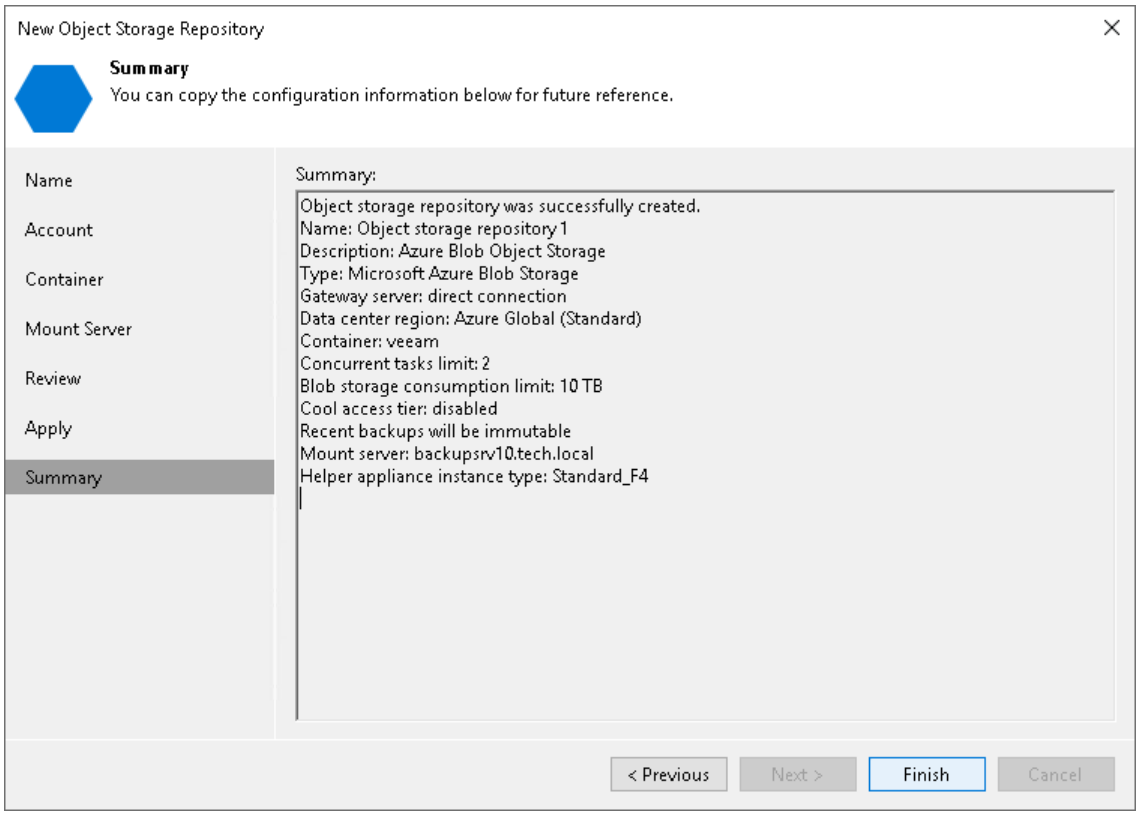


## Adding Azure Archive Storage

This section describes how to add Microsoft Azure Archive Storage to the backup infrastructure as a backup repository. You can only use this repository as an archive extent of the scale-out backup repository. For more information, see Archive Tier. For more information about Azure Archive Storage, see Microsoft Docs.

To add Microsoft Azure Archive Storage, use the **New Object Storage Repository** wizard.

## Step 1. Launch New Object Storage Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **Hyperscalers** > **Microsoft Azure Storage** > **Azure Archive Storage**.

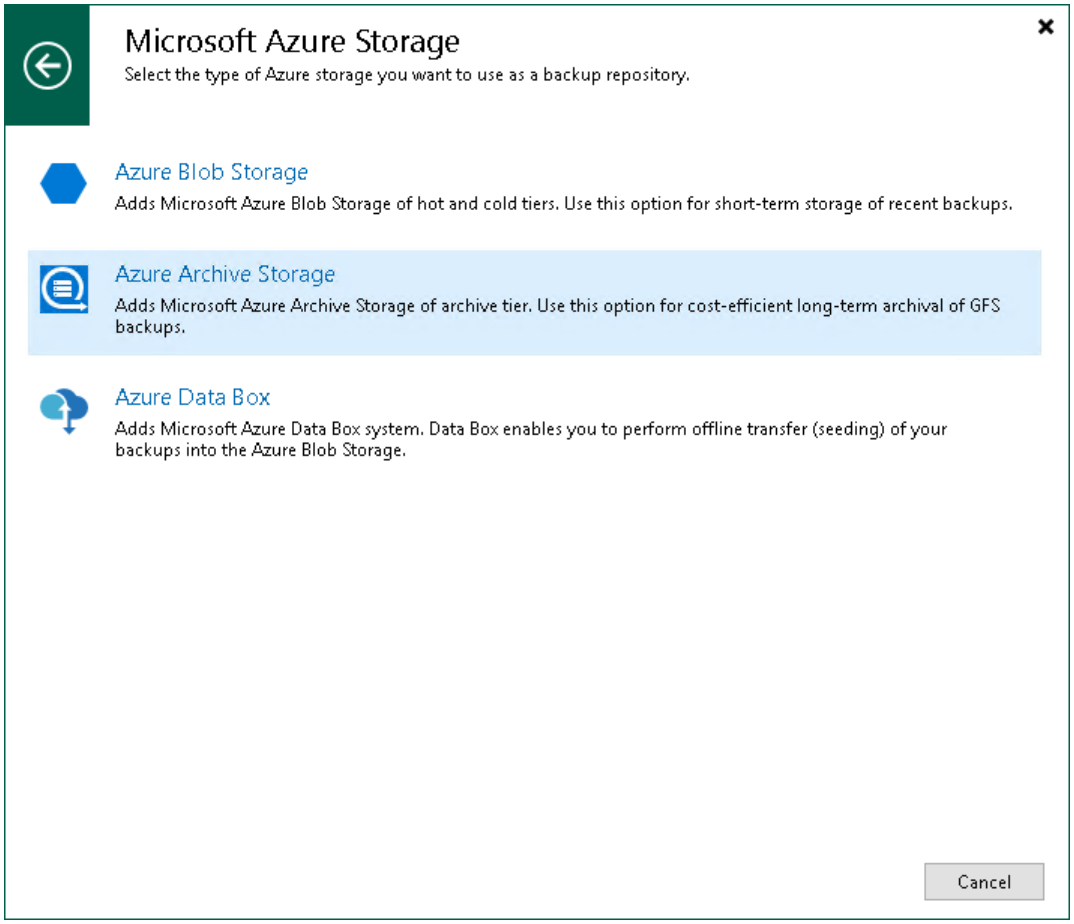

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

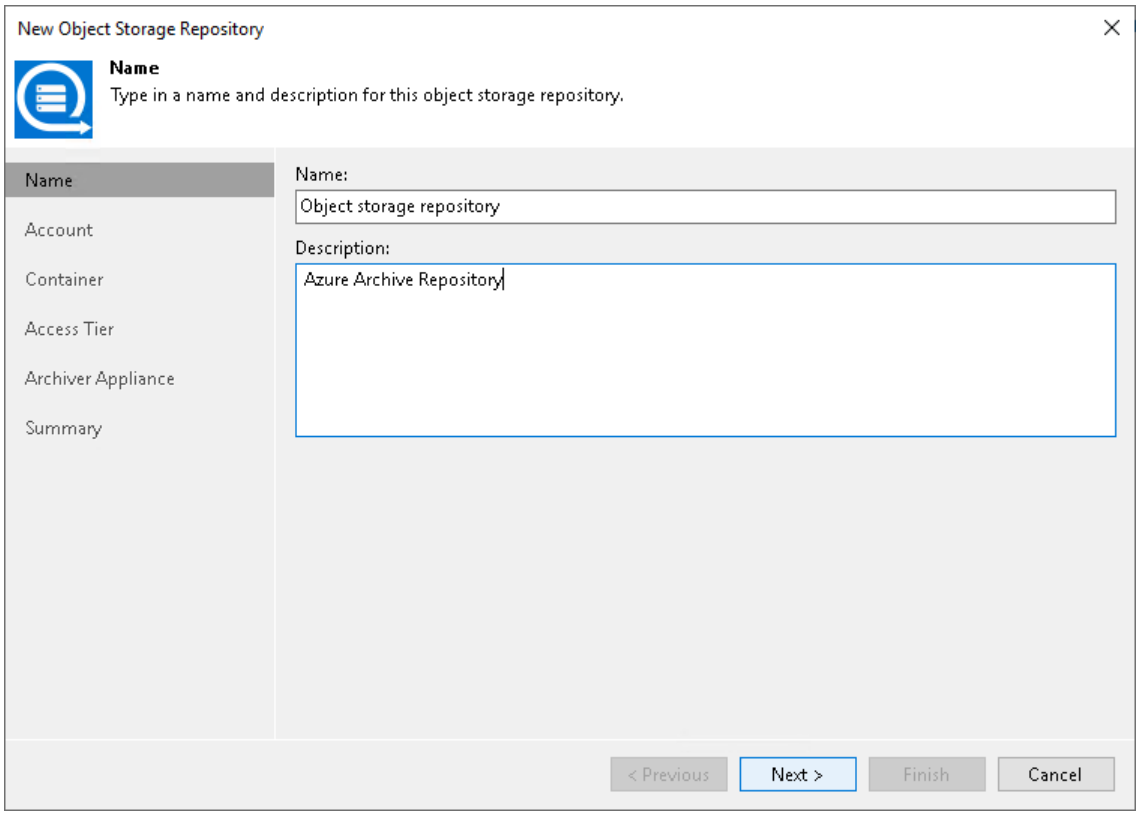

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. From the **Credentials** drop-down list, select user credentials to access your Azure Archive Storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and the credentials record using either your account name and a shared key or specify Microsoft Azure Entra ID storage account. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

   The user account must have permissions listed in section Permissions.

   > **IMPORTANT**
   >
   > Consider the following:
   >
   > - Keep in mind that in case you add a new Microsoft Azure Entra ID storage account, it takes from 30 to 60 seconds to propagate it on the Microsoft Azure side.
   > - Veeam Backup & Replication archive tier feature does not support Microsoft Azure Stack Hub Compute accounts.

2. From the **Region** drop-down list, select the Azure region.
3. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:
   - **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.

- **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

  By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

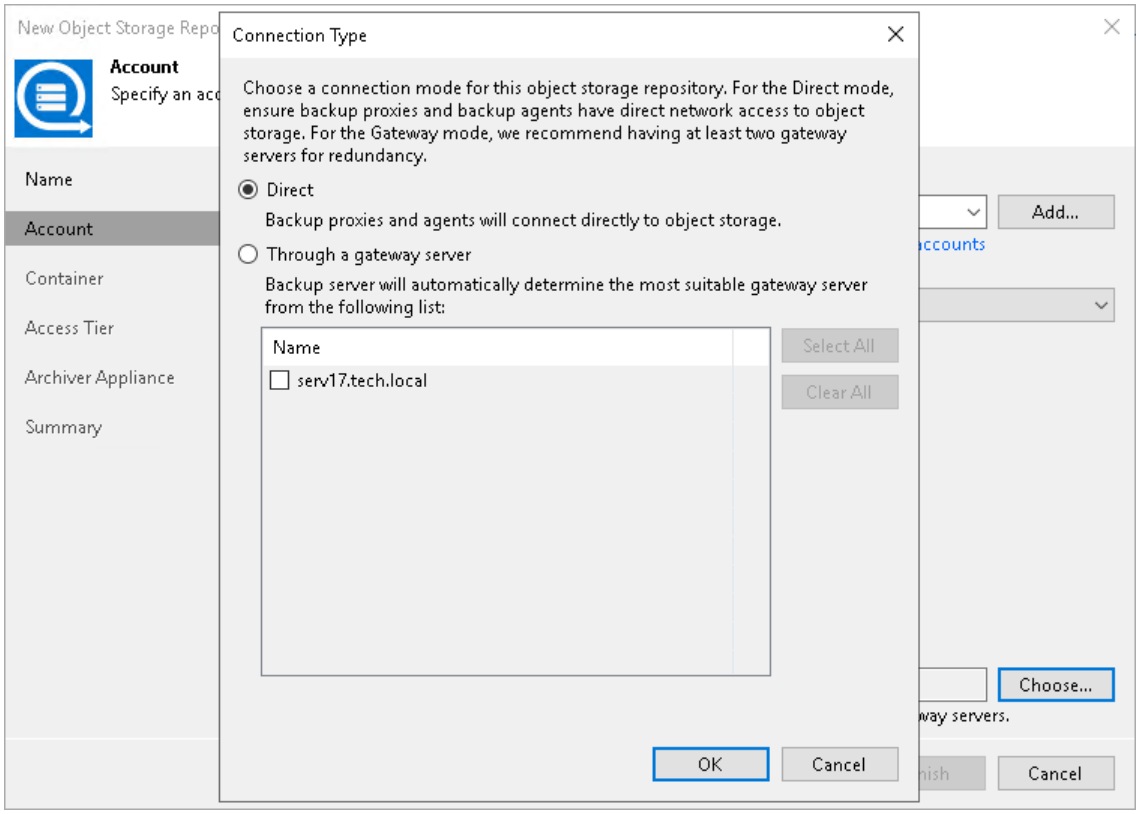

## Step 4. Specify Object Storage Settings

At the **Container** step of the wizard, specify the container and folder that will be used to store data:

1. From the **Container** drop-down list, select a container.

   Make sure that the container where you want to store your backup data was created in advance.

2. To the right of the **Folder** field, click **Browse** and either select an existing folder or click **New Folder**.
3. Select the **Make backups immutable for the entire duration of their retention policy** check box to prohibit deletion of blocks of data from object storage. The immutability period will be equal to the retention period (if any) of the data blocks. All the types of files that are eligible for archive storage can be made immutable. For more information on the immutability feature and the retention policy for each file type, see Immutability for Archive Tier.

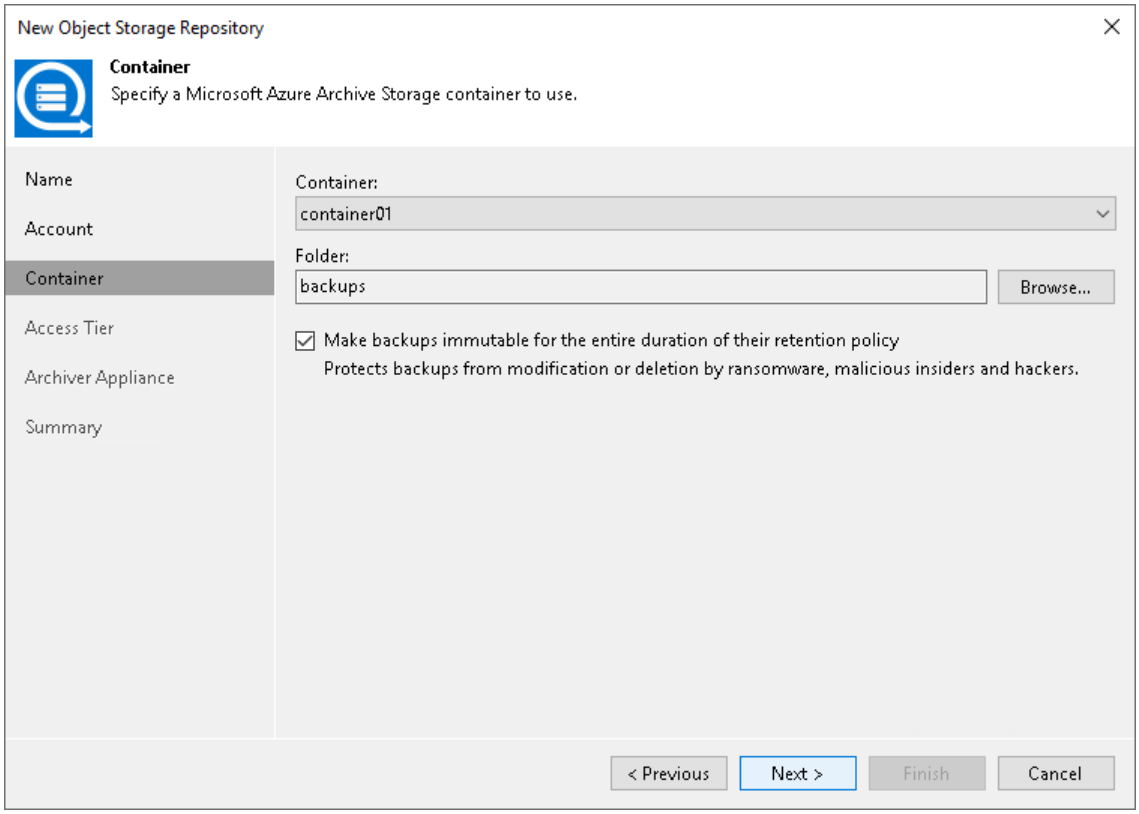

## Step 5. Specify Access Tier Settings

At the **Access Tier Settings** step of the wizard, specify an Azure Storage access tier that will be assigned to blocks in Azure Archive Storage. For more information on Azure Storage access tiers, see Microsoft Docs.

- Select the **Archive** option to assign the *archive* access tier to data blocks. Use this option if you plan to access your data rarely and store it at least for 180 days.
- Select the **Cold** option to assign the *cold* access tier to data blocks. Use this option if you plan to access your data frequently and store it at least for 90 days.

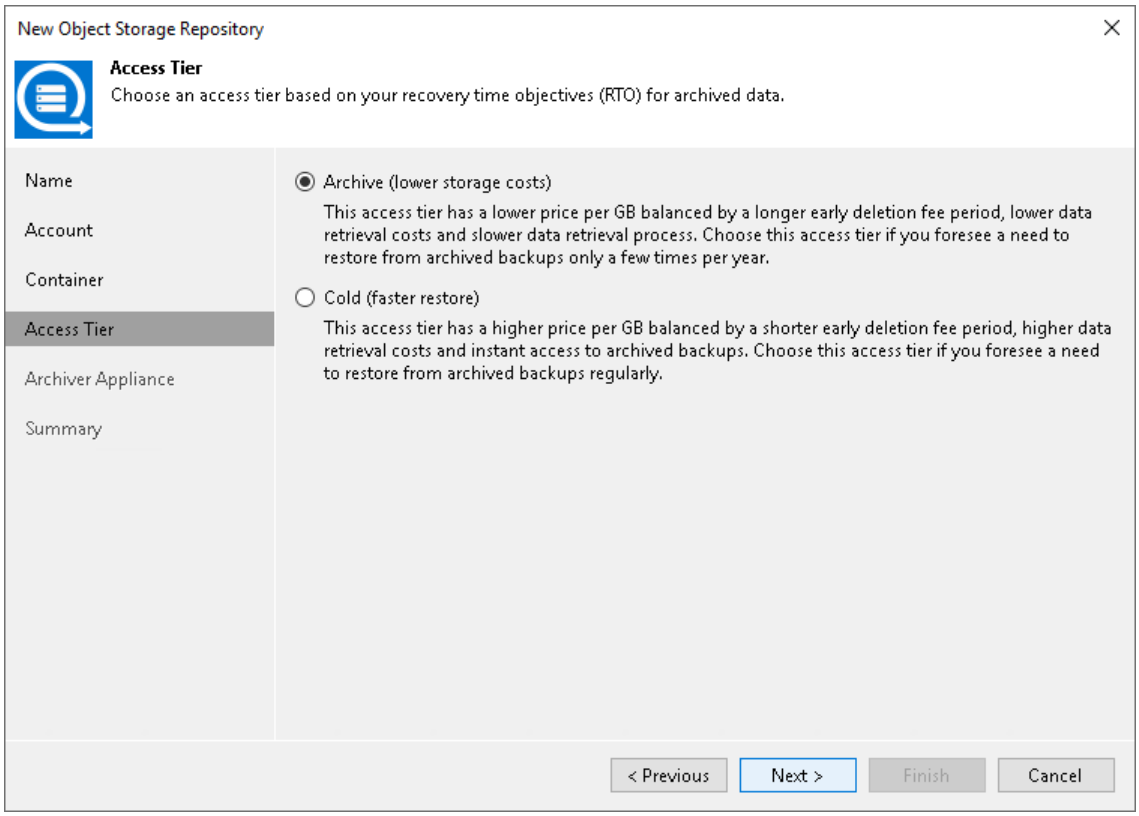

## Step 6. Specify Archiver Appliance

At the **Archiver Appliance** step of the wizard, you can specify archiver appliance settings. An archiver appliance is an auxiliary instance that is necessary to transfer data from Azure Blob storage to Azure Archive Storage. For more information, see the Archiver Appliances section.

> **NOTE**
>
> Veeam Backup & Replication must be able to connect to the machine that you will use as an archiver appliance. Therefore, if your backup server is not located within Microsoft Azure, you must configure public IP addresses for the virtual network in which the appliance resides. For more information on configuring the Azure virtual network, see Microsoft Docs.

To configure and set up an archiver appliance, you must add a Microsoft Azure account to Veeam Backup & Replication. For that, do the following:

1. From the **Subscription** drop-down list, select your Microsoft Azure subscription credentials. If you have not set up credentials beforehand, click **Add**. You will be prompted to the Adding Microsoft Azure Compute Accounts wizard. Follow the wizard to add your account. Before adding your Microsoft Azure account, check the prerequisites.
2. You can use the default settings or customize the archiver appliance. Click **Customize** and in the **Helper Appliance Settings** window and specify the following settings:
   a. From the **Size** drop-down list, select the size of the appliance. For details on the proxy types used by Veeam Backup & Replication, see this Veeam KB article.
   b. From the **Resource group** drop-down list, select a resource group that will be associated with the archiver appliance.
   c. From the **Virtual network** drop-down list, select a network to which the archiver appliance must be connected. If you want to configure Veeam Backup & Replication to connect with Azure Blob Storage Account private endpoints, see this Veeam KB article.

      > **IMPORTANT**
      >
      > Veeam Backup & Replication creates a default network security group within a virtual network with the inbound rules that allow connection using the 443 and 22 ports from everywhere (0.0.0.0/0).

   d. From the **Subnet** drop-down list, select the subnet for the archiver appliance.
   e. In the **Redirector port** field, specify the port that Veeam Backup & Replication will use to route requests between the archiver appliance and backup infrastructure components.

f. Click **OK**.

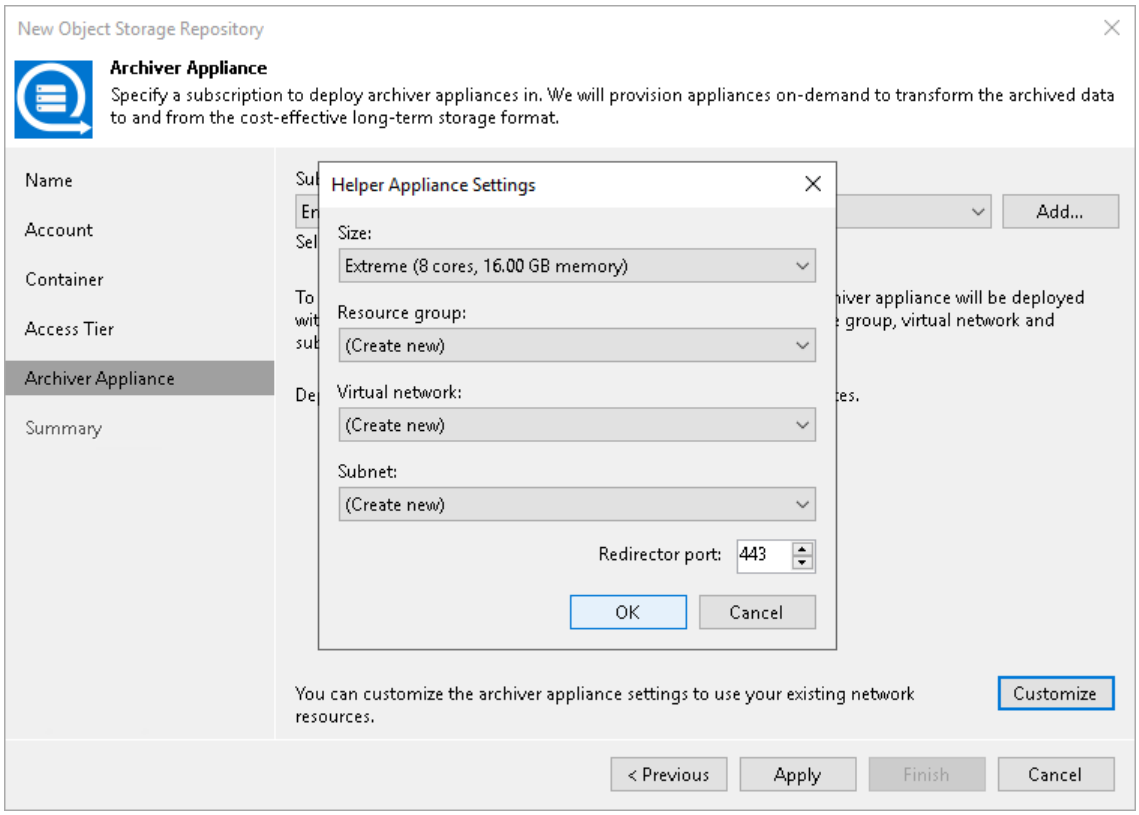
New Object Storage Repository
Archiver Appliance
Specify a subscription to deploy archiver appliances in. We will provision appliances on-demand to transform the archived data to and from the cost-effective long-term storage format.
Name
Account
Container
Access Tier
Archiver Appliance
Summary
Helper Appliance Settings
Size:
Extreme (8 cores, 16.00 GB memory)
Resource group:
(Create new)
Virtual network:
(Create new)
Subnet:
(Create new)
Redirector port:
443
OK
Cancel
Add...
You can customize the archiver appliance settings to use your existing network resources.
Customize
< Previous
Apply
Finish
Cancel

### Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

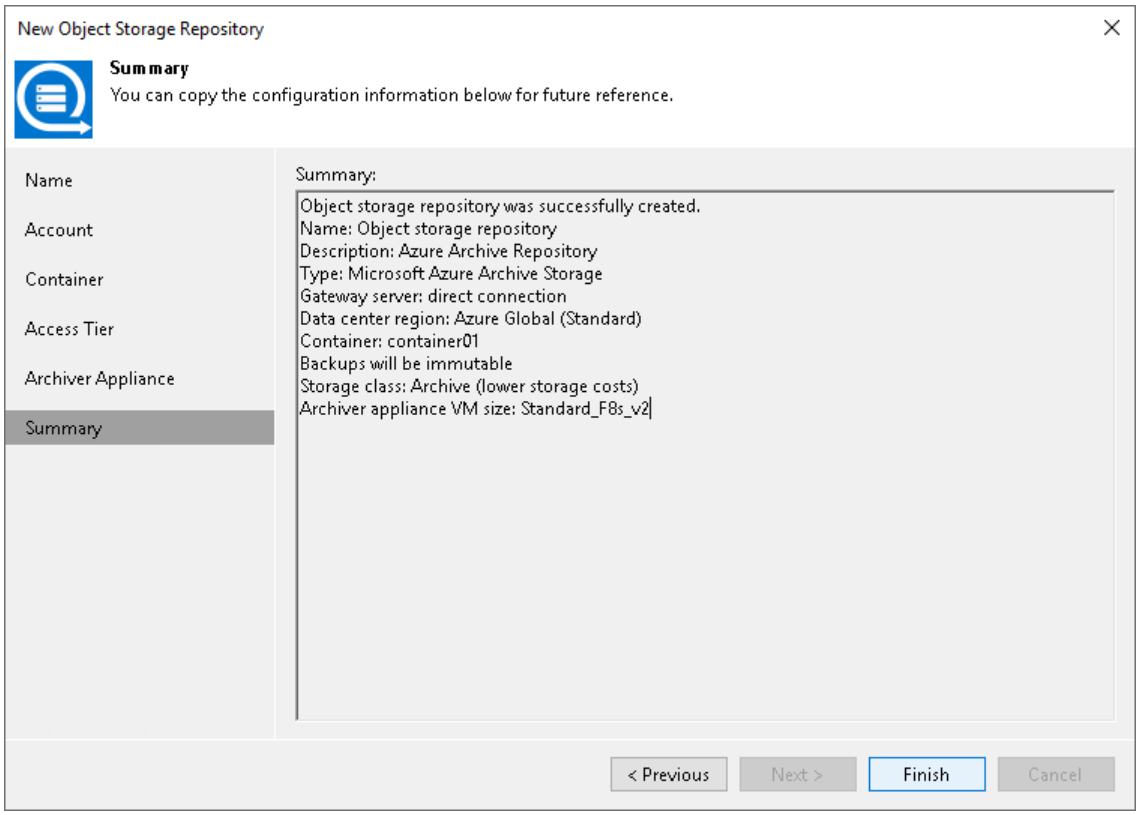


### Adding Azure Data Box Storage

Azure Data Box is a physical device which you can request for a short period of time from Microsoft. You can temporarily attach it to the backup infrastructure and use it as an object storage. For more information about ordering Azure Data Box and preparing to use it, see Microsoft Docs.

This device may become useful when you need to offload a significant number of backup files occupying storage space on your extents, as offloading data to the Azure Data Box device is much faster than transferring the same amount of data directly to Azure object storage. Once you have offloaded backups to Azure Data Box, you need to ship the device back to Microsoft for further data synchronization with your Azure storage account, as described in section Seeding Backups to Azure Blob Storage.

## General Considerations and Limitations

Consider the following:

- Veeam Backup & Replication supports only those Azure Data Box devices that are capable of reading and writing data using REST API.
- The *Azure Data Box disk type* is not supported.
- Direct data copy to the archive tier from Azure Data Box is not supported. When you place your order, do not enable **Copy to archive** option on the **Data destination** step.

- When you configure a scale-out backup repository with Azure Data Box as the capacity extent, it is recommended to only use the copy policy. This way you keep copy of the data on your local storage which helps you reduce the risk of data loss if the device is damaged during shipping. It will also ensure that the backup data is available for restore operations while Azure Data Box is in shipment.
- We do not recommend offloading immutable data to the Azure Data Box device. If you do so, it will become mutable on the Azure Data Box device and will remain mutable after seeding.
- You cannot seed data from the Azure Data Box device to Veeam Data Cloud Vault.
- You can add only one Azure Data Box device to a scale-out backup repository.
- Veeam Cloud Connect service providers can use Azure Data Box only as a capacity extent of a scale-out backup repository.
- You cannot back up data using Veeam Agent backup job or policy to Azure Data Box devices.

For information about other limitations for Microsoft Azure Data Box storage, see Microsoft Docs.

To add Microsoft Azure Blob storage, use the **New Object Storage Repository** wizard.

### Before You Begin

Before you add Microsoft Azure Data Box to the backup infrastructure, complete the following steps:

- Configure Name Resolutions
- Download and Install SSL Certificate

## Configuring Name Resolutions

To resolve a DNS name of your Azure Data Box device, make sure to add the following DNS addresses to the *HOST* file that is located on both the Veeam Backup & Replication server and the gateway server (if any additional gateway server is used):

- *<ip_address> <mydataboxno>.microsoftdatabox.com*
- *<ip_address> <storageaccountname>.blob.<mydataboxno>.microsoftdatabox.com*

For more information on how to learn the values for *<ip_address>*, *<mydataboxno>* and *<storageaccountname>*, see Microsoft Docs.

As an *<ip_address>*, use any of the IP addresses listed under the **Data N** section. You can also use the address specified under the **MGMT** section, but due to its slow connection rate (limited to 1 GbE), using such an address is not recommended.

The following is an example of the *HOST* file.

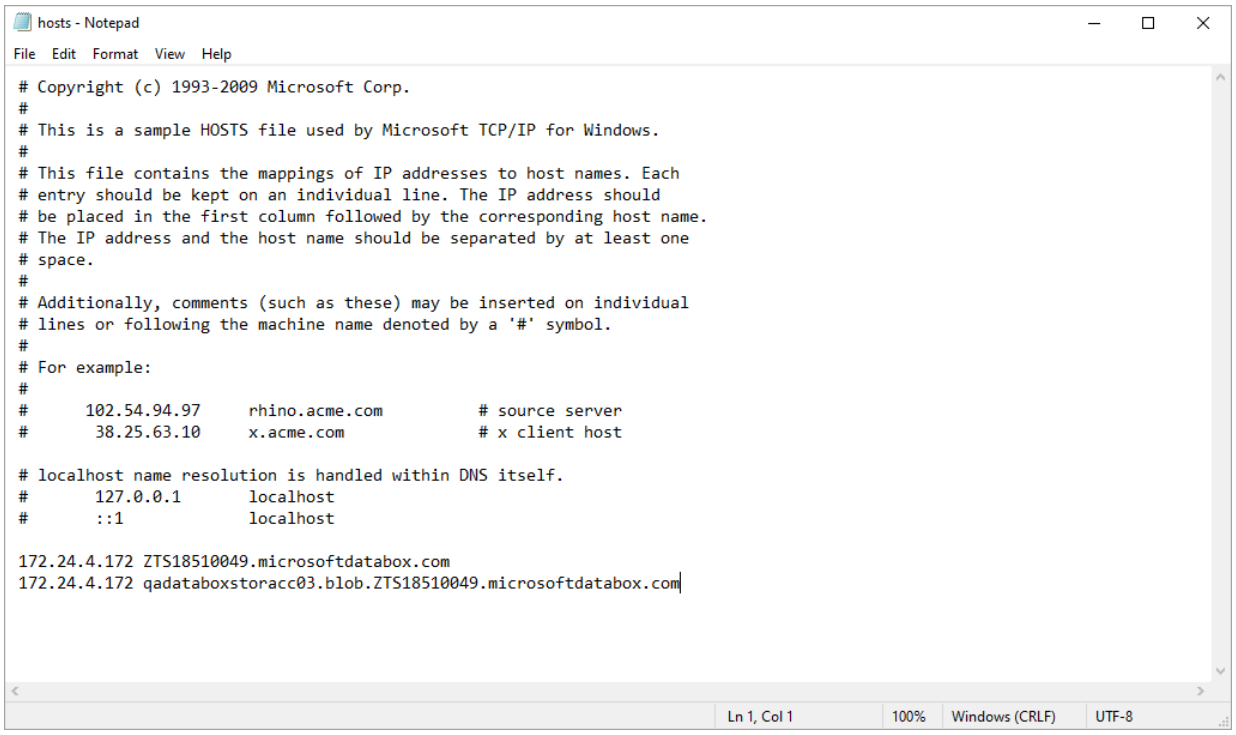


**NOTE**

Consider the following:

- Make sure to configure name resolutions on each server that may be used as a gateway.
- Alternatively, you can create a *microsoftdatabox.com* DNS zone with necessary records on your DNS server if you prefer not to modify the *HOST* file.

# Downloading and Installing SSL Certificate

To establish a secure connection to the Azure Data Box device, make sure to download and import an SSL certificate and install it on both the Veeam Backup & Replication server and the gateway server (if any additional gateway server is used).

For more information on how to download an SSL certificate, see Microsoft Docs.

When installing a certificate, do the following:

1. In the **Certificate Import Wizard** dialog box, select **Local Machine**.
2. In the **Certificate Store** step, select **Place all certificates in the following store** and click **Browse**.

3. Select **Trusted Root Certification Authorities**.

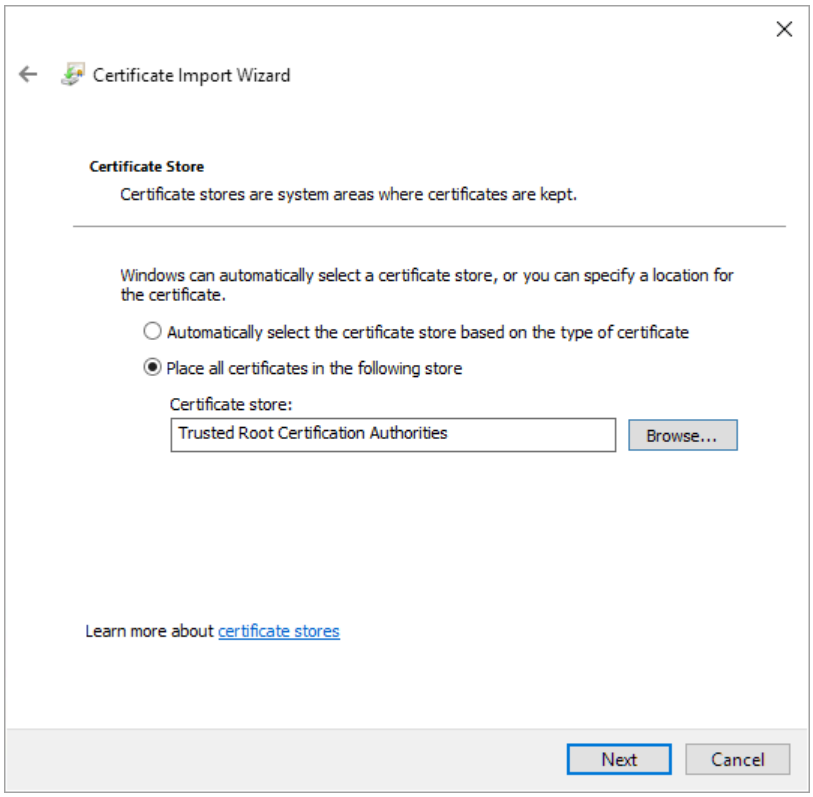


# Configuring registry settings

For information how to configure registry settings for Microsoft Azure Data Box device, see this Veeam KB article.

# Sizing Gateway Server

Consider that Veeam Backup & Replication supports Azure Data Box devices that are capable of reading/writing data using REST API only; the Azure Data Box Disk type is not supported.

As with any other object storage, REST API performance depends on scale. As Azure Data Box is a single endpoint, the individual throughput of this REST API may be limited. The block size used in Veeam Backup & Replication capacity tier for object storage offload matches that of the source job. The default object size will be a compressed 1 MB block, resulting in objects of around 512 KB in size.

The speed of data offload to Azure Data Box devices may reach about 300 MB/s. To achieve this speed, we recommend using a separate gateway server with 8 CPU cores.

## Step 1. Launch New Object Storage Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object Object storage** > **Hyperscalers** > **Microsoft Azure Storage** > **Azure Data Box**.

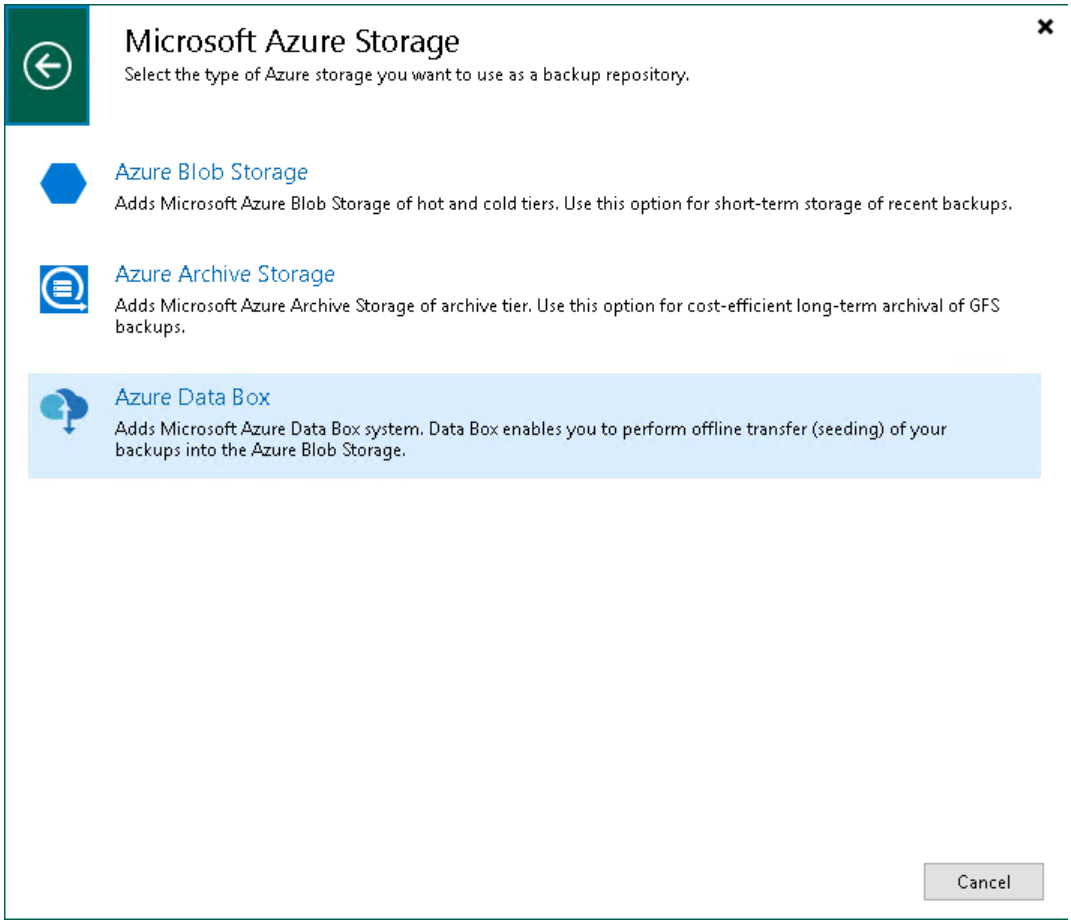

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

If you want to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

> **TIP**
>
> Set the maximum number of concurrent task to a reasonable number to avoid overloading if you plan to upload significant amount of backup chains to the device.

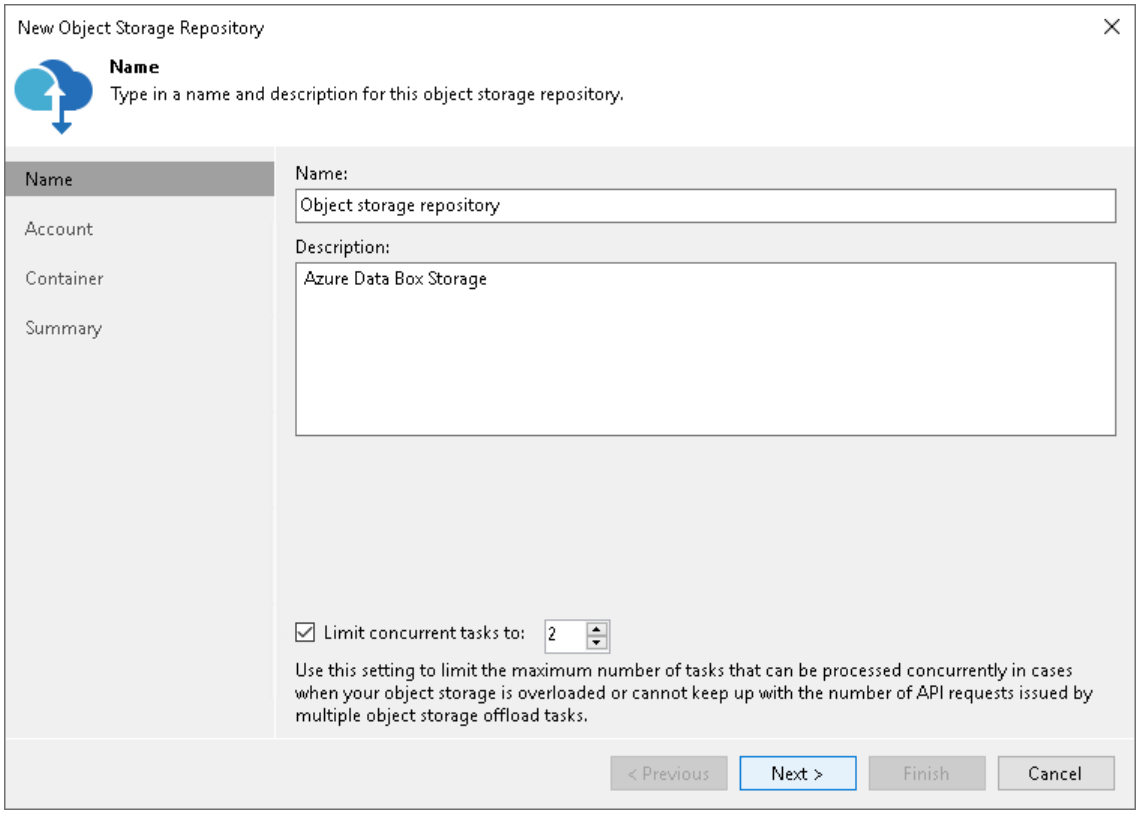

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. In the **Service endpoint** field, specify a service endpoint address of your Azure Data Box device.
2. From the **Credentials** drop-down list, select user credentials to access your Azure Data Box storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

   For more information on where to find connection parameters of your Azure Data Box device, see Getting Data Box Connection Parameters.
3. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:
   - **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.
   - **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

     By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

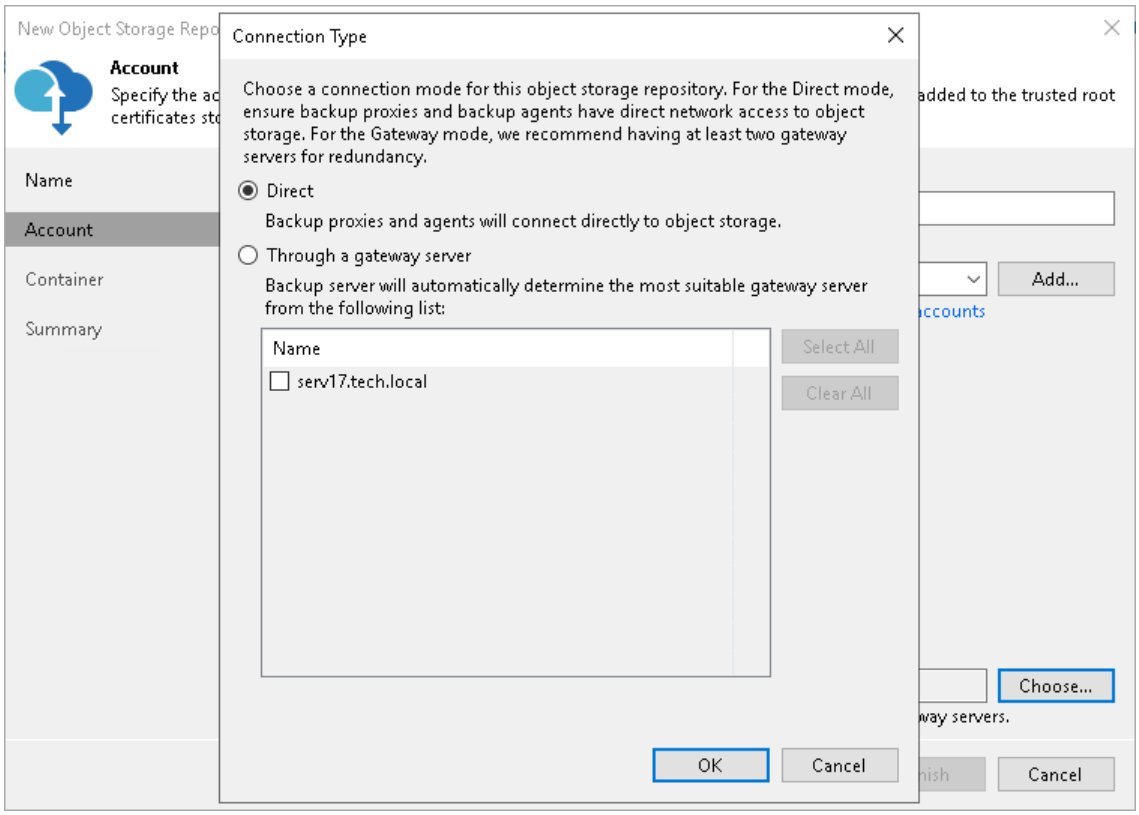

# Getting Data Box Connection Parameters

To find connection parameters of your Azure Data Box device, do the following:

1. Open the Microsoft Azure Data Box portal.
2. In the navigation pane, click **Connect and Copy**.
3. Under the **Access Settings** column of the storage account that you want to use, click **REST** and in the **Access storage account and upload data** dialog box, copy the following:
   a. Under **Storage Account Name**, copy the Azure storage account name.
   b. Under **Key**, copy the storage account key.
   c. Under **Blob Service Endpoint**, copy the service endpoint address that starts exactly after the *blob* word.

      For example, if the complete service endpoint address is *https://qadataboxstoracc03.blob.ZTS18510049.microsoftdatabox.com*, then you will need to copy everything that starts from *ZTS* only. That is, *ZTS18510049.microsoftdatabox.com*. Make sure not to copy the last slash ("/") symbol.

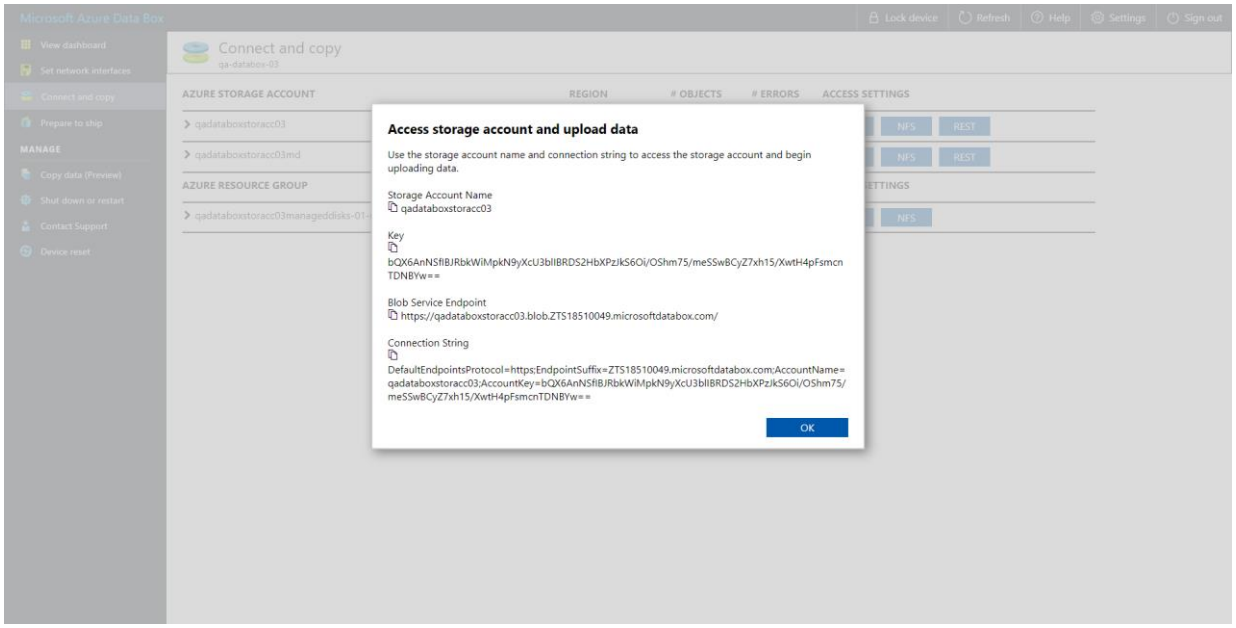

## Step 4. Specify Object Storage Settings

At the **Container** step of the wizard, specify the container and folder that will be used to store data:

1. From the **Container** drop-down list, select a container.

   To create a container, use Microsoft Azure Storage Explorer. For more information on how to connect to the Azure Data Box device using Microsoft Azure Storage Explorer, see Microsoft Docs.

2. In the **Select Folder** field, select a cloud folder to which you want to map your object storage repository.

   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

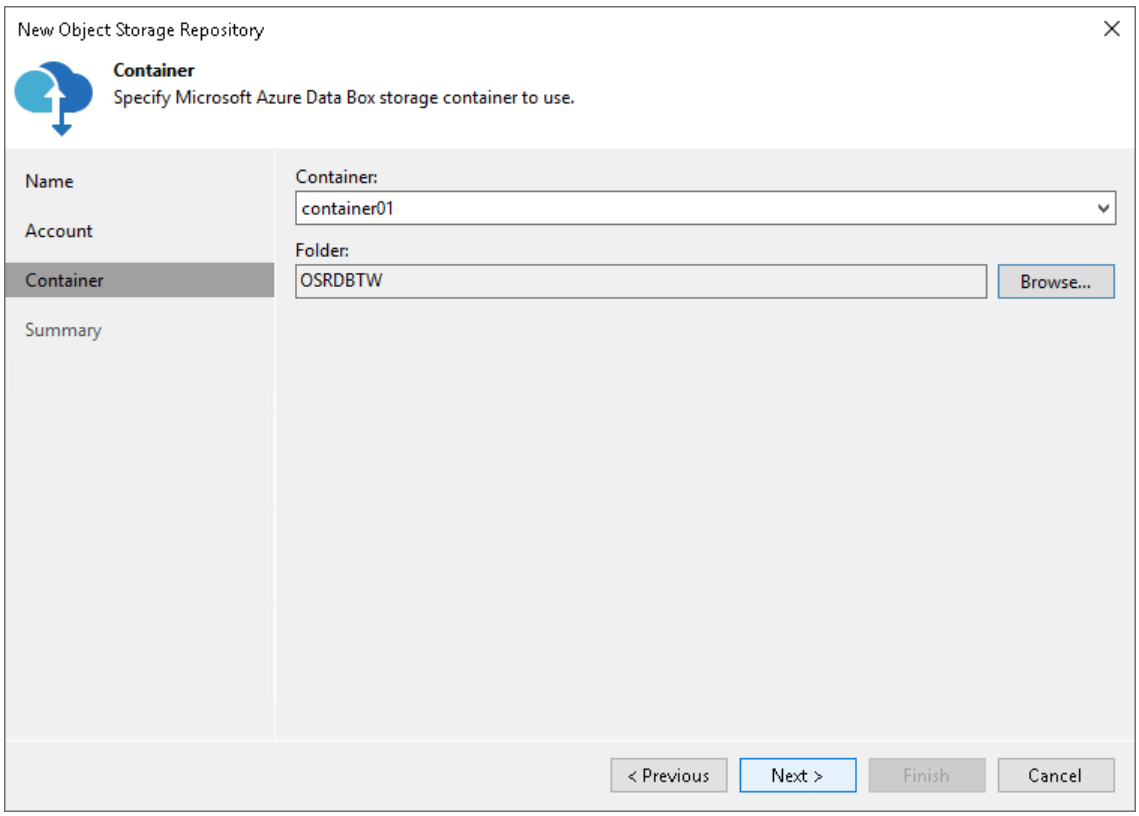

### Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

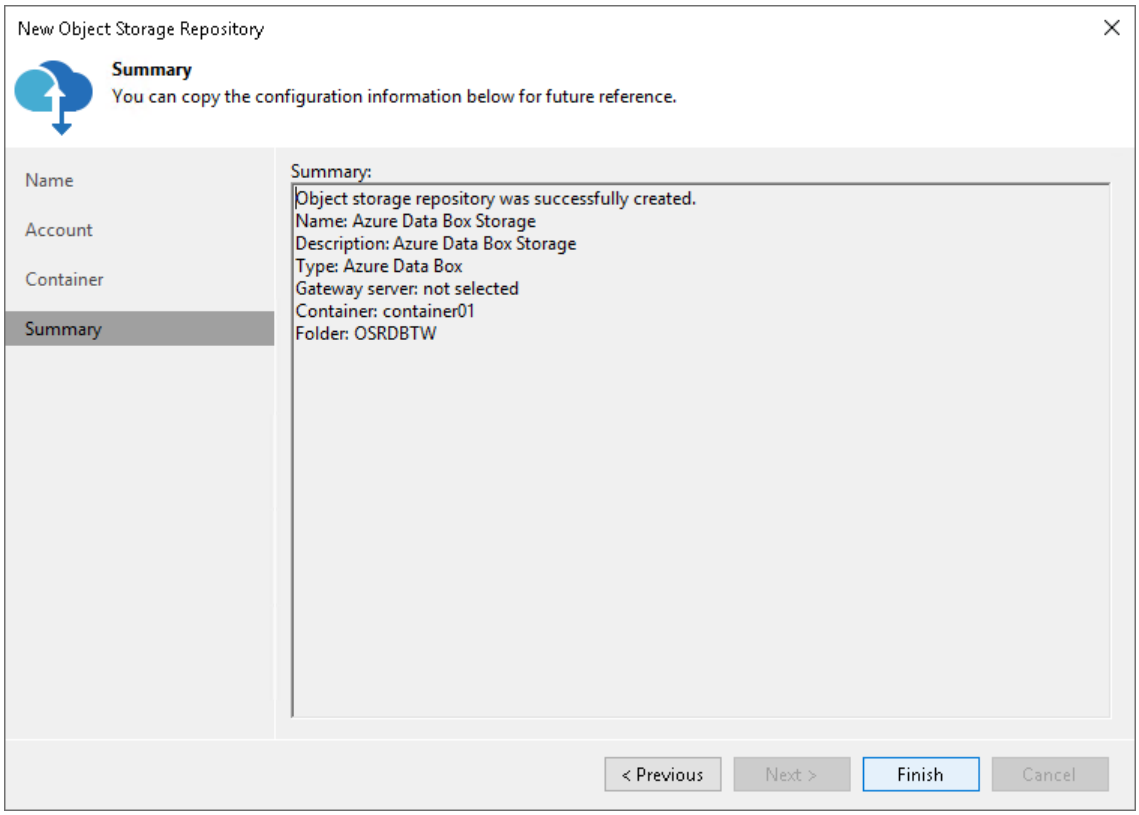


## Adding 11:11 Cloud Object Storage

For more information about 11:11 Cloud Object Storage, see this 11:11 Systems article.

Before you add an 11:11 Cloud Object Storage to the backup infrastructure, check prerequisites. After that, use the **New Object Storage Repository wizard**.

## Step 1. Launch New Object Storage Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **Additional Providers** > **11:11 Cloud Object Storage**.

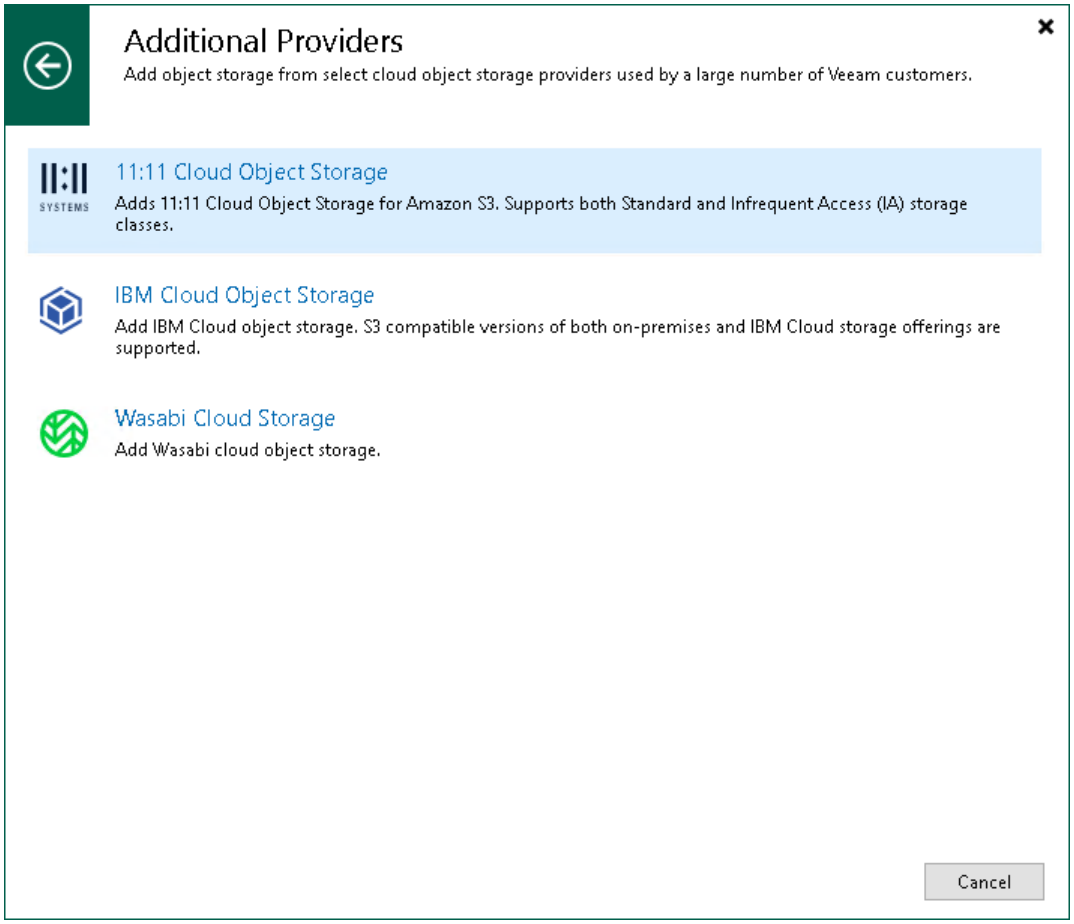

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

If you want to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

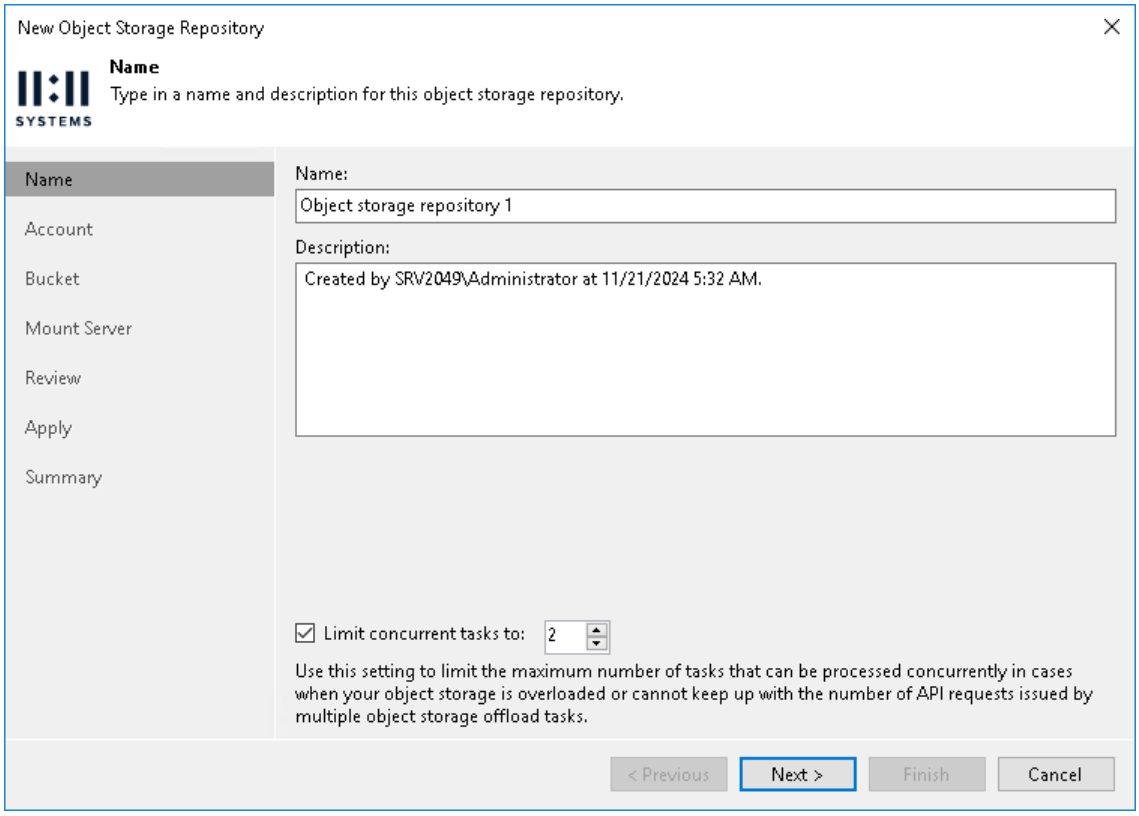

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. In the **Service point** field, specify a service point address of your 11:11 Cloud Object Storage.
2. In the **Region** field, specify a region.
3. From the **Credentials** drop-down list, select user credentials to access your 11:11 Cloud Object Storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.
4. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:
   - **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.
   - **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

     By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

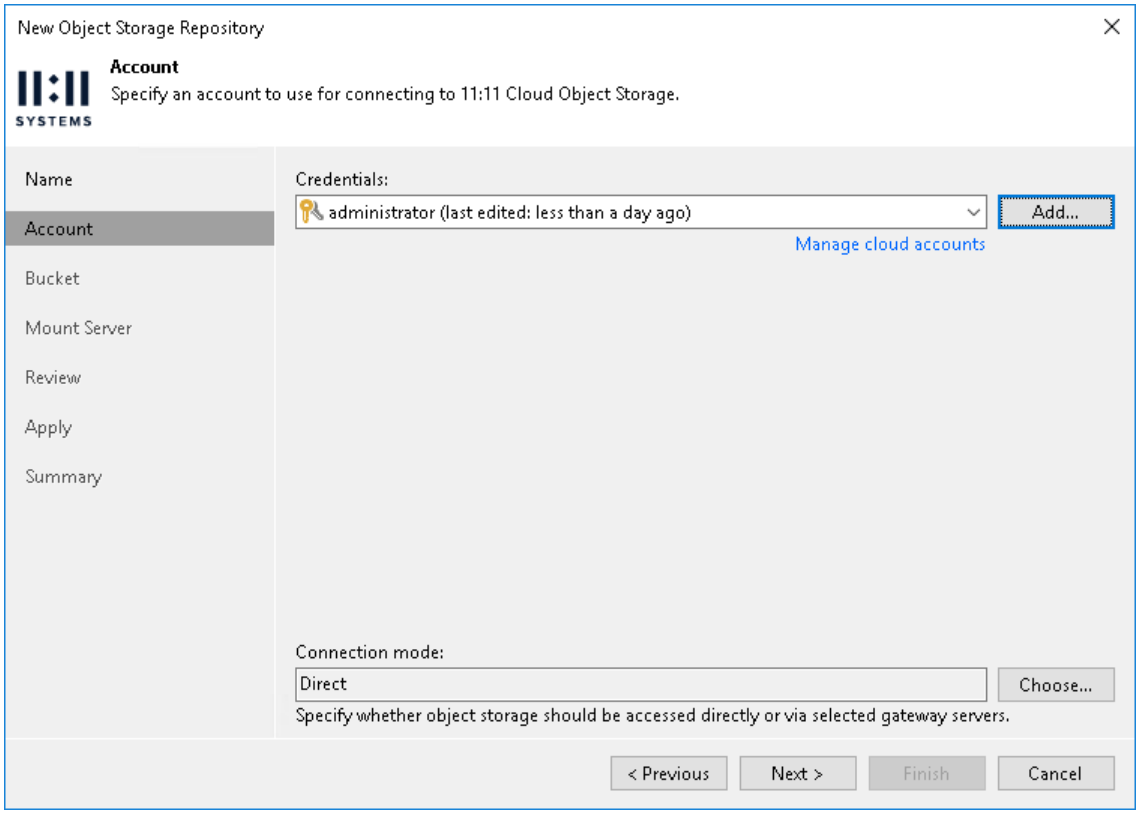

## Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify the bucket and folder that will be used to store data:

1. In the **Bucket** field, enter a name of the bucket or click **Browse** to get the necessary bucket.

   Note that you must create the bucket where you want to store your backup data beforehand.

2. In the **Folder** field, enter a cloud folder name to which you want to map your object storage repository. Alternatively, click **Browse** and either select an existing folder or click **New Folder**.

3. Select the **Limit object storage consumption** to check box to define a soft limit for your object storage consumption. If this limit is exceeded during a job run, Veeam Backup & Replication will complete the job. However, a new job will not be able to start unless you remove the extra data that exceeds the limit or change the soft limit settings. Provide the value in TB or PB.

4. Select the **Make recent backups immutable for** check box to prohibit deletion of blocks of data from object storage. Specify the immutability period. For more information, see Immutability for Object Storage Repositories.

5. If you plan to access your backup data in an infrequent manner, select the **Use infrequent access storage class** check box to mark each block as *S3 Standard-IA (S3 Standard Infrequent Access)*. If you so not check this box, the blocks will be marked as *S3 Standard*.

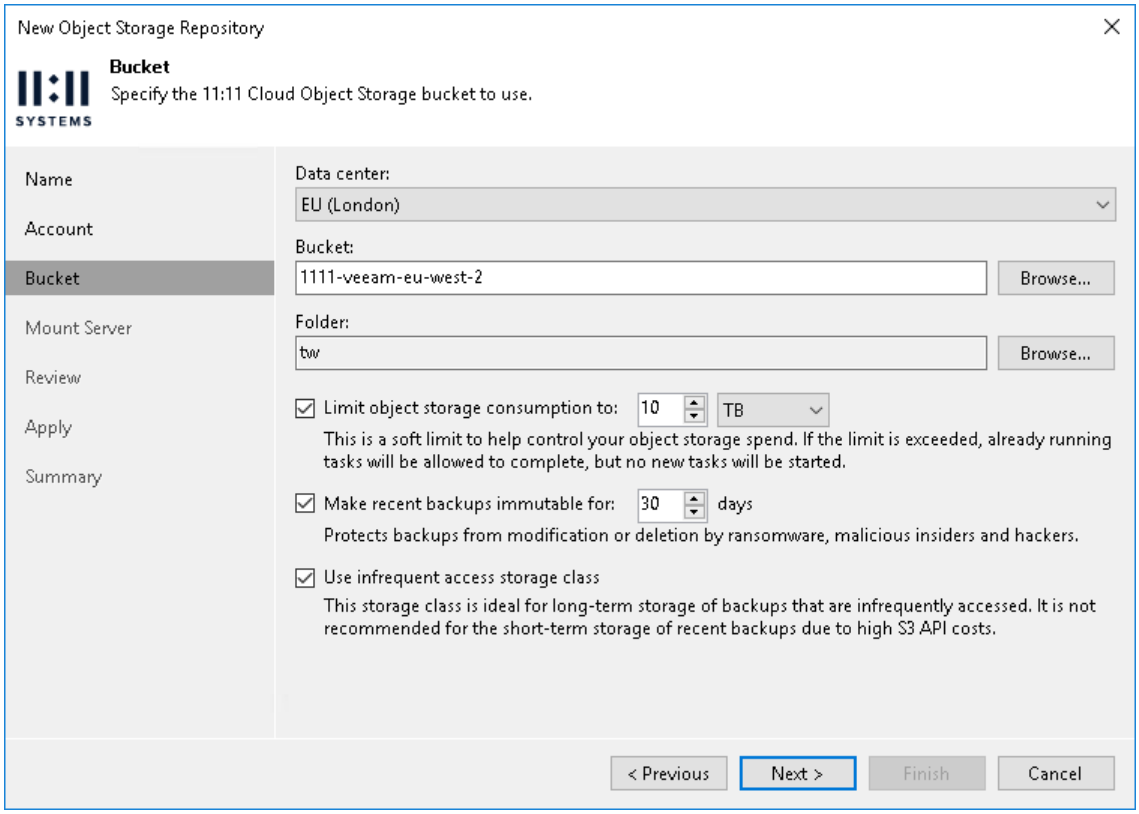

### Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for restore operations, and configure a helper appliance. The helper appliance is a Windows-based or Linux-based virtual or physical server, added to the backup infrastructure, that Veeam Backup & Replication uses to perform a health check of backup files and apply retention to unstructured data backup files. For more information, see Health Check for Object Storage Repositories and Helper Appliance in Unstructured Data Backup.

> **IMPORTANT**
>
> Consider the following:
>
> - If you do not configure a helper appliance, Veeam Backup & Replication will use local resources to perform the health check and apply retention to unstructured data backup files. It will consume more cloud resources and can result in additional costs.
> - To perform the health check, you must enable this option when you configure a job. For more information, see Health Check for Backup Files.

# Specifying Mount Server Settings

To specify the mount server settings, do the following:

1. From the **Mount Server** drop-down list, select a server that you want to use as a mount server. Veeam Backup & Replication uses this server during restore operations to mount VM disks directly from objects located in object storage repositories. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, click **Add New** on the right to open the **New Windows Server wizard**. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder to keep cache that is created during mount operations.

3. Select the **Enable vPower NFS service on the mount server** check box to allow the Veeam vPower NFS Service access the object storage repository. Veeam Backup & Replication will enable the Veeam vPower NFS Service on the necessary mount server. For more information, see Veeam vPower NFS Service.

4. Click **Ports** to customize network ports used by the Veeam vPower NFS Service. In the **vPower NFS Port Settings** window, specify the following settings:

   - Next to the **Mount Port** section, specify the port that the Veeam vPower NFS Service will use to mount the vPower NFS datastore to the ESXi host.
   - Next to the **vPower NFS** port section, specify the port that the Veeam vPower NFS Service will use to connect to the target NFS share.

   For information on ports used by default, see Ports.

5. To specify the helper appliance settings, click **Configure**. From the **Managed server** drop-down list, select a server that you want to use as the helper appliance.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

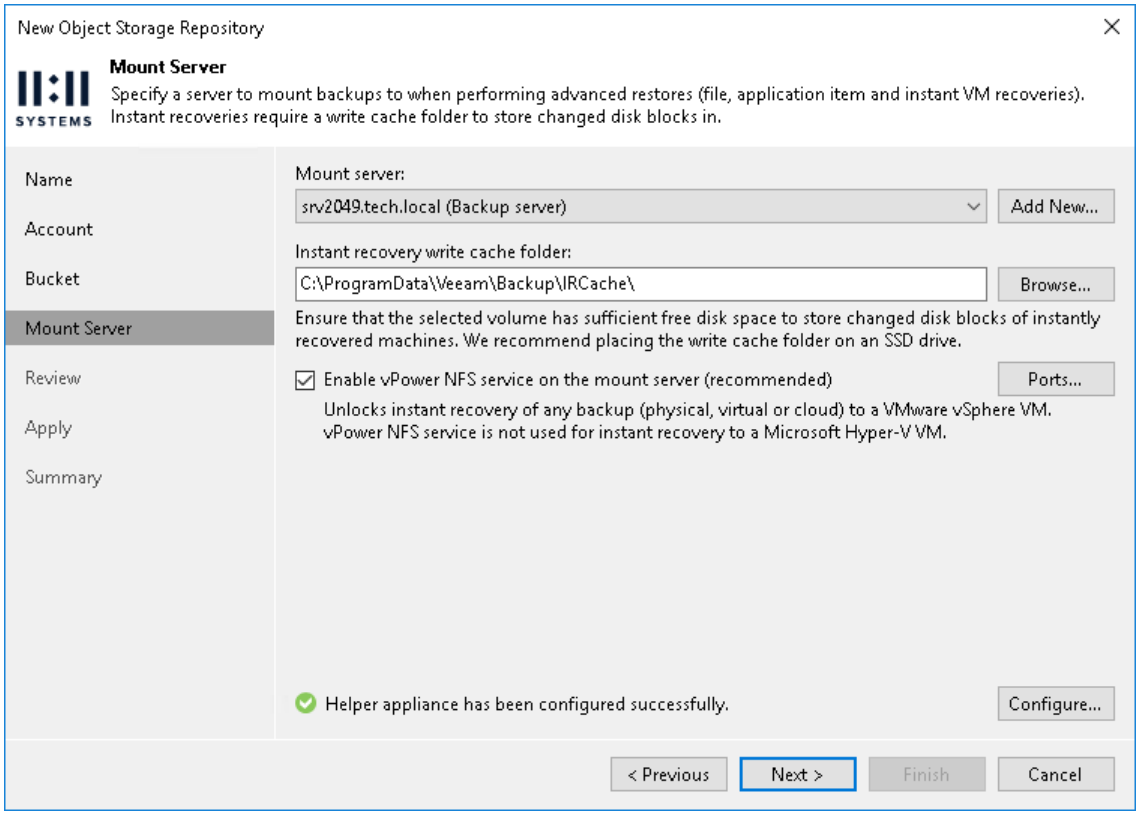
New Object Storage Repository
Mount Server
Specify a server to mount backups to when performing advanced restores (file, application item and instant VM recoveries).
Instant recoveries require a write cache folder to store changed disk blocks in.
Name
Account
Bucket
Mount Server
Review
Apply
Summary
Mount server:
srv2049.tech.local (Backup server)
Add New...
Instant recovery write cache folder:
C:\ProgramData\Veeam\Backup\IRCache\
Browse...
Ensure that the selected volume has sufficient free disk space to store changed disk blocks of instantly recovered machines. We recommend placing the write cache folder on an SSD drive.
Enable vPower NFS service on the mount server (recommended)
Ports...
Unlocks instant recovery of any backup (physical, virtual or cloud) to a VMware vSphere VM.
vPower NFS service is not used for instant recovery to a Microsoft Hyper-V VM.
Helper appliance has been configured successfully.
Configure...
< Previous
Next >
Finish
Cancel

## Step 6. Review Components

At the **Review** step of the wizard, review what components will be processed on the mount server server and their status.

If the backup repository contains backups, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Object Storage (Imported)** node.

If the backup repository contains guest file system index files, select the **Import guest file system index data to the catalog** check box. Veeam Backup & Replication will import index files together with backup files, and you will be able to search for guest OS files inside imported backups. For more information, see the Guest OS File Restore section of the Veeam Backup Enterprise Manager Guide.

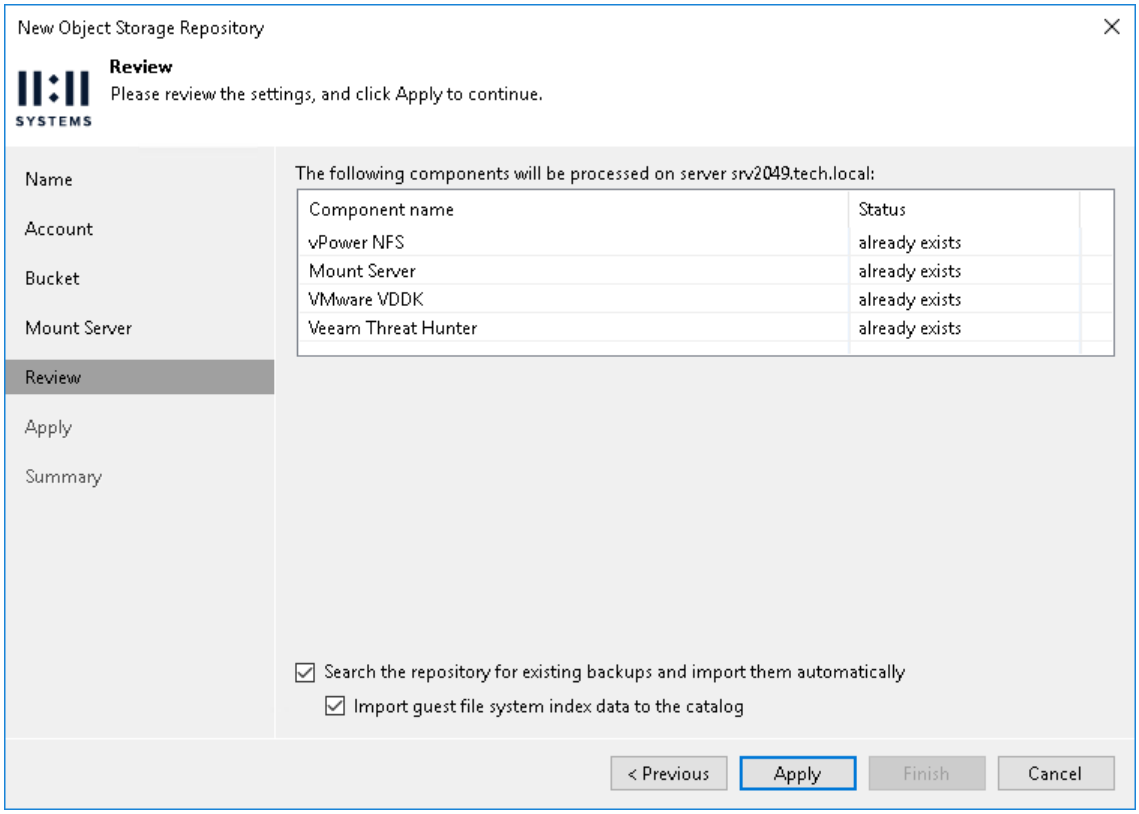

## Step 7. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to complete saving your settings to the configuration database and create backup infrastructure objects.

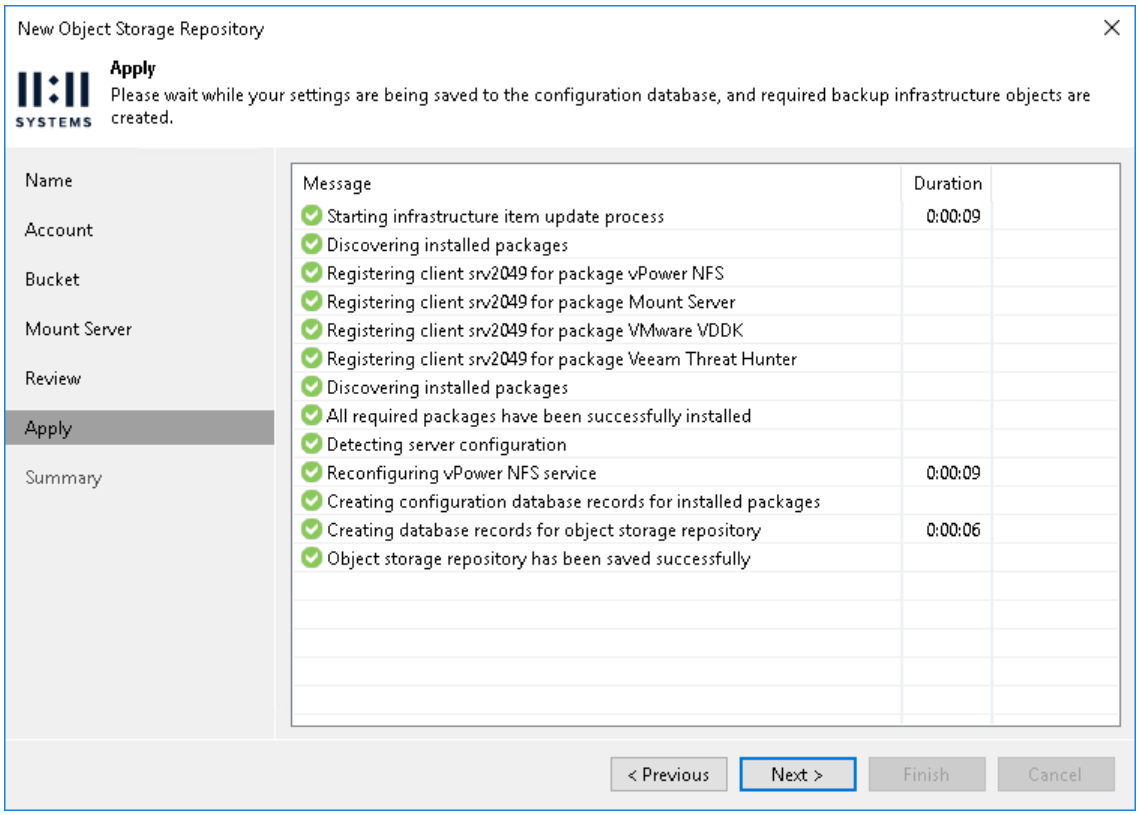

### Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

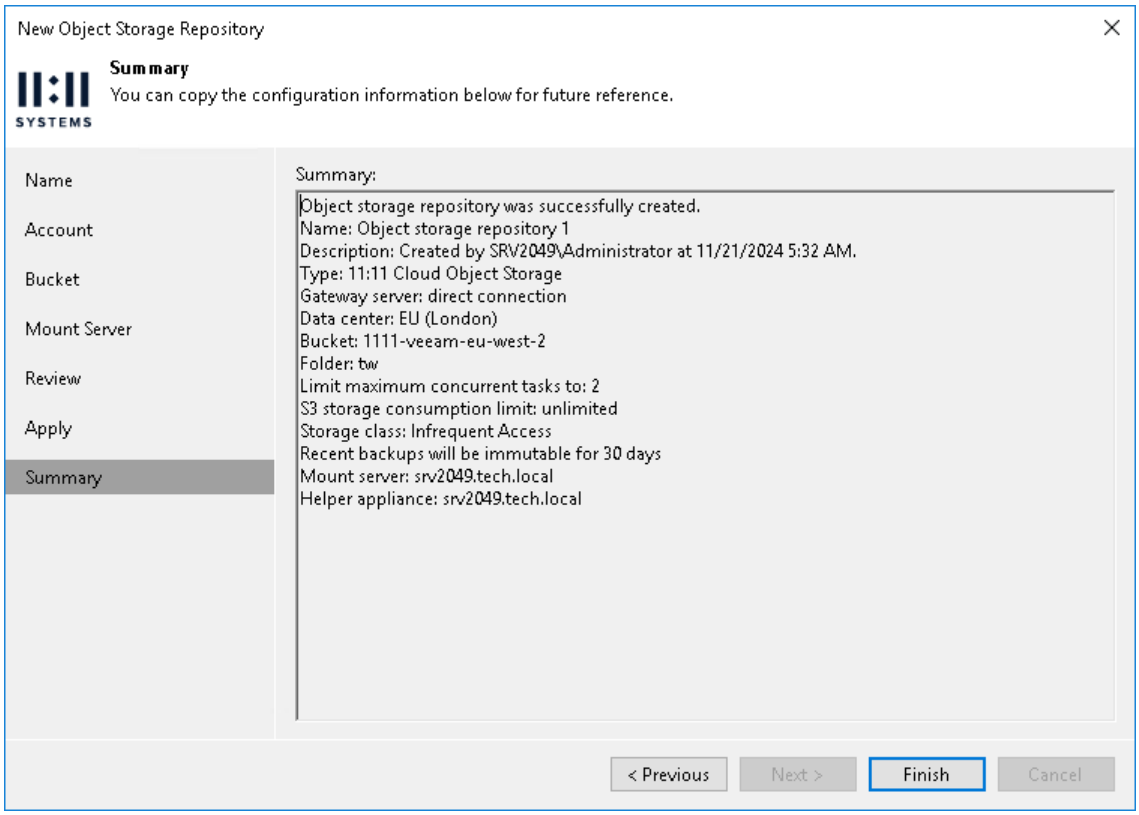


## Adding IBM Cloud Object Storage

For more information about IBM Cloud object storage, see this IBM article.

Before you add an IBM Cloud object storage to the backup infrastructure, check prerequisites. After that, use the **New Object Storage Repository wizard**.

## Step 1. Launch New Object Storage Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **Additional Providers** > **IBM Cloud Object Storage**.

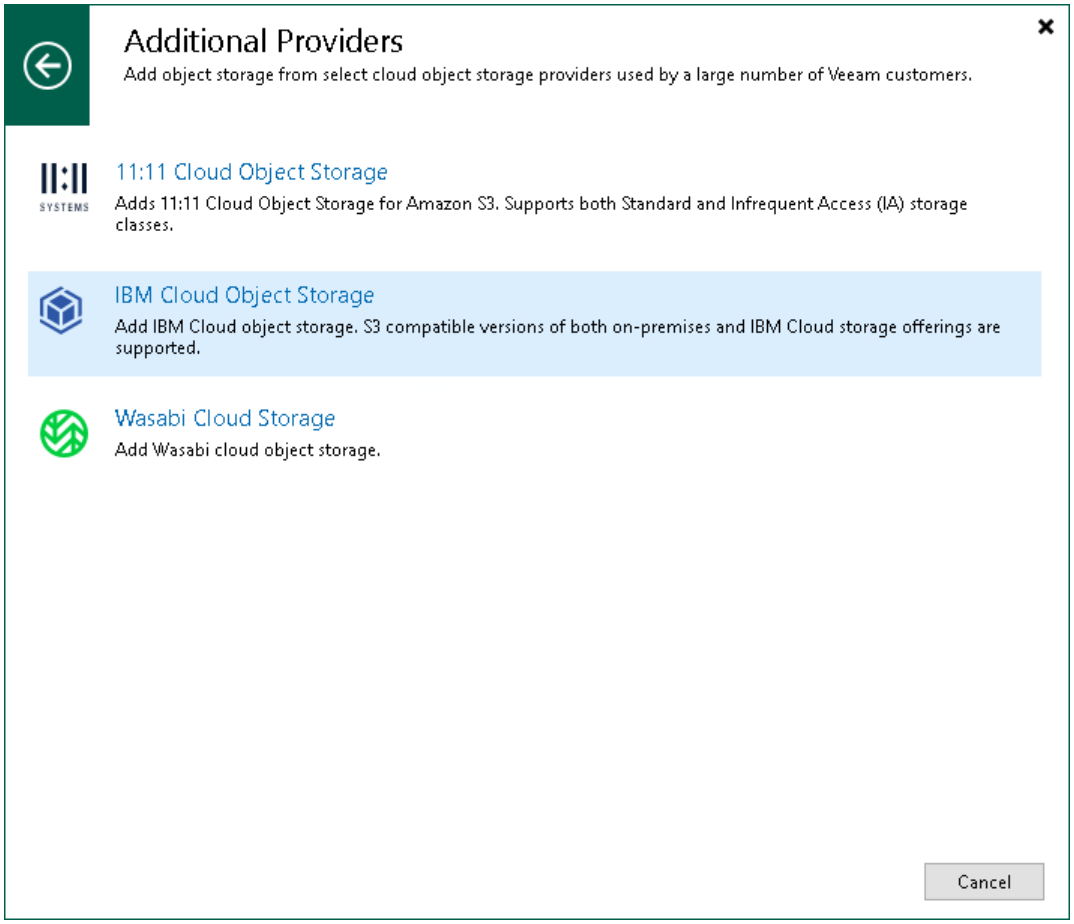

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

If you want to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

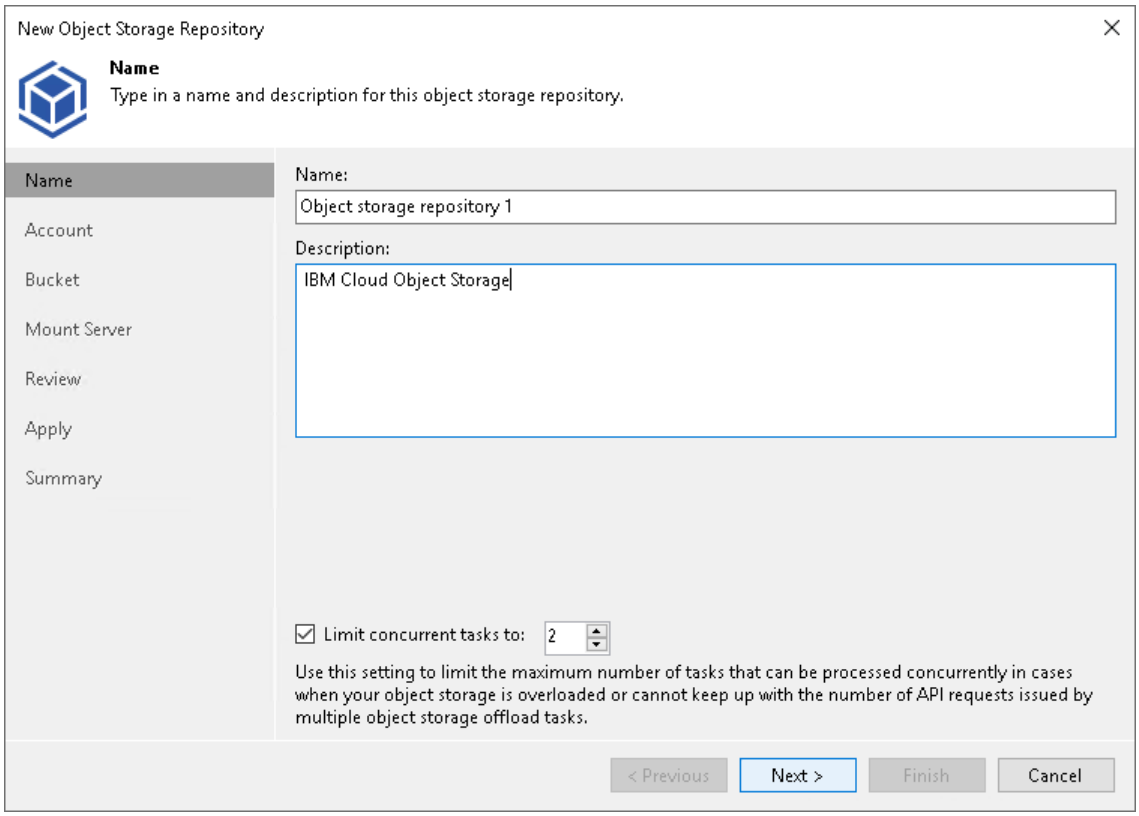

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. In the **Service point** field, specify a service point address of your IBM cloud object storage.
2. In the **Region** field, specify a region.
3. From the **Credentials** drop-down list, select user credentials to access your IBM cloud object storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.
4. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:
   - **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.
   - **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

     By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

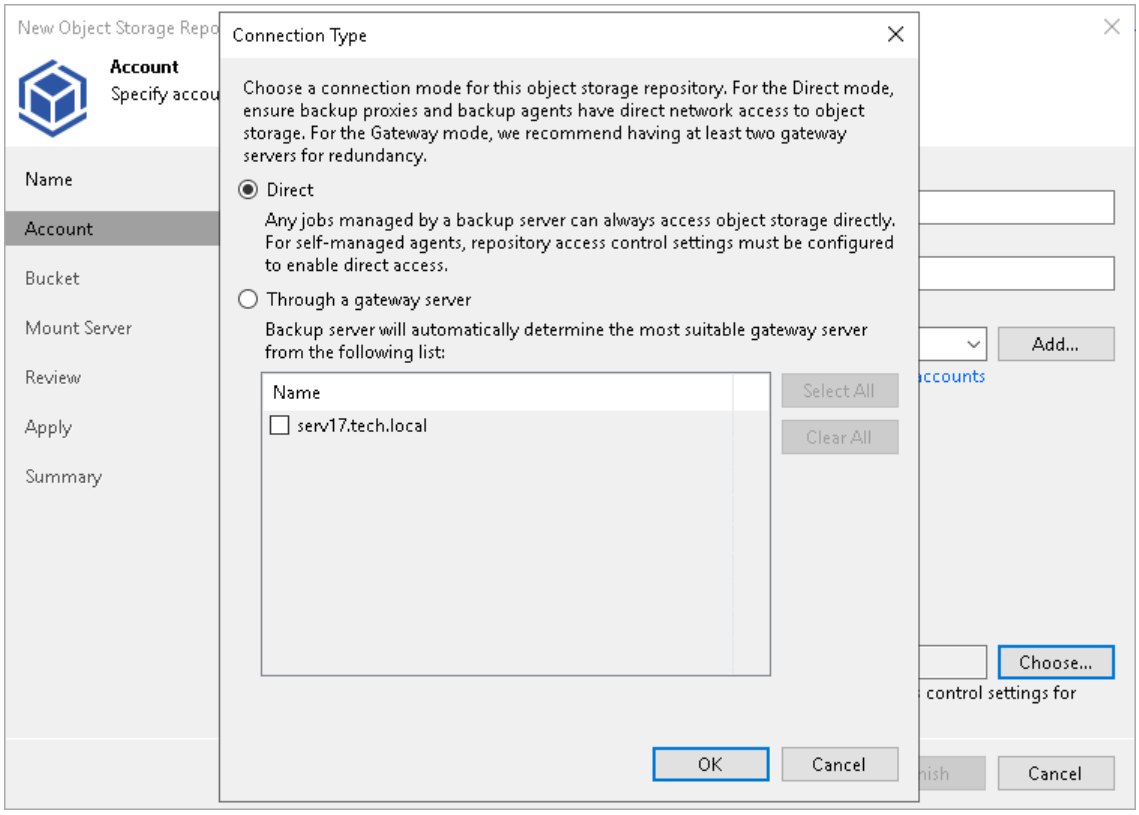

## Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify the bucket and folder that will be used to store data:

1. In the **Bucket** field, enter a name of the bucket or click **Browse** to get the necessary bucket.

   Note that you must create the bucket where you want to store your backup data beforehand.

2. In the **Folder** field, enter a cloud folder name to which you want to map your object storage repository. Alternatively, click **Browse** and either select an existing folder or click **New Folder**.

3. Select the **Limit object storage consumption** to check box to define a soft limit for your object storage consumption. If this limit is exceeded during a job run, Veeam Backup & Replication will complete the job. However, a new job will not be able to start unless you remove the extra data that exceeds the limit or change the soft limit settings. Provide the value in TB or PB.

4. Select the **Make recent backups immutable for** check box to prohibit deletion of blocks of data from object storage. Specify the immutability period. For more information, see Immutability for Object Storage Repositories.

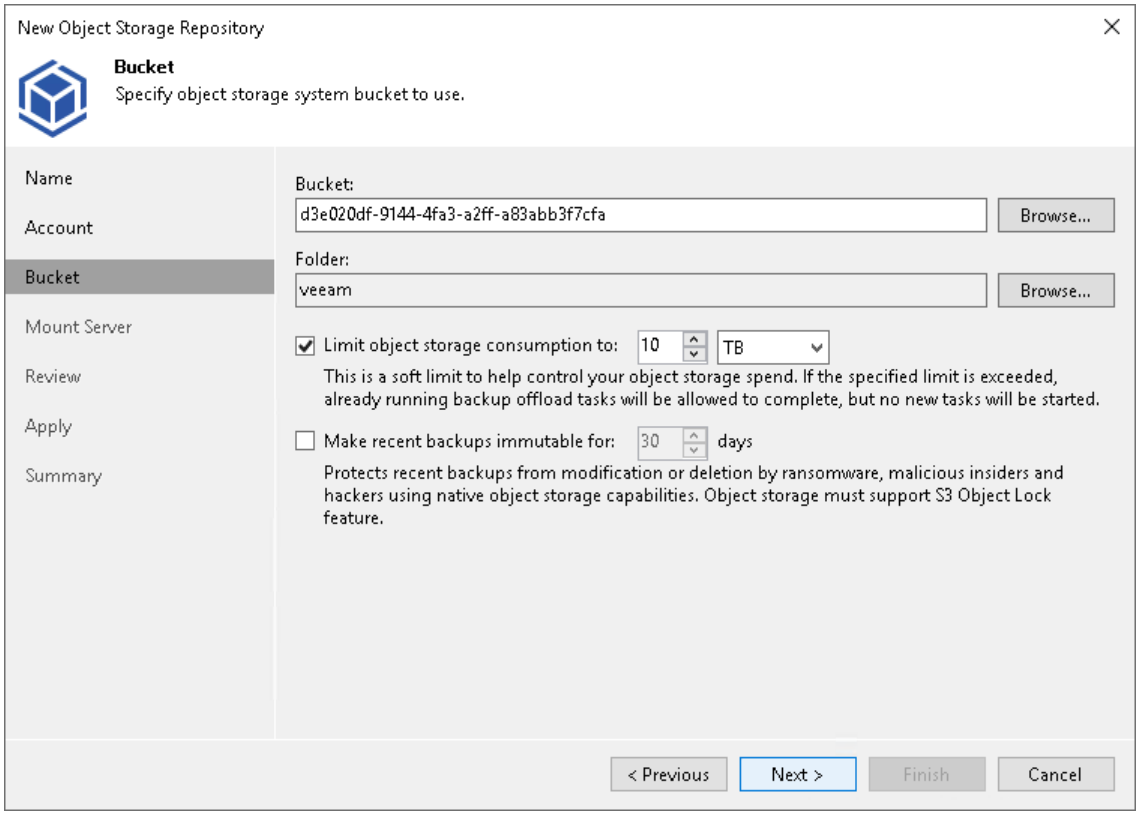

### Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for restore operations, and configure a helper appliance. The helper appliance is a Windows-based or Linux-based virtual or physical server, added to the backup infrastructure, that Veeam Backup & Replication uses to perform a health check of backup files and apply retention to unstructured data backup files. For more information, see Health Check for Object Storage Repositories and Helper Appliance in Unstructured Data Backup.

> **IMPORTANT**
>
> Consider the following:
>
> - If you do not configure a helper appliance, Veeam Backup & Replication will use local resources to perform the health check and apply retention to unstructured data backup files. It will consume more cloud resources and can result in additional costs.
> - To perform the health check, you must enable this option when you configure a job. For more information, see Health Check for Backup Files.

## Specifying Mount Server Settings

To specify the mount server settings, do the following:

1. From the **Mount Server** drop-down list, select a server that you want to use as a mount server. Veeam Backup & Replication uses this server during restore operations to mount VM disks directly from objects located in object storage repositories. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, click **Add New** on the right to open the **New Windows Server wizard**. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder to keep cache that is created during mount operations.

3. Select the **Enable vPower NFS service on the mount server** check box to allow the Veeam vPower NFS Service access the object storage repository. Veeam Backup & Replication will enable the Veeam vPower NFS Service on the necessary mount server. For more information, see Veeam vPower NFS Service.

4. Click **Ports** to customize network ports used by the Veeam vPower NFS Service. In the **vPower NFS Port Settings** window, specify the following settings:

   - Next to the **Mount Port** section, specify the port that the Veeam vPower NFS Service will use to mount the vPower NFS datastore to the ESXi host.
   - Next to the **vPower NFS** port section, specify the port that the Veeam vPower NFS Service will use to connect to the target NFS share.

   For information on ports used by default, see Ports.

5. To specify the helper appliance settings, click **Configure**. From the **Managed server** drop-down list, select a server that you want to use as the helper appliance.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

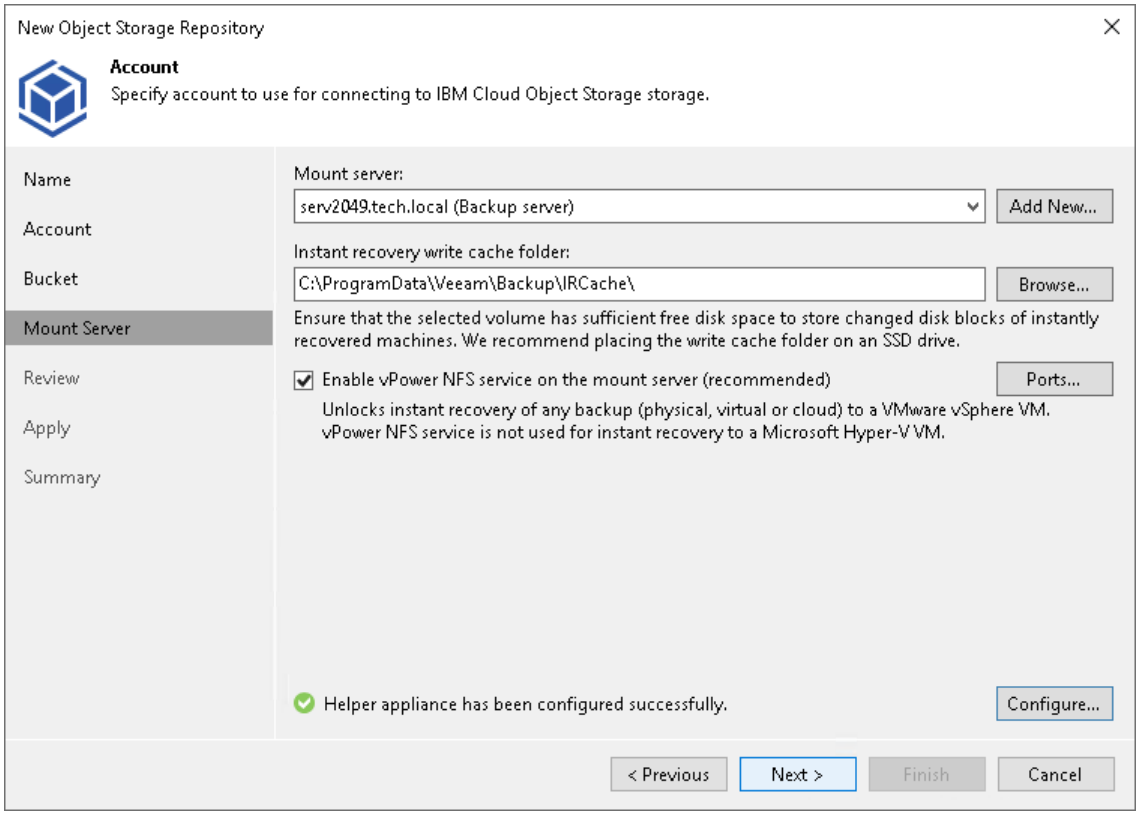
New Object Storage Repository
Account
Specify account to use for connecting to IBM Cloud Object Storage storage.
Name
Account
Bucket
Mount Server
Review
Apply
Summary
Mount server:
serv2049.tech.local (Backup server)
Add New...
Instant recovery write cache folder:
C:\ProgramData\Veeam\Backup\IRCache\
Browse...
Ensure that the selected volume has sufficient free disk space to store changed disk blocks of instantly recovered machines. We recommend placing the write cache folder on an SSD drive.
Enable vPower NFS service on the mount server (recommended)
Ports...
Unlocks instant recovery of any backup (physical, virtual or cloud) to a VMware vSphere VM.
vPower NFS service is not used for instant recovery to a Microsoft Hyper-V VM.
Helper appliance has been configured successfully.
Configure...
< Previous
Next >
Finish
Cancel

## Step 6. Review Components

At the **Review** step of the wizard, review what components will be processed on the mount server server and their status.

If the backup repository contains backups, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Object Storage (Imported)** node.

If the backup repository contains guest file system index files, select the **Import guest file system index data to the catalog** check box. Veeam Backup & Replication will import index files together with backup files, and you will be able to search for guest OS files inside imported backups. For more information, see the Guest OS File Restore section of the Veeam Backup Enterprise Manager Guide.

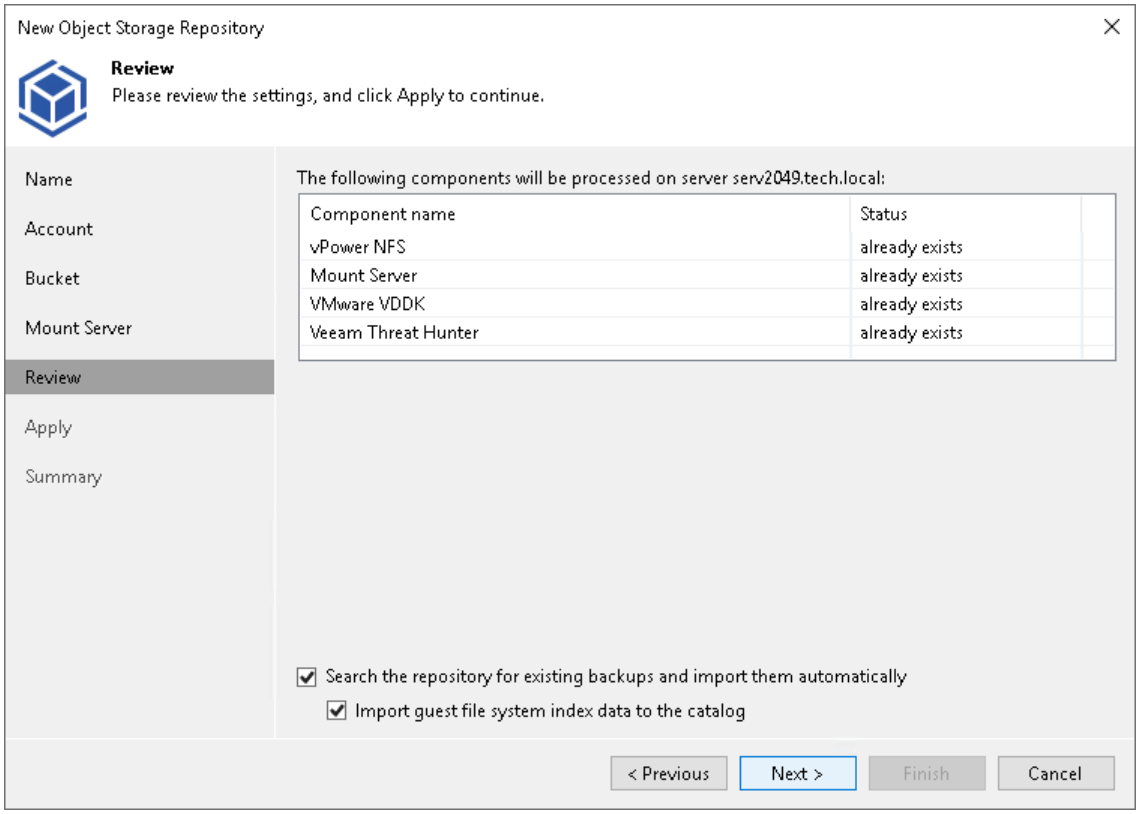

## Step 7. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to complete saving your settings to the configuration database and create backup infrastructure objects.

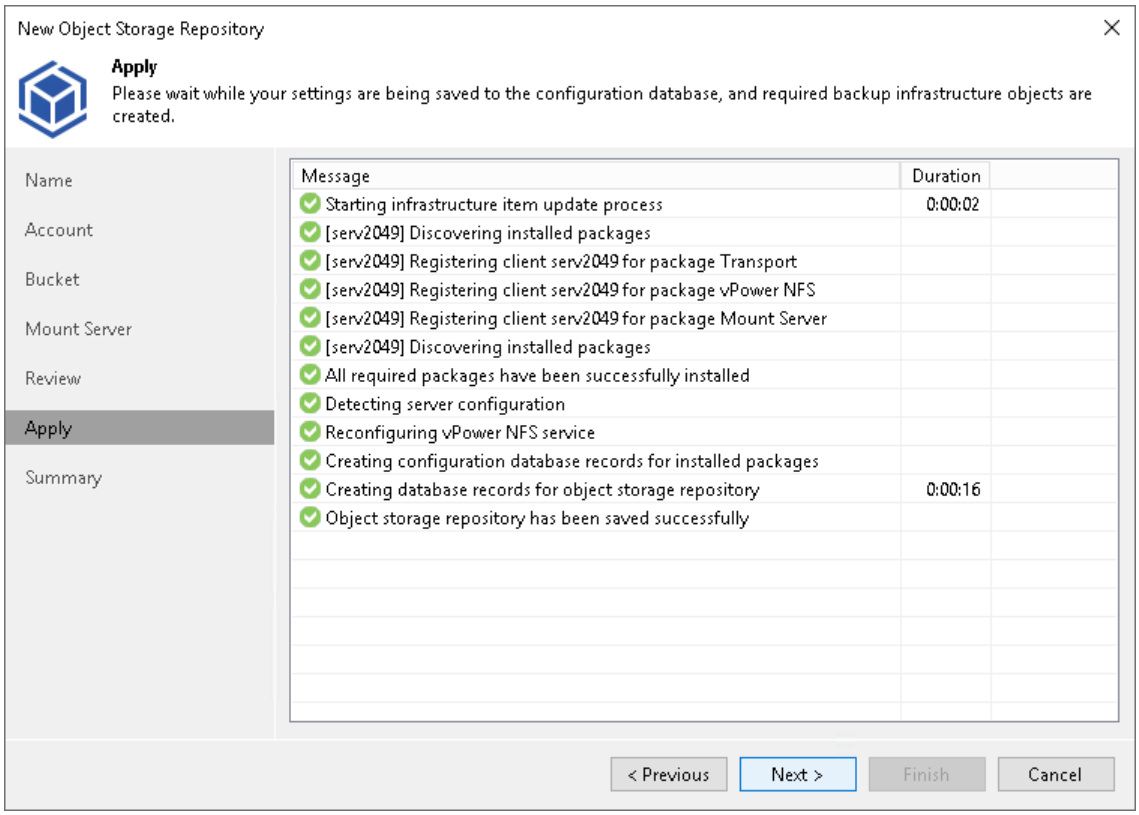

### Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

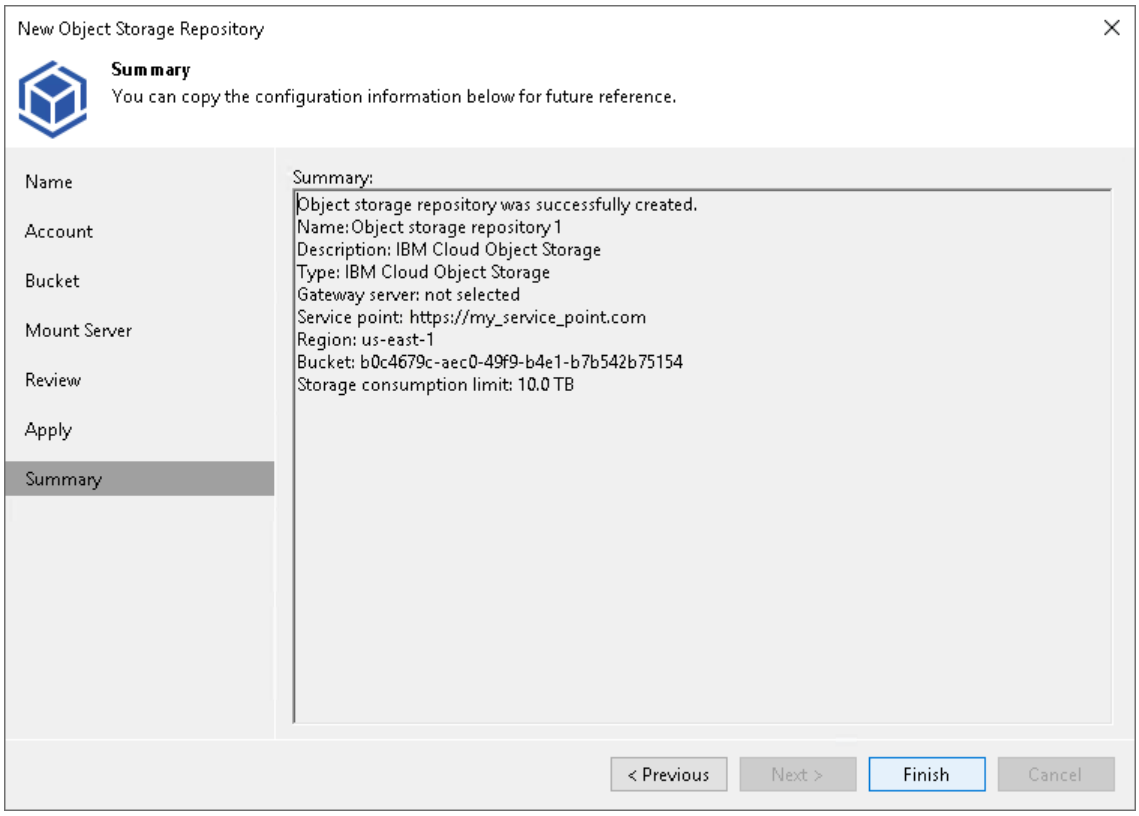


## Adding Wasabi Cloud Object Storage

This section describes how to add Wasabi Cloud object storage to the backup infrastructure. For more information about Wasabi Cloud object storage, see Wasabi documentation.

Before you add a Wasabi Cloud object storage to the backup infrastructure, check prerequisites. After that, use the **New Object Repository wizard**.

## Step 1. Launch New Object Storage Repository Wizard

To launch the **New Object Storage Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add backup repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select **Object storage** > **Additional Providers** > **Wasabi Cloud Storage**.

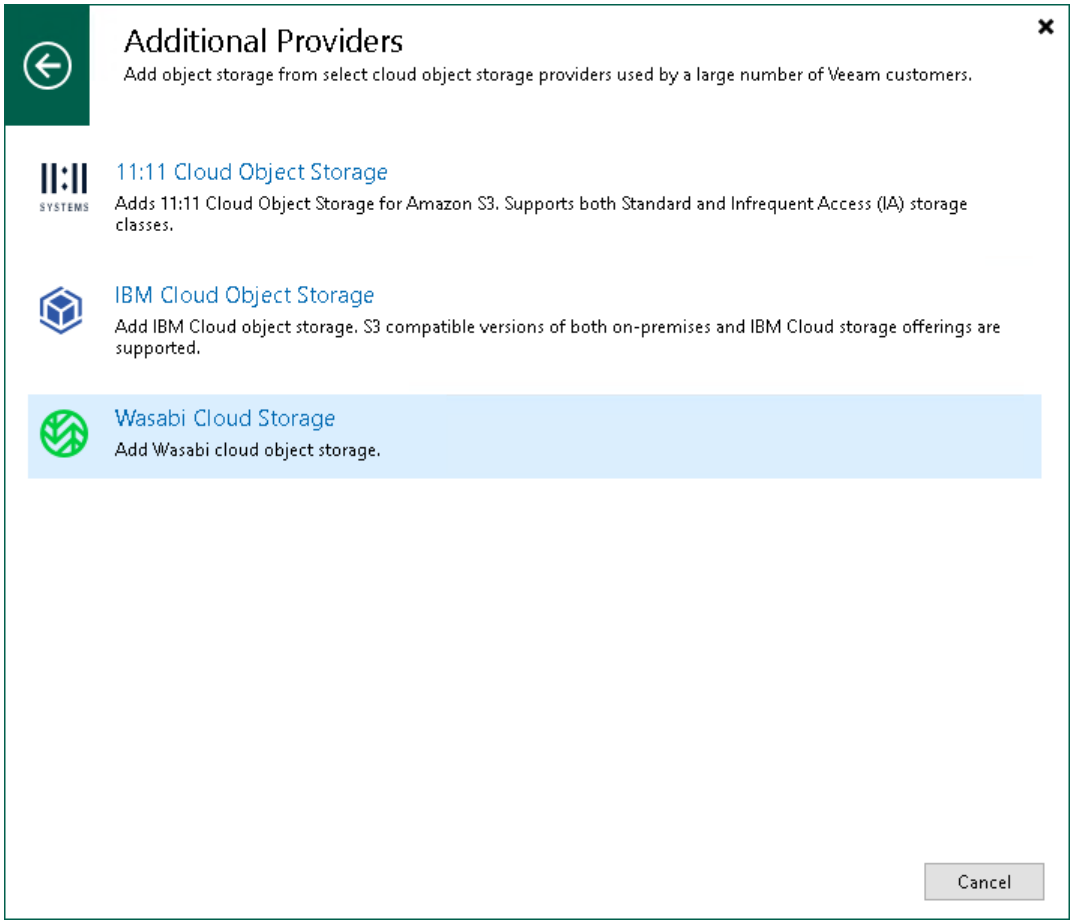

## Step 2. Specify Object Storage Name

At the **Name** step of the wizard, use the **Name** and **Description** fields to enter a name for a new object storage repository and to provide a description for future reference.

If you want to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

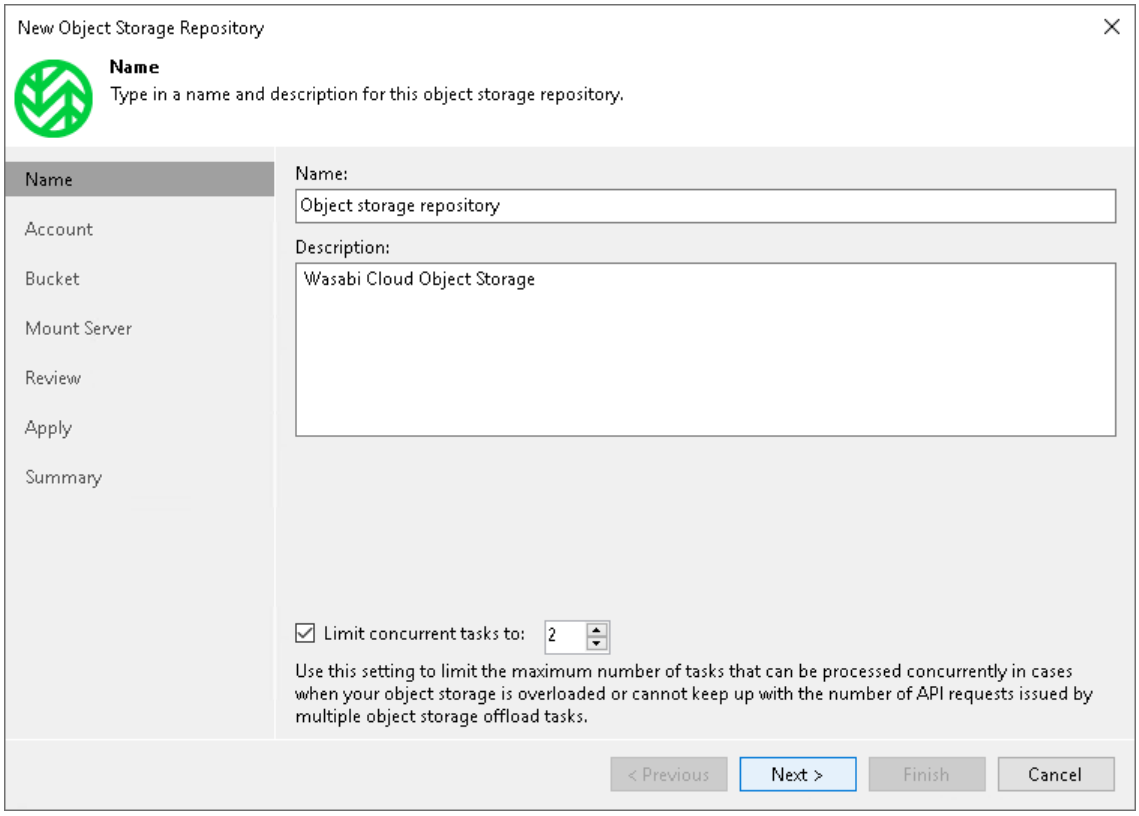

## Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. In the **Region** field, the region is specified by default.
2. From the **Credentials** drop-down list, select user credentials to access your Wasabi cloud object storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.
3. Next to the **Connection mode** field, click **Choose** and specify how Veeam Backup & Replication will transfer data to the object storage repository:
   - **Direct** – select this option if you want to instantly move data of processed VMs or file shares to object storage repositories. Before you select this option, check the following Considerations and Limitations.
   - **Through gateway server** – select this option if you want Veeam Backup & Replication to use gateway servers to transfer data from processed machines or file shares to object storage repositories. From the **Name** list, select gateway servers that you want to use for data transfer operations.

     By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. You can choose any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. Note that you must add the server to the backup infrastructure beforehand. Before you add the server, check the following Considerations and Limitations. For more information on how to add a server, see Virtualization Servers and Hosts.

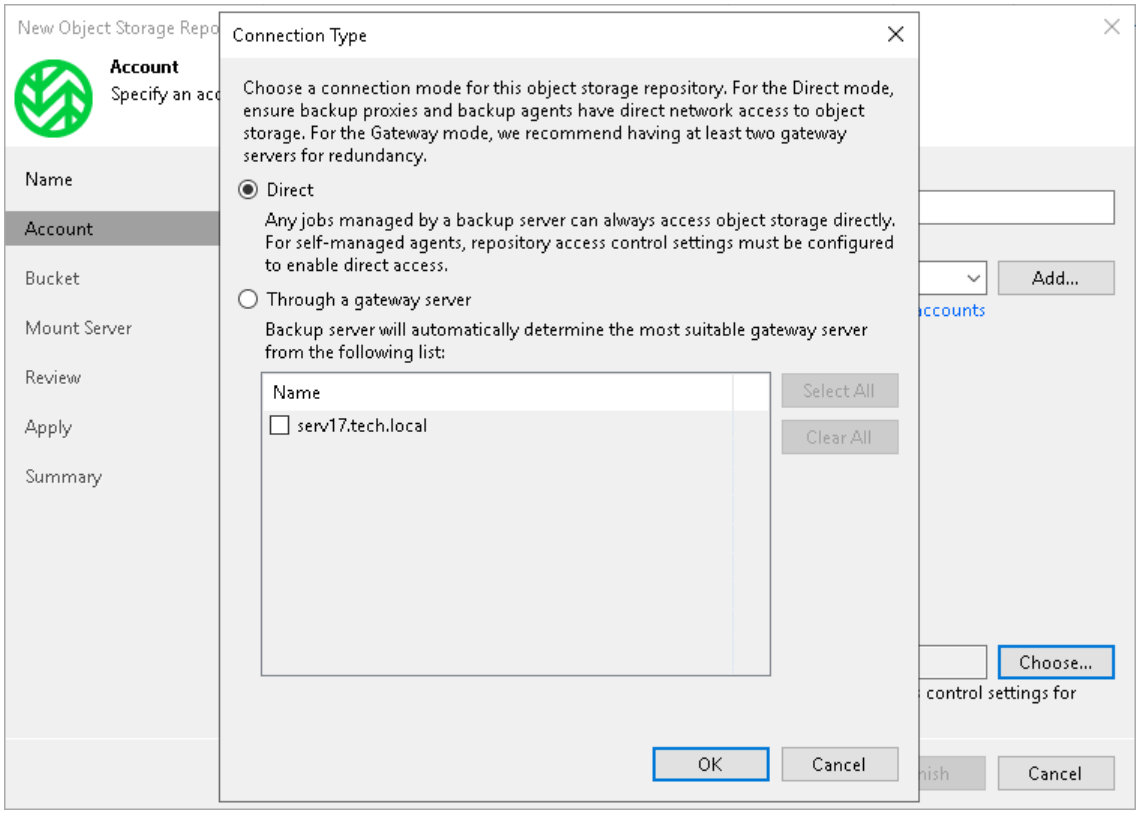

## Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify the bucket and folder that will be used to store data:

1. In the **Bucket** field, enter a name of the bucket or click **Browse** to get the necessary bucket.

   Note that you must create the bucket where you want to store your backup data beforehand.

2. In the **Folder** field, enter a cloud folder name to which you want to map your object storage repository. Alternatively, click **Browse** and either select an existing folder or click **New Folder**.

3. Select the **Limit object storage consumption** to check box to define a soft limit for your object storage consumption. If this limit is exceeded during a job run, Veeam Backup & Replication will allow to complete the job. However, a new job will not be able to start unless you remove the extra data that exceeds the limit or change the soft limit settings. Provide the value in TB or PB.

4. Select the **Make recent backups immutable for** check box to prohibit deletion of blocks of data from object storage. Specify the immutability period. For more information, see Immutability for Scale-Out Backup Repositories.

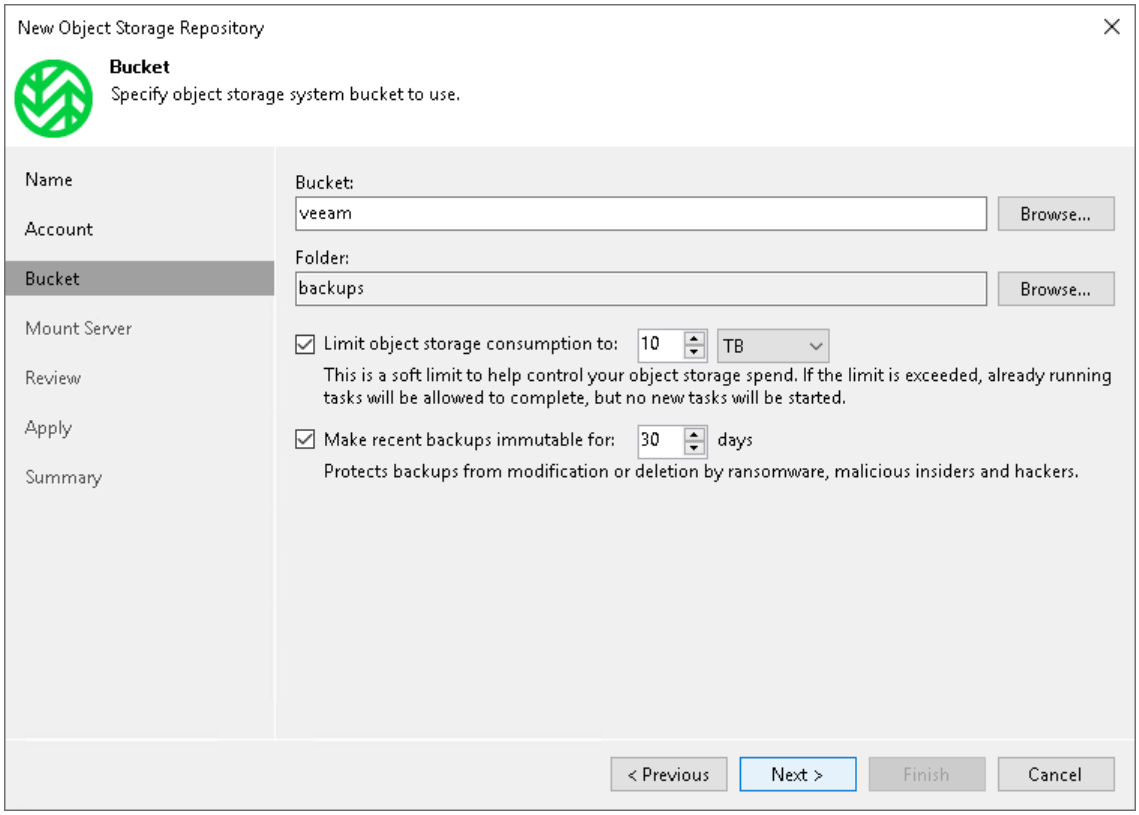

### Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for restore operations, and configure a helper appliance. The helper appliance is a Windows-based or Linux-based virtual or physical server, added to the backup infrastructure, that Veeam Backup & Replication uses to perform a health check of backup files and apply retention to unstructured data backup files. For more information, see Health Check for Object Storage Repositories and Helper Appliance in Unstructured Data Backup.

> **IMPORTANT**
>
> Consider the following:
>
> - If you do not configure a helper appliance, Veeam Backup & Replication will use local resources to perform the health check and apply retention to unstructured data backup files. It will consume more cloud resources and can result in additional costs.
> - To perform the health check, you must enable this option when you configure a job. For more information, see Health Check for Backup Files.

# Specifying Mount Server Settings

To specify the mount server settings, do the following:

1. From the **Mount Server** drop-down list, select a server that you want to use as a mount server. Veeam Backup & Replication uses this server during restore operations to mount VM disks directly from objects located in object storage repositories. For more information, see Mount Servers.

   The **Mount Server** list contains only Microsoft Windows servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, click **Add New** on the right to open the **New Windows Server wizard**. For more information, see Adding Microsoft Windows Servers.

2. In the **Instant recovery write cache folder** field, specify a folder to keep cache that is created during mount operations.

3. Select the **Enable vPower NFS service on the mount server** check box to allow the Veeam vPower NFS Service access the object storage repository. Veeam Backup & Replication will enable the Veeam vPower NFS Service on the necessary mount server. For more information, see Veeam vPower NFS Service.

4. Click **Ports** to customize network ports used by the Veeam vPower NFS Service. In the **vPower NFS Port Settings** window, specify the following settings:

   - Next to the **Mount Port** section, specify the port that the Veeam vPower NFS Service will use to mount the vPower NFS datastore to the ESXi host.
   - Next to the **vPower NFS** port section, specify the port that the Veeam vPower NFS Service will use to connect to the target NFS share.

   For information on ports used by default, see Ports.

5. To specify the helper appliance settings, click **Configure**. From the **Managed server** drop-down list, select a server that you want to use as the helper appliance.

> **IMPORTANT**
>
> Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.

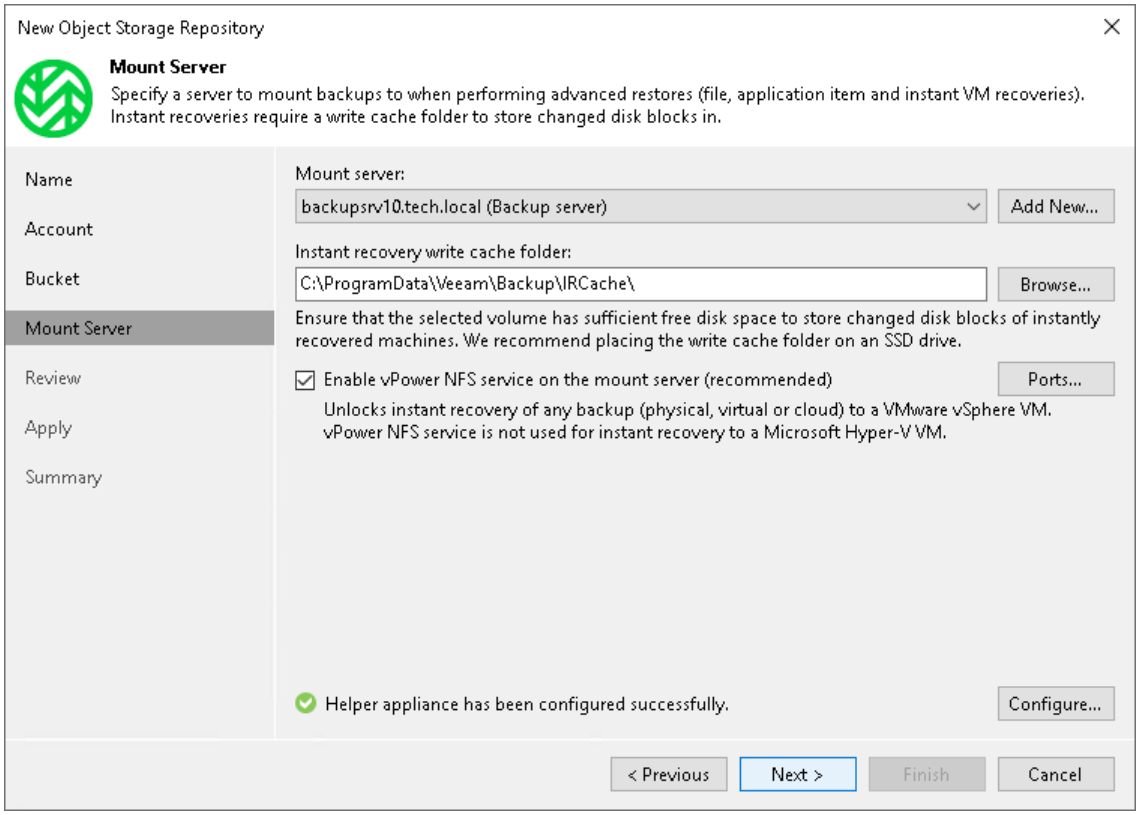
New Object Storage Repository
Mount Server
Specify a server to mount backups to when performing advanced restores (file, application item and instant VM recoveries). Instant recoveries require a write cache folder to store changed disk blocks in.
Name
Account
Bucket
Mount Server
Review
Apply
Summary
Mount server:
backupsrv10.tech.local (Backup server)
Add New...
Instant recovery write cache folder:
C:\ProgramData\Veeam\Backup\IRCache\
Browse...
Ensure that the selected volume has sufficient free disk space to store changed disk blocks of instantly recovered machines. We recommend placing the write cache folder on an SSD drive.
Enable vPower NFS service on the mount server (recommended)
Ports...
Unlocks instant recovery of any backup (physical, virtual or cloud) to a VMware vSphere VM. vPower NFS service is not used for instant recovery to a Microsoft Hyper-V VM.
Helper appliance has been configured successfully.
Configure...
< Previous
Next >
Finish
Cancel

## Step 6. Review Components

At the **Review** step of the wizard, review what components will be processed on the mount server server and their status.

If the backup repository contains backups, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Backups** > **Object Storage (Imported)** node.

If the backup repository contains guest file system index files, select the **Import guest file system index data to the catalog** check box. Veeam Backup & Replication will import index files together with backup files, and you will be able to search for guest OS files inside imported backups. For more information, see the Guest OS File Restore section of the Veeam Backup Enterprise Manager Guide.

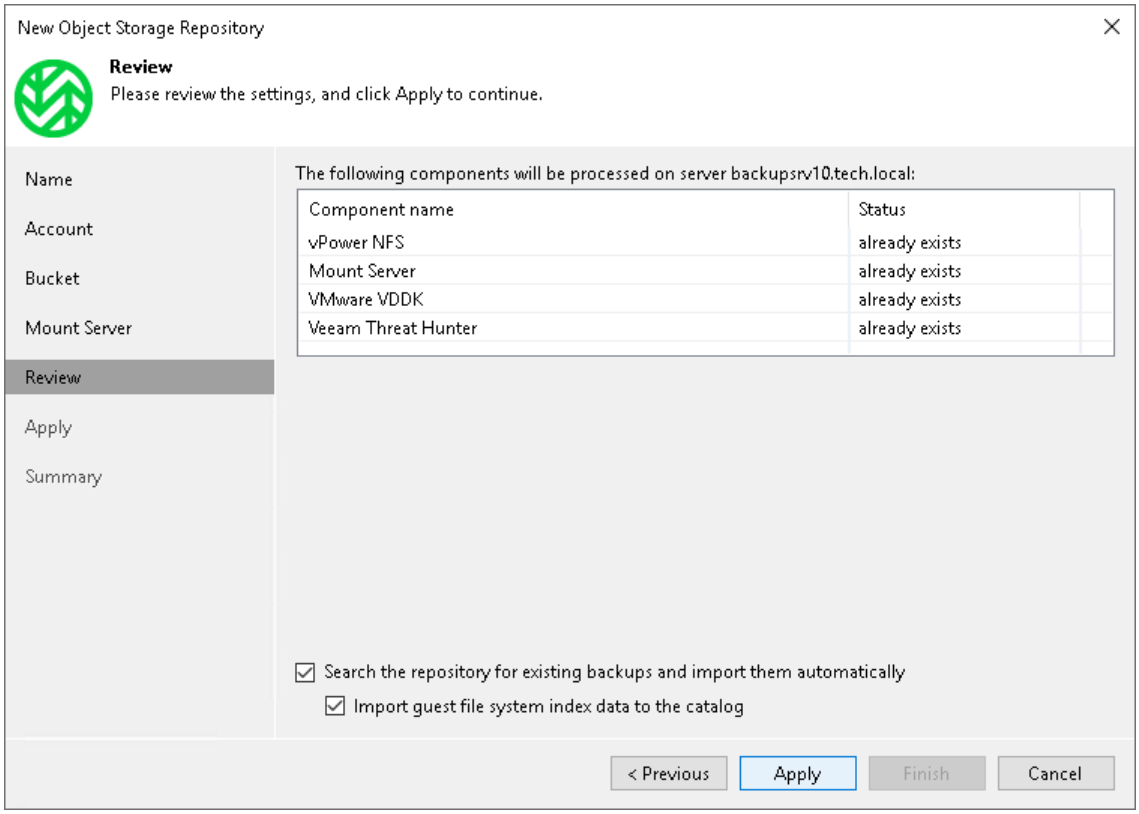

## Step 7. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to complete saving your settings to the configuration database and create backup infrastructure objects.

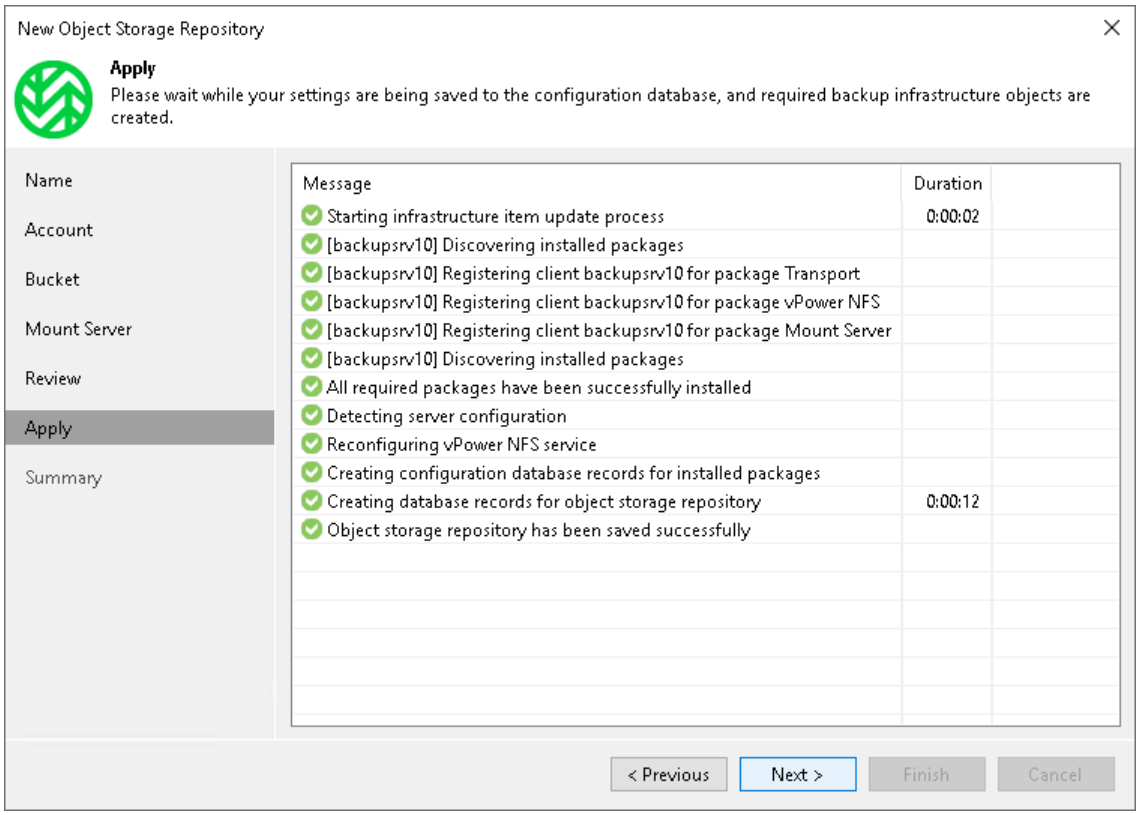

### Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the newly created object storage repository and click **Finish**.

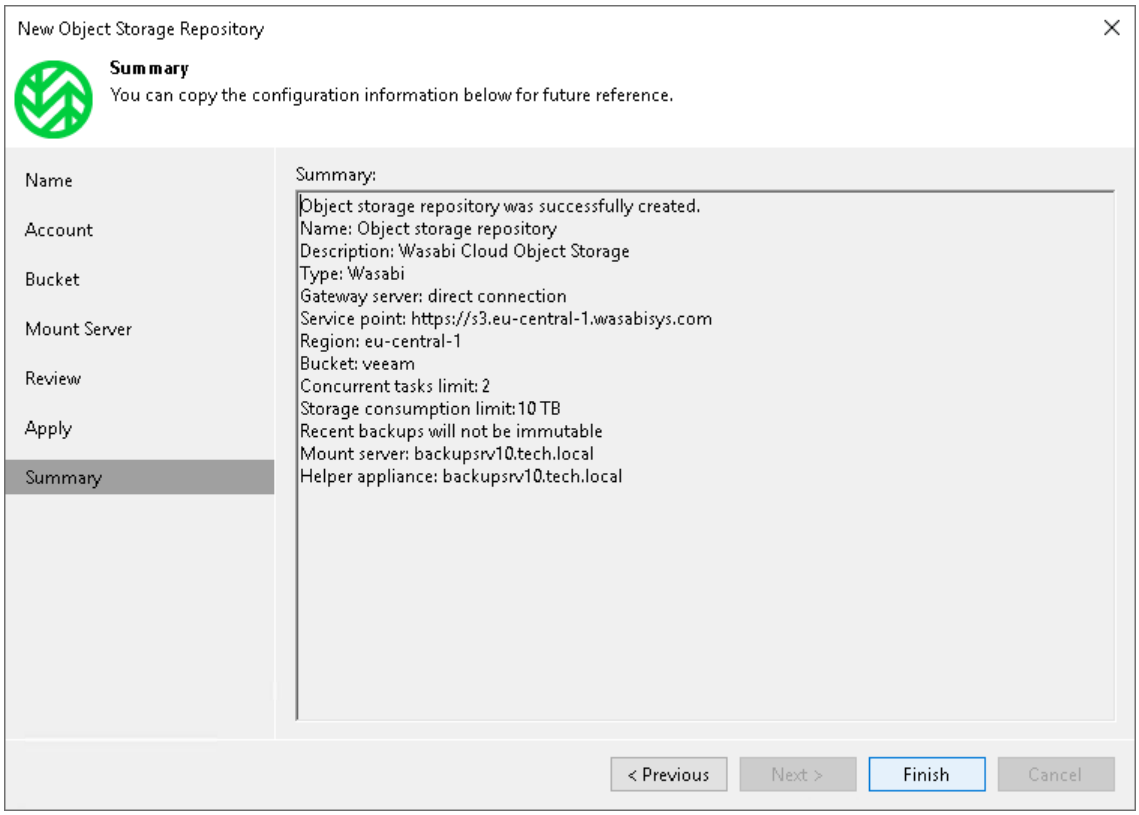


## Working with Veeam Smart Object Storage API (SOSAPI)

Veeam Backup & Replication offers Veeam Smart Object Storage application programming interface (SOSAPI) functionality to interact with S3 compatible object storage repositories. Object storage vendors can leverage SOSAPI to decrease impact on the production environment, significantly improve recovery point objective (RPO) and optimize the interaction between object storage and Veeam Backup & Replication. SOSAPI relies on Amazon S3 API protocol, uses its authorization mechanisms and existing S3 API calls. This approach does not require to set up a complex network infrastructure - there is no need to create additional ports, firewall rules or install any plug-ins. Instead, vendors need to handle the API requests to interact with S3 compatible and S3 compatible with Data Archiving object storage repositories.

Among others, SOSAPI functionality suggests taking advantage of the following options:

- **Capacity reports**

  Veeam Backup & Replication sends a warning if space capacity of an object storage repository reaches its limits. Also, you can check the amount of free space available in the object storage repository using Veeam Backup & Replication UI. For more information, see Editing Settings of Backup Repositories.

- **Enhanced scale-out backup repository placement policy**

  In case you use multiple extents within the scale-out backup repository, Veeam Backup & Replication will select the best extent to keep the backup data evenly.

- **Load balancing on smart entity level**

  A vendor can implement internal load balancing for SOSAPI-enabled repositories.

This option defines that Veeam Backup & Replication will be able to perform read and write task with the necessary repositories using a set of vendor-defined IP addresses instead of only one endpoint address.

# Working with SOSAPI

To start working with SOSAPI, vendors must add to their object storage infrastructure a set of XML configuration files. Veeam Backup & Replication will automatically enable SOSAPI functionality supported by the vendor and provide necessary capabilities to customers. After you add the S3 compatible or S3 compatible with Data Archiving object storage repositories with SOSAPI functionality to the backup infrastructure, it will show up as the S3-integrated type in working area.

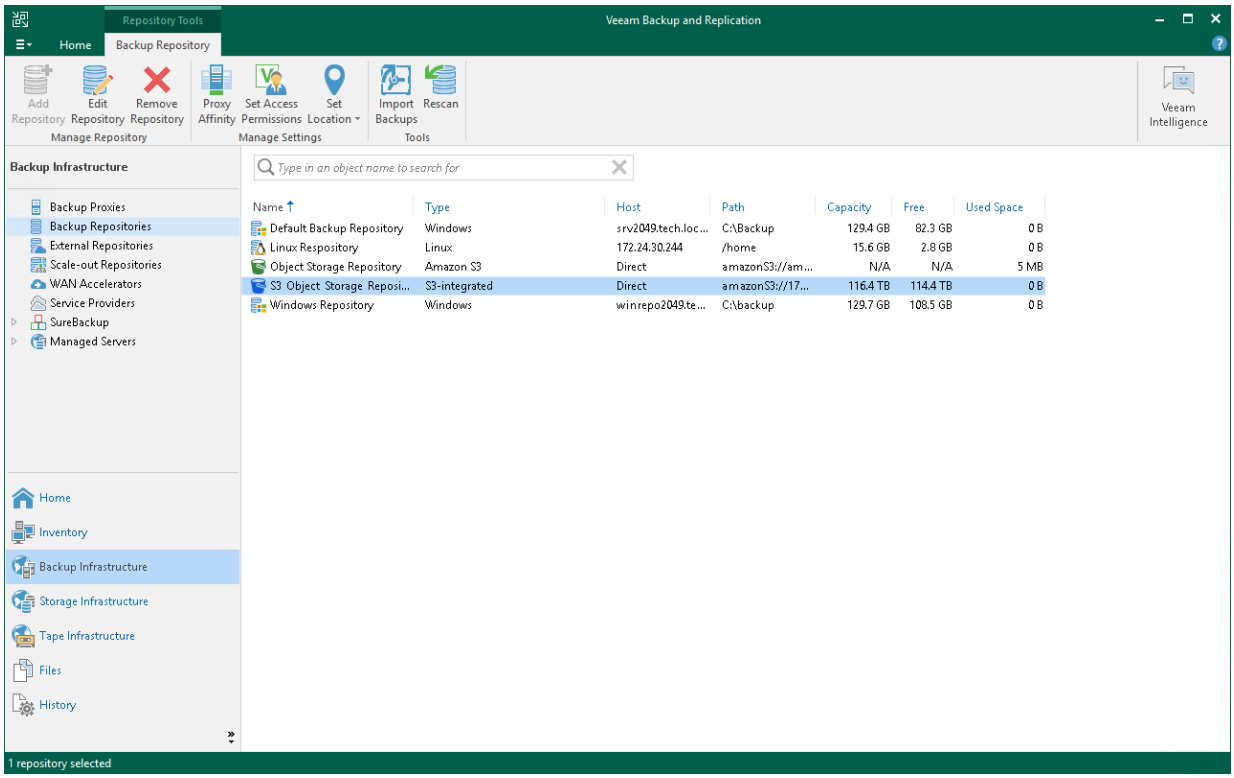


# How SOSAPI Works

To interact with object storage, SOSAPI uses API requests. It sends API requests to S3 compatible object storage repository and gets the necessary information in a set of XML files. These files contain details on the system, object storage repository capacity and correct storage usage, object storage capabilities and a state of the backup processing. Therefore, the SOSAPI functionality allows you to maintain the efficient state of the object storage repository and protect your backup infrastructure from data loss.

You can use SOSAPI functionality for advanced interaction with backups created by the following types of jobs and policies:

- Backups of VMware vSphere VMs created by Veeam Backup & Replication.
- Backups of Microsoft Hyper-V VMs created by Veeam Backup & Replication.
- Backups of file shares created by Veeam Backup & Replication.
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows and managed by Veeam Backup & Replication.
- Backups of virtual and physical machines created by Veeam Agent for Linux and managed by Veeam Backup & Replication.

- Backups of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV.
- Backups of oVirt VMs created by Veeam Backup for OLVM and RHV.
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE.
- Backups stored in S3 compatible object storage repositories of Service Providers.

> **NOTE**
>
> This option works only if a tenant connects to the object storage through a gateway server. For more information, see the Backup to Object Storage section in the Veeam Cloud Connect Guide.

# Managing Object Storage Repositories

You can manage your object storage repositories in various ways: edit settings of an object storage repository or remove an object storage repository.

## Editing Settings of Object Storage Repository

After you have added an object storage repository, you may want to edit its settings.

To edit object storage settings, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Repositories**.
3. In the working area, select an object storage repository and click **Edit Repository** on the ribbon or right-click an object storage repository and select **Properties**.
4. Follow the steps of the **Edit Object Storage Repository** wizard and edit settings as required.

   Note that some settings cannot be modified and will remain disabled while being edited.

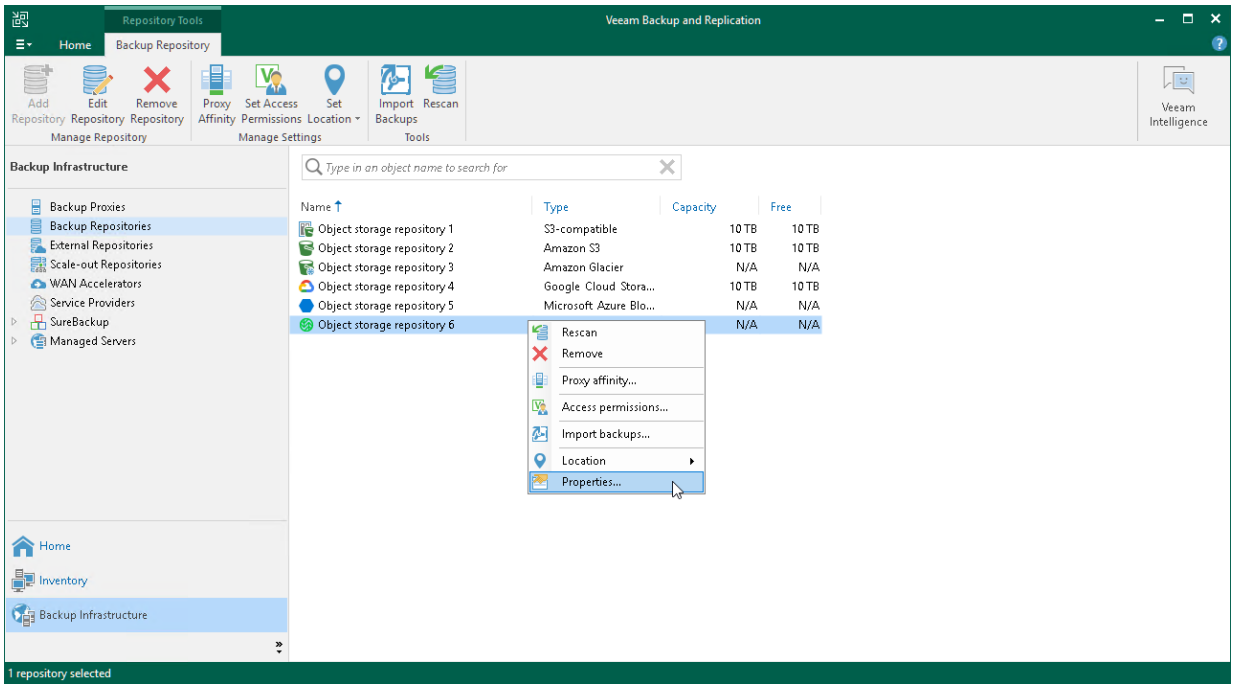

## Removing Object Storage Repository

You can remove any object storage repository from the application scope if you no longer need it.

# Considerations and Limitations

Before you remove an object storage repository, consider the following limitations:

- An object storage repository cannot be removed if it is part of a scale-out backup repository.

  To remove such a repository, you must first exclude an object storage repository from the scale-out backup repository configuration. For more information, see Removing Performance Extents from Scale-Out Repositories and Excluding Capacity Extent from Scale-Out Repositories.

- An object storage repository cannot be removed if backups located in this repository was imported, as described in section Importing Object Storage Backups.

  To remove such a repository, you must first detach object storage, as described in section Detaching Object Storage Backups.

- When an object storage repository is being removed from the environment, the data remains completely unaffected.

  To learn how to remove data, see Removing Backups from Capacity or Archive Tier.

# Removing Object Storage Repository

To remove an object storage repository, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Repositories**.
3. In the working area, select an object storage repository and click **Remove Repository** on the ribbon or right-click an object storage repository and select **Remove**.

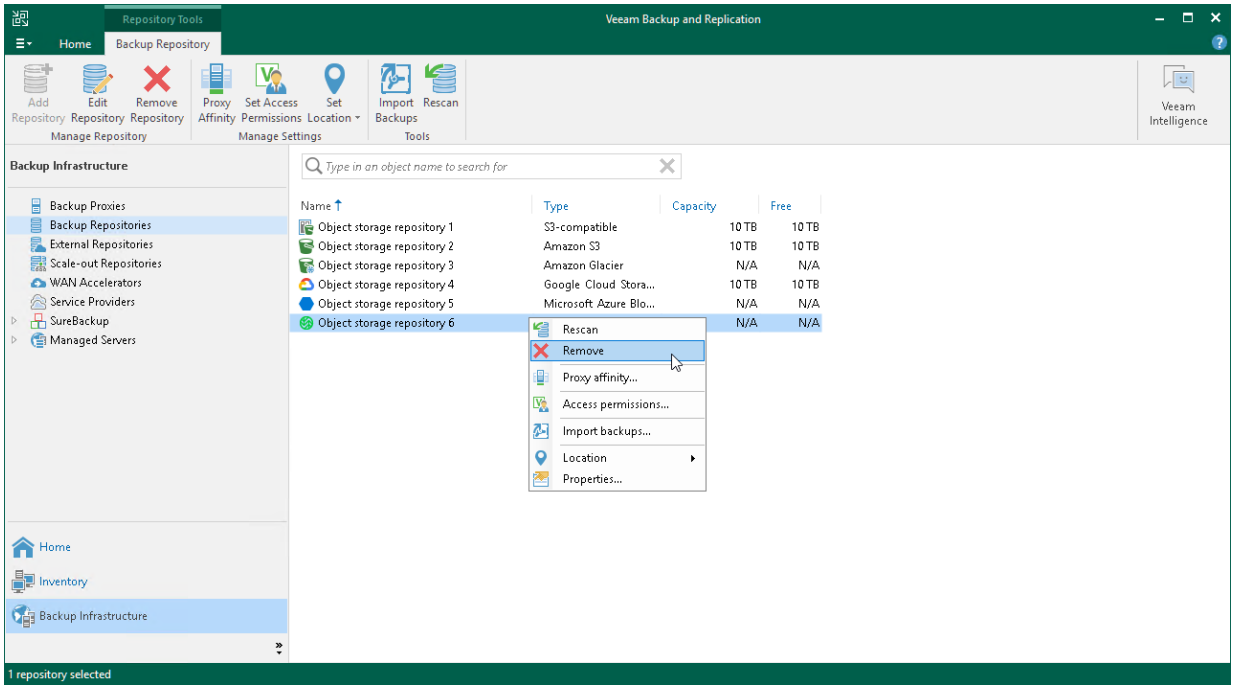

# Managing Object Storage Backups

You can manage your object storage backups in various ways: import, detach or remove object storage backups or seed them to Azure Data Box or AWS Snowball Edge storage devices.

## Importing Object Storage Backups

In case a disaster strikes and your Veeam Backup & Replication with a scale-out backup repository becomes unavailable, you can import into another Veeam Backup & Replication infrastructure backups located in object storage repositories added as performance, capacity or archive extents of the scale-out backup repository. Once you add the necessary object storage repositories in Veeam Backup & Replication infrastructure, you can import backups and they will become available for data recovery operations.

> **NOTE**
>
> If you want to import backups located in standalone object storage repositories or that are added as performance extents of a scale-out backup repository, use the Rescan Backup Repositories option.

Consider the following:

- The **Import Backups** option is available only if the object storage is not a part of the scale-out backup repository.
- Before you start importing backups, make sure to add the object storage repository that stores data you want to import.
- If you have imported backups from the object storage repository, you will not be able to use it as a standalone repository when you configure a job or as an extent of a scale-out backups repository. To be able to do it, you will need to detach this object storage.
- The **Import Backups** option is applicable only to the Capacity Tier backups and Archive Tier backups. It does not support the unstructured data backups.

To import backups, do the following:

1. Launch the Import Backups wizard.
2. Specify password.
3. Wait for import.
4. Finish working with the wizard.

## Step 1. Launch Import Wizard

To launch the **Import Backups** wizard, open the **Backup Infrastructure** view and in the inventory pane select the **Backup Repositories** node. In the preview pane, select necessary object storage and select **Import Backups**. Alternatively, right-click necessary object storage and select **Import backups**.

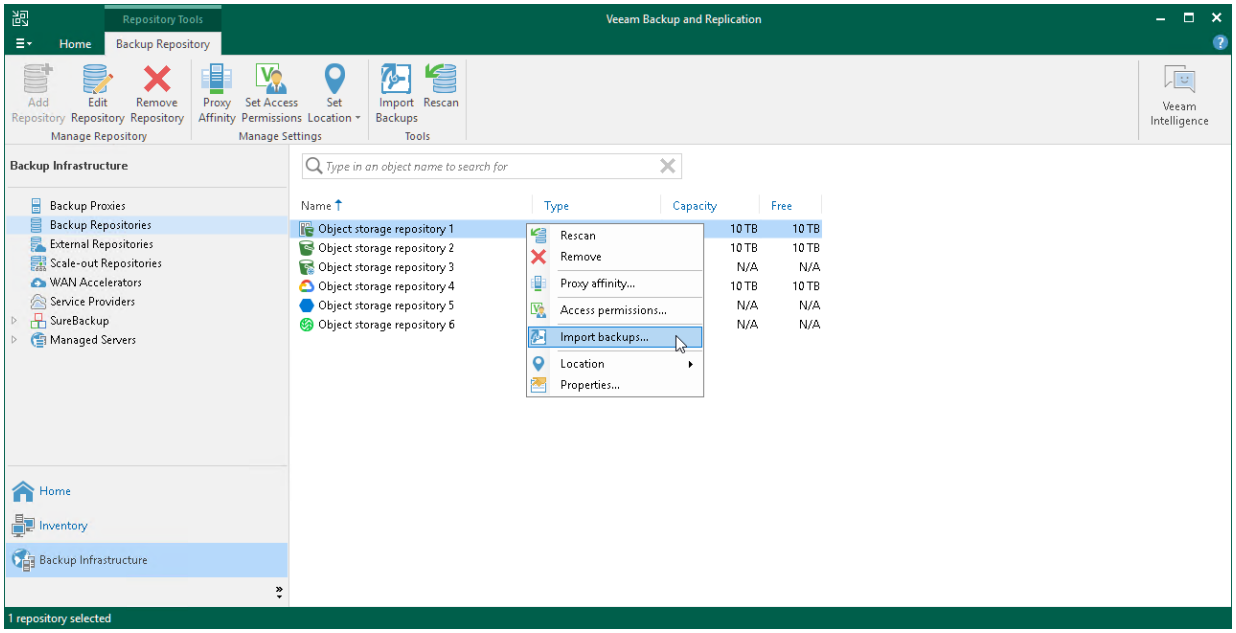

## Step 2. Specify Password

This step is only available when importing encrypted backups.

At the **Password** step of the wizard, in the **Password** field, specify the password that was used to encrypt data during offload or copy sessions.

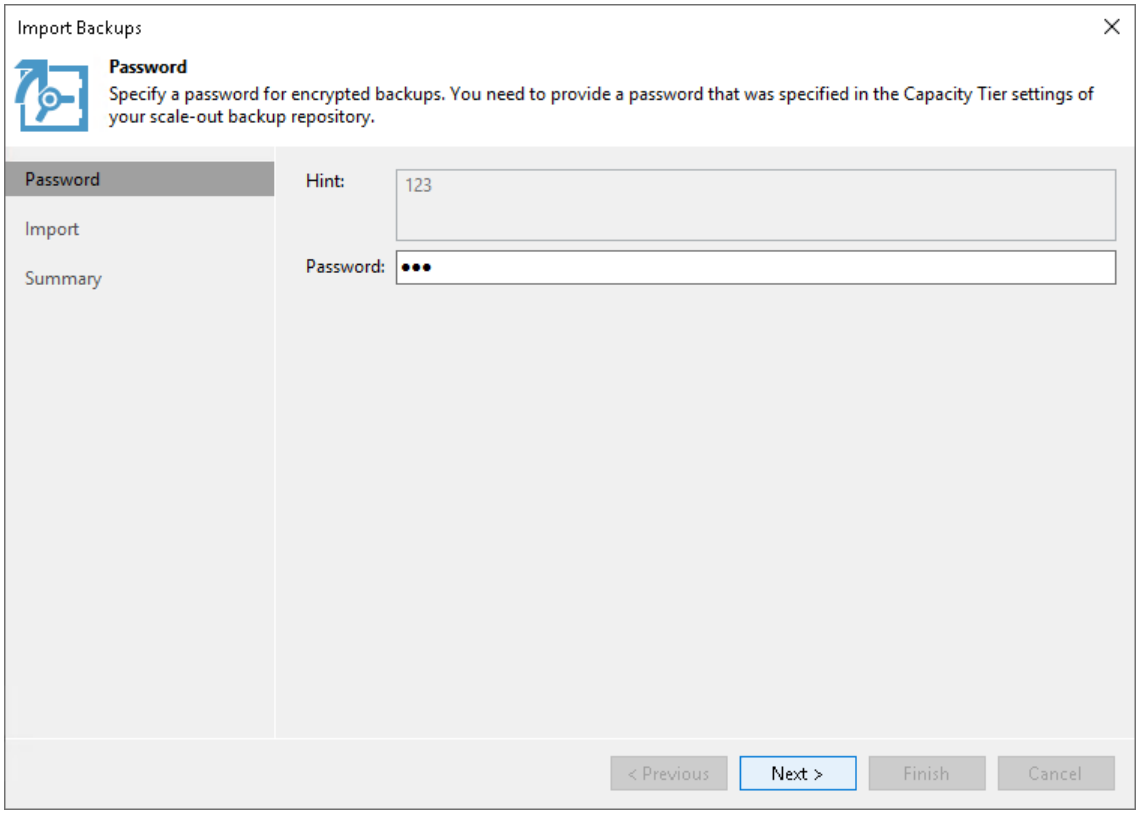

## Step 3. Wait for Import

At the **Import** step of the wizard, wait until Veeam Backup & Replication prepares a temporary database to which information about backups will be added upon import.

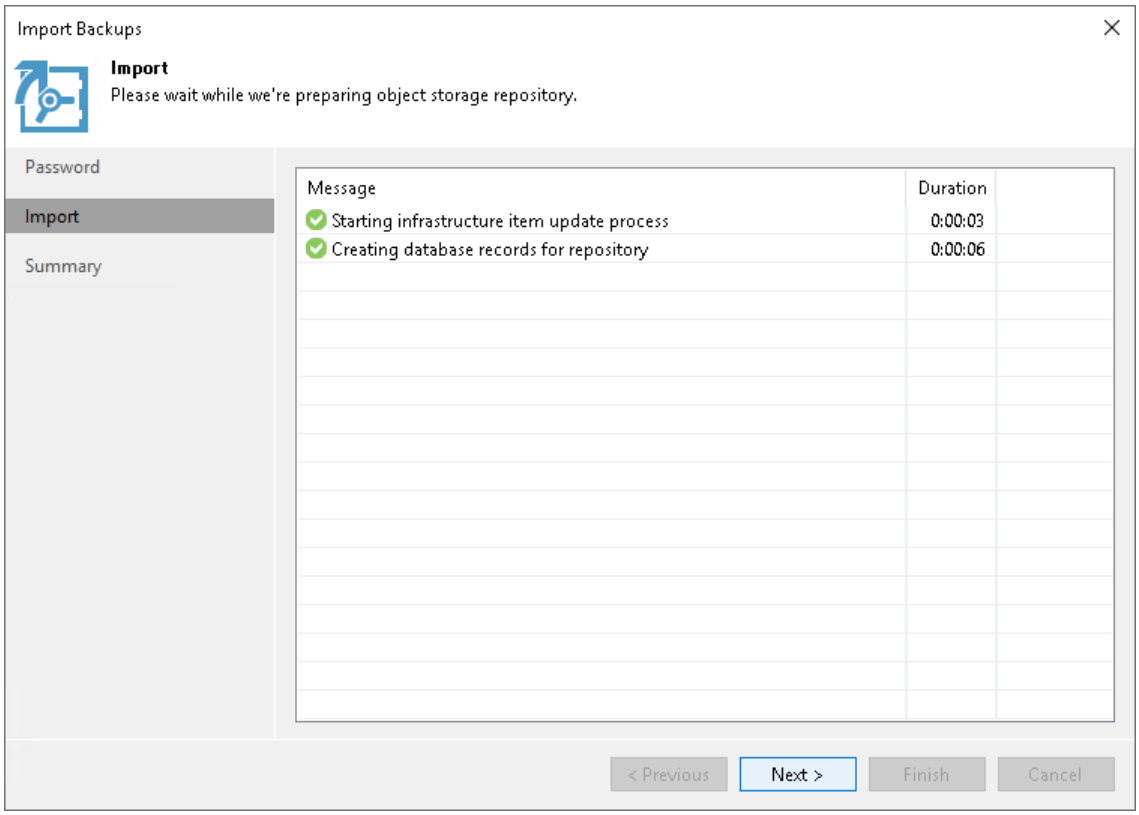

## Step 4. Finish Working with Wizard

At the **Summary** step of the wizard:

1. Review details of the import operation.
2. Click **Finish** to run a **Configuration Database Resynchronize** session.

   During this session, all existing backups will be imported into the Veeam Backup & Replication console.

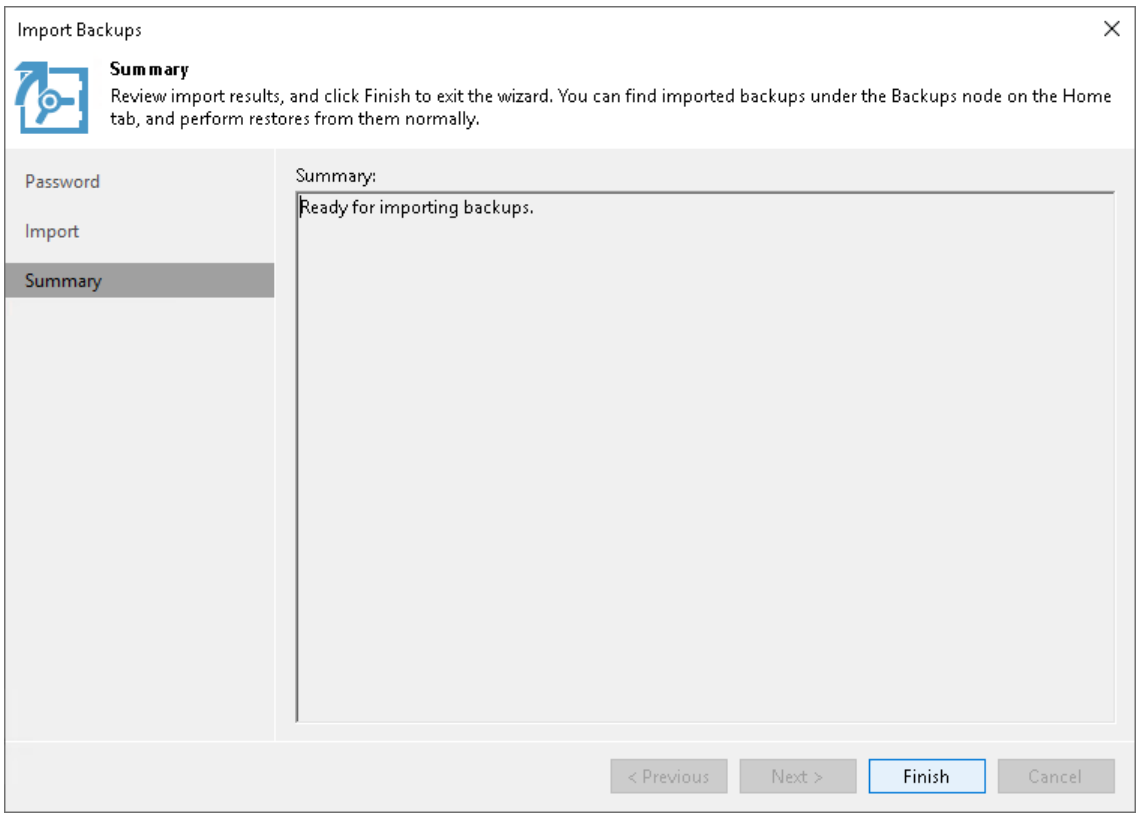


## Viewing Imported Backups

To view imported backups, do the following:

1. Open the **Home** view.
2. In the navigation pane, select the **Backups** > **Capacity Tier (Imported)** or **Archive Tier (Imported)** node.

3. In the preview pane, review imported backups.

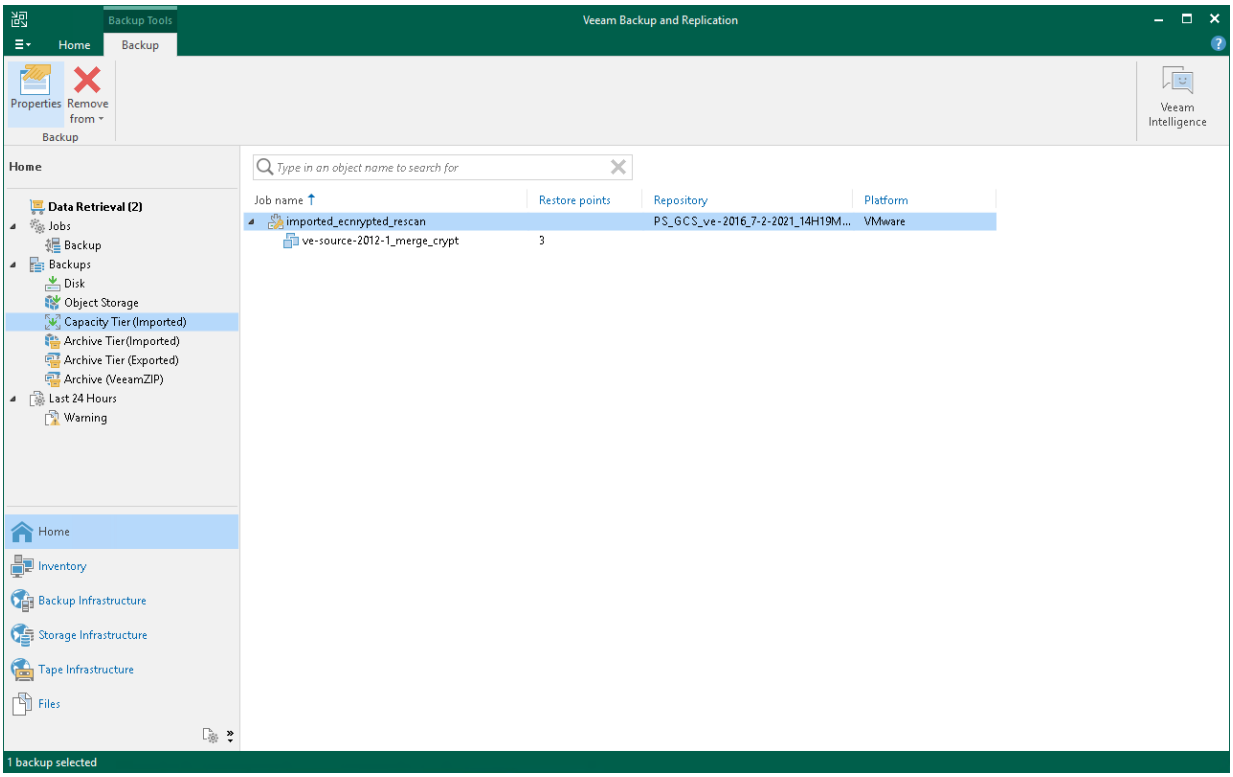


# Detaching Object Storage Backups

You can detach imported object storage after you have finished working with it. After object storage is detached, the imported backups become unavailable.

To detach imported object storage, do the following:

1. Open the **Backup Infrastructure** view.
2. In the navigation pane, select the **Backup Repositories** node.

3. In the preview pane, select object storage the backups of which have been imported and click **Detach Repository** on the toolbar or right-click object storage and select **Detach Repository**.

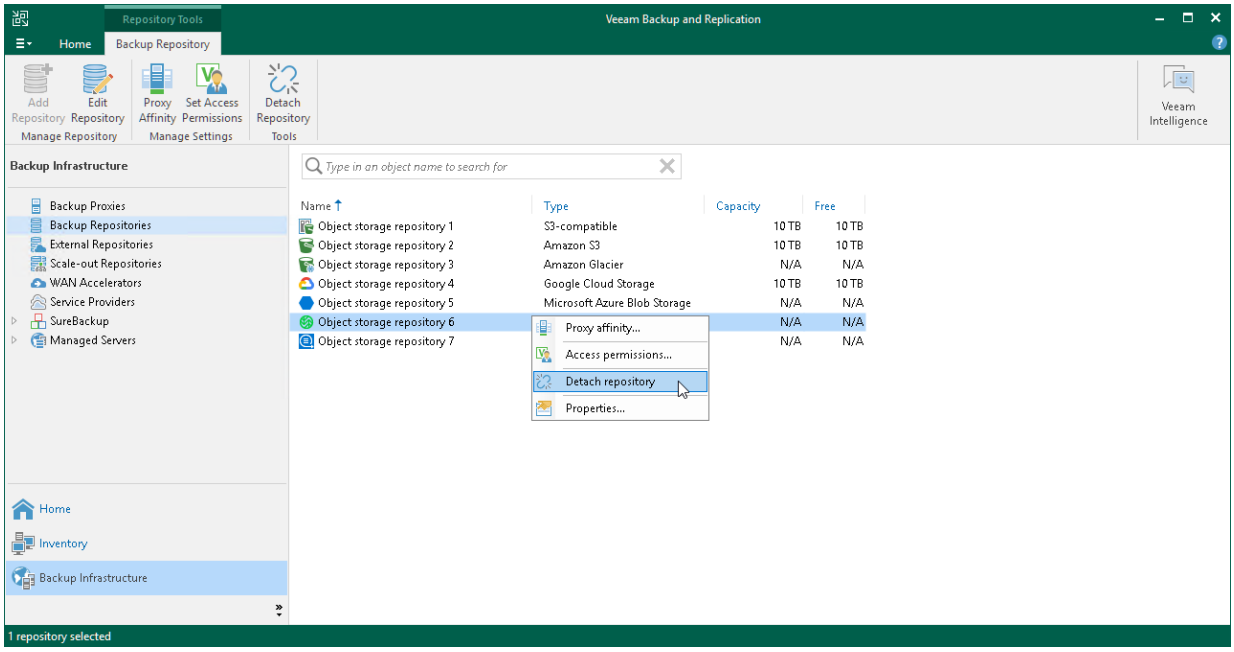


# Seeding Backups to Azure Blob Storage

Seeding backups is an operation which can help in case you want to keep backups in object storage repositories and need to transfer backups in bulk to object storage. In this case, you can utilize the capabilities of the Azure Data Box device that will help you to transfer data from the local backup repository to Azure Blob storage and synchronize with the storage. Azure Data Box is a physical storage device that emulates a cloud storage endpoint and serves as an intermediate between the local backup repository and Azure Blob storage. With the help of Azure Data Box you will not need to create a full backup file and transfer it through the network which may result in high load on the network. Instead, you will transport all necessary data using Azure Data Box.

To seed data to Azure Blob storage, you need to order Azure Data Box from Microsoft, connect to it, copy backups from the local repository and ship it back to Microsoft. After the backup data is uploaded to your Azure Blob storage, you can perform initial sync and store your backups in Azure Blob storage later on.

For more information on ordering Azure Data Box, see Azure Storage Documentation.

> **IMPORTANT**
>
> Consider the following limitations:
>
> - We do not recommend offloading immutable data to the Azure Data Box device. If you do so, it will become mutable on the Azure Data Box device and will remain mutable after seeding.
> - You cannot seed data from the Azure Data Box device to Veeam Data Cloud Vault.

To seed backups to Azure Blob Storage do the following:

1. Review considerations and limitations.
2. Set up Azure Data Box.
3. Add Azure Data Box to the backup infrastructure and then offload the data.
4. Prepare Azure Data Box for shipping.
5. Synchronize the seeded data.

# Setting up Azure Data Box

Before you add Azure Data Box to Veeam Backup & Replication, you need to set up the storage:

1. Install Azure Data Box. For more information, see Azure Storage Documentation.
2. Connect to Azure Data Box through REST APIs. For more information, see Azure Storage Documentation.

   After you perform the steps in the Azure Storage Documentation, you will have a storage account name, access key, and blob service endpoint that you can use to add Azure Data Box to the backup infrastructure.

# Adding Azure Data Box and Transferring the Data

To transfer backups to Azure Data Box, do the following:

1. Add Azure Data Box as an Azure Data Box storage.
2. Add Azure Data Box storage as a standalone object storage repository, the capacity tier or performance tier in a scale-out backup repository.

   **IMPORTANT**

   When you configure a scale-out backup repository with Azure Blob storage as the capacity extent, make sure to select the **Copy backups to object storage as soon as they are created** check box as described in section Add Capacity Tier. In this case, you will keep copy of the data on your local storage which will help you to reduce the risk of data loss if the device is damaged during shipping. It will also ensure that the backup data is available for restore operations while Azure Blob storage is in shipment.

   If you use the other type of backup repository, make sure to keep a copy of your backup in your local storage.

# Preparing Azure Data Box for Shipping

After you have completed copying backup data to the Azure Data Box, you need to ship this device back to Microsoft so that this data can be synchronized with your Azure storage account.

Before shipping Azure Data Box back to Microsoft, do the following:

1. Put an object storage repository that is associated with Azure Data Box into the Maintenance mode, as described in section Switching to Maintenance Mode.

   **NOTE**

   In case you have added Azure Data Box as a standalone object storage repository, make sure that you stop all backup job activities to prevent backups from modification.

2. Perform the steps described in Microsoft Docs.

After shipping the device, wait for Microsoft to receive it and copy your data into your Azure storage account. You will be notified of successful data upload to your Azure storage account.

# Synchronizing the Seeded Data

To connect Veeam Backup & Replication to your Azure storage account, do the following:

1. Add an Azure Blob storage, as described in section Adding Azure Blob Storage. At the Specify Object Storage Settings step, select the same container and folder that you have selected for your Azure Data Box device.
2. Depending on the way you have added Azure Data Box, do one of the following:
   - If you have added Azure Data Box as an object storage repository, as this Azure Data Box as the Azure Blob storage, rescan and map the jobs to the newly synced backups.
   - If you have added Azure Data Box as an extent of a scale-out backup repository, edit settings of this scale-out backup repository. Change your Azure Data Box object storage to the Azure Blob storage.
3. Enable importing backups from Azure Blob storage and wait for the synchronization to complete.
4. Remove Azure Data Box object storage from the backup infrastructure, as described in section Removing Object Storage Repository.

> **NOTE**
>
> Keep in mind that if you delete individual VMs from your backup in the period between mailing Azure Data Box and adding Azure Blob Storage, you may need to run Active Full instead of Incremental job after the final Azure synchronization. To avoid this, do not delete any VMs manually and disable the **Retention Policy for Deleted Items** option in the job settings for that period. After successful synchronization, you can re-enable this option and work as usual.

## Seeding Backups to Amazon S3 Storage

Seeding backups is an operation which can be helpful in case you want to start using object storage repositories and need to transfer a bulk amount of data to object storage. In this case, you can utilize the capabilities of the AWS Snowball Edge device that will help you to transfer data from the local backup repository to Amazon S3 storage and synchronize with the storage. AWS Snowball Edge is a physical storage device that emulates a cloud storage endpoint and serves as an intermediate between the local backup repository and Amazon S3 storage. With the help of AWS Snowball Edge you will not need to create a full backup file and transfer it through the internet which can result in high load on the network. Instead you will transport all necessary data using AWS Snowball Edge.

To seed data to Amazon S3 storage, you need to order AWS Snowball Edge from AWS, connect to it, copy backups from the local repository and ship it back to AWS. After the backup data is uploaded to your Amazon S3 storage, you can perform initial sync and store your backups in Amazon S3 storage later on.

For more information on AWS Snowball Edge, see AWS Documentation.

> **NOTE**
>
> Before you seed backups to AWS Snowball Edge, make sure to review considerations and limitations.

To seed backups to AWS Snowball Edge do the following:

1. Set up AWS Snowball Edge.
2. Add AWS Snowball Edge to the backup infrastructure and then offload the data.
3. Prepare AWS Snowball Edge for shipping.

4. Synchronize the seeded data.

# Setting up AWS Snowball Edge Device

Before you add AWS Snowball Edge to Veeam Backup & Replication, you need to order the device, connect it to your local network, download and install the Snowball edge client and configure the device. For detailed instructions on these steps, see AWS Documentation.

After that, you will have a service point address, access key and a secret key that you can use to add AWS Snowball Edge to the backup infrastructure.

# Adding AWS Snowball Edge and Transferring the Data

To transfer backups to AWS Snowball Edge, do the following:

1. Add AWS Snowball Edge as AWS Snowball Edge Storage.
2. Add AWS Snowball Edge as a standalone object storage repository, the capacity tier or performance tier in a scale-out backup repository.

   **IMPORTANT**

   When you configure a scale-out backup repository with AWS Snowball Edge as the capacity extent, make sure to select the **Copy backups to object storage as soon as they are created** check box as described in section Add Capacity Tier. In this case, you will keep copy of the data on your local storage which will help you to reduce the risk of data loss if the device is damaged during shipping. It will also ensure that the backup data is available for restore operations while AWS Snowball Edge is in shipment.

   If you use the other type of backup repository, make sure to keep a copy of your backup in your local storage.

# Preparing AWS Snowball Edge for Shipping

After you have completed copying backup data to the AWS Snowball Edge device, you need to ship this device back to Amazon so that this data can be synchronized with your Amazon S3 storage.

Before shipping AWS Snowball Edge back to AWS, do the following:

1. Put an object storage repository that is associated with AWS Snowball Edge into the Maintenance mode, as described in section Switching to Maintenance Mode.

   **NOTE**

   In case you have added AWS Snowball Edge as a standalone object storage repository, make sure that you stop all backup job activities to prevent backups from modification.

2. Power off AWS Snowball Edge. For more information, see AWS Documentation.
3. Ship AWS Snowball Edge back to AWS. For more information, see AWS Documentation.

After mailing the device, wait until it is accepted by AWS and notification of successful data upload to your storage account is received (if you subscribed for such notification).

# Synchronizing the Seeded Data

To connect Veeam Backup & Replication to your Amazon S3 storage account, do the following:

1. Add an Amazon S3 storage, as described in section Adding Amazon S3 Storage. At the Specify Object Storage Settings step, select the same bucket and folder that you have selected for your AWS Snowball Edge device.
2. Depending on the way you have added AWS Snowball Edge, do one of the following:
   - If you have added AWS Snowball Edge as an object storage repository, as this AWS Snowball Edge as the Amazon S3 storage, rescan and map the jobs to the newly synced backups.
   - If you have added AWS Snowball Edge as an extent of a scale-out backup repository, edit settings of this scale-out backup repository. Change your AWS Snowball Edge object storage to the Amazon S3 storage.
3. Enable importing backups from Amazon S3 storage and wait for the synchronization to complete.
4. Remove the AWS Snowball Edge object storage from the backup infrastructure, as described in section Removing Object Storage Repository.

# Managing Backup Repositories

You can manage your backup repositories and data stored in them in various ways: edit settings of the backup repositories, set up access permissions, rescan backup repositories or remove backup repositories.

## In This Section


- Editing Settings of Backup Repositories
- Access Permissions
- Encrypting Standalone Application Backups in Backup Repositories
- Rescanning Backup Repositories
- Removing Backup Repositories
- Fast Clone
- Proxy Affinity


## Editing Settings of Backup Repositories

You can edit settings of backup repositories that you have added to the backup infrastructure.

**NOTE**

When editing backup repository settings, you cannot change the selected repository server. Therefore, the **Repository server** field at the **Server** step is grayed out.

**TIP**

The following column headers show information on the repository storage settings:

- The **Capacity** column header — specifies a full size of a storage location where you keep your backups.
- The **Free** column header — specifies the amount of free space on a storage location. It considers a size of all data that is already added to this storage location without using Veeam Backup & Replication. For example, if the storage location capacity is 100 GB, and you added to this location data that occupies 50 GB, the Free column header will display 50 GB.
- The **Used Space** column header — specifies the amount of space occupied by backups created by Veeam Backup & Replication. Note that it does not consider data compression and deduplication that is not performed by means of Veeam Backup & Replication. For example, if Veeam Backup & Replication creates a backup that occupies 10 GB, and a size of this backup decreases to 9 GB after compression or deduplication performed by a storage appliance, the Used Space column header will display 10 GB.
- The **Membership** column header — specifies a scale-out backup repository to which repositories are added as performance extents. This column is hidden by default. To display it, right-click any column header and select **Membership** from the drop-down menu.

To edit settings of a backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Repositories** node.

3. In the working area, select the backup repository and click **Edit Repository** on the ribbon or right-click the backup repository and select **Properties**.
4. Edit the backup repository settings as required. Note that you cannot change the selected repository server and path to the folder used for storing backups.

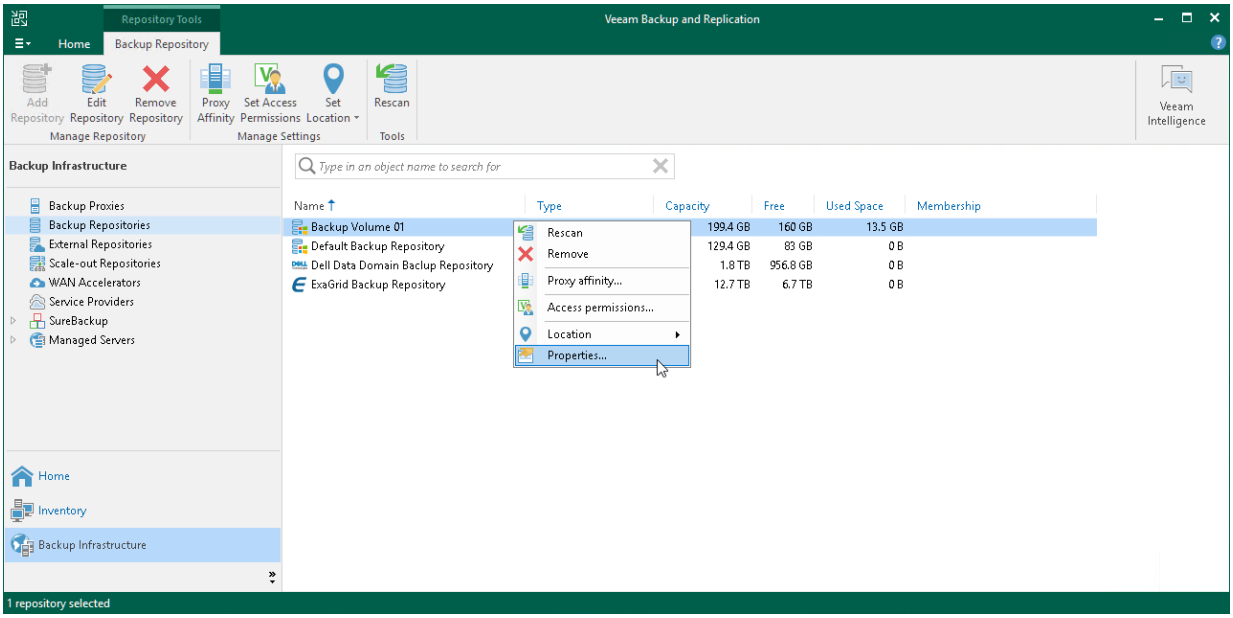


# Editing Access Permissions

If you want to store in the backup repository backups of virtual and physical machines created by Veeam Backup & Replication additional solutions, for example, Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Plug-ins for Enterprise Applications and so on, you must set up access permissions to backup repositories.

Access permissions are granted to security principals such as users and AD groups by the backup administrator who works with Veeam Backup & Replication. Users with granted access permissions can target backup jobs created by additional solutions at this backup repository and perform restore from backups located in this backup repository.

> **NOTE**
>
> Consider the following:
>
> - If you plan to create backups in a Veeam backup repository with Veeam Agent backup jobs configured in Veeam Backup & Replication, you do not need to grant access permissions on the backup repository to users. In the Veeam Agent management scenario, to establish a connection between the backup server and protected computers, Veeam Backup & Replication uses a TLS certificate. To learn more, see the Configuring Security Settings section in the Veeam Agent Management Guide.
> - If you plan to create backups in a Veeam backup repository with Veeam Backup for Nutanix AHV, you do not need to grant access permissions when configuring repositories, you must do that when configuring Nutanix AHV backup appliances.

Right after installation, access permissions on the default backup repository are set to *Allow to everyone* for testing and evaluation purposes. If necessary, you can change these settings.

After you create a new backup repository, access permissions on this repository are set to *Deny to everyone*. To allow users to store backups in the backup repository, you must grant users with access permissions to this repository.

# Managing Permissions of Backup Repositories

To grant access permissions to a security principal:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click one of the following nodes:
   - The **Backup Repositories** node — if you want to grant access permissions on a regular backup repository.
   - The **Scale-out Repositories** node — if you want to grant access permissions on a scale-out backup repository.
3. In the working area, select the necessary backup repository and click **Set Access Permissions** on the ribbon or right-click the backup repository and select **Access permissions**. If you do not see the **Set Access Permissions** button on the ribbon or the **Access permissions** command is not available in the shortcut menu, press and hold the [Ctrl] key, right-click the backup repository and select **Access permissions**.

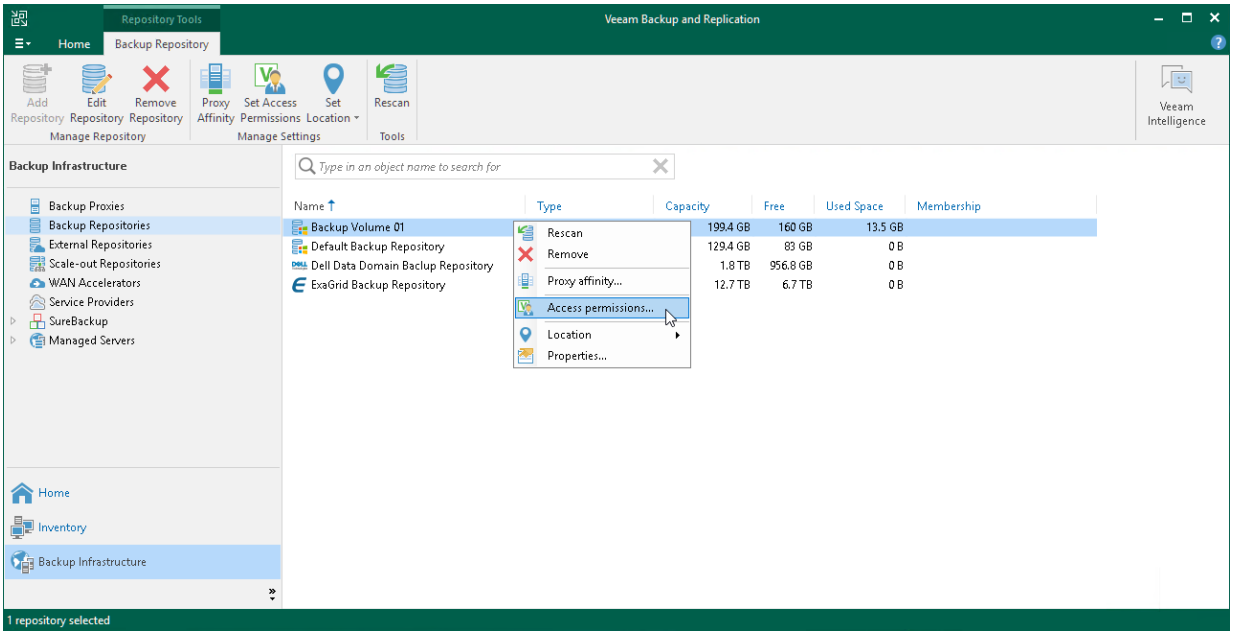


4. In the **Standalone applications** window, specify to whom you want to grant access permissions on this backup repository:
   - **Allow to everyone** — select this option if you want all users to be able to store backups in this backup repository. Setting access permissions to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). Note, however, this scenario is recommended for demo environments only.

- **Allow to the following accounts or groups only** – select this option if you want only specific users to be able to store backups in this backup repository. Click **Add** to add the necessary users and groups to the list.

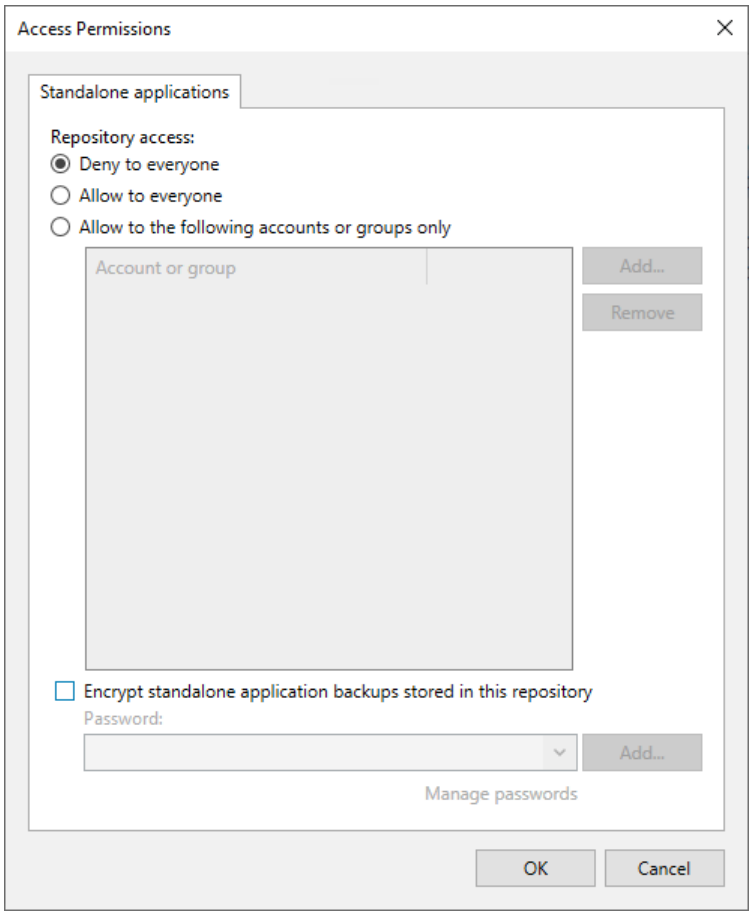


# Managing Permissions for S3 Compatible Object Storage

If you plan to use S3 compatible object storage as an object storage repository, you must set up access permissions to the object storage. These permissions are used if you keep in object storage repositories backups created by Veeam Agent or by Veeam Cloud Connect tenant. For more information, see Backup to Object Storage in the Veeam Agent Management Guide and Backup to Object Storage in the Veeam Cloud Connect Guide.

To manage permissions for S3 compatible object storage, perform the following:

1. In Veeam Backup & Replication, open the **Backup Infrastructure** view.
2. In the inventory pane, click the **Backup Repositories** node.
3. In the working area, select the necessary S3 compatible object storage and click **Set Access Permissions** on the ribbon or right-click the backup repository and select **Access permissions**. If you do not see the **Set Access Permissions** button on the ribbon or the **Access permissions** command is not available in the shortcut menu, press and hold the [Ctrl] key, right-click the backup repository and select **Access permissions**.

4. On the **Security** tab, specify how Veeam Agent or a tenat will access an S3 compatible object storage repository:
   - **Agents share credentials to object storage repository** – select this option if you use the direct connection mode to access the S3 compatible object storage repository. In this case, Veeam Agent will use credentials that you specified when added the S3 compatible object storage repository to the backup infrastructure.

     > **IMPORTANT**
     >
     > This option is not secure since Veeam Backup & Replication will have access permissions on the bucket where you keep your folders with backups.

   - **Provided by the backup server** – select this option if you want to access the S3 compatible object storage repository through a gateway server. By default, Veeam Backup & Replication uses this option.

     > **NOTE**
     >
     > If you use the direct mode to connect to the S3 compatible object storage repository, the access permissions will override the repository settings and the object storage will be accessed indirectly. In this case, Veeam Backup & Replication will use a mount server to access the S3 compatible object storage repository.

   - **Provided by IAM/STS object storage capabilities** – select this option if you use the direct connection mode to access the S3 compatible object storage repository. In this case, Veeam Backup & Replication will create service tokens that Veeam Agent or a tenant will use to access the S3 compatible object storage repository.

     To specify credentials, do the following:

     i. In the **Identity and access management (IAM) endpoint** field, specify the endpoint of your S3 compatible object storage repository.

ii. In the **Security token service (STS) endpoint** field, specify the security token.

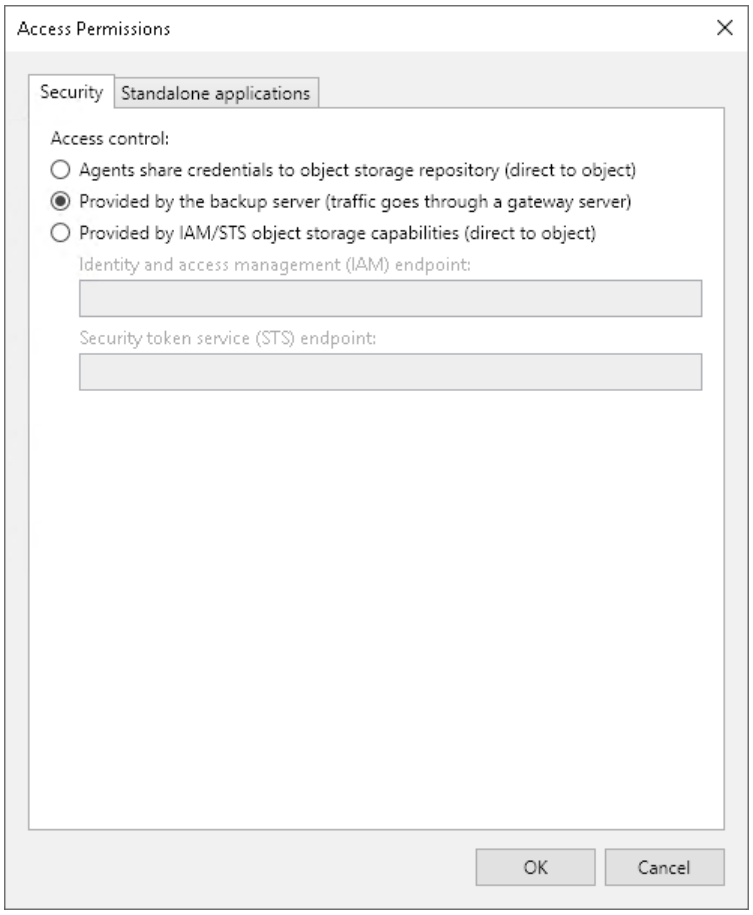


# Encrypting Standalone Application Backups in Backup Repositories

At the storage encryption level, you can encrypt backup files created by the following standalone applications:

- Veeam Agents operating in the standalone mode
- Veeam Backup for Nutanix AHV
- Veeam Backup for Proxmox VE
- Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization
- Veeam Kasten for Kubernetes

Veeam Backup & Replication encrypts files at the standalone application side in the following way:

| Backup Repository | Encryption |
|---|---|
| Regular backup repository:<br>• Microsoft Windows/Linux-based repository<br>• NFS file share<br>• SMB (CIFS) file share<br>• Object storage repository<br>• Deduplicating storage appliance | Backup files are encrypted before uploading to the backup repository. To encrypt backup files, a standalone application uses encryption password specified in the **Access Permissions** settings. For more details, see Enabling Encryption for Standalone Application Backups. |
| Scale-out repository | Backup files are encrypted before uploading to the scale-out backup repository.<br><br>To encrypt backup files before uploading to a performance tier, a standalone application uses encryption password specified in the **Access Permissions** settings for the scale-out backup repository. For more details, see Enabling Encryption for Standalone Application Backups.<br><br>To encrypt backup files before uploading to a capacity or archive tier, a standalone application uses encryption password specified in the capacity tier settings. For more details, see Adding Scale-Out Backup Repositories.<br><br>If you enable encryption both in the capacity tier settings and in the **Access Permissions** settings for the scale-out backup repository, backup files will be encrypted twice before uploading to a capacity or archive tier. |

# Requirements and Limitations

Consider the following:

- For backup files created by Veeam Agents operating in the standalone mode, data encryption with KMS keys is not supported. For more information, see Key Management System Keys.
- To encrypt backup files created by Veeam Agents operating in the managed mode, you must configure encryption in the backup job settings. For example, to learn how to encrypt backup files created by managed Veeam Agent for Microsoft Windows, see the Storage Settings section in the Veeam Agent Management Guide.
- To encrypt backup files created by other platforms and applications, you must configure encryption in the backup job settings. For more information, see Job Encryption.
- Data encryption has a negative effect on the deduplication ratio. If you want to achieve a higher deduplication ratio, use only deduplicating storage appliance encryption. For more details, see Data Encryption and Deduplication.

# Enabling Encryption for Standalone Application Backups

To encrypt standalone application backups, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Repositories** or **Scale-out Repositories** node.
3. In the working area, select the necessary backup repository and click **Set Access Permissions** on the ribbon or right-click the backup repository and select **Access permissions**. If you do not see the **Set Access Permissions** button on the ribbon or the **Access permissions** command is not available in the shortcut menu, press and hold the [Ctrl] key, right-click the backup repository and select **Access permissions**.

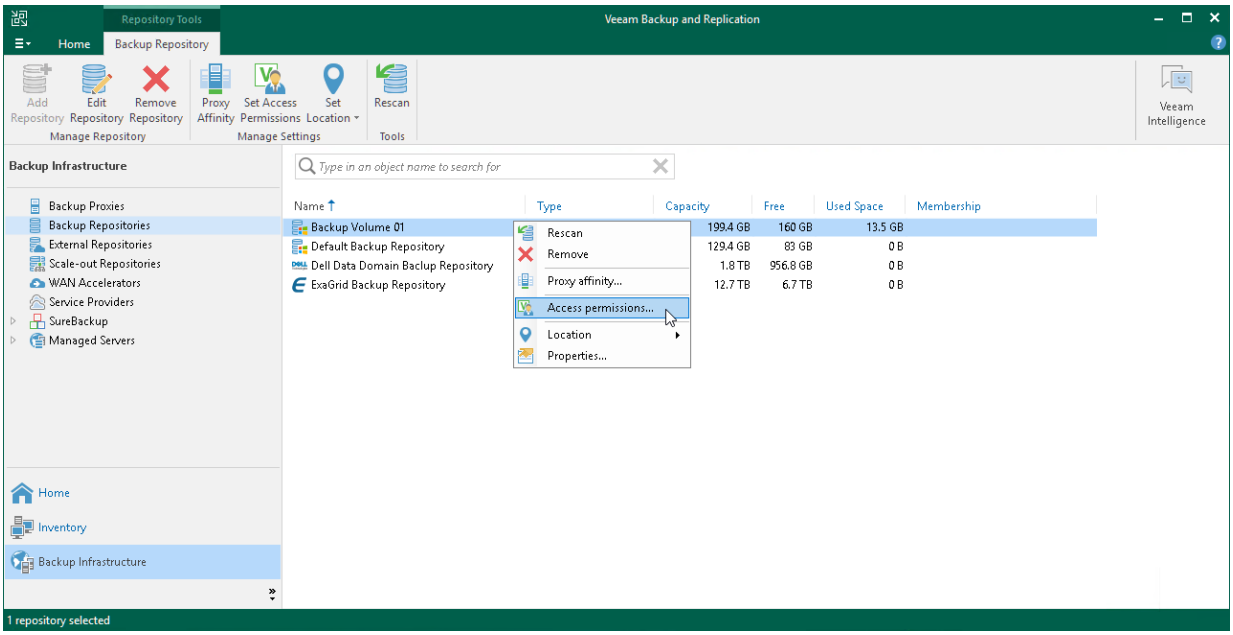

4. In the **Standalone applications** window, select the **Encrypt standalone application backups stored in this repository** check box and choose the necessary password from the field under the check box. If you have not specified a password beforehand, click **Add** on the right or the **Manage passwords** link to add a new password.

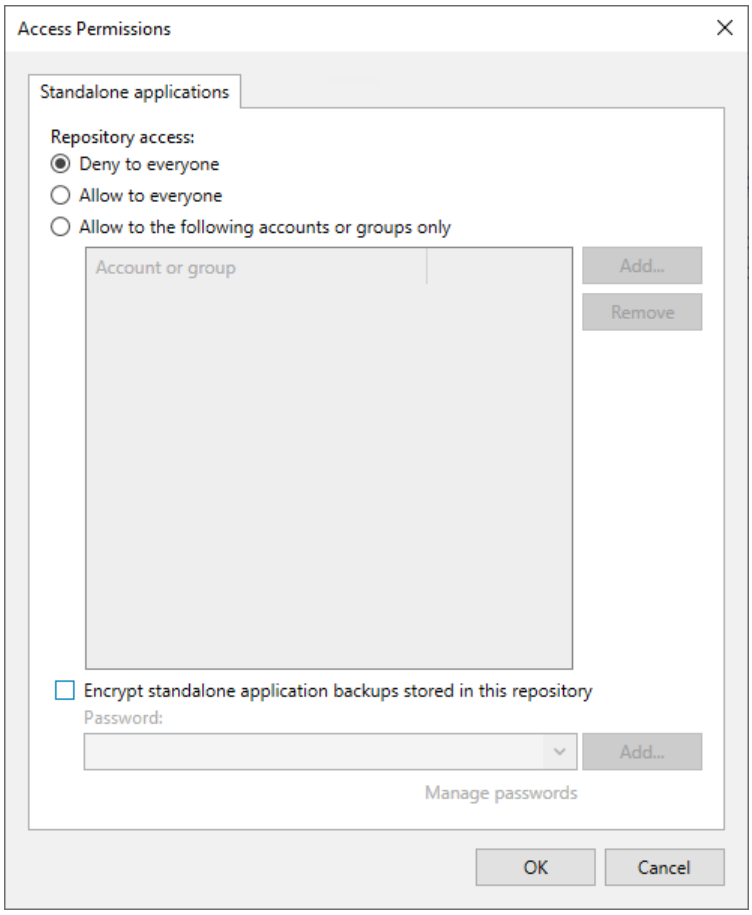


# Rescanning Backup Repositories

You can rescan a backup repository configured in the backup infrastructure. Backup repository rescan may be required, for example, if you have archived backups from a backup repository to tape and deleted backup files in the backup repository. Or you have copied backups to the backup repository manually and want to work with them in Veeam Backup & Replication.

During the rescan operation, Veeam Backup & Replication gathers information about backups that are currently available in the backup repository and updates the list of backups in the configuration database. After the rescan operation, backups that were not in this configuration database will be shown on the **Home** view in the **Backups** > **Disk (Imported)** node. If backups are encrypted, they will be shown in the **Backups** > **Disk (Encrypted)** node.

> **IMPORTANT**
>
> Consider the following:
>
> - It is recommended that you stop or disable all jobs before performing the rescan. Veeam Backup & Replication skips from scanning backups created by active jobs.
> - Veeam Backup & Replication will not be able to import backups automatically while performing a rescan if VBM files are not available. In this case you will have to import backups manually using the VBK files. For more information, see Importing Backups Manually.

To rescan a backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Repositories** node.
3. In the working area, select the backup repository and click **Rescan Repository** on the ribbon or right-click the backup repository and select **Rescan repository**.

> **TIP**
>
> By default, Veeam Backup & Replication skips from rescan files whose paths contain `$RECYCLE.BIN` or `.Trash-<digits>`. You can exclude other file paths with a registry value. For more information, contact Veeam Customer Support.

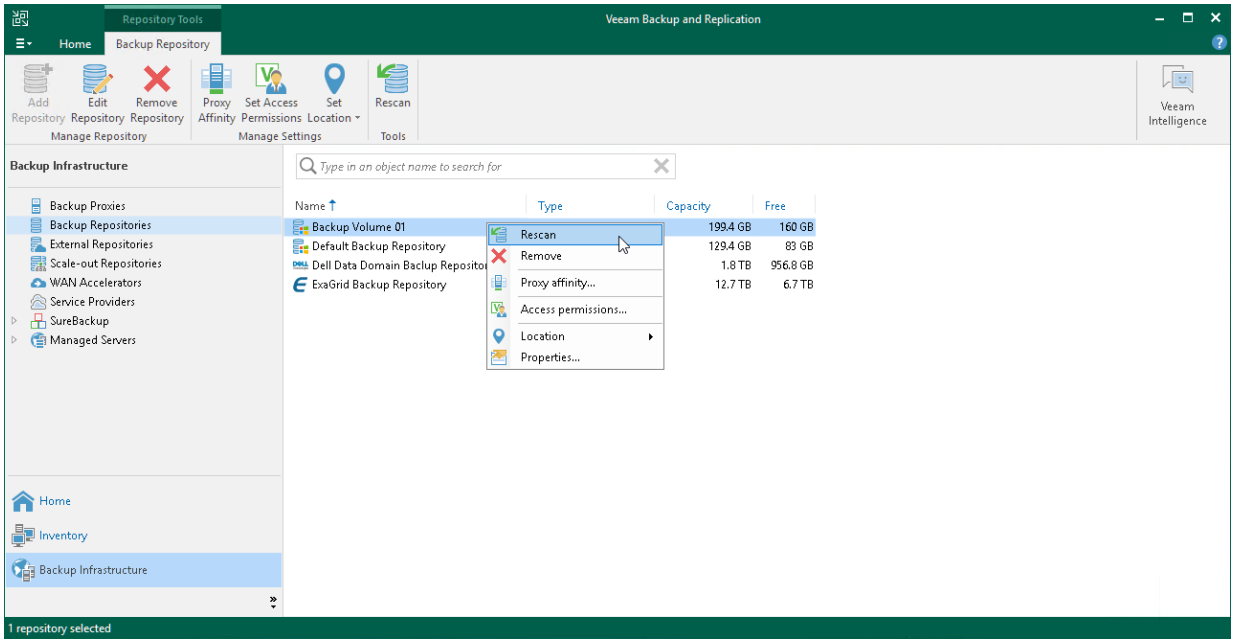


# Removing Backup Repositories

You can permanently remove a backup repository from the backup infrastructure. When you remove a backup repository, Veeam Backup & Replication unassigns the backup repository role from the server and this server is no longer used as a backup repository. The actual server remains in the backup infrastructure.

Veeam Backup & Replication does not remove backup files and other data stored in the backup repository. You can re-connect the backup repository at any time and import backups from this backup repository to Veeam Backup & Replication.

You cannot remove a backup repository that is used by any job. To remove such backup repository, you first need to delete a reference to this backup repository in the job settings.

To remove a backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Repositories** node.

3. In the working area, select the backup repository and click **Remove Repository** on the ribbon or right-click the backup repository and select **Remove**.

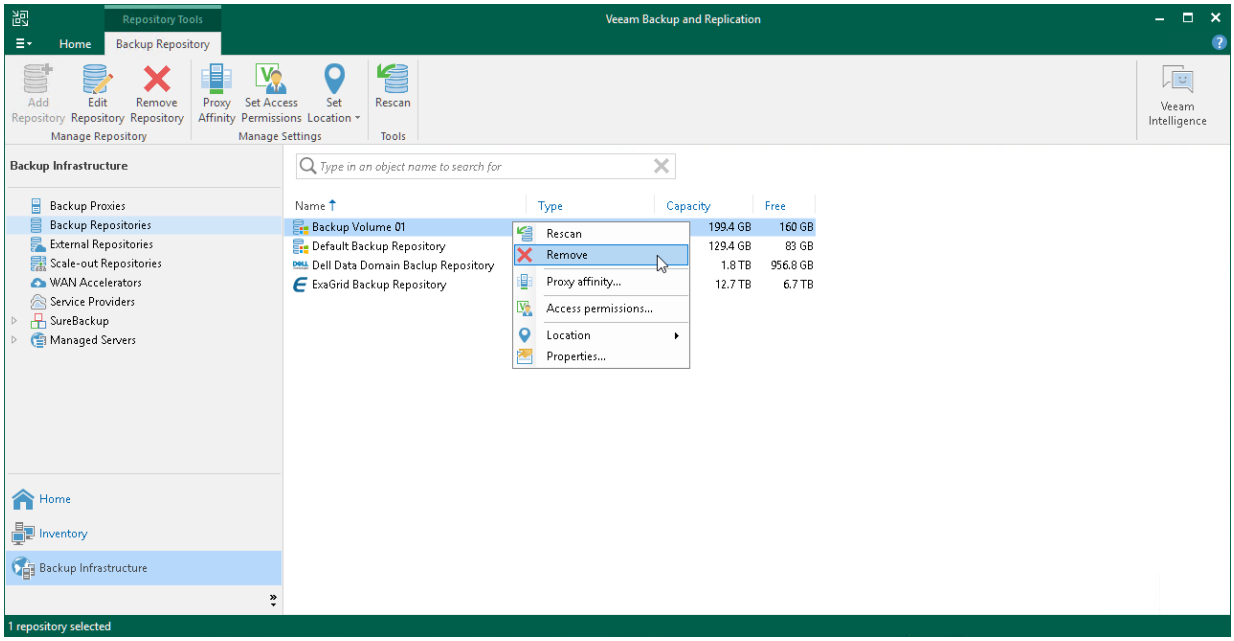


# Fast Clone

Fast Clone is the Veeam Backup & Replication technology that helps create quick file copies. Fast Clone increases the speed of synthetic backup creation and transformation, reduces disk space requirements and decreases the load on storage devices.

With this technology, Veeam Backup & Replication references existing data blocks on volumes instead of copying data blocks between files. Data blocks are copied only when files are modified.

Veeam Backup & Replication supports Fast Clone for the following types of backup repositories:

- Linux server
- Microsoft Windows server
- SMB share
- Dell Data Domain
- ExaGrid
- Fujitsu ETERNUS CS800
- HPE StoreOnce
- Infinidat InfiniGuard
- Quantum DXi

Depending on the repository type, Fast Clone uses different technologies and has different requirements and limitations. For more information, see Fast Clone for Deduplicating Storage Appliances, Fast Clone for Linux Repositories and Fast Clone for Microsoft Windows and SMB Repositories.

Fast Clone requires that the starting and ending file offsets are aligned to cluster boundaries. For this reason, Veeam Backup & Replication automatically enables the **Align backup file data blocks** option for backup repositories that support Fast Clone.

Veeam Backup & Replication uses Fast Clone for the following operations:

- In **backup jobs**: to merge backup files, create synthetic full backups (including GFS backups), transform reverse incremental backups and compact full backup files.
- In **backup copy jobs**: to merge backup files, create GFS backups (synthetic method) and compact full backup files.
- To migrate (evacuate) backups between performance extents.
- To migrate (evacuate) Veeam Cloud Connect tenant backups between performance extents.
- To rebalance SOBR extents.
- To download backups from capacity extents.
- To move backups to another repository.
- To copy backups to another repository.
- To export backups into standalone .VBK files.
- To upgrade the backup chain format for backup copies created in the periodic copy mode.

# Fast Clone for Deduplicating Storage Appliances

For deduplicating storage appliances (ExaGrid, Quantum DXi and so on), Fast Clone is based on the reflink technology and must be supported by vendors.

# Fast Clone for Linux Repositories

For Linux backup repositories, Fast Clone is based on the reflink technology.

## Requirements for Linux Repositories

To use Fast Clone, Veeam Backup & Replication requires that Linux backup repositories meet the following conditions:

- File system is XFS.
- The Linux distribution is supported. For more information, see Backup Repository (XFS integration).
- The Linux kernel supports reflinks.
- Cyclic redundancy check (CRC) is enabled.
- The minimum supported data block size is 1 KB. The maximum supported block size is 4KB.

## Configuring a Linux Repository

To configure a Linux backup repository for work with Fast Clone, perform the following steps:

1. Format the disk where backups will be stored using the following XFS volume format string:

```
mkfs.xfs -b size=4096 -m reflink=1,crc=1 /dev/sda1
```

Consider that:

- **size=4096** sets file system block size to 4096 bytes.
- **reflink=1** enables reflinking for the XFS instance (disabled by default).
- **crc=1** enables checksums, required for reflink=1 (enabled by default).

2. Mount the disk to a specific directory:

```
mkdir /backups
mount /dev/sda1 /backups
```

3. Run the following command to ensure that the disk is properly configured:

```
df -hT
```

4. To permanently mount the disk, add the entry to the `/etc/fstab` configuration file. It is recommended to use the UUID to specify the disk.

```
# <file system>    <mount point>       <type>   <options>   <dump> <pass>
UUID=<UUID>        /backups            xfs      defaults    0       0
```

To get the UUID, run the following command:

```
blkid /dev/sda1
```

### Limitations

If you have manually moved backup chains to a Linux backup repository with Fast Clone support, you must create active full backups for these chains after the move to activate Fast Clone. You can also schedule the backup file compact operation instead of active full backup. Note that creation of active full backups is not required if you use the Veeam copy or move features.

## Fast Clone for Microsoft Windows and SMB Repositories

For Microsoft Windows and SMB backup repositories, Fast Clone is based on ReFS block cloning technology of Microsoft. For more information on block cloning, see Microsoft Docs.

By default, Veeam Backup & Replication uses Fast Clone for all Microsoft Windows and SMB backup repositories that meet the requirements. You can disable this option with a registry value. For more information, contact Veeam Customer Support.

### Requirements for Microsoft Windows and SMB Repositories

#### Microsoft Windows Backup Repository

To use Fast Clone, Veeam Backup & Replication requires that Microsoft Windows backup repositories meet the following conditions:

- OS is Microsoft Windows Server 2016 (or later) or Microsoft Windows 10 Pro for Workstations (or later).

- File system is ReFS 3.1 (or later).

> **NOTE**
>
> All ReFS supported configurations must use Windows Server Catalog certified hardware. For other requirements, limitations and known issues, see this Veeam KB article.

**Shared Folder Backup Repository**

To use Fast Clone, Veeam Backup & Replication requires that SMB backup repositories support FSCTL_DUPLICATE_EXTENTS_TO_FILE and FSCTL_SET_INTEGRITY_INFORMATION. SMB shares configured on Microsoft Windows machines must also support the SMB 3.1.1 protocol and the ReFS 3.1 (or later) file system.

Depending on the type of the performed job, Veeam Backup & Replication also imposes the following requirements on backup infrastructure components.

| Type of Job | Requirements to backup infrastructure components |
|---|---|
| **Backup job** | Protocol: SMB 3.1.1<br><br>OS: Microsoft Windows Server 2016 (or later) or Microsoft Windows 10 Pro for Workstations (or later) on the following backup infrastructure components:<br>• If a gateway is selected manually: Gateway server.<br>• If a gateway is selected automatically: Mount server associated with the backup repository, or backup server. For reverse incremental backup chains, Microsoft Windows Server 2016 (or later) or Windows 10 Pro for Workstations (or later) must additionally be installed on backup proxies assigned for the job. |
| **Backup copy job** | Protocol: SMB 3.1.1<br><br>OS: Microsoft Windows Server 2016 (or later) or Microsoft Windows 10 Pro for Workstations (or later) on the following backup infrastructure components:<br>• If a gateway is selected manually: Gateway server.<br>• If a gateway is selected automatically:<br>[For direct data transport path] Mount server associated with the backup repository, or backup server.<br>[For data transport path over WAN accelerators] Microsoft Windows Server 2016 (or later) or Microsoft Windows 10 Pro for Workstations (or later) on the target WAN accelerator. |

## Limitations

The following limitations apply when Veeam Backup & Replication uses Fast Clone for Microsoft Windows or SMB backup repositories:

- Veeam Backup & Replication does not use Fast Clone for backup repositories configured with Veeam Backup & Replication 9.5 or an earlier version. After upgrade, such backup repositories will work as backup repositories without Fast Clone support. To leverage Fast Clone, edit settings of such backup repositories and complete the **Edit Backup Repository** wizard without changing settings.

- After you have enabled Fast Clone for existing repositories as described in the previous paragraph or have manually moved backup chains to backup repositories with Fast Clone support, you must create active full backups for backup chains stored in / moved to the repositories to activate Fast Clone. You can also schedule the backup file compact operation instead of performing active full backup.
- Due to Microsoft limitations, the source and destination files must be stored on the same ReFS volume. For more information, see Restrictions and Remarks at Microsoft Docs.

  If you add a backup repository with Fast Clone support as an extent to a scale-out backup repository, make sure that you enable the Data Locality placement policy for this scale-out backup repository. If backup files are stored on different extents, Fast Clone will not be used.
- Veeam Backup & Replication automatically aligns data blocks at a 4KB or 64 KB block boundary depending on the volume configuration or SMB share used storage.

  We recommend that you use ReFS volume formatted with 64 KB cluster size to provide better performance with large data volumes.
- When you copy data from a ReFS volume to another location, the file system downloads cloned data blocks. For this reason, copied data occupy more space in the target location than it used to occupy in the source location. This can happen, for example, if you evacuate an extent that supports block cloning from a scale-out backup repository and migrate VM backup data to another extent: copied data will require more space than it originally took.
- If you plan to assign the role of a backup repository to Microsoft Windows Server 2016 version 1709 (or later) or Microsoft Windows 10 Pro for Workstations (or later), consider the following limitations:
  - Fast Clone and Windows data deduplication cannot be used simultaneously. Thus, if you target a backup job to a repository supporting Fast Clone and enable Windows data deduplication, the Fast Clone technology will not be used for this job.
  - If you target a backup job to an SMB (CIFS) ReFS repository and enable Windows data deduplication, the job will fail. Veeam Backup & Replication does not support such scenario.

# Proxy Affinity

By default, Veeam Backup & Replication assigns backup proxies and repositories for jobs or tasks independently of each other. If you need to bind backup proxies to specific backup repositories and use them together, you can define proxy affinity settings. Proxy affinity determines what backup proxies are eligible to access a specific backup repository and read/write data from/to this backup repository.

Proxy affinity lets you control assignment of resources in the backup infrastructure and reduce administration overhead. For example, in case of a geographically distributed infrastructure, you can restrict a backup repository in the local site from communicating with backup proxies in a remote site. Or you can configure proxy affinity rules based on a connection speed between backup proxies and backup repositories.

Proxy affinity settings are specified at the level of a backup repository. By default, Veeam Backup & Replication lets all backup proxies in the backup infrastructure access the backup repository. Using proxy affinity settings, you can define a list of backup proxies that can access this backup repository.

Proxy affinity can be set up for the following types of backup repositories:

- Backup repositories
- Scale-out backup repositories
- Cloud repositories (proxy affinity settings are configured on the tenant side)

Proxy affinity rules are applied to the following types of jobs and tasks that engage backup proxies and repositories:

- Backup jobs, including VMware Cloud Director backup and backup jobs from storage snapshots on primary and target storage arrays
- VeeamZIP
- VM copy
- Entire VM restore
- Hard disk restore

Proxy affinity rules are not applied to replication jobs.

Proxy affinity rules are not restrictive. You can think of affinity rules as a priority list. If backup proxies from the proxy affinity list cannot be used for some reason, for example, these backup proxies are inaccessible, Veeam Backup & Replication automatically fails over to the regular processing mode. It picks the most appropriate backup proxy from the list of proxies selected for the job or task and records this action in the job statistics.

When you target a job at a backup repository for which proxy affinity settings are configured, you must make sure that you assign a backup proxy from the proxy affinity list for job or task processing. If you assign a backup proxy that is not bound to this backup repository, Veeam Backup & Replication will display a warning. For job processing, Veeam Backup & Replication will use the backup proxy that you define in the job settings, which may result in degraded job performance.

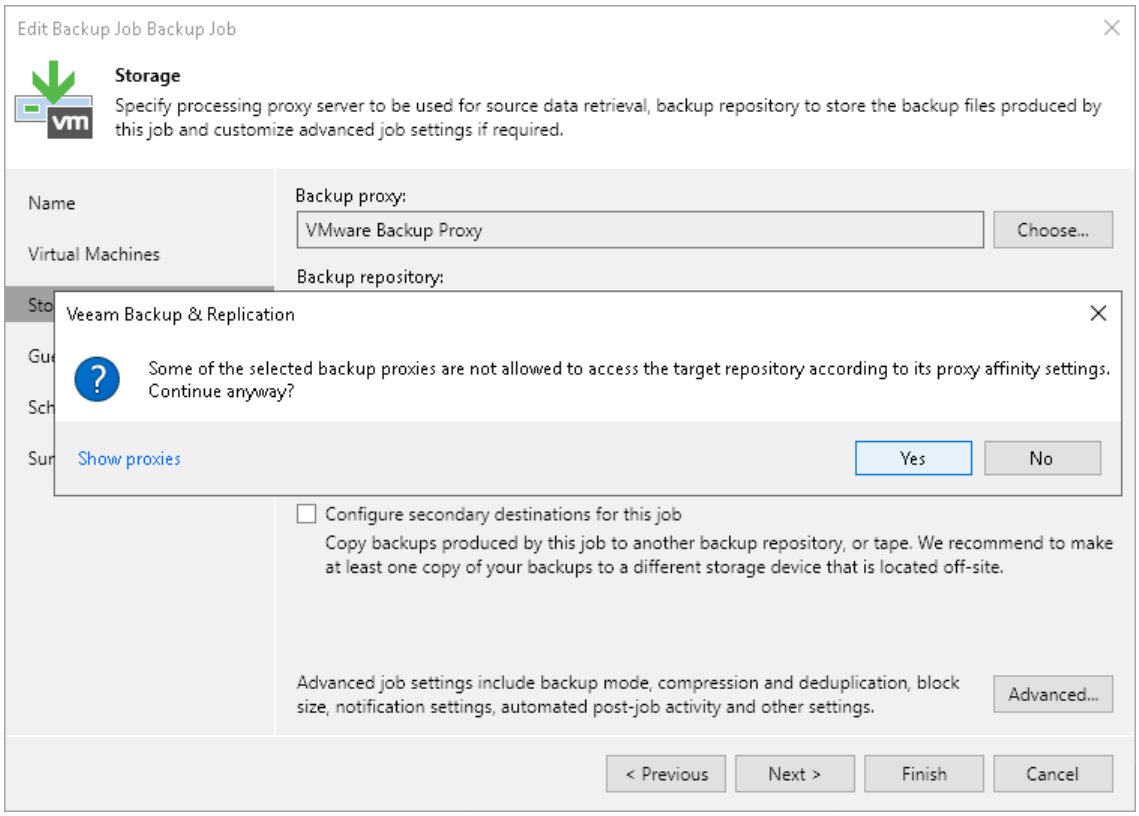


# Proxy Affinity for Scale-Out Backup Repositories

In case of a scale-out backup repository, you can configure proxy affinity settings at the extent level. Proxy affinity settings cannot be configured at the scale-out backup repository level.

Extent selection rules have a higher priority than proxy affinity rules. Veeam Backup & Replication first selects an extent and then picks a backup proxy according to the proxy affinity rules specified for this extent.

For example, you have 2 backup proxies: *Backup Proxy 1* and *Backup Proxy 2*. You create a backup job and target it at a scale-out backup repository configured in the following way:

- Scale-out backup repository policy is set to Data Locality.
- Scale-out backup repository has 2 extents: *Extent 1* has 100 GB of free space and is bound to *Backup Proxy 1*; *Extent 2* has 1 TB of free space and is bound to *Backup Proxy 2*.

In the backup job settings, you define that *Backup Proxy 1* must be used for job processing.

When you run the backup job, Veeam Backup & Replication will store backup files to *Extent 2* since it has more free space. For job processing, it will pick *Backup Proxy 1* and will display a message in the job statistics that requirements of proxy affinity rules cannot be met.

In case of restore from a scale-out backup repository, backup files may be located on different extents. In this case, Veeam Backup & Replication picks a backup proxy according to the following priority rules (starting from the most preferable one):

1. Backup proxy is added to the affinity list for all extents.
2. Backup proxy is added to the affinity list for the extent where the full backup file is stored.
3. Backup proxy is added to the affinity list for at least one extent.

## Specifying Proxy Affinity Settings

For every backup repository, you can configure proxy affinity settings — define a list of backup proxies that can work with this backup repository. Proxy affinity binds backup proxies to specific backup repositories. When transporting data to/from the backup repository, Veeam Backup & Replication tries to pick a VMware backup proxy from the proxy affinity list whenever it is possible.

To configure a proxy affinity list for a backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Repositories**.
3. In the working area, select the backup repository and click **Proxy Affinity** on the ribbon or right-click the backup repository and select **Proxy affinity**.

4. In the **Proxy Affinity** window, select **Prefer the following backup proxies for this repository** and select check boxes next to the backup proxies that you want to bind to the backup repository.

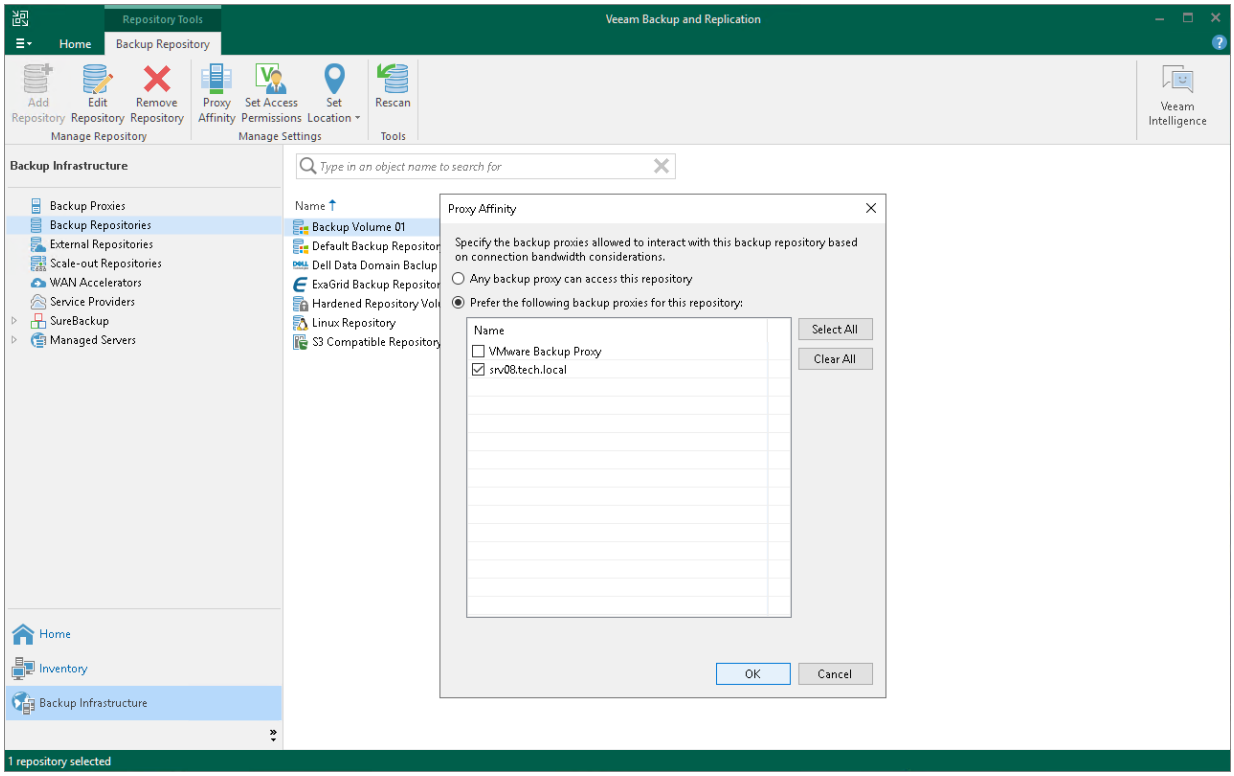

# Scale-Out Backup Repositories

A scale-out backup repository is a repository system with horizontal scaling support for multi-tier storage of data. A scale-out backup repository consists of one or more backup repositories or object storage repositories called performance tier, and can be expanded with object storage repositories for long-term and archive storage: capacity tier and archive tier. All the storage devices and systems inside the scale-out backup repository are joined into a system, with their capacities summarized.

**NOTE**

Consider the following:

- Scale-out backup repository is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.
- If you configure a scale-out backup repository and then downgrade to the Standard license, you will not be able to run jobs targeted at the scale-out backup repository. However, you will be able to perform restore from the scale-out backup repository.

This feature provides the following benefits:

- It provides a convenient way of managing the backup storage.
- A scale-out backup repository can be expanded at any moment: if the extents of your scale-out backup repository run out of space, you can add a new extent to the existing scale-out backup repository. For example, if backup data grows and the backup repository reaches the storage limit, you can add a new storage system to the scale-out backup repository. The free space on this storage system will be added to the capacity of the scale-out backup repository. As a result, you will not have to move backups to a backup repository of a larger size.
- It supports any backup target supported by Veeam: Windows or Linux servers with local or DAS storage, network shares, deduplicating storage appliances, object storage repositories. All the features of any storage device or system are preserved.
- It allows you to set up granular performance policy. For more information, see Backup File Placement for Performance Tier.
- It provides practically unlimited cloud-based storage capacity: you can instruct Veeam Backup & Replication to offload data from extents to the cloud for long-term storage. For details, see Capacity Tier.

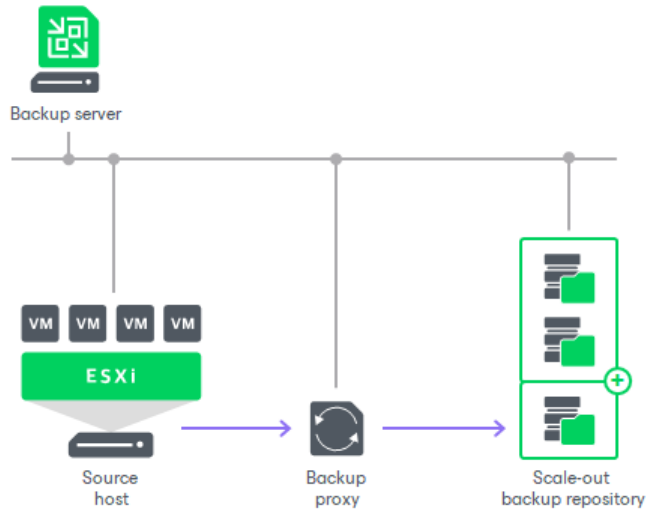

A scale-out backup repository can comprise different tiers, or logical levels of storage.

- Performance tier is the level used for the fast access to the data. It consists of one or more backup repositories or object storage repositories called performance extents.

  For more information, see Performance Tier.

- Capacity tier is an additional level for storing data that needs to be accessed less frequently. However, you still can restore your data directly from it. The capacity tier consists of cloud-based or on-premises object storage repositories called capacity extent.

  For more information, see Capacity Tier.

- Archive tier is an additional level for archive storage of infrequently accessed data. Applicable data can be transported either from the performance or capacity tier. For restore from the archive tier, data must undergo preparation process.

  For more information, see Archive Tier.

You can use a scale-out backup repository with almost all types of jobs and tasks that Veeam Backup & Replication runs. For a list of unsupported scenarios, see Limitations for Scale-Out Backup Repositories.

Backup files stored in the scale-out repository can be used for all types of restores, replication from backup and backup copy jobs. You can verify such backups with SureBackup jobs. The scale-out backup repository can be used as a staging backup repository for restore from tape media. Files restored from the tape media are placed to the extents according to data placement policy configured for the scale-out backup repository. For more information, see Backup File Placement for Performance Tier.

To deploy a scale-out backup repository, you must configure one or more backup repositories or object storage repositories and add them to a scale-out backup repository as extents. You can use the following types of repositories:

- Microsoft Windows backup repositories
- Linux backup repositories
- SMB (CIFS) shared folders
- NFS shared folders
- Deduplicating storage appliances
- Object Storage Repositories

# Limitations for Scale-Out Backup Repositories

## Limitations for Jobs

You cannot use a scale-out backup repository as a target for the following types of jobs:

- Configuration backup job.
- Replication jobs (including replica seeding).
- VM copy jobs.
- File backup jobs (in case your scale-out backup repository consists of object storage extents).
- Object storage backup job (in case your scale-out backup repository consists of object storage extents).
- Scale-out backup repositories are not supported as a backup destination for cloud machines.
- Veeam Agent for Microsoft Windows or Veeam Agent for Linux backup jobs that back up cloud machines (that is, Amazon EC2 instances or Microsoft Azure virtual machines) in managed mode. For more information, see Veeam Agent Management Guide.

## General Limitations

Scale-out backup repositories have the following limitations:

- You cannot add a backup repository as an extent to the scale-out backup repository if any job of an unsupported type is targeted at this backup repository or if the backup repository contains data produced by jobs of unsupported types (for example, replica metadata). To add such backup repository as an extent, you must first target unsupported jobs to another backup repository and remove the job data from the backup repository in question.
- Scale-out backup repositories do not support rotated drives. If you enable the **This repository is backed by rotated hard drives** setting on an extent, Veeam Backup & Replication will ignore this setting and will work with such repository as with a standard extent.
- Backup repositories that you plan to add as extents to the scale-out backup repository must be added to the backup infrastructure. For more information, see Backup Repositories.
- If a backup repository is added as an extent to the scale-out backup repository, you cannot use it as a regular backup repository. You cannot target jobs to this backup repository. Instead, you have to target jobs to the configured scale-out backup repository.
- You cannot add a scale-out backup repository as an extent to another scale-out backup repository.
- You cannot add a backup repository as an extent if this backup repository is already added as an extent to another scale-out backup repository.
- You cannot use the same object storage repository as an extent of two or more different scale-out backup repositories at the same time.
- You cannot add a backup repository as an extent if this backup repository is already used as a backup destination by VMware Cloud Director organizations.
- You cannot add a backup repository in which some activity is being performed (for example, a backup job or restore task) as an extent to the scale-out backup repository.

- You cannot assign the role of a cache repository for file backup or object storage backup to a scale-out backup repository and its extents. To learn more about the cache repository, see the Backup Infrastructure for Unstructured Data Backup section.
- You can add only one AWS Snowball Edge storage and Azure Data Box device to a scale-out backup repository.
- Data located in object storage repositories is organized into a separate backup chain for every machine in a job.
- We do not recommend adding more than 16 active extents to one tier within one scale-out backup repository, otherwise you may have performance issues.
- For some S3 compatible object storage repositories, we do not recommend that you use more than one bucket for each scale-out backup repository. If you use multiple folders within one bucket for several scale-out backup repositories, it will slow down data processing since the metadata generated by S3 compatible is handled for each bucket. To check requirements for your object storage, contact your S3 compatible vendor.
- If you want to use S3-compatible repositories by different providers in the same tier, one of this repository must be in the sealed mode. For example, if you want to use both Wasabi Cloud object storage and IBM Cloud object storage within the same tier, Wasabi Cloud object storage must be in sealed mode.
- If you use immutability and have several extents in your performance tier, you must enable it for all active extents within this tier. If you want to use a mixed immutability configuration (with some extents immutable and some mutable), you must switch either all mutable or all immutable extents to the sealed mode.
- If you apply the **Forget** or **Remove from disk** options to a missing restore point in a scale-out backup repository, the backup file associated with the missing restore point will be deleted from capacity tier and archive tier on the next offload and archiving job run.
- If you want to use the Extract utility to work with backup files located on any of the extents of your scale-out backup repository, make sure that incremental and full backup files are located on the same extent.
- You can use the Veeam Backup Validator utility only for backups stored in the performance tier which consists of backup repositories (except object storage repositories). Make sure that incremental and full backup files are located on the same extent.
- If you want to add an object storage repository added as a capacity or performance extent to another backup server, you must switch this object storage repository to the Maintenance mode on the initial server. Note that the secondary server will request the ownership over the object storage repository. After you change the ownership, the backup jobs (except for the offload job to capacity tier) will fail on the primary backup server.
- To let Veeam Backup & Replication automatically import backups during rescan of a scale-out backup repository, names of VBM files and paths to VBM files (starting from the backup repository root to VBM files, not including the root itself) must contain only allowed characters:
  - Alphanumeric characters: a-zA-Z0-9
  - Special characters: _-.+=@^

  Names of VBM file and paths to VBM files must not contain spaces.

  If a name of a VBM file or a path to a VBM file contains prohibited characters, Veeam Backup & Replication will fail to import such backup during rescan of the scale-out backup repository. To import such backup, you can replace prohibited characters with the underscore character, for example: `C:\My Repository\Backup_Job\Backup_Job.vbm`. You do not need to rename the actual backup files.
- Veeam Backup & Replication does not split one backup file across multiple extents.

- If a repository is used as a cloud repository, you cannot add it as an extent of a scale-out backup repository. For more information, see Veeam Cloud Connect.
- [For Nutanix AHV VM backups] Due to specifics of backup jobs for AHV VMs, Veeam Backup for Nutanix AHV always creates a separate backup chain for each VM added to a backup job. Thus, even if you clear the **Use per-machine backup files** check box in the advanced settings of a scale-out backup repository, backups of multiple AHV VMs are not stored in a single backup file.

# Limitations for Specific Tier

Keep in mind that every tier of a scale-out backup repository has its own limitations. For more information on the limitations for every type of tier, check the following sections:

- Limitations for Performance Tier
- Limitations for Capacity Tier
- Limitations for Archive Tier

# License-Based Limitations

Enterprise edition of Veeam Backup & Replication has the following limitations for scale-out backup repositories:

- You can create 2 scale-out backup repositories.
- For each scale-out backup repository, you can have either backup repositories or object storage repositories added as a performance extent or a capacity extent. You can have 3 active, and 1 inactive (that is, put to the Maintenance mode) performance or capacity extents. You can put an active extent to the Maintenance mode if it has no free space, and you want to evacuate backup data from it.
- If you add 4 performance extents and do not put any of them to the Maintenance mode, the jobs targeted at the scale-out backup repository will fail.

> **NOTE**
>
> Veeam Universal License and Enterprise Plus of Veeam Backup & Replication editions have no limitations on the number of scale-out backup repositories or performance extents and capacity extents.

# Limitations for Veeam Solutions

For more information on limitations for a specific Veeam solution that utilizes scale-out backup repositories functionality, see the following sections of the necessary guide:

- Veeam Plug-in for SAP HANA — to check limitations for an SAP-certified backup and recovery solution that allows you to back up and restore SAP HANA databases.
- Veeam Plug-in for Oracle RMAN — to check limitations for an Oracle-certified backup and recovery solution that allows you to back up and restore Oracle databases.
- Veeam Plug-in for SAP on Oracle — to check limitations for an SAP-certified backup and recovery solution that allows you to back up and restore Oracle databases to which an SAP application is connected.
- Veeam Plug-in for Microsoft SQL Server — to check limitations for a backup and recovery solution that allows you to back up and restore Microsoft SQL Server databases.

- Veeam Plug-in for IBM Db2 — to check limitations for a backup and recovery solution that allows you to back up and restore IBM Db2 databases.
- Veeam Cloud Connect Guide — to check limitations for service providers and tenants.
- Veeam Agent Management Guide — to check limitations for Veeam Agent management solution.
- Veeam Agent for Microsoft Windows — to check limitations for data protection and disaster recovery solution for physical and virtual machines running Windows-based operating systems.

# Immutability for Scale-Out Backup Repositories

Veeam Backup & Replication allows you to prohibit deletion of data from object storage repositories added as the extents of the scale-out backup repository by making that data temporarily immutable. It is done for increased security: immutability protects your data against loss as a result of attacks, malware activity or other injurious actions.

You can enable the immutability feature for any tier of scale-out backup repository.

To learn how immutability works with the performance tier of the scale-out backup repository, see Immutability for Performance Tier.

To learn how immutability works with the capacity tier of the scale-out backup repository, see Immutability for Capacity Tier.

To learn how immutability works with the archive tier of the scale-out backup repository, see Immutability for Archive Tier.

> **NOTE**
>
> Before you use an object storage repository as an immutable extent, you must configure immutability for this object storage beforehand. For more information, see Enabling immutability for object storage repositories.

## GFS Backups Immutability Period

The immutability period for backups with GFS flags depends on multiple settings of a scale-out backup repository. The following table shows how the immutability period is counted depending on the scale-out backup repository configuration.

| Scale-Out backup Repository Settings | Performance Tier Immutable | Capacity Tier Immutable | Archive Tier Immutable |
|---|---|---|---|
| **Perfomance Tier Only** | The immutability period is the same as the GFS retention period[1].<br><br>**Note**: The actual retention depends on the Block Generation settings. For more information, see Block Generation. | – | – |
| **Performance Tier + Capacity Tier with Copy Policy** | The immutability period is the same as the GFS retention period[1].<br><br>**Note**: The actual retention depends on the Block Generation settings. For more information, see Block Generation. | The immutability period is the same as the immutability period of an object storage repository set as the capacity extent. | – |

| Scale-Out backup Repository Settings | Performance Tier Immutable | Capacity Tier Immutable | Archive Tier Immutable |
|---|---|---|---|
| **Performance Tier + Capacity Tier with Move Policy** | The immutability period is the same as the immutability period of an object storage repository set as the performance extent. | The immutability period is the same as the immutability period of an object storage repository set as the capacity extent. | — |
| **Performance Tier + Capacity Tier with Copy Policy + Archive Tier** | The immutability period is the same as the immutability period of an object storage repository set as the performance extent. | The immutability period is the same as the immutability period of an object storage repository set as the capacity extent. | The immutability period is the same as the GFS retention period[1]. |
| **Performance Tier + Capacity Tier with Move Policy + Archive Tier** | The immutability period is the same as the immutability period of an object storage repository set as the performance extent. | The immutability period is the same as the immutability period of an object storage repository set as the capacity extent. | The immutability period is the same as the GFS retention period[1]. |
| **Performance Tier + Archive Tier** | The immutability period is the same as the immutability period of an object storage repository set as the performance extent. | — | The immutability period is the same as the GFS retention period[1]. |

[1] Valid only in case the GFS retention period is longer than the immutability period for object storage repository.

# Performance Tier

Performance tier of a scale-out backup repository is the level used for fast access to the data.

The performance tier of a scale-out backup repository can comprise one or more performance extents. A performance extent is a backup repository or an object storage repository added to the scale-out backup repository. The list of the performance extents is displayed at the Add Performance Extents step of the **New Scale-out Backup Repository** wizard. You can use the performance tier with almost all types of jobs and tasks that Veeam Backup & Replication runs. For a list of unsupported scenarios, see Limitations for Performance Tier.

After you add a backup repository or an object storage repository to the scale-out backup repository, they no longer exist as individual backup repositories.

When a backup repository or an object storage repository is added as a performance extent, some of its original settings are kept, and some are not. The following settings are kept, or inherited:

- Number of tasks that can be performed simultaneously
- Read and write data rate limit
- Data decompression settings
- Block alignment settings
- Connection mode (using gateway server or direct)
- Object storage consumption (for object storage repositories)

The following settings are not inherited:

- Rotated drive settings. Rotated drive settings are ignored and cannot be configured at the level of the scale-out backup repository.
- Per-machine backup file settings. Per-machine settings can be configured at the level of the scale-out backup repository.

## Supported Types of Repositories

You can use the following types of repositories as performance extents:

- Microsoft Windows backup repositories
- Linux backup repositories
- SMB (CIFS) shared folders
- NFS shared folders
- Deduplicating storage appliances
- Object Storage Repositories

## Limitations for Performance Tier

Consider the following limitations for performance tier:

- Backup repositories that you plan to add as performance extents to the scale-out backup repository must be added to the backup infrastructure. For more information, see Backup Repositories.

- The same limitations that are specific for certain types of backup repositories apply to the performance extents. For example, if you add Dell Data Domain as a performance extent to the scale-out backup repository, you will not be able to create a backup chain longer than 120 points in this scale-out backup repository.
- The performance extents of the scale-out backup repository should be located in the same site. Technically, you can add to the scale-out backup repository the performance extents that reside in different sites. However, in this case Veeam Backup & Replication will have to access VM backup files on storage devices in different locations, and the backup performance will degrade.
- Within a scale-out backup repository, the mount server of a performance extent will act as a gateway server of the capacity extent if all of the following is true:
  a. You use SMB share/NFS share/deduplicating storage appliances as performance extents of your scale-out backup repository.
  b. You have chosen **Automatic selection for the gateway server** at the Specify Shared Folder Settings step of the **New backup repository** wizard.
  c. For the object storage that you use as the capacity extent, you have not selected to connect to object storage using a gateway server at the **Account** step of the **New Object Repository** wizard.
- If you want to add object storage repositories of another provider as performance extents, you must switch the object storage of the former extent to the sealed mode. For example, if you used Amazon S3 object storage as performance extents and you want to add Microsoft Azure Blob as a new extent, you must first seal the Amazon S3 object storage.
- If you want to use S3-compatible repositories by different providers as performance extents, one of this repository must be in the sealed mode. For example, if you want to use both Wasabi Cloud object storage and IBM Cloud object storage within the same performance tier, Wasabi Cloud object storage must be in sealed mode.
- If the object storage contains offloaded backup data, you will be offered to synchronize this data with your performance extents. For more information, see Synchronizing Capacity Tier Data.
- You cannot have a mixed configuration and use different types of repositories within one performance tier. For example, a performance tier cannot consist of a Windows-based backup repository and the object storage repository added as performance extents.
- Data located in object storage repositories is organized into a separate backup chain for every machine in a job.
- If you use immutability and have several extents in your performance tier, you must enable it for all active extents within this tier. If you want to use a mixed immutability configuration (with some extents immutable and some mutable), you must switch either all mutable or all immutable extents to the sealed mode.
- You cannot add as performance extents one Wasabi bucket as an S3 Compatible Object Storage and another Wasabi bucket as a Wasabi Cloud Object Storage to the same type of tier (either performance tier or capacity tier) of a scale-out backup repository.
- You cannot use the same object storage repository as a performance extent and as a capacity extent.
- Veeam Cloud Connect service providers can not use Azure Data Box and AWS Snowball Edge storage as a performance extent of a scale-out backup repository.
- You cannot use direct backup object storage repositories as performance extents to keep backups created with Veeam Plug-ins for Enterprise Applications.

# Extent Structure of Performance Tier with Object Storage Repositories

If your performance tier consists of object storage repositories, Veeam Backup & Replication creates and maintains the following structure of directories after backups are moved to the performance tier extents.

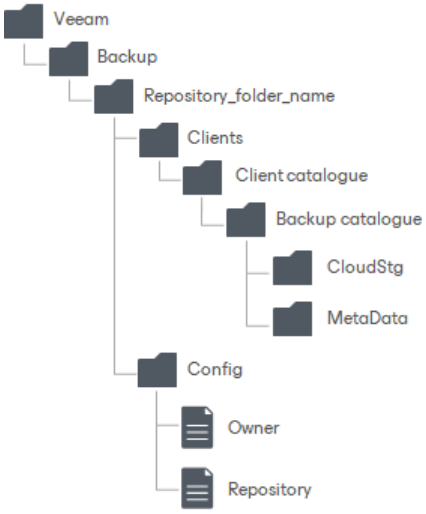


| Directory | Description |
|---|---|
| **Veeam/Backup/** | Standard folders created by Veeam Backup & Replication. |
| **Repository_folder_name** | Contains information on the repository name and the repository ID. |
| **Clients** | Contains backups. |
| **Client catalogue** | Contains information on solutions that create backups to this repository. |
| **Backup catalogue** | Contains backup ID. |
| **CloudStg** | Contains data blocks. |
| **MetaData** | Contains metadata. |
| **Config** | Contains information on the object storage repository infrastructure. |
| **Owner** | Contains information on a repository owner. |
| **Repository** | Contains information on a repository. |

# Backup File Placement for Performance Tier

Veeam Backup & Replication stores backup files on all performance extents of the scale-out backup repository.

When you configure a scale-out backup repository, you must set the backup file placement policy for backup repositories. The backup file placement policy describes how backup files are distributed between extents. You can choose one of the following policies:

- Data locality
- Performance

You can also select an extent for backup file placement, which has its nuances if you set the *Performance* policy for the scale-out backup repository:

- Extent Selection
- Extent Selection for Backup Repositories with Performance Policy
- Extent Selection for Object Storage Repositories Added as Performance Extents

Keep in mind that at the beginning of a job, Veeam Backup & Replication retrieves the actual free space on a scale-out backup repository and estimates the size of a restore point to be created. Then the estimated restore point size is subtracted from the actual free space to determine the estimated free space. This estimated free space is then used by all further concurrently running jobs targeted at the same scale-out backup repository. For details, see Backup Size Estimation.

If you do not select the **Strict placement policy enforcement** check box when you specify backup placement policy, and if the necessary extent is not accessible, Veeam Backup & Replication will disregard the policy limitations and attempt to place the backup file to the extent that has enough free space for the backup file.

For example, you have set the *Performance* policy for the scale-out backup repository and specified that full backup files must be stored on *Extent 1* and incremental backup files must be stored on *Extent 2*. If before an incremental backup job session *Extent 2* goes offline, the new incremental backup file will be placed to *Extent 1*.

> **NOTE**
>
> Veeam Backup & Replication places backups of Microsoft SQL transaction logs, Oracle archived logs and PostgreSQL WAL files to the extent configured for storing incremental backup files. If such extent is not accessible, Veeam Backup & Replication will attempt to place log backups to any other extent that has enough free space.

# Data Locality

If you set the *Data locality* policy for a scale-out backup repository, all backup files that belong to the same backup chain are stored on the same extent of the scale-out backup repository.

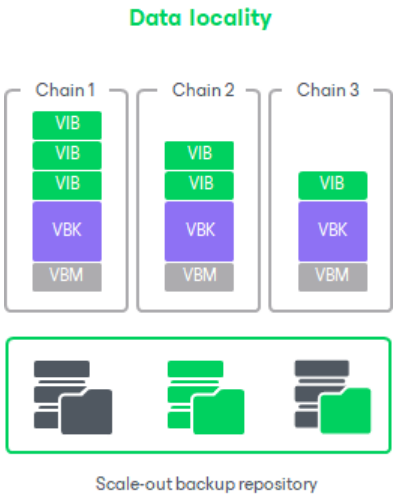


The *Data locality* policy does not put any limitations to backup chains. A new backup chain can be stored on the same extent or another extent. For example, if you create an active full backup, Veeam Backup & Replication may store the full backup file to another extent, and all dependent incremental backup files will be stored together with this full backup file.

**NOTE**

If you use a deduplicating storage appliance as an extent to the scale-out backup repository, Veeam Backup & Replication will attempt to place a new full backup (active or synthetic) to the extent where the full backup from the previous backup chain resides. Such behavior will help increase the data deduplication ratio. However, if you use ExaGrid appliances as extents of scale-out backup repository, Veeam Backup & Replication will place a new backup to the extent with the highest storage capacity.

**TIP**

If you plan to use Fast Clone on performance extents with volumes formatted with ReFS or XFS, select *Data locality*.

# Performance

If you set the *Performance* policy for a scale-out backup repository, full backup files and incremental backup files that belong to the same backup chain are stored on different extents of the scale-out backup repository. If necessary, you can explicitly specify on which extents full backup files and incremental backup files must be stored.

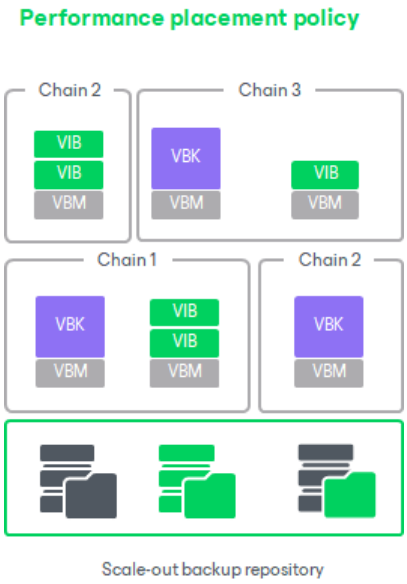


The *Performance* policy can improve performance of transformation processes if you use raw data devices as extents. When Veeam Backup & Replication performs transformation, it needs to access a number of backup files in the backup repository. If these files are located on different storage devices, the I/O load on the devices hosting backup files will be lower.

If you set the *Performance* policy, you must make sure that the network connection between extents is fast and reliable. You must also make sure all extents are online when the backup job, backup copy job or a restore task starts. If any extent hosting backup files in the current backup chain is not available, the backup chain will be broken, and Veeam Backup & Replication will not be able to complete the task. To avoid data loss in this situation, you can enable the **Perform full backup when required extent is offline** option for the scale-out backup repository. With this option enabled, Veeam Backup & Replication will create a full backup instead of incremental backup if some files are missing from the backup chain.

# Extent Selection

To select an extent for backup file placement, Veeam Backup & Replication checks the following conditions:

1. Availability of extents on which backup files reside. If some extent with backup files from the current backup chain is not accessible, Veeam Backup & Replication will trigger a full backup instead of incremental (if this option is enabled). For more information, see Adding Backup Repository Extents.
2. Backup placement policy set for the scale-out backup repository.
3. Load control settings — maximum number of tasks that the extent can process simultaneously.

> **IMPORTANT**
>
> If a Service Provider uses an extent as a cloud repository, Veeam Backup & Replication compares the percentage of active tasks over the extent task limit and selects the extent with the least load over the extent. For example, if the A extent tasks limit is 5 tasks, and currently 3 tasks are active, and the B extent tasks limit is 30 tasks, and 10 tasks are currently running, Veeam Backup & Replication selects the B extent. If there are several extents with equal percentage values, Veeam Backup & Replication checks a disc space and selects the best suitable extent based on this check. For example, if there are extents with the same percentage of running tasks, Veeam Backup & Replication compares their disk space and selects the extent with the least disk space occupied.

4. Amount of free space available on the extent — the backup file is placed to the extent with the most amount of free space.
5. Availability of files from the current backup chain — extents that host incremental backup files from the current backup chain (or current VM) have a higher priority than extents that do not host such files.

## Extent Selection for Backup Repositories with Performance Policy

If you set the *Performance* policy for the scale-out backup repository, Veeam Backup & Replication always stores full backup files and incremental backup files that belong to the same backup chain on different extents. To choose the extent to which a backup file can be stored, Veeam Backup & Replication applies this policy and policies mentioned above.

For example, a scale-out backup repository has 2 extents that have 100 GB and 200 GB of free space. You set the *Performance* policy for the scale-out backup repository and define that all types of backup files (full and incremental) can be placed on both extents.

When a backup job runs, Veeam Backup & Replication picks the target extent in the following manner:

1. During the first job session, Veeam Backup & Replication checks to which extent a full backup file can be stored. As both extents can host the full backup file, Veeam Backup & Replication checks which extent has more free space, and picks the extent that has 200 GB of free space.
2. During incremental job session, Veeam Backup & Replication checks to which extent an incremental backup file can be stored. As both extents can host the incremental backup file, Veeam Backup & Replication picks the extent that does not store the full backup file — the extent that has 100 GB of free space.

## Extent Selection for Object Storage Repositories Added as Performance Extents

If you add an object storage repository as a performance extent, neither data locality nor performance policies are applied to these extents. Instead, Veeam Backup & Replication checks the following conditions to distribute data between extents.

1. During the first job session, Veeam Backup & Replication checks availability of extents. In case, an extent is set to Sealed or Maintenance mode, these extents will be skipped from backup file placement.
2. After that, Veeam Backup & Replication checks the amount of backup chains in available extents and selects the extent with a minimal amount of backup chains.
3. If some extents have storage space limitations, Veeam Backup & Replication calculates the average storage of the backup chains located in all extents and select the extent that contains a minimum amount of backup chains. If all extents have the same number of backup chains, Veeam Backup & Replication selects the extent that has more free space.

4. During subsequent job sessions, Veeam Backup & Replication moves backup chains to the extent that was specified before.

> **NOTE**
>
> If an extent is set to Sealed mode or Maintenance mode, Veeam Backup & Replication will not target a backup chain to this extent. In this case, Veeam Backup & Replication starts a new backup chain on a new extent and follows the same algorithm as during the first job session.
> Veeam Backup & Replication will remove backup chains from the Sealed extents according to retention policy.

# Backup Size Estimation

At the beginning of a job, Veeam Backup & Replication retrieves the actual free space on a scale-out backup repository and estimates the size of a restore point to be created. Then the estimated restore point size is subtracted from the actual free space to determine the estimated free space. This estimated free space is then used by all further concurrently running jobs targeted at the same scale-out backup repository.

Veeam Backup & Replication assumes that the following amount of space is required for backup files:

- The size of the first full backup file is equal to 50% of source VM data.
- The size of further full backup files is equal to 100% of the previous full backup file size.
- The size of the first incremental backup file is equal to 10% of the previous full backup file size.
- The size of further incremental backup files in the backup chain is equal to 100% of the previous incremental backup file size.

  In case of reverse incremental backup chains, during incremental job sessions Veeam Backup & Replication allocates 10% of the previous full backup file size on the extent where a rollback file is placed and additional 10% on the extent where the full backup file resides.

This mechanism is also applied to backup files created by backup copy jobs.

> **NOTE**
>
> Consider the following:
>
> - On every extent of a scale-out backup repository, Veeam Backup & Replication reserves 1% of storage space to guarantee correct update of backup metadata files (VBM) and success of merge operations.
> - Make sure that you have enough free space on the extent where the full backup file resides. Veeam Backup & Replication requires 10% of the size of the full backup file to perform merge operations in the backup chain. If the disk space is low, merge operations may fail.
> - The actual free space value is only captured at the start of a job targeted at the scale-out backup repository while no other jobs actively use the same scale-out backup repository. For more information, see this Veeam KB article.

# Immutability for Performance Tier

Veeam Backup & Replication allows you to prohibit deletion of data from the repositories added as a performance extent by making that data temporarily immutable. It is done for increased security: immutability protects your data from loss as a result of attacks, malware activity or any other injurious actions.

You can enable immutability for the following types of repositories:

- Amazon, S3-compatible and Azure object storage repositories. For more information, see Immutability for Object Storage Repositories.
- Hardened Repository. For more information, see Hardened Repository.
- HPE StoreOnce. For more information, see Immutability for HPE StoreOnce.
- Dell Data Domain. For more information, see Immutability for Dell Data Domain.

The immutable data within the performance extent cannot be subject to the following operations:

- Manual removal of data, as described in section Deleting Backups from Scale-Out Backup Repositories.
- Removal of data by the retention policy, as described in section Retention Policy.
- Removal of data using any cloud service provider tools.
- Removal of data by the cloud service provider technical support department.
- Removal of data by the **Remove deleted items data after** option, as described in section Maintenance Settings.

# Immutability for Backup Object Storage Repositories Added as Performance Extents

Immutability settings for the following types of backups files depends on a configuration of a scale-out backup repository:

- Backups with GFS flags.
- Backups created by VeeamZIP jobs.
- Exported backup files.

Consider the following scenarios:

- If retention policy configured for these backup files exceeds immutability settings, Veeam Backup & Replication applies retention that is defined for these types of backups. Immutability settings defined for an object storage repository are ignored.
- If you do not add a capacity tier or archive tier to your scale-out backup repository, Veeam Backup & Replication follows retention policy configured for these files in case it exceeds immutability settings. Immutability settings defined for an object storage repository are ignored.
- If you add a capacity tier to your scale-out backup repository and select move backups to capacity tier, immutability is set according to the object storage repository settings. Retention policy for these backups files is ignored. Note that data blocks are deleted from the performance tier once the immutability period for the performance tier ends.

## Block Generation

To reduce I/O operations and associated costs, Veeam Backup & Replication will add several days to the immutability expiration date. This period is called Block Generation. You do not have to configure it, the Block Generation setting is applied automatically.

Depending on the type of the object storage repository, Veeam Backup & Replication will add the following values for the default generation period:

- 30 days — for Amazon S3 object storage and IBM Cloud object storage.
- 10 days — for all other types of object storage repositories.

For example, if you set your immutability period to 30 days for your object storage repository, Veeam Backup & Replication will add 10 days to specific objects to reduce I/O operations with the data blocks over time. Thus, you will have immutability set for 30 days + 10 days of Block Generation set for data blocks in your object storage repositories.

# How Block Generation Works

When the first data block (a full backup) arrives, its immutability period by default is set to 30 + 10 = 40 days. The first full backup starts its generation, that will be appended with the incremental backups. All the incremental backups within the generation (that is, within the 10-days period) will have the same immutability expiration date as the full backup. For instance, a data block that was offloaded on day 9 will have the same immutability expiration date as a data block offloaded on day 1. Thus we ensure that the immutability period for all the data blocks within a generation is no less than 30 days.

To maintain the backup consistency, Veeam Backup & Replication can extend immutability expiration for all data blocks in the backup chain and assign these blocks to a new generation.

For example, within one forward incremental backup chain, a full backup file can not be removed before an incremental backup file. On the other hand, an incremental backup file makes no sense without relevant full backup file. So the immutability period is extended for all data blocks in the backup chain.

With 10 days of immutability automatically added, it means there is no need to extend the immutability period for the incremental backups in forward chain and for the unchanged blocks of current full backups in reverse chain offloaded within those 10 days. On the 11th day, though, the immutability period will have to be extended (to ensure that the immutability period for all the data blocks within a generation is no less than 30 days).

> **NOTE**
>
> Consider the following:
>
> - For data blocks located in object storage repositories, Veeam Backup & Replication extends immutability period for every data block of every backup file in the whole backup chains, even in inactive part.
> - Veeam Backup & Replication will not extended immutability for the data blocks that are not used in any exiting backup files.

# Transferring Backups to Performance Tier

Prior to Veeam Backup & Replication version 12, it was possible to keep data only on object storage repositories added as capacity extents of scale-out backup repository. Starting from Veeam Backup & Replication version 12, it is possible to add an object storage repository as a performance extent and back up data directly to this extent.

In case you have an object storage repository that is added as a capacity extent and want to use it as a performance extent, you can transfer existing backup from the capacity extent to the performance extent.

To do this, you must perform the followings steps:

1. Add a new object storage repository to the backup infrastructure.

2. Download a necessary backup chain to the existing backup repository added as a performance extent. For more information, see Downloading Data from Capacity Tier.
3. Move backups from the backup repository to a new object storage repository. For more information, see Moving Backups.
4. Add the object storage as a performance extent to a new scale-out backup repository.

# Removing Performance Extents from Scale-Out Repositories

You can remove a performance extent from the scale-out backup repository, for example, if you do not want to store backup files on the underlying storage anymore. When you remove a performance extent, Veeam Backup & Replication puts the relevant backup repository to the Maintenance mode and unassigns the performance extent role from it. The backup repository ceases to exist as a part of the scale-out backup repository and becomes an individual backup repository.

If the performance extent contains backup files, it is strongly recommended that you perform the following actions before you remove the extent:

1. Put the performance extent to the Maintenance mode.
2. Evacuate backups from the extent. For more information, see Evacuating Backups from Extents.

In this case, backup files will be moved to other performance extents of the scale-out backup repository, and the backup chain will remain consistent. If you do not evacuate backups from the performance extent, the backup chain may get broken as some restore points will be missing from it.

To remove a performance extent from the scale-out backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Scale-out Repositories**.
3. In the working area, select the scale-out backup repository and click **Edit Scale-out Repository** on the ribbon or right-click the scale-out backup repository and select **Properties**.
4. Move to the **Performance tier** step of the wizard.
5. In the **Extents** list, select the performance extent and click **Remove**.

If the performance extent contains backup files, Veeam Backup & Replication will suggest evacuating them. To evacuate files, click **No**, close the wizard and evacuate backup files. For more information, see Evacuating Backups from Extents.

If you do not want to evacuate the backup files, click **Yes** and proceed with the wizard.

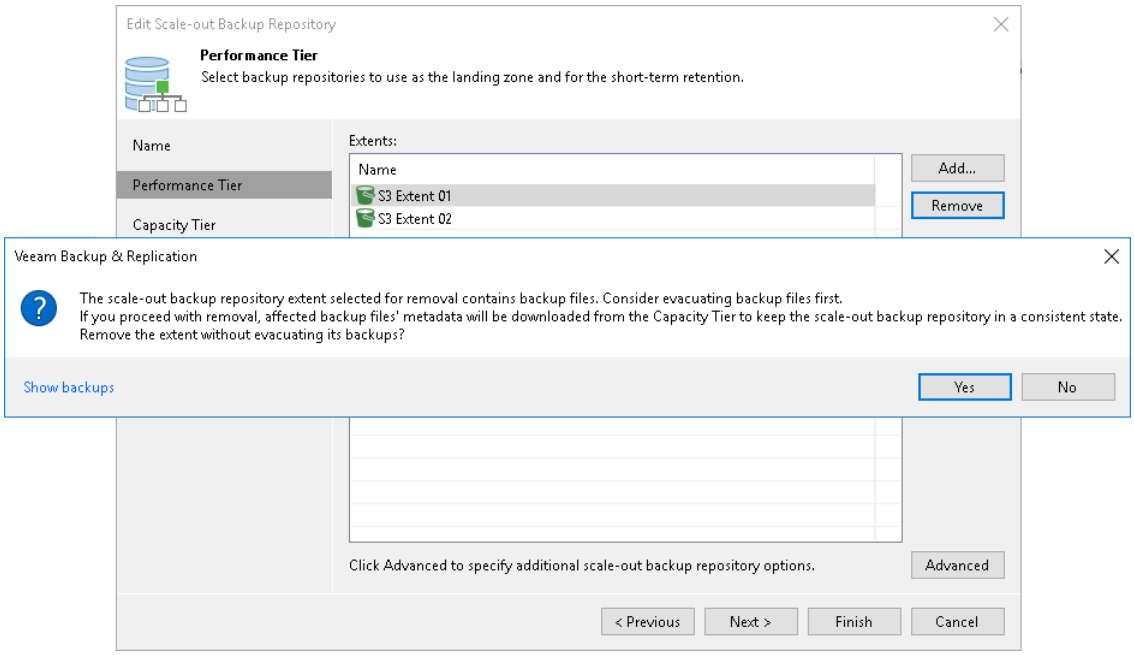

# Capacity Tier

Capacity tier is an additional tier of storage that can be attached to a scale-out backup repository. Data from the scale-out backup repository performance extents can be transported to the capacity tier for long-term storage. You can use the capacity tier with almost all types of jobs and tasks that Veeam Backup & Replication runs. For a list of unsupported scenarios, see Limitations for Capacity Tier.

This feature is most useful in the following cases:

- You are running out of storage space.
- Your organization policies allow you to store only a certain amount of data on your extents, while the outdated data should be stored elsewhere.
- You seek to store data on several sites to ensure its safety in case of a disaster.

With capacity tier, you can perform the following operations:

- Move inactive backup chains to capacity extents, as described in section Moving Backups to Capacity Tier and Manually Moving Backups to Capacity Tier.
- Copy new backup files as soon as these files are created, as described in section Copying Backups to Capacity Tier.
- Download data that was moved from capacity extents back to the performance extents, as described in section How Downloading from Capacity Tier Works.
- Restore your data. For more information, see Restore from Capacity Tier. In particular, you can promptly restore data from the capacity tier in case of disaster without creating a scale-out backup repository anew. For more information about this feature, see Importing Object Storage Backups.

## Supported Types of Object Storage Repositories

The capacity tier consists of multiple capacity extents. The capacity extent can be either a cloud-based object storage repository or on-premises object storage repository, such as:

- S3-compatible object storage repository
- Amazon S3
- AWS Snowball Edge
- Microsoft Azure Blob storage
- Microsoft Azure Data Box
- IBM Cloud Object storage
- Google Cloud Object storage
- Wasabi Cloud Object storage
- Veeam Data Cloud Vault

Before an object storage repository can be configured as the capacity extent, it must be added to Veeam Backup & Replication. For more information, see Adding Object Storage Repositories.

The capacity extents are displayed in the scale-out backup repository wizard, on the Capacity Tier step.

For information on configuring capacity tier and synchronizing capacity tier data, see Add Capacity Tier.

# Limitations for Capacity Tier

# General Limitations

Consider the following limitations for the capacity tier:

- If you want to add object storage repositories of another provider as capacity extents, you must switch the object storage of the former extent to the sealed mode. For example, if you used Amazon S3 object storage as capacity extents and you want to add Microsoft Azure Blob as a new extent, you must first seal the Amazon S3 object storage.
- If you want to use S3-compatible repositories by different providers as capacity extents, one of this repository must be in the sealed mode. For example, if you want to use both Wasabi Cloud object storage and IBM Cloud object storage within the same capacity tier, Wasabi Cloud object storage must be in sealed mode.
- You cannot add more than one instance of Azure Databox or Snowball Edge AWS object storage repository as a capacity extent.
- You cannot add the object storage repository that contains imported backups as a capacity extent. For more information, see Importing Object Storage Backups.
- You cannot copy transaction log backups to the capacity tier.
- You cannot use the capacity tier as a target for file backup jobs and object storage backup jobs.
- You cannot add the object storage repository that contains imported backups as a capacity extent. For more information, see Importing Object Storage Backups.
- Before you start using the capacity tier, make sure to check the pricing plans of your cloud storage provider to avoid additional costs for offloading and downloading backup data.
- Within a scale-out backup repository, the mount server of a performance extent will act as a gateway server of the capacity extent if all of the following is true:
  a. You use SMB share/NFS share/deduplicating storage appliances as performance extents of your scale-out backup repository.
  b. You have chosen **Automatic selection for the gateway server** at the Specify Shared Folder Settings step of the **New backup repository** wizard.
  c. For the object storage that you use as the capacity extent, you have not selected to connect to object storage using a gateway server at the **Account** step of the **New Object Repository** wizard.
- Downloading or restoring data from capacity tier does not reuse the existing blocks present on performance tier if performance tier consists of object storage repositories.
- Full restore points downloaded from capacity tier to XFS performance extents do not use Fast Clone.
- If you use immutability and have several extents in your capacity tier, you must enable it for all active extents within this tier. If you want to use a mixed immutability configuration (with some extents immutable and some mutable), you must switch either all mutable or all immutable extents to the sealed mode.
- You cannot use the same object storage repository as a performance extent and as a capacity extent.
- You cannot simultaneously store capacity tier and unstructured data backups in the same object storage repository. You must either add this object storage repository as the capacity extent or use it as the backup repository for the unstructured backup job.

- If you use the forever forward incremental backup method to create your backups, Veeam Backup & Replication will ignore the move policy and will apply the copy policy to move backups to capacity tier.
- You must enable the capacity tier encryption for Veeam Data Cloud Vault if you want to use it as a capacity extent.
- If you add S3 compatible repositories with multiple buckets as capacity extents, Veeam Backup & Replication will ignore this setup and store backups in the master bucket.

# Limitations for Veeam Solutions

For more information on limitations for a specific Veeam solution that utilizes capacity tier functionality, see the following sections of the necessary guide:

- Veeam Plug-in for SAP HANA — to check limitations for an SAP-certified backup and recovery solution that allows you to back up and restore SAP HANA databases.
- Veeam Plug-in for Oracle RMAN — to check limitations for an Oracle-certified backup and recovery solution that allows you to back up and restore Oracle databases.
- Veeam Plug-in for SAP on Oracle — to check limitations for an SAP-certified backup and recovery solution that allows you to back up and restore Oracle databases to which an SAP application is connected.
- Veeam Plug-in for Microsoft SQL Server — to check limitations for a backup and recovery solution that allows you to back up and restore Microsoft SQL Server databases.
- Veeam Plug-in for IBM Db2 — to check limitations for a backup and recovery solution that allows you to back up and restore IBM Db2 databases.

# Capacity Extent Structure

When data is transferred to the capacity extent, Veeam Backup & Replication creates and maintains the following structure of directories:

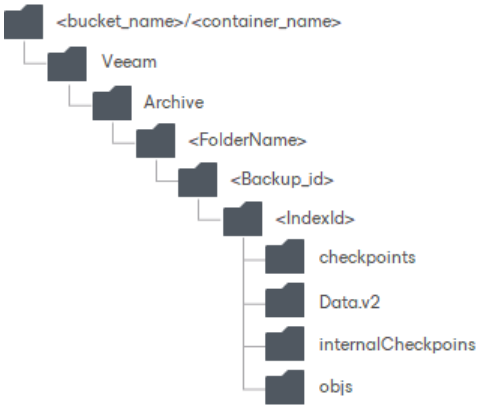


| Directory | Description | Misc |
|---|---|---|
| **<bucket_name> or <container_name>** | A bucket or container name.<br>Buckets and containers must be created in advance. | N/A |

| Directory | Description | Misc |
|---|---|---|
| **Veeam/Archive/** | Standard folders created by Veeam Backup & Replication. | |
| **<FolderName>** | A repository folder that you create when adding a new capacity extent. | |
| **<Backup_id>** | Contains objects in a backup. | These folders are automatically removed during data removal. |
| **<Index_Id>** | An identifier of an object in a backup.<br>• If a backup was created using the per-machine method, then each VM will be placed to its own directory.<br>• If a backup was created as a single-file backup, then all the VMs will be placed to a unique directory. | |
| **Checkpoints** | Contains meta information about the state of offloaded backup chains. Such meta information is updated upon each successful offload session. | |
| **Data.v2** | Contains data blocks and metadata related to these blocks. | |
| **internalCheckpoints** | Auxilliary checkpoints. Contain information on backup files and metadata. | |
| **objs** | Contains meta information and other auxiliary data. | |

# Backup File Placement for Capacity Tier

Data distribution between capacity extents depends on whether it is the first or a subsequent move or copy of a backup chain. To select an extent for backup file placement, Veeam Backup & Replication checks the following conditions:

1. During the first job session, Veeam Backup & Replication checks availability of extents. In case, an extent is set to Sealed or Maintenance mode, these extents will be skipped from backup file placement.
2. After that, Veeam Backup & Replication checks the amount of backup chains in available extents and selects the extent with a minimal amount of backup chains.
3. If some extents have storage space limitations, Veeam Backup & Replication calculates the average storage of the backup chains located in all extents and select the extent that contains a minimum amount of backup chains. If all extents have the same number of backup chains, Veeam Backup & Replication selects the extent that has more free space.

4. During subsequent job sessions, Veeam Backup & Replication moves backup chains to the extent that was specified before.

> **NOTE**
>
> If an extent is set to Sealed mode or Maintenance mode, Veeam Backup & Replication will not move a backup chain to this extent. In this case, a new extent is selected and Veeam Backup & Replication follows the same algorithm as during the first job session. The backup chain located in an unavailable extent will be moved again to another extent. Veeam Backup & Replication will remove backup chains from the unavailable extents according to retention policy.

# Immutability for Capacity Tier

Veeam Backup & Replication allows you to prohibit deletion of data from capacity extents by making that data temporarily immutable. It is done for increased security: immutability protects your data from loss as a result of attacks, malware activity or any other injurious actions.

You can enable immutability for data stored in Amazon, S3-compatible and Azure object storage repositories used as capacity extents of the scale-out backup repository. After you enable immutability, Veeam Backup & Replication will prohibit data deletion from capacity tier until the immutability expiration date comes.

Backups are immutable only during the immutability period set for the object storage repository even if their retention policy allows for longer storage. Immutability period ignores retention policies set for the following types of backups:

- Backups with GFS flags
- Backups created by VeeamZIP jobs
- Exported backup files

The immutable data within the capacity extents cannot be subject to the following operations:

- Manual removal of data, as described in section Removing Backups from Capacity or Archive Tier.
- Removal of data by the retention policy, as described in section Retention Policy.
- Removal of data using any cloud service provider tools.
- Removal of data by the cloud service provider technical support department.
- Removal of data by the **Remove deleted items data after** option, as described in section Maintenance Settings.

# Enabling Immutability

To enable immutability, you must do the following:

1. Configure the necessary settings when you create an S3 bucket or an Azure container.

   For more information, see Enabling Immutability.

2. Enable the immutability option when you add an object storage repository to the backup infrastructure at the **Container** step (for Azure object storage repository) or **Bucket** step (for Amazon S3 or S3 compatible object storage repositories) of the new **Object Storage Repository** wizard.

## Block Generation

To reduce I/O operations and associated costs, Veeam Backup & Replication will add several days to the immutability expiration date. This period is called Block Generation. You do not have to configure it, the Block Generation setting is applied automatically.

Depending on the type of the object storage repository, Veeam Backup & Replication will add the following values for the default generation period:

- 30 days — for Amazon S3 object storage and IBM Cloud object storage.
- 10 days — for all other types of object storage repositories.

For example, if you set your immutability period to 30 days for your object storage repository, Veeam Backup & Replication will add 10 days to specific objects to reduce I/O operations with the data blocks over time. Thus, you will have immutability set for 30 days + 10 days of Block Generation set for data blocks in your object storage repositories.

# How Block Generation Works

Block Generation works in the following way. When the first data block (a full backup) arrives, its immutability period is set to 30 + 10 = 40 days. The first full backup starts its generation, that will be appended with the incremental backups. All the incremental backups within the generation (that is, within the 10-days period) will have the same immutability expiration date as the full backup. For instance, a data block that was offloaded on day 9 will have the same immutability expiration date as a data block offloaded on day 1. Thus we ensure that the immutability period for all the data blocks within a generation is no less than 30 days.

Consider this example: within one forward incremental backup chain, a full backup file can not be removed before an incremental backup file. On the other hand, an incremental backup file makes no sense without relevant full backup file. So the immutability period is extended for all data blocks in the backup chain.

The Block Generation period was introduced to reduce the number of requests to the capacity tier, thereby saving traffic and reducing costs that can be incurred by your storage provider. With 10 days of immutability automatically added, it means there is no need to extend the immutability period for the incremental backups in forward chain and for the unchanged blocks of current full backups in reverse chain offloaded within those 10 days. On the 11th day, though, the immutability period will have to be extended (to ensure that the immutability period for all the data blocks within a generation is no less than 30 days).

The immutable blocks of data may be reused during the offload. Veeam Backup & Replication continues to keep reused or dependent blocks of data locked by continuously assigning them to new generations, thereby extending their immutability expiration period. This is valid both for forward and reverse incremental backup chains.

The extension of immutability works differently in different cases:

# Forward incremental backup chain

- New full backup file in the new generation:
  - Immutability is extended for the data blocks that are being reused from the old backup chain.
  - Immutability is set anew for new blocks of the new full backup file.
- New incremental backup file in the new generation:
  - Immutability is extended for all the data blocks from the current backup chain.
  - Immutability is set anew for the new blocks of the latest incremental backup file.

# Reverse incremental backup chain

- New full backup file in the new generation:
  - Immutability is extended for the data blocks that are being reused from the previous full backup file.
  - Immutability is set anew for new blocks of the new full backup file.
- Current full backup file in the new generation:
  - Immutability is extended for all the data blocks of this full backup file that are already stored in the capacity tier.
  - Immutability is set anew for the new blocks of this full backup file.

# Encryption for Capacity Tier

Veeam Backup & Replication allows you to encrypt archived data. This helps you protect the data from unauthorized access.

You can enable data encryption in the following ways:

- When you create a backup or backup copy job.
- When you add a capacity extent to your scale-out backup repository.

To get benefits of both encryption levels, you can use job-level and the capacity tier encryption within the same object storage. Both encryption levels prevents unauthorized access to your data, but the capacity tier encryption allows you to encrypt backup chain metadata and restore points. After you enable encryption on the capacity tier level, Veeam Backup & Replication encrypts the entire collection of blocks along with the metadata. Data encrypted by the backup job stays as is on the capacity tier, Thus, this backup has 2 levels of encryption: on the job level an on the capacity tier level.

To encrypt data, Veeam Backup & Replication uses the following scheme:

1. The user specifies the password in the archiving tier settings.
2. Veeam Backup & Replication generates an encryption key and uses the password to encrypt this key.
3. To encrypt the backups in the capacity tier, Veeam Backup & Replication uses this encryption key.

To access the data in the archive extent you need to provide the password that is used to decrypt the encryption keys.

**IMPORTANT**

Consider the following:

- When you disable encryption, the encryption key is added to the file system of your backup repository and all encrypted backups located on the capacity tier become unencrypted. If you enable encryption again, Veeam Backup & Replication will encrypt only new data with a new encryption key.
- If you enable encryption for the capacity extent that already contains backups, it will not automatically encrypt these backups. If a backup job creates an active full or synthetic full backup, it will consist of encrypted and unencrypted data blocks after offload to the capacity tier. This backup will remain in the capacity tier in this state until new encrypted data blocks completely replace the unencrypted blocks.
- If you enable encryption after you have already offloaded data to capacity tier, Veeam Backup & Replication will not encrypt previously offloaded backup chains.

# Health Check for Capacity Tier

A health check is an operation that allows you to ensure that the restore point is consistent, and you will be able to restore data from this restore point. For data located in the capacity tier, Veeam Backup & Replication offers a special health check mechanism that differs from the standard health check in the following ways:

- The health check for capacity tier verifies metadata for the whole backup, not just the latest restore point.
- The health check for capacity tier verifies that data blocks are present on capacity extents of a scale-out backup repository. It does not read data from data blocks. Instead, it lists data blocks to make sure all blocks in the object storage are available for rebuilding every restore point in the backup chain.

If some blocks are missing, Veeam Backup & Replication marks these blocks and restore points associated with them as corrupted. It will prevent Veeam Backup & Replication from reusing the corrupted blocks in certain operations. Instead, Veeam Backup & Replication will upload these blocks to the capacity extent again during the offload session.

If health check detects metadata corruption, Veeam Backup & Replication will mark all backup chain as corrupted. In this case, you will not be able to restore from this backup chain and Veeam Backup & Replication will prevent all offload sessions for the corrupted backup chain. To add a new backup chain to the capacity tier, you must either delete the corrupted backup chain from capacity tier or detach the backup from a backup job.

**NOTE**

By default, Veeam Backup & Replication uses local resources to perform the health check. The type of the local resource depends on the connection mode of the object storage repository added as a capacity extent:

- For direct connection mode — Veeam Backup & Replication will use a mount server to perform the health check.
- For gateway server connection mode — Veeam Backup & Replication will use a gateway server to perform the health check.

By default, the health check runs monthly every last Saturday and runs during the offload session after backup cleanup completes. You can change the schedule and run the health check weekly or monthly on specific days in the scale-out backup repository settings.

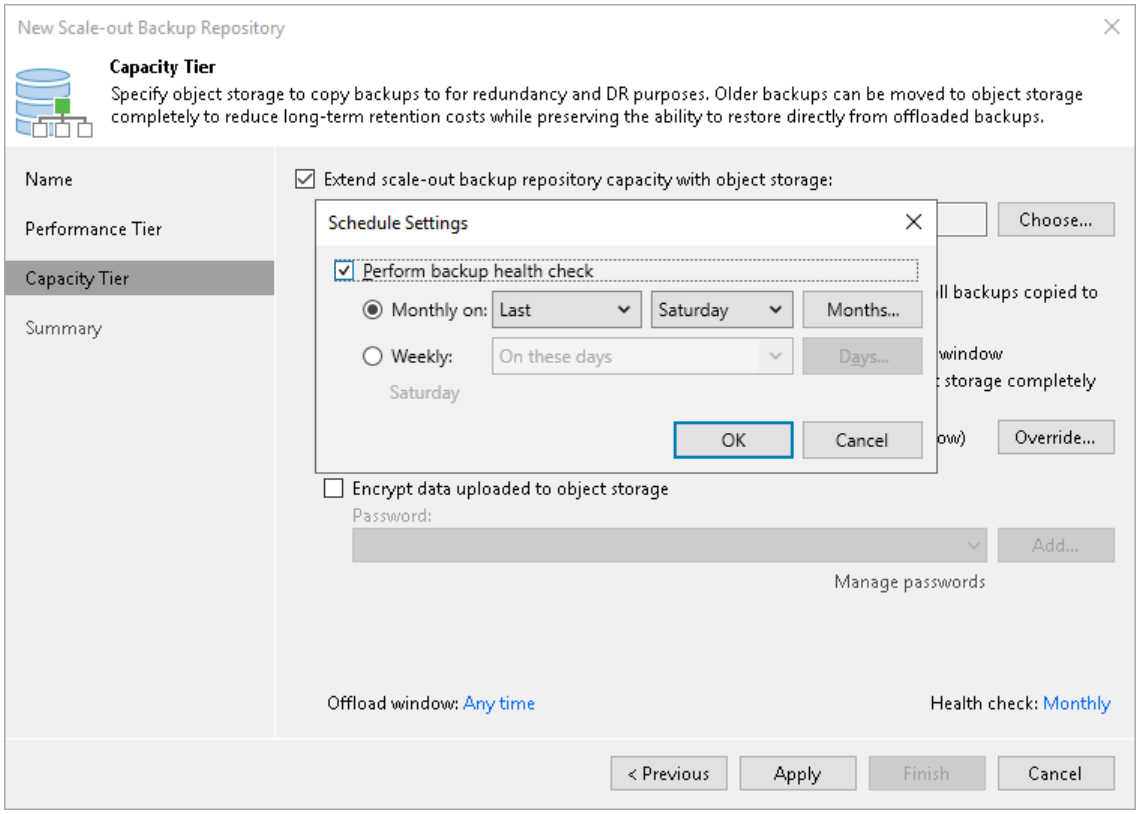


# How Health Check for Capacity Tier Works

Veeam Backup & Replication performs the health check of a backup in the following way:

1. Veeam Backup & Replication starts the health check of the whole backup during the first offload session on the scheduled date. Veeam Backup & Replication checks if the metadata of the backup is consistent and no metadata is missing. Veeam Backup & Replication also checks if all data blocks for every restore point are available on the capacity extent. Veeam Backup & Replication does not read data from data blocks.
2. If Veeam Backup & Replication does not find any corrupted data, the health check completes successfully.

   If Veeam Backup & Replication detects corrupted data, the health check completes with an error. Depending on the detected data inconsistency, Veeam Backup & Replication behaves in one of the following ways:

   - If the health check detects corrupted metadata, the offload session fails. Veeam Backup & Replication will mark all backup chain as corrupted and prevents all offload sessions for the corrupted backup chain. To add a new backup chain to the capacity tier, you must either delete the corrupted backup chain from capacity tier or detach the backup from a backup job.
   - If the health check detects missing data blocks, Veeam Backup & Replication marks these blocks and restore points associated with them as corrupted. Veeam Backup & Replication will upload these blocks to the capacity extent again during the offload session.

   You can view the health check result in the restore point statistics. If the health check finds corrupted data, it will display information on it, as well as list all restore points that share the corrupted data blocks.

> **NOTE**
>
> For immutable backups, Veeam Backup & Replication performs the health check only for valid restore points according to the retention policy. Immutable data associated with removed restore points can still remain in the repository depending on the immutability period, but Veeam Backup & Replication will not perform the health check for such data.

# Managing Capacity Tier

You can manage your capacity tier and the offloaded data in the following ways:

- Download previously offloaded data from the capacity extent back to the performance extents.
- Migrate data to another capacity tier.
- Reduce the number of cost-based operations incurred by your cloud storage provider and decrease the amount of traffic being sent over the network when moving or copying data to the capacity tier. For more information, see Indexes.
- Configure the retention policy.
- Exclude the capacity extents from the scale-out backup repository scope.

## Data Transfer

Veeam Backup & Replication allows you to transfer data to and from capacity extents, as well as move data from capacity extents to archive extents.

You can perform the following data transfer operations:

- Copy policy: the backups are copied to capacity extents automatically.
- Move policy: the inactive backup chains can be transferred to the capacity extents.
- Previously offloaded data can also be downloaded from the capacity tier back to the performance extents.

To manage data transfer to and from capacity extents, Veeam Backup & Replication uses system sessions.

The following types of backup files can be moved or copied to capacity extents:

- Regular backups (except transaction logs)
- Veeam backups for Amazon, Google and Microsoft Azure (using backup copy jobs)
- Backups created with Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for Unix or Veeam Agent for Mac
- Backups created with Veeam Plug-ins for Enterprise Applications (Oracle RMAN, SAP HANA, SAP on Oracle)
- Backups created with Veeam Backup for Nutanix AHV
- Backups created with Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization
- Backups created with VeeamZIP jobs
- Backups exported by Kasten policies
- Exported backup files

- Orphaned backups

> **NOTE**
>
> Consider the following:
>
> - Imported backups can be copied or moved to capacity extents only manually.
> - You can move only inactive backup chains to capacity extents.
> - If you use the forever forward incremental backup method to create your backups, Veeam Backup & Replication will ignore the move policy and will apply the copy policy since this backup method will always produce an active backup chain. If you want to make the backup chain inactive, you must create either a synthetic full or an active full backup.

# In This Section

- Copying Backups to Capacity Tier
- Moving Backups to Capacity Tier
- Manually Moving Backups to Capacity Tier
- Downloading Data from Capacity Tier
- Viewing Capacity Tier Sessions Statistics

## Copying Backups to Capacity Tier

Veeam Backup & Replication allows you to copy backups from the performance extents of your scale-out backup repository to capacity extents as soon as these backups are created.

To enable data copy, make sure to select the **Copy backups to object storage as soon as they are created** option, as described in section Add Capacity Tier.

To copy data to capacity extents, Veeam Backup & Replication performs the following steps:

1. Veeam Backup & Replication verifies that performance extents are available and are not in the Maintenance mode.
2. After a backup (or backup copy) job that is targeted to a scale-out backup repository finishes, Veeam Backup & Replication initiates a copy session.

   A complete name of each copy session consists of the backup (or backup copy) job name plus the *Offload* postfix. That is, if your backup (or backup copy) job name is *Amazon*, the copy session name will be *Amazon Offload*.
3. During a copy session, Veeam Backup & Replication extracts data blocks and metadata from each new backup file (.VBK, .VIB, .VRB) or data blocks created on performance extents and copies these blocks and metadata to capacity extents.

Backup files with metadata are created as described in section Moving Backups to Capacity Tier.

Having such replica gives you the ability to quickly restore data as of the latest state in case of trouble with any backup files, any unexpected failure of any of your performance extents, or even of the entire scale-out backup repository, as described in section Restore Scenarios.

# Combining Copy and Move Operations

You can combine both the **Copy backups to object storage as soon as they are created** option and the **Move backups to object storage as they age out of the operational restores window** option. In such a scenario, a copy session will be copying newly created backups right upon creation.

Once the backup chain becomes inactive (for example, sealed) and exceeds the operational restore window, data blocks will be removed from each associated backup file. Such a behavior mimics data movement, but instead of moving data that was already copied, Veeam Backup & Replication purges associated data blocks from the performance extents, thereby saving traffic and reducing costs that may be incurred by your storage provider for performing read/write operations.

The following figure shows an example in which both options are enabled, suggesting that each backup file has been copied to object storage upon its creation.

The backup chain on the left becomes inactive after a new full backup file is created and consists of one .VBK file and five .VIB files. In this case, only the first four backup files (represented as grey blocks) in this inactive backup chain exceed the operational restore window. Veeam Backup & Replication will wait until .VIB files that were created on Thursday and Friday fall out of operational restore window. After that, Veeam Backup & Replication will remove blocks of data from these .VBK and .VIB files that belong to the inactive backup chain.

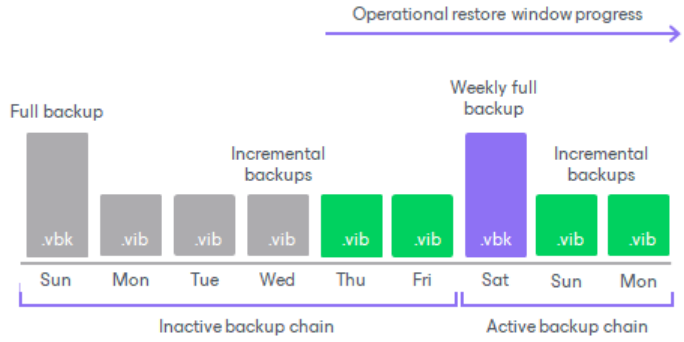


After copy is complete, the new **Capacity Tier** node appears in the **Home** view, under the **Backups** node and shows backups that have been copied to the capacity extent.

> **NOTE**
>
> The copy is not performed during prohibited hours specified in the scale-out backup repository backup window configuration. You can configure the backup window at the Add Capacity Tier step of the **New Scale-out Backup Repository** wizard.

## Moving Backups to Capacity Tier

To collect backup files that belong to inactive backup chains from the performance extents and move them to the capacity tier, Veeam Backup & Replication uses an offload session which is executed automatically every 4 hours.

To enable data movement, make sure to select the **Move backups to object storage as they age out of the operational restores window** option, as described in the Add Capacity Tier section.

A complete name of each offload session is built up of the scale-out backup repository name plus the *Offload* postfix. That is, if your scale-out backup repository name is *Amazon*, the offload session name will be *Amazon Offload*.

The offload session manages the following:

- Validation Process
- Data Transfer

# Validation Process

Before your data can safely be moved to the capacity tier, Veeam Backup & Replication performs the following mandatory verifications and required actions:

- Verifies whether data that is about to be moved belongs to an inactive backup chain.

  For more information, see Backup Chain Detection.

- Verifies whether performance extents are available and have not been put into the Maintenance mode.

  > **IMPORTANT**
  >
  > Veeam Backup & Replication does not offload data from Linux-based performance extents that have internet access through HTTP(S) proxy. All Linux-based performance extents configured in your scale-out backup repository must have direct access to the internet.

- Verifies whether the capacity extents have not been put into the Maintenance or the Sealed mode.

  For more information, see Switching to Maintenance Mode and Switching to Sealed Mode.

- Verifies whether configuration parameters that define how and when inactive backup chains must be moved to capacity extents are met.

  Such parameters are configured as described in section Add Capacity Tier.

- Synchronizes the backup chain state between the performance and capacity extents to maintain retention policies.

  For more information, see Retention Policy.

# Data Transfer

After the validation process is complete, the *SOBR Offload* session collects backup files that have passed verification. Such verified backup files are collected from all the performance extents added to a scale-out backup repository.

The performance extent A has an inactive backup chain consisting of one VBK file and three VIB files, that is, 4 restore points in total. Each of these files comprises metadata and the actual blocks of data. During the offload session, Veeam Backup & Replication will collect the actual blocks of data from all the backup files (VBK and VIB) and offload these blocks to the object storage repository.

Each offloaded block may be of different size, which is defined during configuring storage optimization. The offloaded blocks are placed to the blocks directory in your capacity extents.

Such approach will be applied to all inactive backup chains that satisfy validation criteria.

After offload is complete, the new **Capacity Tier** node appears in the **Home** view, under the **Backups** node and shows backups that have been moved to capacity extents.

> **NOTE**
>
> The offload is not performed during prohibited hours specified in the scale-out backup repository backup window configuration. You can configure the backup window at the Add Capacity Tier step of the **New Scale-out Backup Repository** wizard.

# Offload Session Statistics

The offload session results are saved to the configuration database and available for viewing, as described in section Viewing Offload Session Results.

## Backup Chain Detection

Veeam Backup & Replication will transfer to capacity extents only those restore points that belong to inactive backup chains. To ensure a backup chain is inactive, Veeam Backup & Replication verifies its state. This does not apply to the copy policy: all newly created restore points are copied immediately.

When a backup job runs for the first time, Veeam Backup & Replication creates an initial full backup file. It contains complete information about the VMs that are being backed up. With each subsequent backup job session, new incremental backup files are created. They contain only the changes that have occurred since the last backup session.

For forward incremental backup method, the active backup chain is the one that has not yet been sealed with a new full backup file.

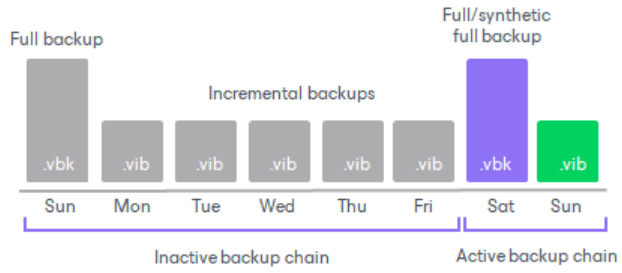


To transform an active backup chain into inactive, a new active full (or synthetic full) backup file must be created for this chain. This can be done either manually, as described in section Performing Active Full Backup. Else, you can configure a schedule according to which new active or synthetic full backups will be created automatically, as described in sections Active Full Backup and Synthetic Full Backup.

Once a new full backup file is created and the offload session is being executed, Veeam Backup & Replication collects all the restore points (full and incremental) that were created prior to the latest active full, verifies that they belong to an inactive chain, and prepares them to be moved to capacity extents. This process is called detect. For more information, see Moving Backups to Capacity Tier.

The same applies to backup chains created by backup copy jobs. Veeam Backup & Replication will transfer to capacity extents only those restore points that belong to an inactive backup copy chain. Note that if you enable backup copy GFS, Veeam Backup & Replication implements the forward incremental retention policy. If you disable backup copy GFS, Veeam Backup & Replication implements the forever forward incremental retention policy.

Note that Veeam Backup & Replication will not transfer to the capacity tier the corrupted restore points and the files dependent on those. For more information on the corrupted restore points, see Health Check for Backup Files.

The same applies to the backup chains created with the reverse-incremental method. In this case, all the .VRB files starting from the third restore point will be considered inactive automatically. Thus, you do not need to create an active full (nor synthetic full) backup manually unless you want to offload all the restore points including the most recent .VBK file and the first two .VRB files.

Consider the following examples:

- Four VRB files are inactive and can be offloaded:

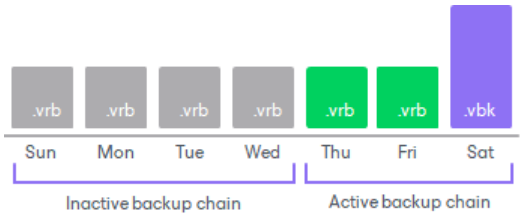


- Six VRB files and a VBK file belong to an inactive chain and can be offloaded:

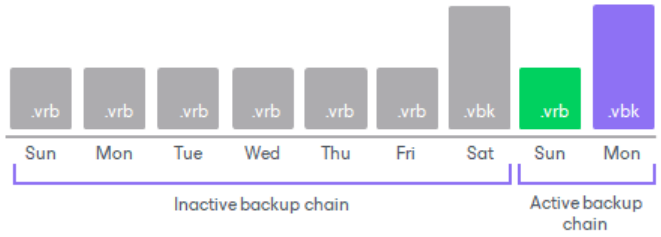


> **NOTE**
>
> A full backup file and the first two incremental backup files (that is, two .VRB files that immediately follow the most recent .VBK file) will not be offloaded until another full backup file is created successfully.

The structure of the backup chains can be different. That depends on whether your backups were created using the per-machine method (for more information, see Backup Chain Formats) or as a single-file backup, with all VMs placed into a single file. The type of the backup chain structure does not affect the offload process.

For more information on how Veeam Backup & Replication creates and manages backup chains, see Backup Chain.

## Manually Moving Backups to Capacity Tier

The **Move to capacity tier** option allows you to manually offload selected backup files to capacity extents.

Consider that backup files you want to offload must belong to an inactive backup chain. For more information, see Backup Chain Detection.

To move backup data to capacity tier, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** > **Disk** node (for backup repositories) or **Backup** > **Object Storage** (for object storage repositories) node.
3. In the working area, right-click a backup job and select **Properties**.
4. In the **Properties** window, right-click a backup file that you want to offload and select **Move to capacity tier**.

   Alternatively, you can use the **Move to capacity tier** control at the bottom.

5. Select the necessary backup file and specify the offload options:
   - **This backup and its dependencies** — to offload the selected incremental backup file (VIB) and all increments related to it.
   - **This backup** — to offload the selected full backup file (VBK) only.
   - **This backup and all dependent increments** — to offload the selected full backup file (VBK) along with its increments.

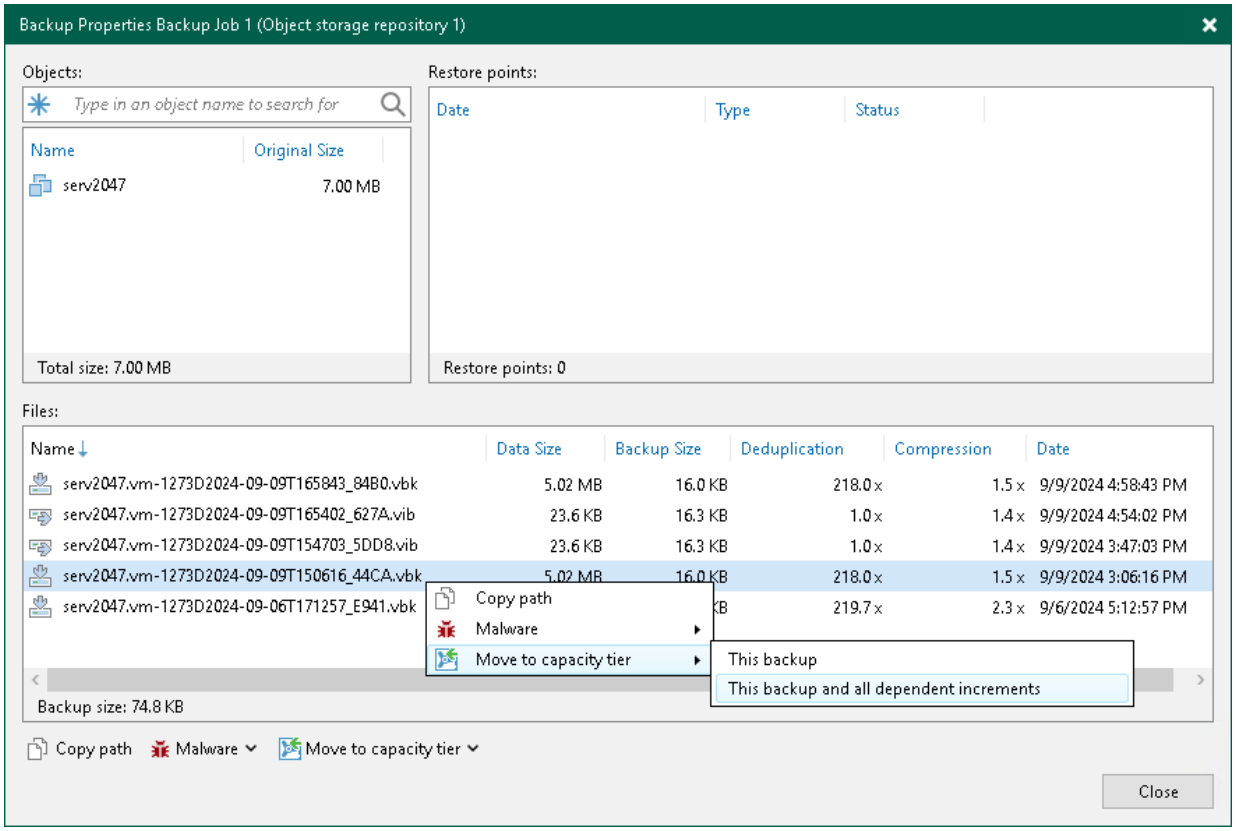


## Downloading Data from Capacity Tier

In Veeam Backup & Replication, you can manually download offloaded backups back to the performance extents. You can download one backup at a time using the **Copy to Performance Tier** option, or get all offloaded backups in bulk using the **Download** feature.

For more information, see How Downloading from Capacity Tier Works.

> **NOTE**
>
> Veeam Backup & Replication uses the Fast Clone technology only if you download all offloaded backups back to the performance extents using the **Download** option. Fast Clone is not supported for the **Copy to Performance Tier** option.

# Downloading Single Backup Chain

To download previously offloaded backup data back to the performance extents, one backup at a time, do the following:

1. Open the **Home** view.

2. In the inventory pane, select the **Backups** or **Backups** > **Capacity Tier** node.
3. In the working area, right-click a backup job and select **Properties**.
4. In the **Properties** window, right-click an offloaded backup file and select **Copy to performance tier**.

   Alternatively, you can use the **Copy to performance tier** control at the bottom.
5. Select one of the following options:
   - For VIB/VBK backup files:
     - **This backup and all dependent increments** — to copy the selected backup along with its increments.
   - For VBK backup files:
     - **This backup** — to copy the full backup only.
     - **This backup and all dependent increments** — to copy the selected backup along with its increments.

**NOTE**

To remove copied blocks from the performance extents, use the **Move to capacity tier** option, as described in section Manually Moving Backups to Capacity Tier.

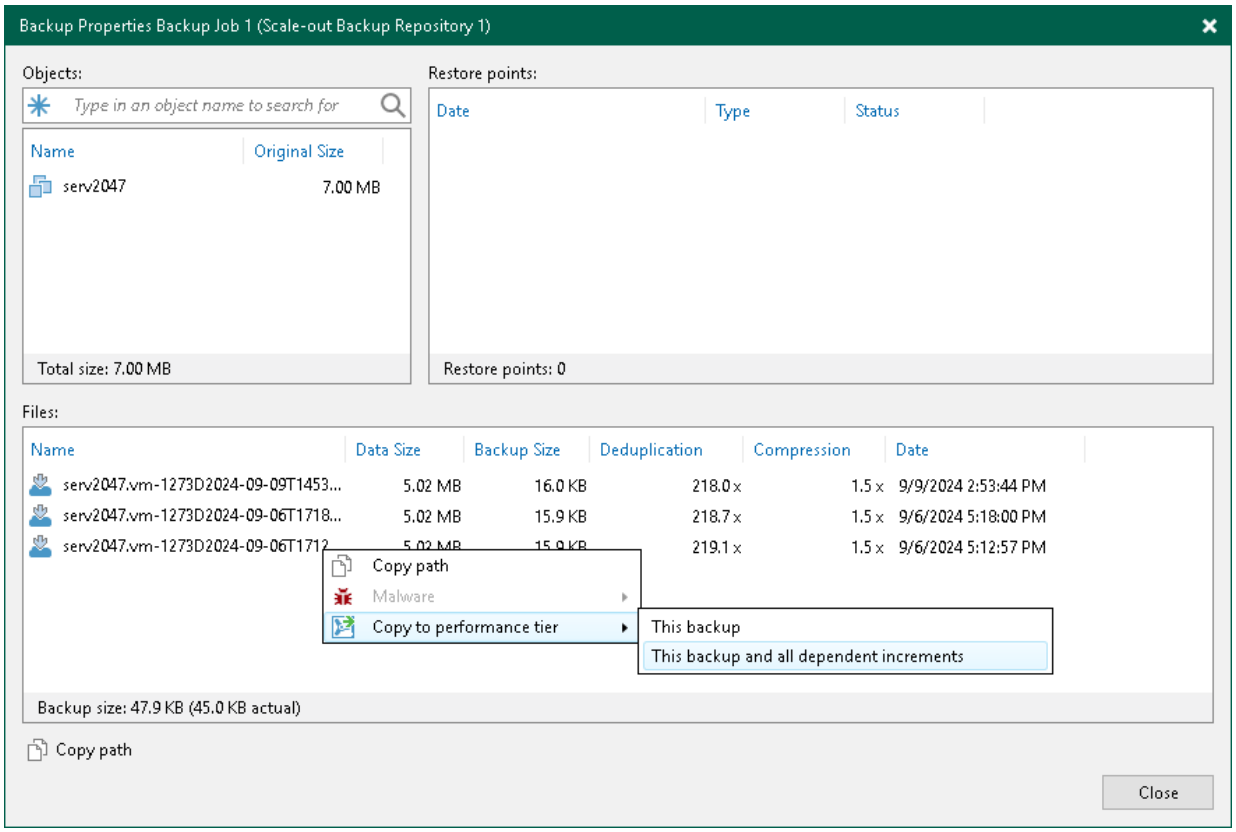


# Downloading All Backups

To download all offloaded backups back to the performance extents, do the following:

1. Open the **Home** view.

2. In the inventory pane, select the **Backups** node and select a backup job.
3. On the ribbon, click **Download** or in the working area, right-click a backup job and select **Download**.
    - If the latest backup chain is already present on your performance extent, you will be asked if you wish to download all other backup files. Click **Yes** if you wish to do so or **No** to cancel the download.
    - If the latest backup chain is not on your performance extent yet, you will be asked which backup files you want to download. Click **All Backups** to download all backup files or **Latest Only** to download just the latest backup chain.

> **NOTE**
>
> To remove downloaded data from the performance extents, use the **Move to capacity tier** option, as described in section Manually Moving Backups to Capacity Tier.

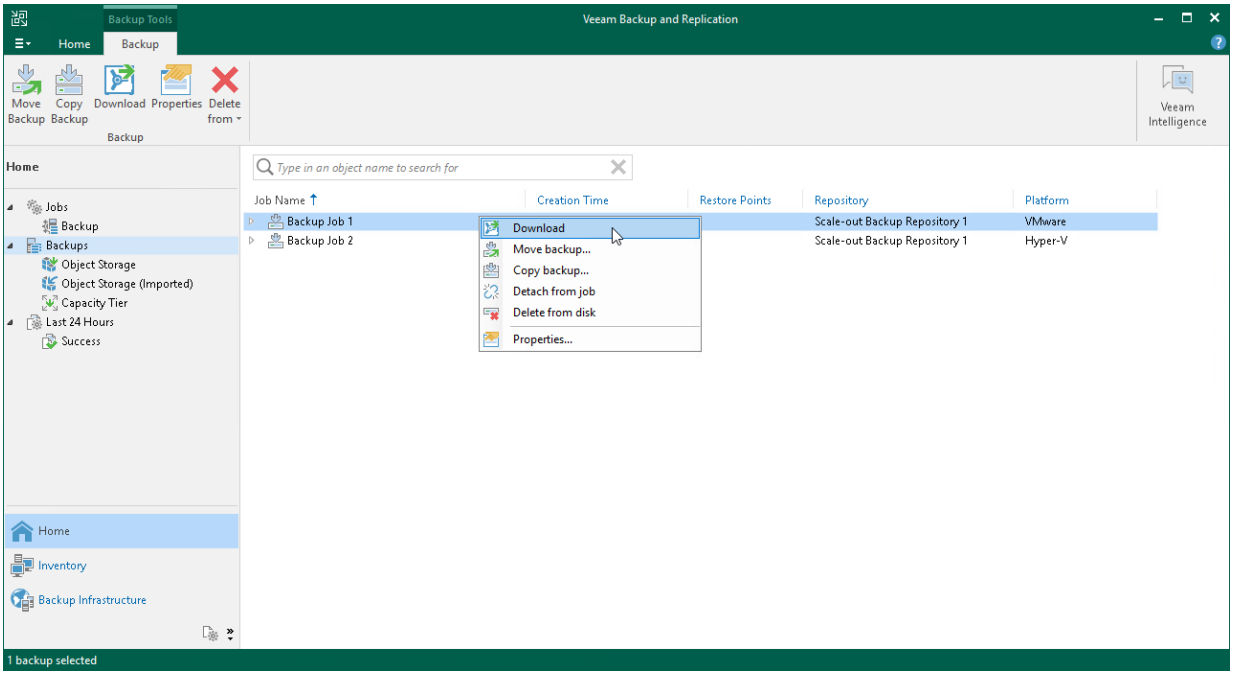


## How Downloading from Capacity Tier Works

To download data from the capacity extent back to the performance extents, Veeam Backup & Replication uses the SOBR Download job.

The SOBR Download job is triggered right after you select the **Copy to performance tier** option; it collects offloaded blocks of data from capacity extents and copies them back to the performance extents. For more information, see Downloading Data from Capacity Tier.

Consider the following:

- Veeam Backup & Replication copies all data blocks directly from object storage repositories added as extents to capacity tier.
- If a performance extent is unable to accommodate copied data due to lack of free storage space, Veeam Backup & Replication will find another extent in the associated scale-out backup repository that has sufficient storage capacity to receive the data. If your scale-out backup repository has no performance extents other than the one running out of space, the copy will not be possible.

- If you have removed any of the performance extents from a scale-out backup repository without evacuating backup files with metadata, the copy will not be possible until the files with metadata are downloaded back to the performance extents in the course of the rescan process. For more information on the rescan process, see Rescanning Scale-Out Repositories.

  Backup files with metadata are created as described in section Moving Backups to Capacity Tier.

- The SOBR download session results are saved to the configuration database and available for viewing, as described in section Viewing Download Session Results.

## Viewing Capacity Tier Sessions Statistics

Veeam Backup & Replication allows you to perform the following actions:

- View Offload Session Results.
- View Copy Session Results.
- View Download Session Results.

## Viewing Offload Session Results

To review offload session results, do the following:

1. Open the **History** view.
2. In the inventory pane, select the **Storage management** node.
3. In the working area, right-click an offload session and select **Statistics**.

For more information, see Moving Backups to Capacity Tier.

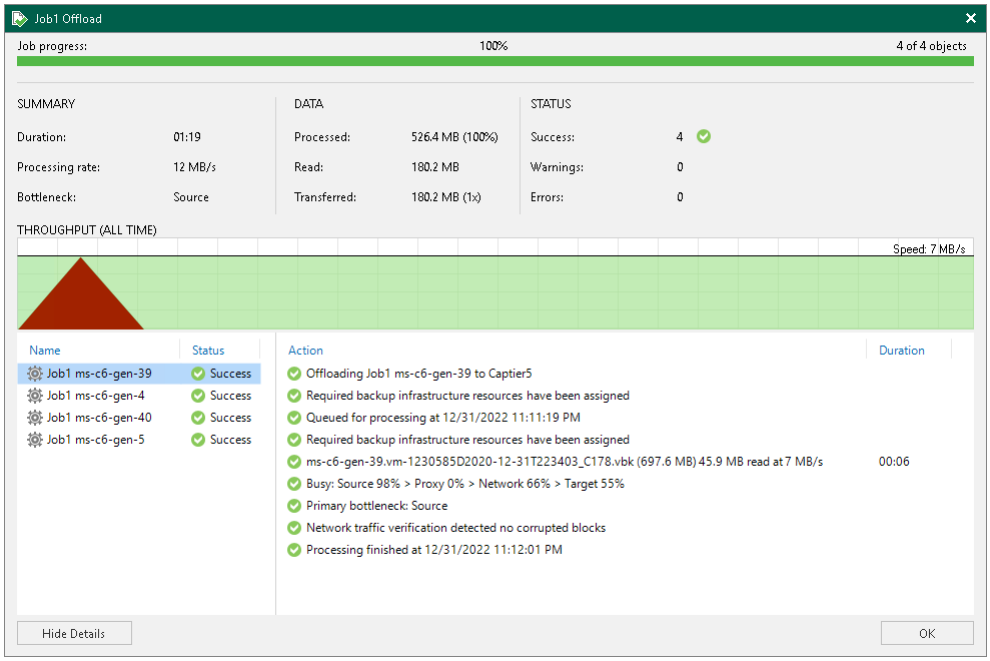


Veeam Backup & Replication displays offload job session statistics for the following counters:

- The **Job progress** bar shows percentage of the offload session completion.
- The **Summary** box shows general information about the offload session:
  - **Duration** – duration of the offload session.

- **Processing rate** – average speed of VM data processing. This counter is a ratio between the amount of data that has actually been read and the offload session duration.
- **Bottleneck** – bottleneck in the data transmission process. To learn more about bottlenecks, see Performance Bottlenecks.

- The **Data** box shows information about processed data:
  - **Processed** – total size of all VM disks processed by the offload session.
  - **Read** – the amount of data read from from the performance tier to the capacity tier.
  - **Transferred** – the amount of data transferred from performance tier to capacity tier.
- The **Status** box shows information about the job results. This box informs how many tasks have completed with the *Success*, *Warning* and *Error* statuses (1 task per 1 VM).
- The pane in the lower-left corner shows a list of objects processed by the offload session.
- The pane in the lower-right corner shows a list of operations performed during the session. To see a list of operations for a specific object, click the object in the pane on the left. To see a list of operations for the whole offload session, click anywhere on the blank area in the left pane.

## Viewing Copy Session Results

To review copy session results, do the following:

1. Open the **History** view.
2. In the inventory pane, select the **Storage management** node.
3. In the working area, right-click a copy session and select **Statistics**.

For more information, see Copying Backups to Capacity Tier.

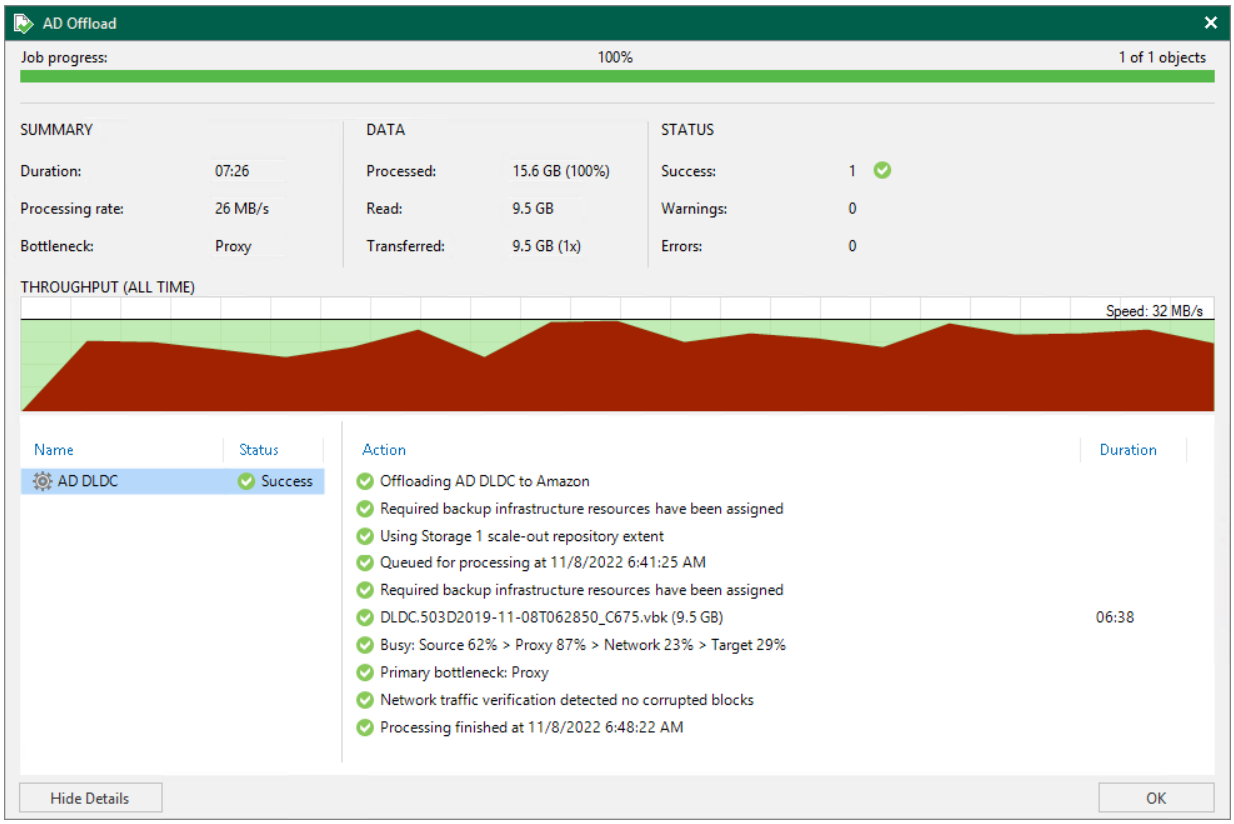

Veeam Backup & Replication displays offload job session statistics for the following counters:

- The **Job progress** bar shows percentage of the copy session completion.
- The **Summary** box shows general information about the copy session:
  - **Duration** – duration of the copy session.
  - **Processing rate** – average speed of VM data processing. This counter is a ratio between the amount of data that has actually been read and the offload session duration.
  - **Bottleneck** – bottleneck in the data transmission process. To learn more about bottlenecks, see Performance Bottlenecks.
- The **Data** box shows information about processed data:
  - **Processed** – total size of all VM disks processed by the copy session.
  - **Read** – the amount of data read from the extents.
  - **Transferred** – the amount of data transferred from the performance tier to the capacity tier.
- The **Status** box shows information about the copy results. This box informs how many tasks have been completed with the *Success*, *Warning* and *Error* statuses (1 task per 1 VM).
- The pane in the lower-left corner shows a list of objects processed by the copy session.
- The pane in the lower-right corner shows a list of operations performed during the session. To see a list of operations for a specific object, click the object in the pane on the left. To see a list of operations for the whole copy session, click anywhere on the blank area in the left pane.

## Viewing Download Session Results

To review **SOBR Download** job session results, do the following:

1. Open the **History** view.
2. In the inventory pane, select the **Storage management** node.
3. In the working area, right-click a **SOBR Download** session and select **Statistics**.

For more information, see How Downloading from Capacity Tier Works.

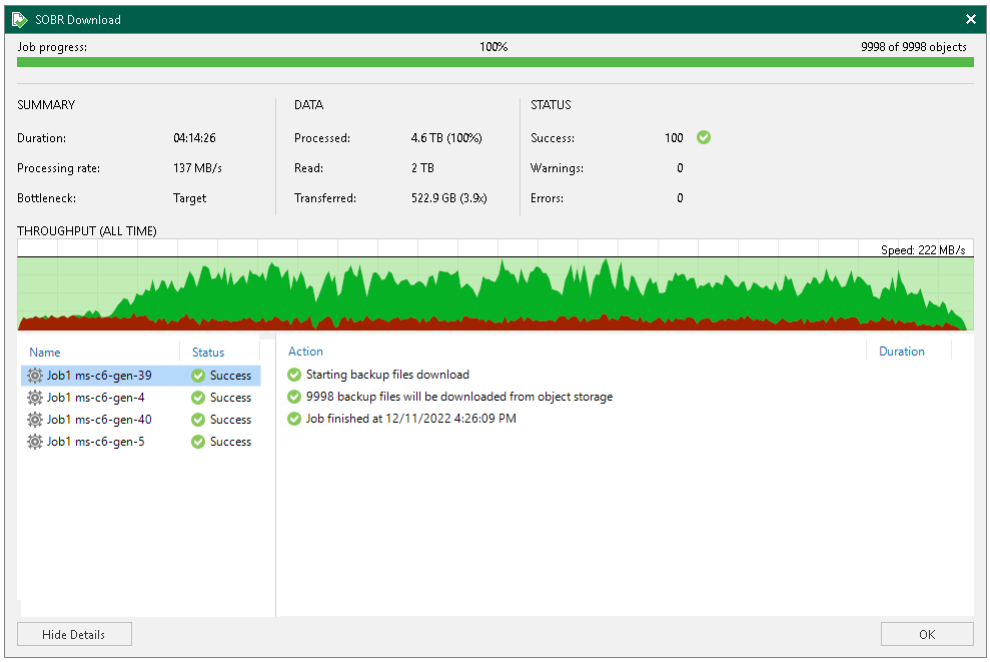


Veeam Backup & Replication displays **SOBR Download** job session statistics for the following counters:

- The **Job progress** bar shows percentage of the job completion.
- The **Summary** box shows general information about the job session:
  - **Duration** – duration of the job session.
  - **Processing rate** – average speed of data processing. This counter is a ratio between the amount of data that has actually been read and the job session duration.
  - **Bottleneck** – bottleneck in the data transmission process. To learn more about bottlenecks, see Performance Bottlenecks.
- The **Data** box shows information about processed data:
  - **Processed** – total size of data blocks being downloaded from object storage repository plus blocks (if any) being taken from the extents of your scale-out backup repository.
  - **Read** – the amount of data read from both the object storage repository and extents of your scale-out backup repository.
  - **Transferred** – the amount of data downloaded from object storage.
- The **Status** box shows information about the job results. This box informs how many tasks have completed with the *Success*, *Warning* and *Error* statuses.
- The pane in the lower-left corner shows a list of objects processed by the job.
- The pane in the lower-right corner shows a list of operations performed during the session. To see a list of operations for a specific object, click the object in the pane on the left. To see a list of operations for the whole job session, click anywhere on the blank area in the left pane.

## Indexes

> **IMPORTANT**
>
> Starting from Veeam Backup & Replication 12, indexes are no longer used. Indexes will be removed after the first offload session once you upgrade to Veeam Backup & Replication version 12.

To reduce the number of cost-based operations incurred by your cloud storage provider and to decrease the amount of traffic being sent over the network when moving or copying data to object storage, Veeam Backup & Replication uses indexes.

Index behavior has the following peculiarities:

- Indexes are created (or updated) during each offload or copy session and consist of hash values of blocks that are being transferred to the capacity tier. These hashes are retrieved from meta information of your backup files (VBK, VIB, or VRB).
- Indexes are stored in the *ArchiveIndex* directory that is located on the Performance Tier.

  On each subsequent offload/copy session, Veeam Backup & Replication reuses these indexes to verify whether new blocks that are about to be transferred to the capacity tier have not been offloaded earlier. Verification is done by comparing existing indexes hashes with that of a block being transferred.
- If backups are created using the per-machine method, indexes are built per backup chain and cannot have any cross references to any other backup chains. If all VM data is backed up to a single storage, indexes are built for the whole backup.
- Indexes are updated every time a backup chain is modified.

  For example, some data may have been removed due to the retention policy threshold, or you may have removed it manually. Both scenarios will modify your indexes upon the next successful offload or copy session to maintain consistency.
- Corrupted indexes can be rebuilt by using the **Rescan** feature, as described in section Rescanning Scale-Out Repositories.

  Once an index is rebuilt, Veeam Backup & Replication will have to wait for 24 hours before it can offload any data again. This is necessary to comply with the eventual consistency model of S3 compatible object storage repositories.

# ArchiveIndex Directory Structure

When Veeam Backup & Replication creates indexes, it also creates and maintains the following directory structure on each extent.

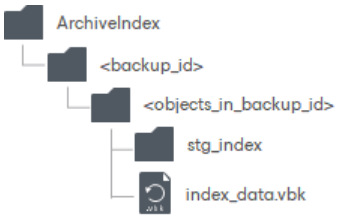


| Directory | Description |
|---|---|
| **ArchiveIndex** | The root directory for keeping indexes. |

| Directory | Description |
| --- | --- |
| **<backup_id>** | Contains objects in a backup file. |
| **<objects_in_backup_id>** | An identifier of an object in a backup file.<br>• If a backup was created using the per-machine method, then each VM will be placed to its own directory.<br>• If a backup was created as a single-file backup, then all the VMs will be placed to a unique directory. |
| **stg_index** | Contains indexes of offloaded backup files (VBK, VIB, or VRB). |
| **index_data.vbk** | Contains meta information on hash values stored in index files. |

## Retention Policy

Retention policy defines the number of restore points to keep on your performance extents and capacity extents is configured at the Specify Backup Storage Settings step of the **New Backup Job** wizard.

You can manage retention policies to remove obsolete restore points both from the performance extents and the capacity tier.

The restore points that fall under the retention policy will be removed both from the performance and capacity extents in the following manner:

- An earliest restore point will be removed from the backup chain on the associated extent.
- Data blocks that correspond to the restore point that is being removed will be purged from the capacity extent upon the next offload or copy session.

  Make sure that the capacity extent has not been put into the Maintenance mode, as this mode prevents synchronization of the state of the backup chain in the performance tier with that of the capacity tier.

  For more information about the move and copy offload sessions, see Moving Backups to Capacity Tier and Copying Backups to Capacity Tier.
- Immutable blocks of data is removed after the immutability period is over.

  When a retention policy encounters immutable copied/moved blocks of data, it removes such blocks from the associated backup files that are located on the performance extents only, informing Veeam Backup & Replication that these blocks no longer exist and must be removed from the capacity tier once mutable.

  For more information about immutability, see Immutability for Scale-Out Backup Repositories.

## Excluding Capacity Extent from Scale-Out Repositories

You can exclude a capacity extent from the scale-out backup repository scope, for example, if you no longer want to use any third party services to store your data.

Consider that after you exclude an object storage repository that is being used as a capacity extent and is storing offloaded backup data, Veeam Backup & Replication automatically puts the excluded object storage repository into the Maintenance mode. Once a repository is in the Maintenance mode, you will not be able to restore your data from it. To switch back to normal, you will have to re-add that repository as a capacity extent and synchronize existing backup chains with your performance extents. After the synchronization is complete, the existing backups will become available as Imported. For more information, see Synchronizing Capacity Tier Data.

> **NOTE**
>
> If you have an archive extent in the same scale-out backup repository, it will also switch to the Maintenance mode when you exclude a capacity extent. Note that this is not applied if you use the same object storage repository by the same provider for performance, capacity and archive tiers, for example, by Amazon S3. In this case, after removing a capacity extent from capacity tier the performance and archive tiers will remain in a scale-out backup repository.

To exclude a capacity extent from the scale-out backup repository scope, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Scale-out Repositories**.
3. In the working area, select a scale-out backup repository and click **Edit Scale-out Repository** on the ribbon, or right-click a scale-out backup repository and select **Properties**.
4. Move to the **Capacity tier** step of the wizard.
5. Clear the **Extend scale-out backup repository capacity with object storage** check box.

   If you have data on the object storage, you will be asked to confirm the action in the dialog box. After that, the object storage repository will be immediately put into the Maintenance mode.

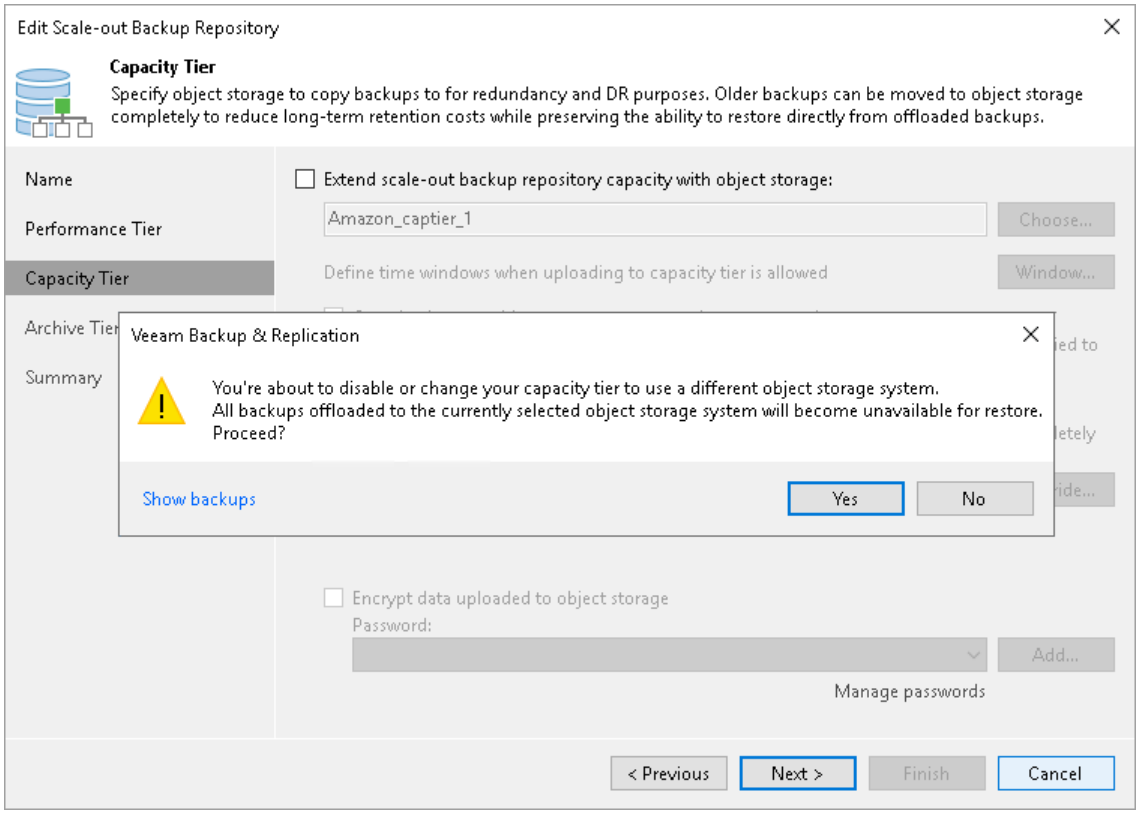

# Restore from Capacity Tier

You can restore your data directly from the capacity tier back to production servers or to Microsoft Azure, Amazon EC2 or Google Cloud platforms. Capacity tier data recovery does not differ from that of a standard backup data recovery and can be performed by using any of the following methods:

- Instant Recovery to VMware vSphere
- Instant Recovery to Microsoft Hyper-V
- Entire VM Restore
- VM Files Restore
- Virtual Disk Restore
- Guest OS File Restore
- Instant Disk Recovery
- Disk Export
- Application Item Restore
- Exporting Backups

Data recovery can also be done directly to Amazon EC2, Microsoft Azure or Google as described in sections Restore to Amazon EC2, Restore to Microsoft Azure and Restore to Google Compute Engine.

> **NOTE**
>
> To increase the speed of restore to Microsoft Azure, make sure that the following requirements are met:
>
> - To restore data, use Azure restore proxy appliances.
> - Azure restore proxy appliances must be located in the same Azure region, where you keep backups and Azure VMs created from backups.

## Restore Scenarios

This section explains possible restore scenarios from the capacity tier.

# Unavailability of Backup Files

If one of the offloaded backup files becomes unavailable on any of the extents in a scale-out backup repository, you can restore it by doing the following:

1. Rescan a scale-out backup repository.

   For more information on how to rescan a scale-out backup repository, see Rescanning Scale-Out Repositories.

2. Copy data from the capacity tier to the performance tier.

   For more information on how to copy data, see Downloading Data from Capacity Tier.

> **NOTE**
>
> The performance or capacity extents must not be in the Maintenance mode.

# Unavailability of Extents

A performance extent in a scale-out backup repository may become unavailable or be in the Maintenance mode. To restore data in such case, you can use any method described in section Data Recovery.

For example, you are restoring a virtual machine consisting of three data blocks, of which two blocks reside on the *Extent 1* and another required block is stored on the *Extent 2* which is unavailable. In such scenario, Veeam Backup & Replication gets two blocks from the *Extent 1* and another required block from the capacity tier.

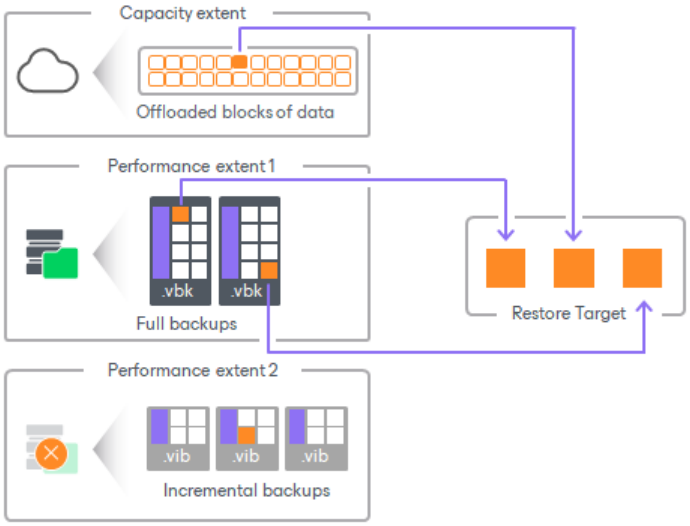


# Unavailability of Scale-Out Backup Repository

To restore data if the entire scale-out backup repository is unavailable, you can use any method described in section Data Recovery.

If the entire scale-out backup repository becomes unavailable, Veeam Backup & Replication restores data from the capacity tier only.

For example, both performance extents that store required backup files to restore a virtual machine are not available. In such a scenario, Veeam Backup & Replication restores data from the capacity tier only.

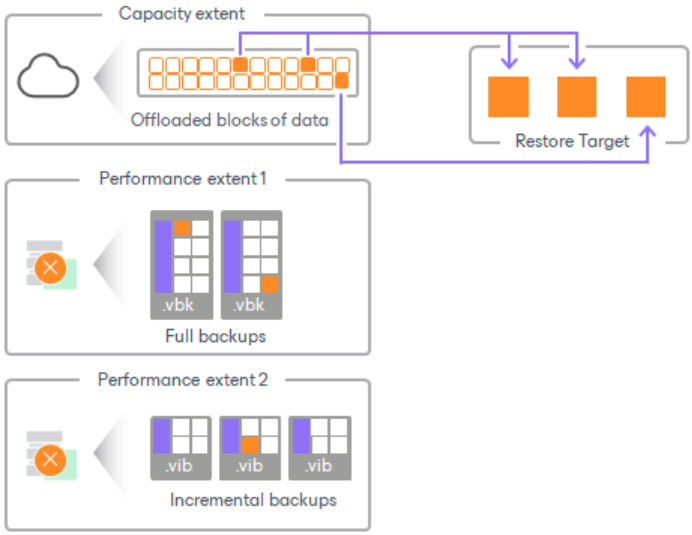

# Unavailability of Backup Server

To get access to backups in the capacity tier in case the entire configuration of the backup server becomes corrupted and your scale-out backup repositories are no longer available, you can:

- Restore the configuration of the backup server from the configuration backup, as described in section Managing Configuration Database.
- Import backups from the capacity tier, as described in section Importing Object Storage Backups.

# Archive Tier

Archive tier is an additional tier of storage that can be attached to a scale-out backup repository. You can transfer data to the archive tier directly from the performance tier or from the capacity tier. To transfer data to the archive tier, the archiving job session runs periodically. Archived data in the archive tier is cheaper than in the performance and capacity tier. However, restoring data from the archive tier is longer and more expensive compared to the capacity tier. Data must be prepared for restore from the archive tier.

This feature is most useful in the following cases:

- You have a lot of rarely (no more than once a quarter) accessed data that has to be stored in an archive.
- You want to save costs and space on storing archived data.

## Supported Types of Object Storage Repositories

The archive tier consists of a single archive extent. One of the following cloud-based object storage repositories providing "cold" data storage can be added as the archive extent:

- Amazon S3 Glacier.
- S3 compatible object storage with data archiving.
- Microsoft Azure Archive Storage.

Before an object storage repository can be configured as the archive extent, it must be added to Veeam Backup & Replication backup infrastructure. For more information, see Adding Amazon S3 Glacier Storage or Adding Azure Archive Storage.

You can add the archive extent to your scale-out backup repository and configure its settings on the Add Archive Tier step of the **New Scale-out Backup Repository** wizard.

## Direct Data Transfer to Archive Tier

Veeam Backup & Replication supports multiple scenarios of direct data transfer to the archive tier.

- Direct data transfer from the performance to the archive tier if your scale-out backup repository has the following configuration:
  - The performance tier consists of direct attached storage (Microsoft Windows Server, Linux Server, Hardened Repository) and the archive tier consists of all supported types of object storage repositories.
  - The performance tier consists of network attached storage (SMB (CIFS) Share, NFS Share) and the archive tier consists of all supported types of object storage repositories.
- Direct data transfer from the performance or the capacity tier to the archive tier if your scale-out backup repository has the following configuration:
  - The performance tier or the capacity tier consists of Veeam Data Cloud Vault and the archive tier consists of Azure Archive Storage.
  - The performance tier or the capacity tier consists of Amazon S3 Storage and the archive tier consists of Amazon S3 Glacier Storage.
  - The performance tier or the capacity tier consists of Azure Blob Storage and the archive tier consists of Azure Archive Storage.

- The performance tier or the capacity tier consists of S3 Compatible Object Storage and the archive tier consists of all supported types of object storage repositories.
- The performance tier or the capacity tier consists of IBM Cloud Object Storage and the archive tier consists of S3 compatible object storage with data archiving.
- The performance tier or the capacity tier consists of Wasabi Cloud Object Storage and the archive tier consists of S3 compatible object storage with data archiving.
- The performance tier or the capacity tier consists of 11:11 Cloud Object Storage and the archive tier consists of either Amazon S3 Glacier Storage or S3 compatible object storage with data archiving.

# Supported Types of Backup Files

The following types of backup files are supported for archive storage:

- Backup files with GFS flags.
- VeeamZIP backup files.
- Exported backup files.
- Orphaned backups with GFS flags.
- Backups created by Veeam Backup for Nutanix AHV.
- Backups created by Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization.
- Backups exported by Kasten policies.

# Limitations for Archive Tier

The archive tier has the following limitations:

- If you transfer data to the archive tier from an object storage repository, the archive tier configuration depends on the object storage provider. For more information, see the Direct Data Transfer to Archive Tier section.
- Veeam Backup & Replication does not support direct data transfer from performance extents that consist of Google Cloud Object Storage to archive extents.
- Migrating data to another archive tier is not supported.
- Imported backups cannot be offloaded to archive tier.
- Incremental backup files cannot be stored in the archive tier.
- Microsoft Azure Archive Storage with the *archive* access tier does not support Azure accounts with the following redundancy options: zone-redundant storage (ZRS), geo-redundant storage (GZRS) and read-access geo-zone-redundant storage (RA-GZRS). For more information, see Microsoft Docs.

# Archiving Job

The process of moving backup data from the performance or capacity tier to the archive tier is called an archiving job. This job runs in a separate session and follows the default schedule of the offload job (this job moves data from the performance tier to the capacity tier).

During the archiving job, new backup data is offloaded to the archive tier, while the outdated backup data is cleaned up from the archive tier during the same archiving session. To clean up outdated data, a cleanup task runs in parallel with the archiving job. After the archiving job finishes, the moved data is deleted from the performance tier. For the capacity tier, it happens during the next offload job.

# How Archiving Job Works

The archiving job works according to the following algorithm:

1. Veeam Backup & Replication gets metadata information about blobs and data blocks that should be archived.
2. If any blocks are available in the archive tier, Veeam Backup & Replication adds them to the archiving job. The archiving job will re-use these blocks within the archiving session.
3. Veeam Backup & Replication gets a list of all blocks that should be archived.
4. Veeam Backup & Replication arranges the blocks from the list into sets. The identical blocks from different tiers are arranged into one set.
5. From each set of identical blocks, Veeam Backup & Replication selects one block that will be archived. If the set contains a block not present in the archived tier, Veeam Backup & Replication adds this block to the blob.

   > **NOTE**
   >
   > Consider the following:
   >
   > - This option depends on whether the data block reuse is enabled. By default, Veeam Backup & Replication enables the data block reuse when you configure the archive tier. To disable this option, modify the settings of your archive tier.
   > - The archiving tier inherits encrypted data from the capacity tier. The blocks encrypted on the capacity level are moved to the archive tier in the encrypted state.

6. Veeam Backup & Replication arranges blocks into blobs according to the blob optimization algorithm.
7. Veeam Backup & Replication generates metadata for blobs that will be archived. This metadata is used to identify the location of each block in the blob.
8. Veeam Backup & Replication uploads blobs with blocks that are not present in the archive tier.
9. Veeam Backup & Replication uploads metadata to the archive tier.

# Blob Optimization

Blob optimization helps organize data efficiently and reduces the storage space required on the archive tier. During the archiving job, Veeam Backup & Replication modifies the blocks of data from the performance/capacity tier and rearranges these data into blobs. The size of these blobs differs from the size of blocks of data on the performance/capacity tier and depends on the storage optimization option of a backup job.

To arrange blocks into blobs, the archiving job analyzes the size of data blocks in the performance/capacity tier. It does not decompress data blocks, but uses already compressed data and arranges them into a blob according to the following rules:

- The blob must not contain more than 512 data blocks.
- The blob must not occupy more than 512 MB.

Depending on the storage optimization option, the size of a blob in the archive tier can be calculated the following way:

- If the storage optimization option is set to 1 or 4 MB, the archiving job rearranges data blocks into a single blob that occupies no more than 512 MB.
- If the storage optimization option is set to 512 KB, the archiving job rearranges data blocks into a single blob that occupies no more than 256 MB.
- If the storage optimization option is set to 256 KB, the archiving job rearranges data blocks into a single blob that occupies no more than 128 MB.

For example, due to the storage optimization option of a backup job, each data block located on the performance/capacity tier occupies no more than 1 MB. During the archiving job session, the job checks the block size and recompiles it into larger blobs. Thus, the total amount of blocks in a blob will be 512, and this blob will occupy no more than 512 MB.

# Considerations and Limitations

You can archive backups if they meet the following conditions:

- The type of a backup file is supported. For details, see Archive Tier.
- Backup files belong to inactive backup chains.
- Backup files have been created N days ago according to **Archive backup files older than N days** setting of the Add Archive Tier step of the wizard.
- If you have selected the **Archive backups only if the remaining retention time is above minimal storage period** check box of the Add Archive Tier step of the wizard:
  a. [For Amazon S3 Glacier storage] Only backup files with retention period no less than 90 days.
  b. [For Microsoft Azure Archive storage] Only backup files with retention period no less than 180 days.

Backup files that meet all the following conditions will be cleaned up:

- Backup files that do not have original files on the performance tier.
- Backup files with the expired or unspecified immutability period. For more information, see Immutability for Archive Tier.

# Archiver Appliances

Archiver appliance is a temporary virtual machine that is located in the Public Cloud (AWS or Microsoft Azure). Veeam Backup & Replication uses it for the following operations:

- To transfer data from the capacity tier to the archive tier.
- To transfer data from performance tier that consist of object storage repositories to the archive tier.
- To perform health check operations of data blocks located in the performance tier.

The template for all the archiver appliances is set up at the **Archiver Appliance** step of the Adding Amazon S3 Glacier or Adding Azure Archive Storage wizard. You can either create the archiver appliance with the default settings, or specify the archiver appliance settings manually: set up the size of the virtual machine and cloud resources where the archiver appliance will be created.

> **NOTE**
>
> Veeam Backup & Replication must be able to connect to the machine that you will use as an archiver appliance. Therefore, if your backup server is not located within the Public Cloud (AWS or Microsoft Azure), you must configure public IP addresses for the Amazon subnet or Azure virtual network in which the appliance resides. For more information on configuring the subnet for Amazon VPC, see AWS Documentation. For more information on configuring the Azure virtual network, see Microsoft Docs.

After the archiving job is finished, all the archiver appliances are automatically deleted. If the job ends prematurely, the archiver appliances will be deleted as well. Also, any archiver appliance can be deleted if there are no more tasks for it.

Note that the archiever appliance is not used for the direct data transfer from the performance tier to the archive tier if your scale-out backup repository has the following configuration:

- Archive tier consists of S3 compatible object storage with data archiving.
- Performance tier consist of a direct attached storage or network attached storage and archive tier consist of Amazon S3 object storage or Microsoft Azure Blob storage.

In these cases, the component that transfers data depends on the connection type of the object storage repository:

- If you use the direct connection type, the data mover that runs on performance extents transfers data to the archive tier.
- If you use the connection through gateway servers, the specified gateway servers transfer data to the archive tier.

# Archive Extent Structure

When data is transferred to the archive extent, Veeam Backup & Replication creates and maintains the following structure of directories:

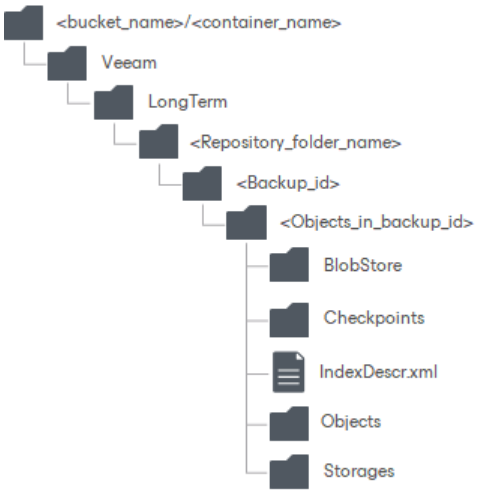


| Directory | Description | Misc |
|---|---|---|
| **<bucket_name> or <container_name>** | A bucket or container name.<br>Buckets and containers must be created in advance. | N/A |

| Directory | Description | Misc |
|---|---|---|
| **Veeam/LongTerm/** | Standard folders created by Veeam Backup & Replication. | |
| **<repository_folder_name>** | A repository folder that you create when adding a new archive extent. | |
| **<backup_id>** | Contains objects in a backup. | These folders are automatically removed during data removal. |
| **<objects_in_backup_id>** | An identifier of an object in a backup. | |
| **IndexDescr** | Contains size of a backup. | |
| **BlobStore** | Contains data blobs created by the archiving session, as described in section Archive Tier Data Transfer. | |
| **Checkpoints** | Contains meta information about the state of archived backup chains. Such meta information is updated upon each successful archiving session. | |
| **Objects** | Contains meta information and other auxiliary data. | |
| **Storages** | Contains a replicated version of archived backup files with metadata. | |

# Immutability for Archive Tier

Veeam Backup & Replication allows you to prohibit deletion of data from the archive extent by making that data temporarily immutable. It is done for increased security: immutability protects your data from loss as a result of attacks, malware activity or any other injurious actions.

You can enable immutability for data stored in Amazon S3 Glacier, S3-compatible with data archiving and Microsoft Azure Archive Storage repositories used as archive extents of the scale-out backup repository. After you enable immutability, Veeam Backup & Replication will prohibit data deletion from the archive tier until the immutability expiration date comes.

When you enable immutability for the archive tier, keep in mind that only the settings of the archive extents will be taken into account. The settings of the capacity extents and of the original data blocks will be ignored.

For Amazon S3 Glacier, S3-compatible with data archiving and Microsoft Azure Archive Storage, all the types of files that are suitable for archive storage can be made immutable:

- Backup files with GFS flags assigned: in case GFS retention is extended in the backup job or backup copy job settings, the immutability period for existing backup files will be prolonged at the end of the archiving session. For more information about GFS retention policy, see Long-Term Retention Policy (GFS).

- VeeamZIP backup files with specified retention (deletion date). For more information, see Creating VeeamZIP Backups.
- Exported backup files with specified retention (deletion date). For more information, see Exporting Backups.

The immutability period of a backup file will be equal to its retention period at the moment of archiving. If the retention period is not specified for VeeamZIP backup files or exported backup files, such files will not be made immutable.

# Enabling Immutability

To enable immutability, you must do the following:

1. Configure the necessary settings when you create an S3 bucket or an Azure container.

   For more information, see Enabling Immutability.

2. Enable the immutability option when you add an object storage repository to the backup infrastructure at the **Container** step (for Azure object storage repository) or **Bucket** step (for Amazon S3 or S3 compatible object storage repositories) of the new **Object Storage Repository** wizard.

# Encryption for Archive Tier

Veeam Backup & Replication allows you to encrypt archived data. This helps you protect the data from unauthorized access.

You can enable data encryption in the following ways:

- When you create a backup or backup copy job.
- When you add an archive extent to your scale-out backup repository.

To get benefits of both encryption levels, you can use job-level and archive tier encryption within the same object storage. Both encryption levels prevents unauthorized access to your data, but the archive tier encryption allows you to encrypt backup chain metadata and restore points. After you enable encryption on the archive tier level, the archiving job encrypts the entire collection of blocks along with the metadata.

To encrypt data, Veeam Backup & Replication uses the following scheme:

1. The user specifies the password in the archiving tier settings.
2. Veeam Backup & Replication generates an encryption key and uses the password to encrypt this key.
3. To encrypt the backups in the archive tier, Veeam Backup & Replication uses this encryption key.

To access the data in the archive extent you need to provide the password that is used to decrypt the encryption keys.

# Encryption Scenarios

When data is moved to the archiving tier, Veeam Backup & Replication applies encryption on the archive tier level. Data encrypted by the backup job stays as is on the archive tier, Thus, this backup has 2 levels of encryption: on the job level an on the archive tier level.

Depending on the configuration of a scale-out backup repository, Veeam Backup & Replication will process backups according the following scenarios:

- If you use the performance tier, Veeam Backup & Replication encrypts data on the archive tier level.

- If you use the performance tier and the capacity tier with encryption, Veeam Backup & Replication decrypts the backups from the capacity tier and encrypts them on the archive tier level.
- If you do not apply encryption on the capacity tier level, Veeam Backup & Replication will encrypt backups on the archive tier level.
- If you apply encryption on the capacity tier level, but do not apply encryption on the archive tier level, Veeam Backup & Replication will decrypt these backups and will move them to the archiving tier unencrypted.

> **IMPORTANT**
>
> When you disable encryption, the encryption key is added to the file system of your backup repository and all encrypted backups located on the archive tier become unencrypted. If you enable encryption again, Veeam Backup & Replication will encrypt only new backups with a new encryption key.

# Managing Archive Tier

You can manage your archive tier and the archived data in the following ways:

- Move outdated backups from the capacity extent or the performance extent to the archive extent.
- Delete outdated backup files.
- Exclude the archive extent from the scale-out backup repository scope.

## Archive Tier Data Transfer

You can move outdated backup files from the capacity tier or performance tier to the archive tier. After that, the transported files are deleted from the capacity tier or performance tier. They can stay in the performance tier or be deleted from there depending on the data transfer policy of the scale-out backup repository:

- In case the copy policy is selected, the original file stays in the performance tier. The copied file is copied from the performance tier to the capacity tier. After the archive job runs, it is moved from the capacity tier to the archive tier.
- In case the move policy is selected, the original file is moved from the performance tier to the capacity tier. After the archive job session, the files are moved from the capacity tier to archive tier.
- In case you transfer data from the performance tier to the archive tier, the original file is moved from the performance tier to archive tier.

## Excluding Archive Extent from Scale-Out Backup Repository

You can exclude an archive extent from the scale-out backup repository scope, for example, if you no longer want to use any third party services to store your data.

Consider that after you exclude an object storage repository that is being used as an archive extent and is storing archived backup data, Veeam Backup & Replication automatically puts the excluded object storage repository into the Maintenance mode. Once a repository is in the Maintenance mode, you will not be able to restore your data from it. To switch back to normal, you will have to re-add that repository as an archive extent and synchronize existing backup chains with data in this scale-out backup repository. After the synchronization is complete, the existing backups will become available as Imported.

To exclude an archive extent from the scale-out backup repository scope, do the following:

1. Open the **Backup Infrastructure** view.

2. In the inventory pane, click **Scale-out Repositories**.
3. In the working area, select a scale-out backup repository and click **Edit Scale-out Repository** on the ribbon, or right-click a scale-out backup repository and select **Properties**.
4. Move to the **Archive tier** step of the wizard.
5. Clear the **Archive GFS full backups to object storage** check box.

   You will be asked to confirm the action in the dialog box. After that, the object storage repository will be immediately put into the Maintenance mode.

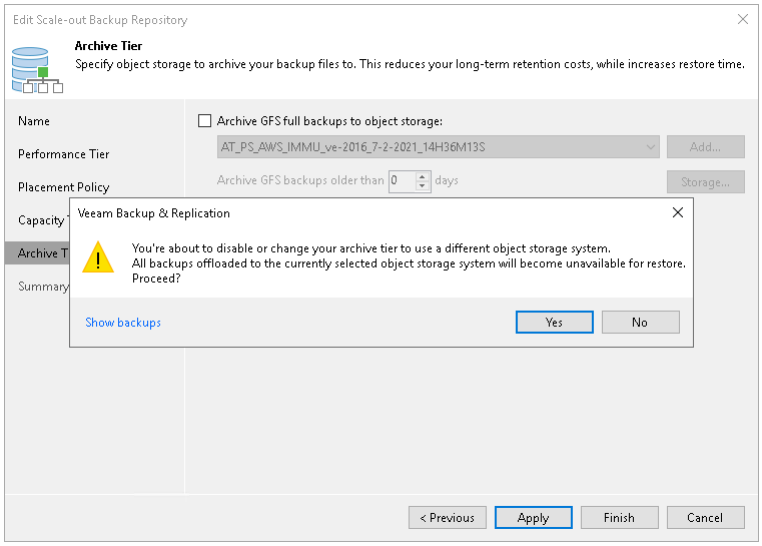


# Restore from Archive Tier

You can restore data from the archive storage. To do so, you have to receive a temporary access to the data first. Thus, the process of restore from the archive tier consists of two consecutive parts:

1. Data retrieval
2. Restoring retrieved data

## Data Retrieval

Data retrieval is the process of receiving temporary access to archived data, so that it can be restored.

The process of retrieving data from the archive tier takes course in a separate retrieval job session. It is completed when a restore point is available for reading and restore. For information on how to launch the retrieval job, see Retrieving Backup Files.

When the retrieval job is complete, the retrieved data will be available for a limited period of time, during which you can restore it. However, you can extend the availability period.

> **NOTE**
>
> If you launch restore job when retrieval job is not over yet, the restore job will be pending until the retrieval job is complete.

Retrieval cost varies depending on the desired speed of the process and on the targeted period of the data accessibility. The set of options is different for different vendors.

# Retrieval Options for Amazon S3 Glacier and S3 Compatible Object Storage with Data Archiving

Amazon and S3 compatible object storage with Data Archiving provide the following options for data retrieval. The indicated time is approximate. For more information on Amazon, see AWS Documentation.

> **NOTE**
>
> Before you start data retrieval from S3 compatible object storage with Data Archiving, verify that the necessary method is supported by your S3 compatible vendor. To get this information, contact your vendor.

- Expedited: the most expensive method. Retrieval takes 1-5 minutes.

  > **NOTE**
  >
  > This option is not available if you assign the Amazon S3 Glacier Deep Archive storage class to data blocks.

- Standard accelerated: retrieval time is faster than the standard option and allows you to retrieve a group of object. Note that it will result in additional costs.

  > **NOTE**
  >
  > To perform the standard accelerate restore, you must set the necessary permissions. For more information, see S3 Batch Retrieval Permissions for Amazon S3 Glacier.

- Standard: retrieval time 3-5 hours for Amazon S3 Glacier and within 12 hours for Amazon S3 Glacier Deep Archive.
- Bulk: the least expensive method. Retrieval time within 5-12 hours for Amazon S3 Glacier and within 48 hours for Amazon S3 Glacier Deep Archive.

## Expiration of the Retrieved Data for Amazon S3 Glacier

Estimated time of data expiration is calculated from the moment of the data retrieval launch. Required number of days is added to that moment.

However, during the retrieval job Veeam Backup & Replication constantly requests the S3 API to move the expiration time forward, until the job is completed and all the data blocks are ready.

For more information on calculating the estimated expiration time, see AWS Documentation.

For information on how to extend the expiration time, see Extending Data Availability.

# Retrieval Options for Azure Archive Storage

Azure provides the following options for data retrieval. The indicated time is approximate. For more information, see Microsoft Docs.

- High Priority: most expensive method. Retrieval may take less than one hour.
- Standard Priority: least expensive method. Retrieval takes up to 15 hours.

### Expiration of the Retrieved Data for Azure Archive Storage

Unlike Amazon S3 Glacier, Azure API does not temporarily change the storage class of a backup file from Archive Tier to Hot. Instead, a temporary copy of a backup file is created in Hot storage class. Since the deletion of this temporary copy is managed by Veeam Backup & Replication, the expiration time is quite accurate (within ten minutes, which is the frequency of the deletion process).

For information on how to extend the expiration time, see Extending Data Availability.

## Retrieving Backup Files

To start the retrieval of archived backup files, use the **Retrieve Backup** wizard.

## Step 1. Launch Retrieve Backup Wizard

To launch the retrieval job, do one of the following:

- Open the **Home** view. In the inventory pane select **Archive Tier**. In the working area, select the backup job whose files you want to retrieve and click **Properties** on the ribbon. In the **Backup Properties** window, click on **Retrieve**.
- Open the **Home** view. In the inventory pane select **Archive Tier**. In the working area, select the VM whose guest OS files you want to restore and on the ribbon select the necessary type of restore. Proceed to the **Select Restore Point** step of the wizard. If this restore point has not been retrieved yet, you will be prompted to launch the **Retrieve Backup** wizard.

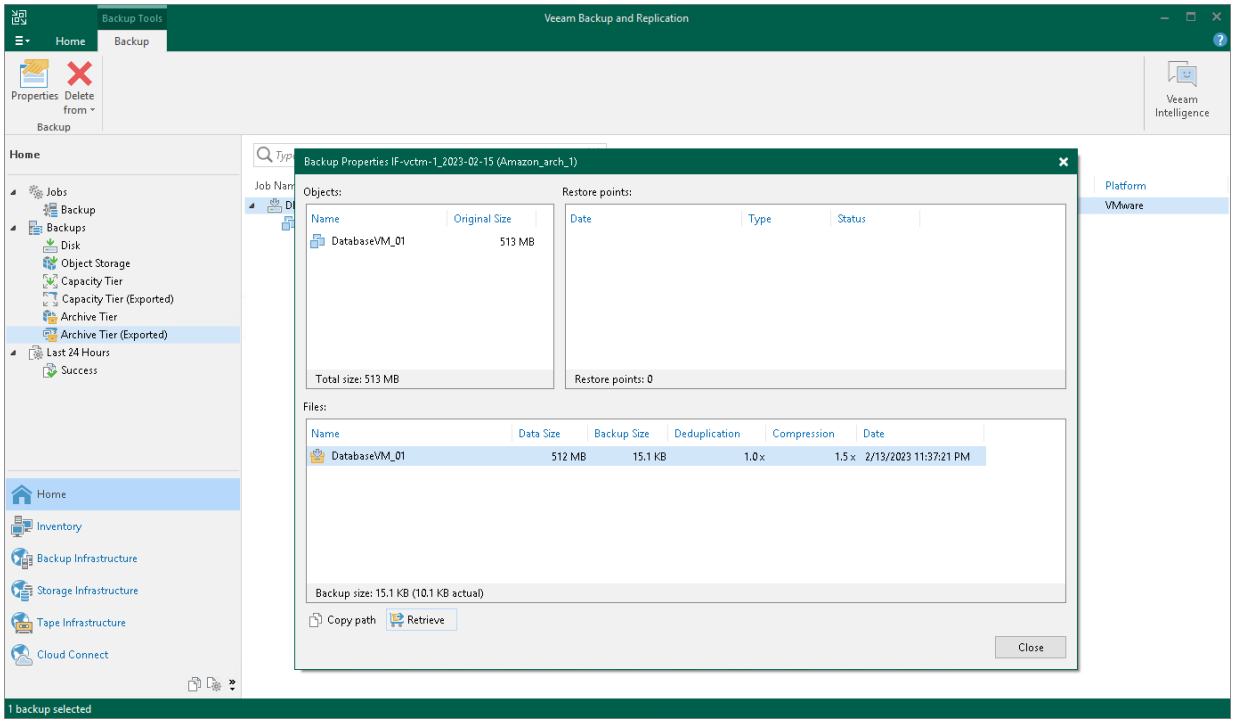

## Step 2. Select Retrieval Mode

At the **Retrieval Mode** step of the wizard, select the desired retrieval option. For information on the retrieval modes for different archive storage options, see Data Retrieval.

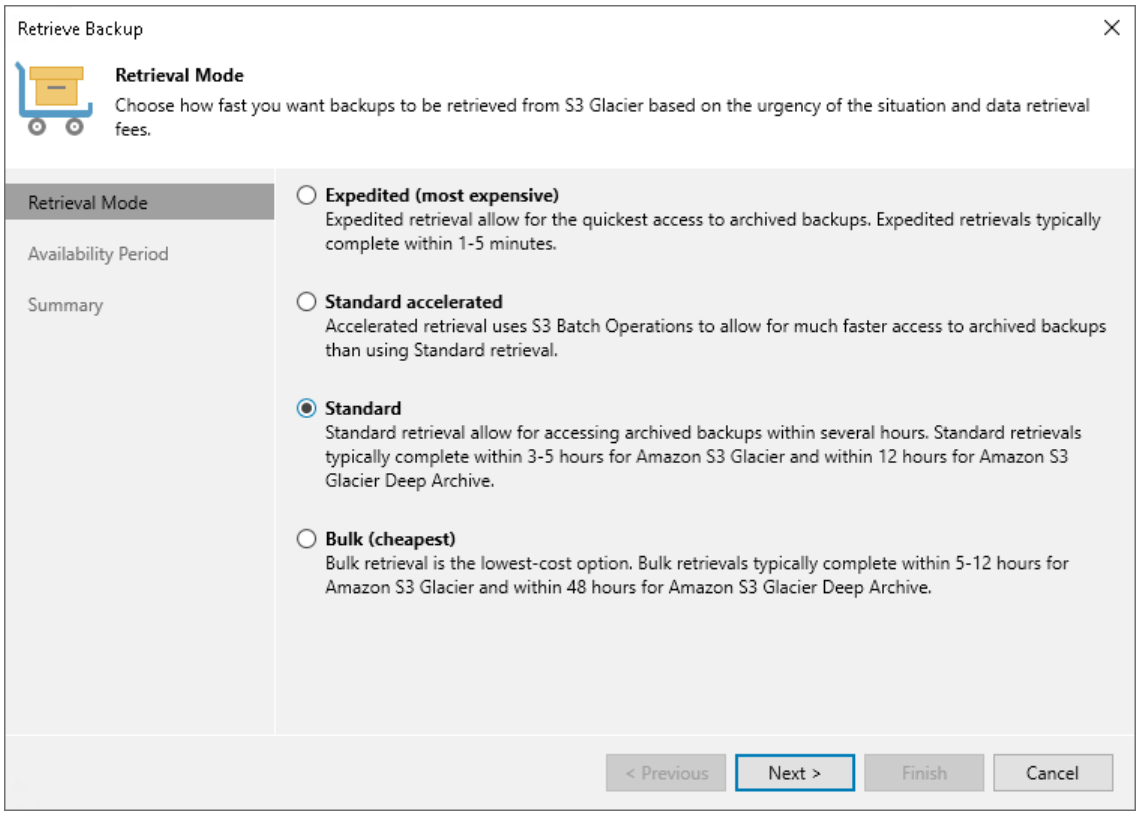

## Step 3. Select Availability Period

At the **Availability Period** step of the wizard, select the desired availability period of the retrieved backup files. During that period you will be able to restore the data.

If you want to receive a notification that the availability period is about to end, select the **Send a notification email N hours before the published data expires** check box and choose the desired time for the notification.

> **TIP**
>
> You will be able to extend that period if necessary. For information on how to extend backup files availability, see Extending Data Availability.

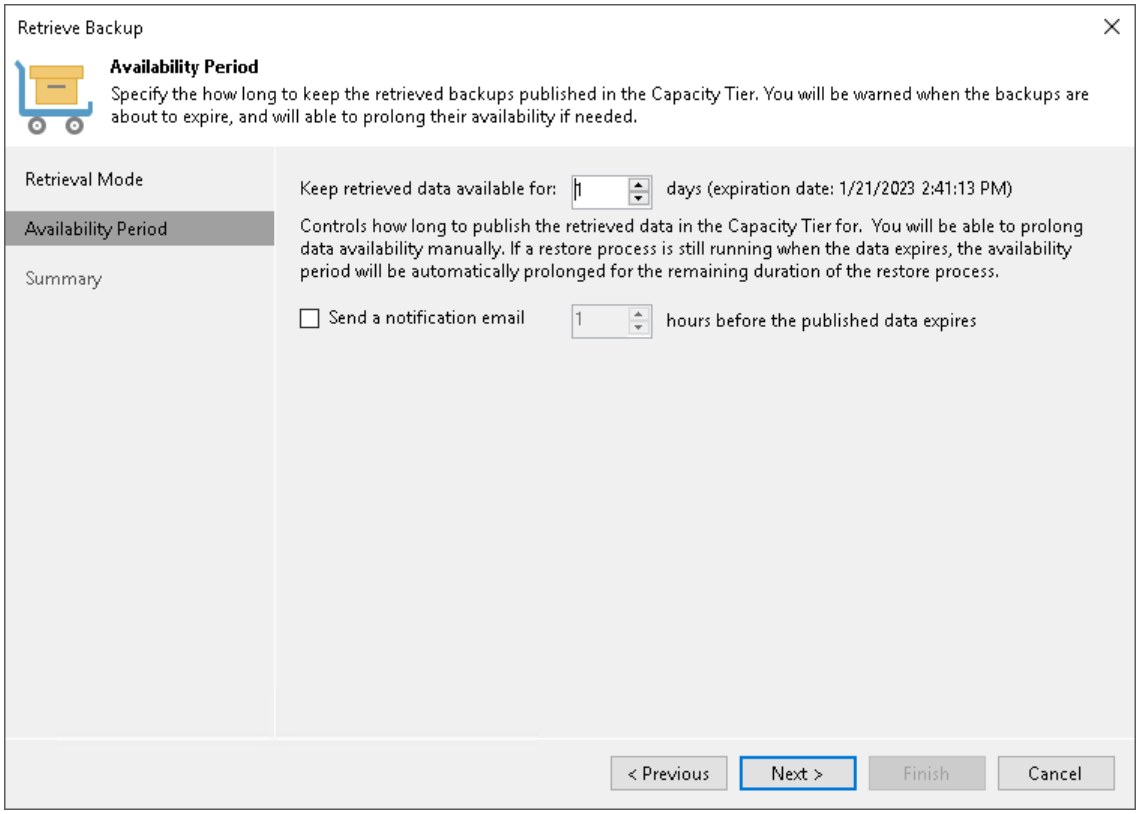

### Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review the retrieval settings. Click **Finish** to exit the wizard.

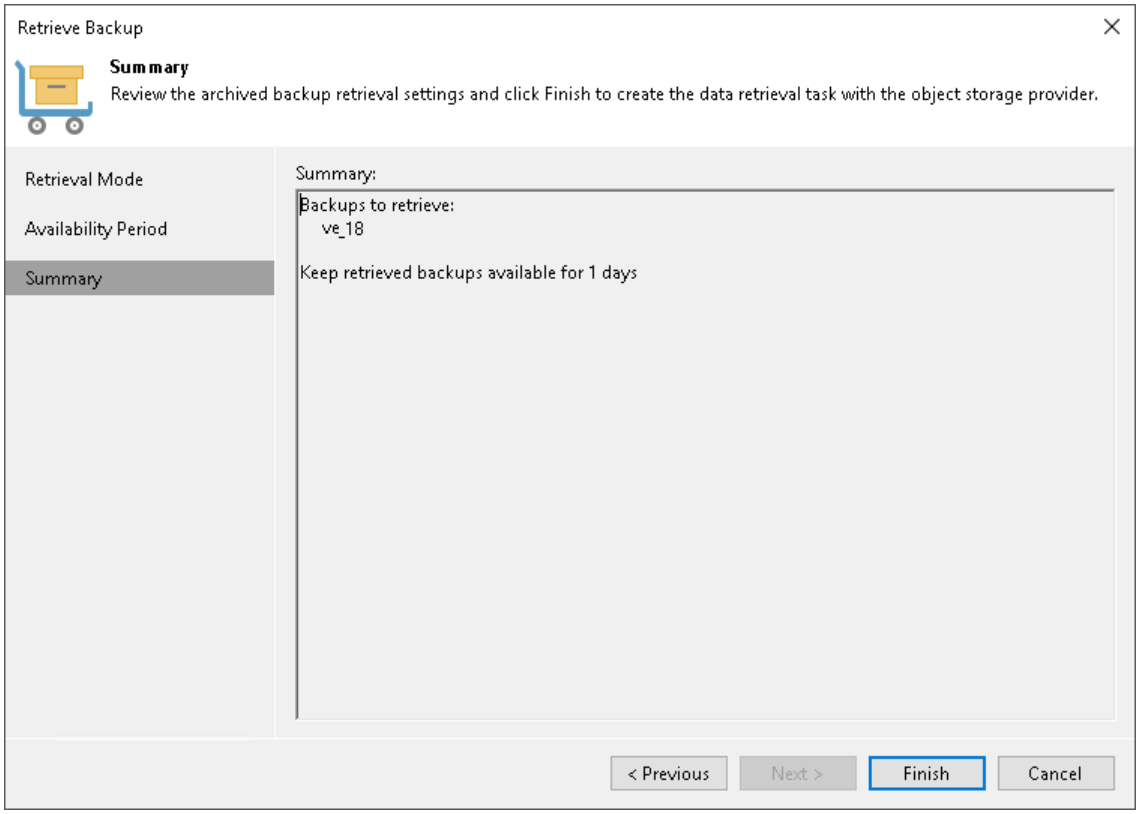


## Extending Data Availability

The availability period of the retrieved data can be prolonged:

- Manually. If the retrieval job still continues at the moment of manual extension, the new extension time is automatically added to the job.
- Automatically. This happens in cases when the expiration time is close, but the restore process still continues.

At the Select Availability Period step of the **Retrieve Backup** wizard, you can request a notification that the data availability is about to expire. If you extend the expiration time, the notification will be postponed according to the expiration time.

To extend the availability period, do the following:

- Open the **Home** view. In the inventory pane select **Data Retrieval**. In the working area, select the retrieved backup file whose availability period you want to extend, and click **Extend Availability** on the ribbon. You will be prompted to select the number of days you want to add to the availability expiration time.

- Open the **Home** view. In the inventory pane select **Archive Tier**. In the working area, select the backup job for whose retrieved files you want to extend availability and click **Properties** on the ribbon. In the **Backup Properties** window, click on **Extend Availability Period.** You will be prompted to select the number of days you want to add to current expiration time.

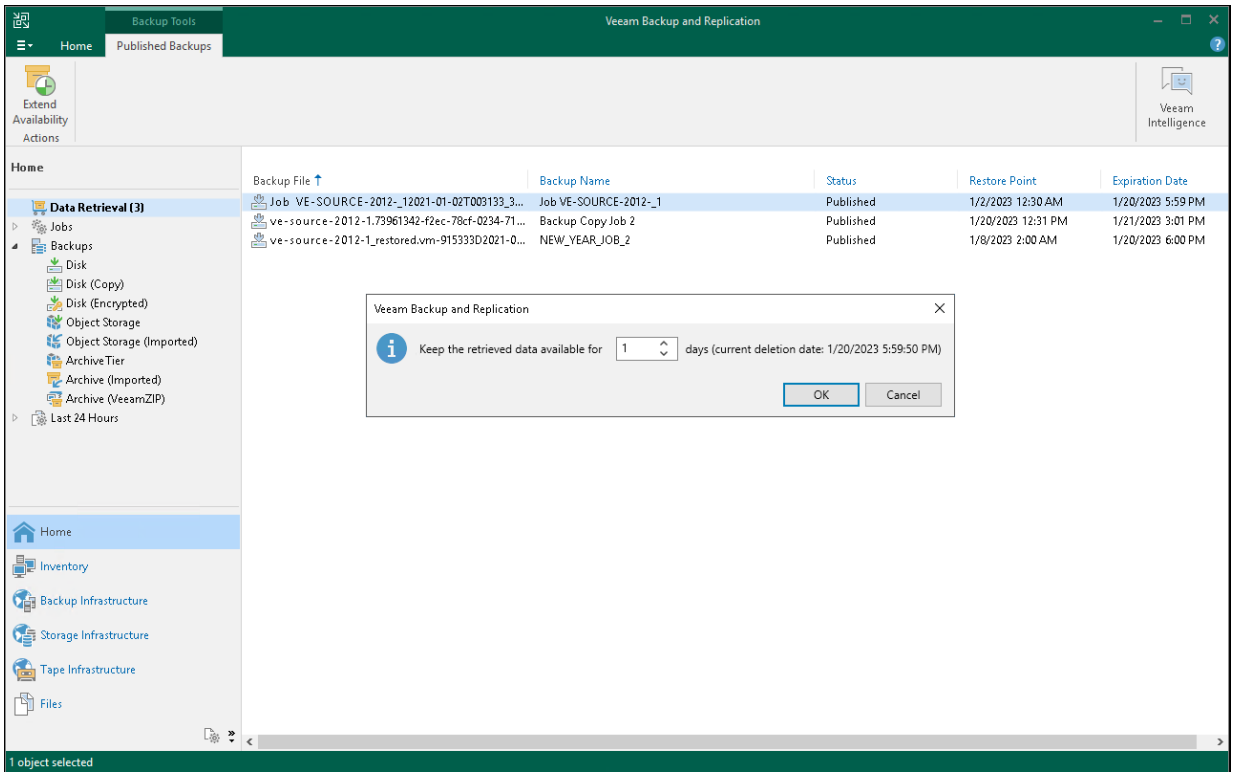


## Restoring Retrieved Data

You can restore your data directly from the archive tier back to production servers or to Microsoft Azure, Amazon EC2 or Google Compute Engine platforms. The retrieved data recovery does not differ from standard backup data recovery and can be performed by using any of the methods documented in the Data Recovery section.

# Adding Scale-Out Backup Repositories

To configure the scale-out backup repository, use the **New Scale-out Backup Repository** wizard.

# Step 1. Launch New Scale-Out Backup Repository Wizard

To launch the **New Scale-out Backup Repository** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane select **Scale-out Repositories** and click **Add Scale-out Repository** on the ribbon.
- Open the **Backup Infrastructure** view. In the inventory pane right-click **Scale-out Repositories** and select **Add Scale-out Backup Repository**.

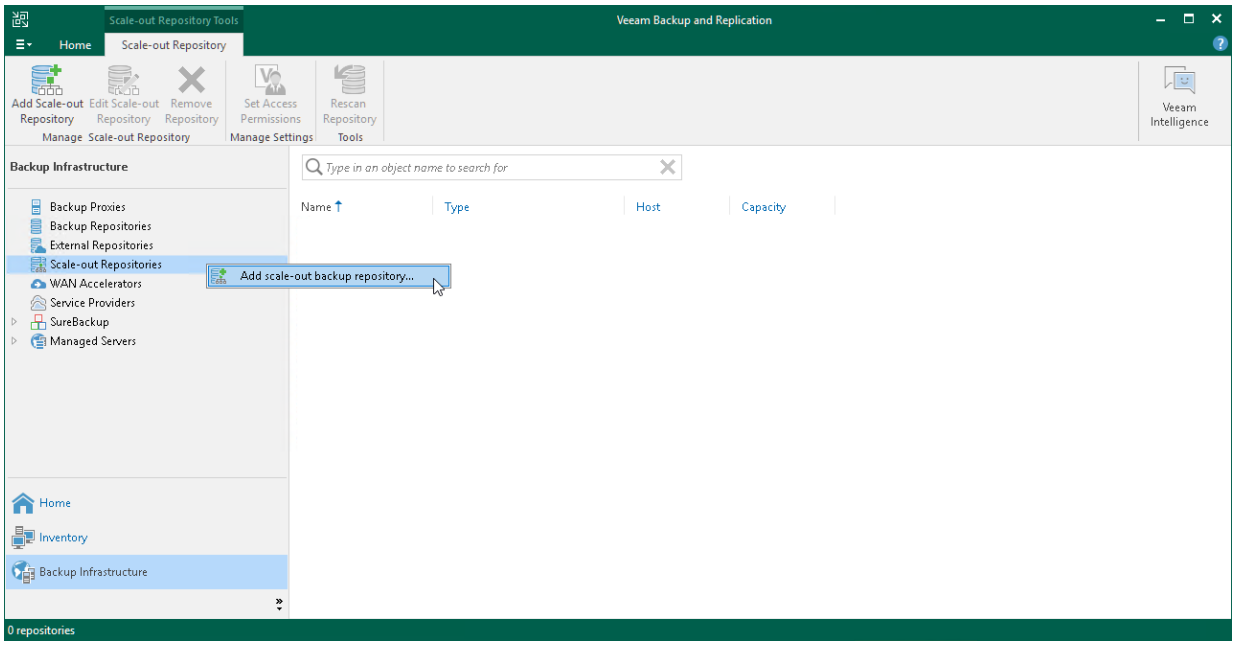

# Step 2. Specify Scale-Out Backup Repository Name

At the **Name** step of the wizard, specify a name and description for the scale-out backup repository.

1. In the **Name** field, specify a name for the scale-out backup repository.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who added the backup repository, date and time when the backup repository was added.

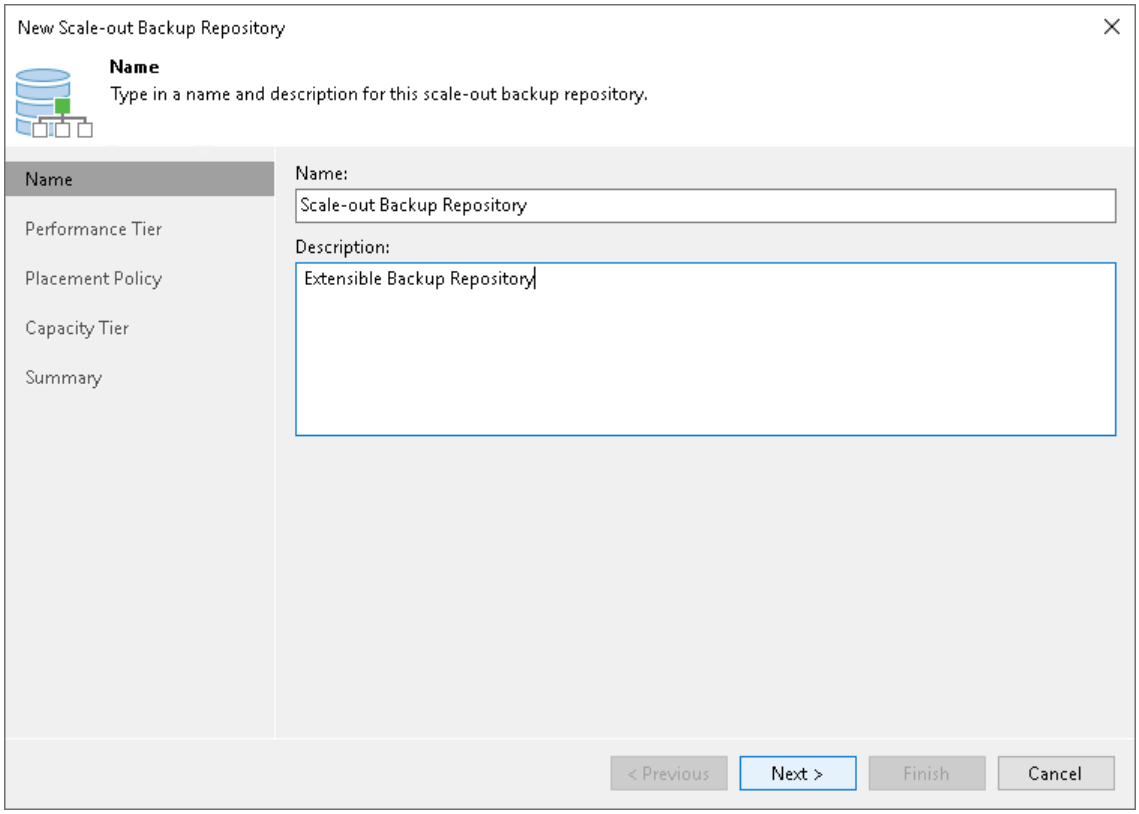

# Step 3. Add Performance Extents

At the **Performance Tier** step of the wizard, specify which backup repositories or object storage repositories you want to add as performance extents, and configure options for the scale-out backup repository.

1. On the right of the **Extents** list, click **Add**.
2. In the **Extents** window, select check boxes next to backup repositories or object storage repositories that you want to add as performance extents.
3. Click **OK**.
4. At the lower right corner of the **Extents** list, click **Advanced**.
5. Specify advanced options for the scale-out backup repository:
   a. If you want to create a separate backup chain for every machine in the job, check that the **Use per-machine backup files** check box is selected . With this option enabled, during one backup job session Veeam Backup & Replication will produce a number of backup files — one per every machine, and will write these files to the backup repository in multiple streams simultaneously. It is recommended that you enable this option to achieve better storage and compute resource utilization, especially if you use as a backup repository a deduplicating storage appliance that supports multiple write streams.

      > **NOTE**
      >
      > For object storage repositories that you add as performance extents, this option is selected by default and you cannot disable it.

b. To preserve the consistency of backup chains in the scale-out backup repository, select the **Perform full backup when required extent is offline** check box. If an extent that contains previous restore points from the current backup chain gets offline, the backup chain will be broken.
Veeam Backup & Replication will not be able to add a new incremental backup file. With this option enabled, Veeam Backup & Replication will create a full backup file instead of an incremental backup file. If you enable this option, you must make sure that you have enough free space in the scale-out backup repository to host a full backup file.

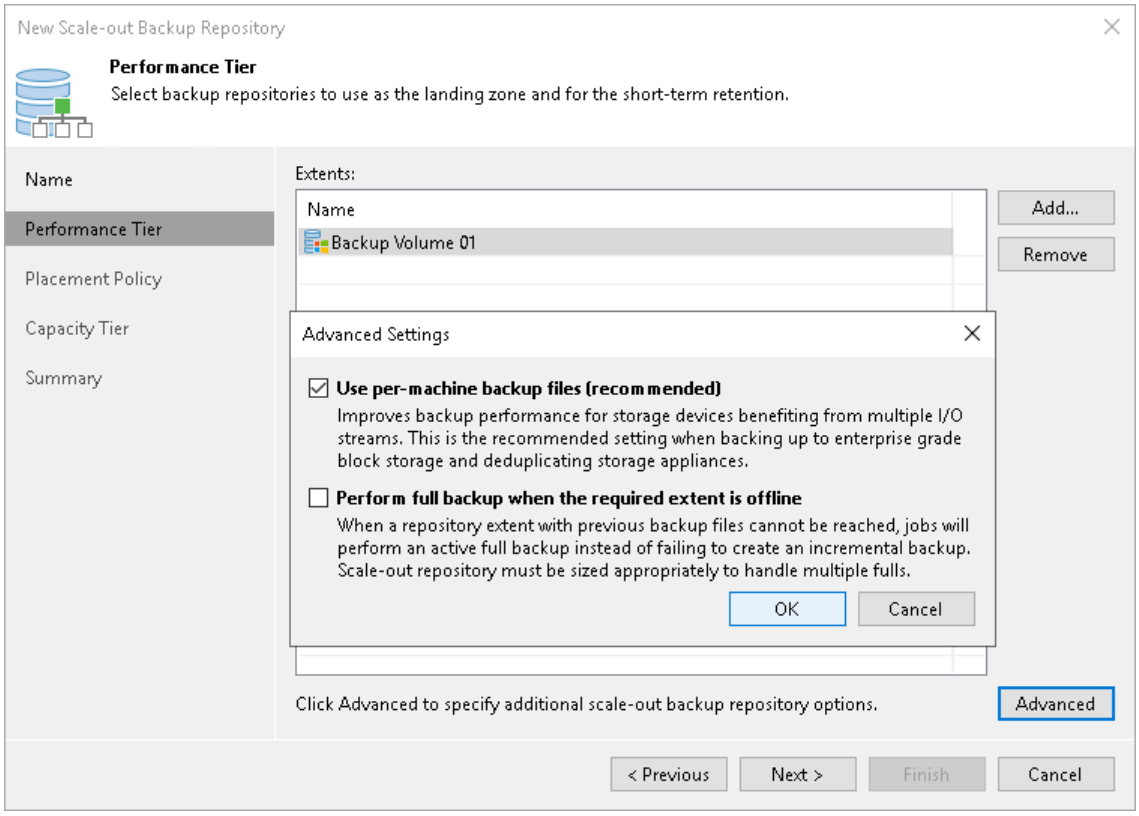

If a backup repository that you add as a performance extent is already used by jobs of supported type or there are backups pointing at the backup repository (for example, independent backups created by VeeamZIP), Veeam Backup & Replication will offer you to update a link to the backup repository in the job properties. Click **Yes** to update the link and target the jobs and backups at the scale-out backup repository. If you click **No**, you will not be able to pass to the next steps of the wizard.

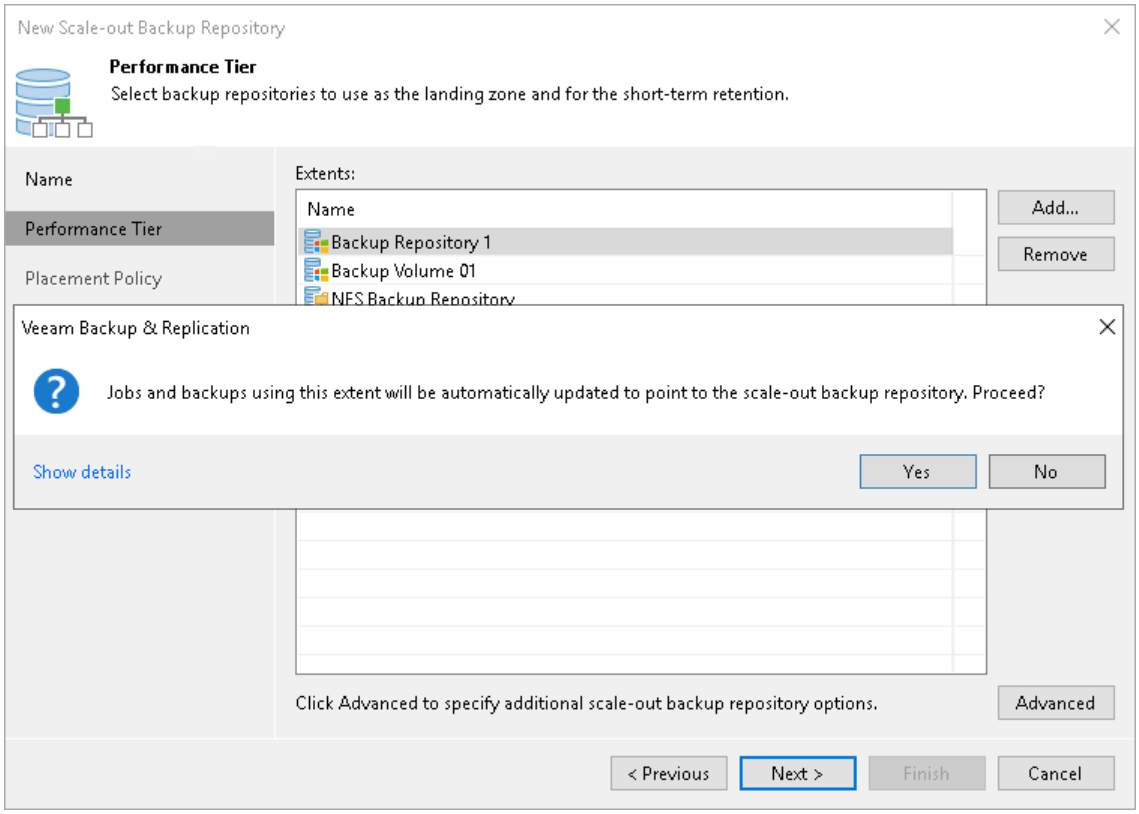

# Step 4. Specify Backup Placement Policy

[This step is not available if you use object storage repositories as performance extents. For more information on how backup file placement policy works for performance extents configured of objects storage repositories, see Backup File Placement]

At the **Policy** step of the wizard, specify how you want to store backup files on performance extents of the scale-out backup repository.

1. Set the backup file placement policy for the scale-out backup repository:
   - Select **Data locality** if you want to store backup files that belong to the same backup chain together. In this case, a full backup file and subsequent incremental backup files will be stored to the same performance extent of the scale-out backup repository. The new backup chain may be stored to the same performance extent or to another performance extent (unless you use a deduplicating storage appliance as a performance extent).
   - Select **Performance** if you want to store full and incremental backup files to different performance extents of the scale-out backup repository. If you set the Performance policy, you must make sure that the network connection is fast and reliable so that Veeam Backup & Replication can access all backup files from the backup chain.

   For more information, see Backup File Placement.
2. If you select the **Performance** policy, you can restrict which types of backup files can be stored on a specific performance extent. For example, if you have added three performance extents to the scale-out backup repository, you may want to store full backup files on one extent and incremental backup files — on the other two extents.
   a. Click **Customize**.
   b. In the **Backup Placement Settings** window, select a performance extent and click **Edit**.
   c. Select a check box next to the type of backup files that you want to store on the extent: **Full backup files** or **Incremental backup files**. By default, Veeam Backup & Replication can store both full and incremental backup files on the same extent.
3. If you select the **Strict placement policy enforcement** check box, Veeam Backup & Replication will not create a backup if it violates the backup placement policy and may result in that a backup job will fail.

> **NOTE**
>
> This option is ignored if you either rebalance extents of your scale-out backup repositories or evacuate data from scale-out backup repositories.

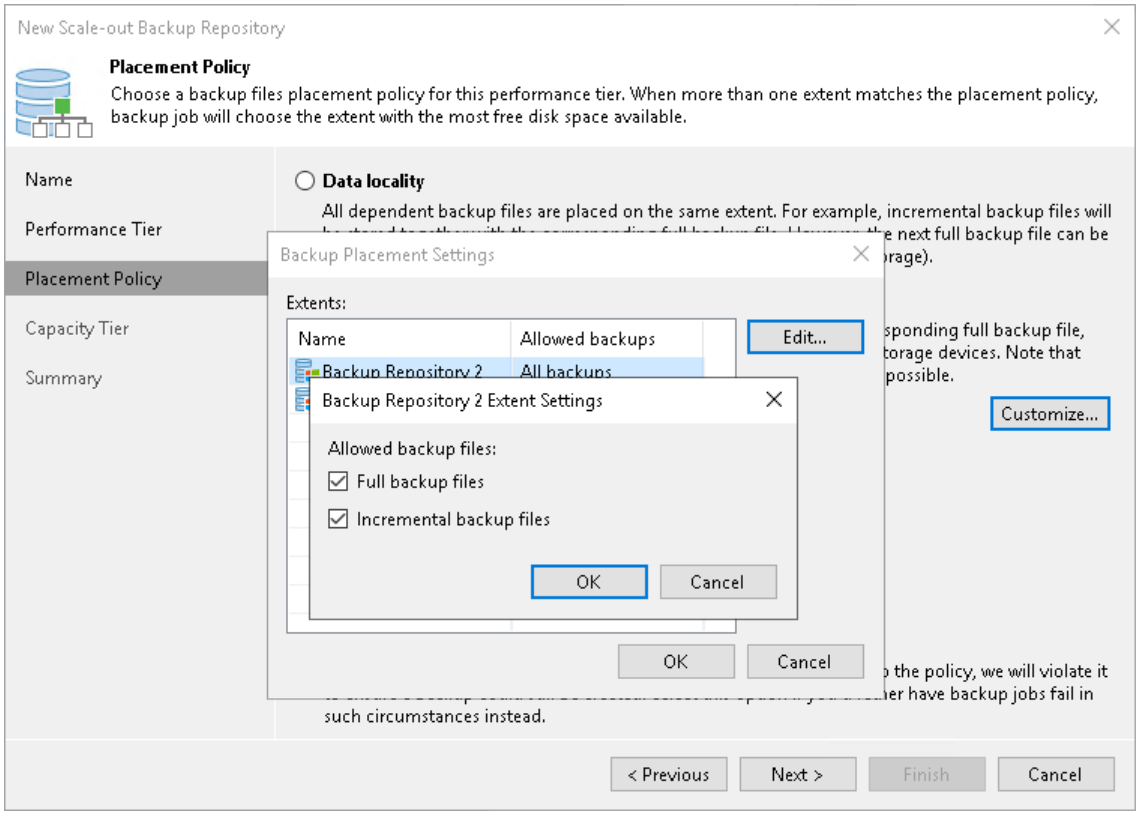
New Scale-out Backup Repository
Placement Policy
Choose a backup files placement policy for this performance tier. When more than one extent matches the placement policy, backup job will choose the extent with the most free disk space available.
Name
Performance Tier
Placement Policy
Capacity Tier
Summary
Data locality
All dependent backup files are placed on the same extent. For example, incremental backup files will
Backup Placement Settings
Extents:
Name
Allowed backups
Edit...
Backup Repository 2
All backups
Backup Repository 2 Extent Settings
Allowed backup files:
Full backup files
Incremental backup files
OK
Cancel
Customize...
possible.
OK
Cancel
such circumstances instead.
< Previous
Next >
Finish
Cancel

# Step 5. Add Capacity Tier

Before you add a capacity tier, check the prerequisites.

At the **Capacity Tier** step of the wizard, select object storage repositories that you want to add as capacity extents. Then specify when to move and copy data.

> **TIP**
>
> If you already have a scale-out backup repository in your backup infrastructure and you want to add a capacity tier, select the scale-out backup repository, click **Edit Scale-out Repository** on the ribbon or right-click the scale-out backup repository and select **Properties**. In the **Edit Scale-out Backup Repository** wizard go to the **Capacity Tier** step and proceed with the following steps.

To configure capacity tier, do the following:

1. Select the **Extend scale-out backup repository capacity with object storage** check box.
2. To select the object storage repositories to which you want to offload your data, click **Choose**.
3. In the **Capacity Tier Extents** window select the check box in front of the necessary object storage repositories. Make sure that these repositories are added to the backup infrastructure in advance. In case object storage repositories are not added, click **Add** and follow the steps of the Adding Object Storage Repository wizard.
4. Click **Offload window** and specify when it is allowed or prohibited to move or copy data to object storage.
5. Select the **Copy backups to object storage as soon as they are created** check box to copy new backups as soon as they are created, as described in section Copying Backups to Capacity Tier.

   When selecting this option, you will be asked whether to copy all backup files that you may already have on any of the performance extents, or only those that have been created recently.

   If you select **Latest**, only backup files that belong to the last active backup chain will be copied from each of the performance extents. If you select **All**, Veeam Backup & Replication will copy all backup files that belong to all backup chains located on any of the specified extents.
6. Select the **Move backups to object storage as they age out of the operational restore window** check box to move inactive backup chains to the capacity extent, as described in section Moving Backups to Capacity Tier.

   In the **Move backup files older than X days** field, specify the operational restore window to define a period after which inactive backup chains on your performance extents will be considered outdated and, therefore, should be moved to the capacity extent. Consider that "0" is an acceptable value, which you can specify to offload inactive backup chains on the same day they are created.

   To override behavior of moving old backups, click **Override**, select the **Move oldest backup files sooner if scale-out backup repository is reaching capacity** check box and define a threshold in percent to force data transfer if a scale-out backup repository has reached the specified threshold.
7. To offload data encrypted, select **Encrypt data uploaded to object storage** and provide a strong password. With this option selected, the entire collection of blocks along with the metadata will be encrypted while being offloaded.If you have not created the password beforehand, click **Add** or use the **Manage passwords** link to specify a new password.

   > **NOTE**
   >
   > You must enable capacity tier encryption if you want to use Veeam Data Cloud Vault as a capacity extent.

8. If you want to specify a schedule for a health check, click the *Monthly* link and define the schedule settings. For more information, see Health Check for Capacity Tier.

> **TIP**
>
> You can combine both the **Copy backups to object storage as soon as they are created** option and the **Move backups to object storage as they age out of the operational restores window** option, as described in section Copying Backups to Capacity Tier.

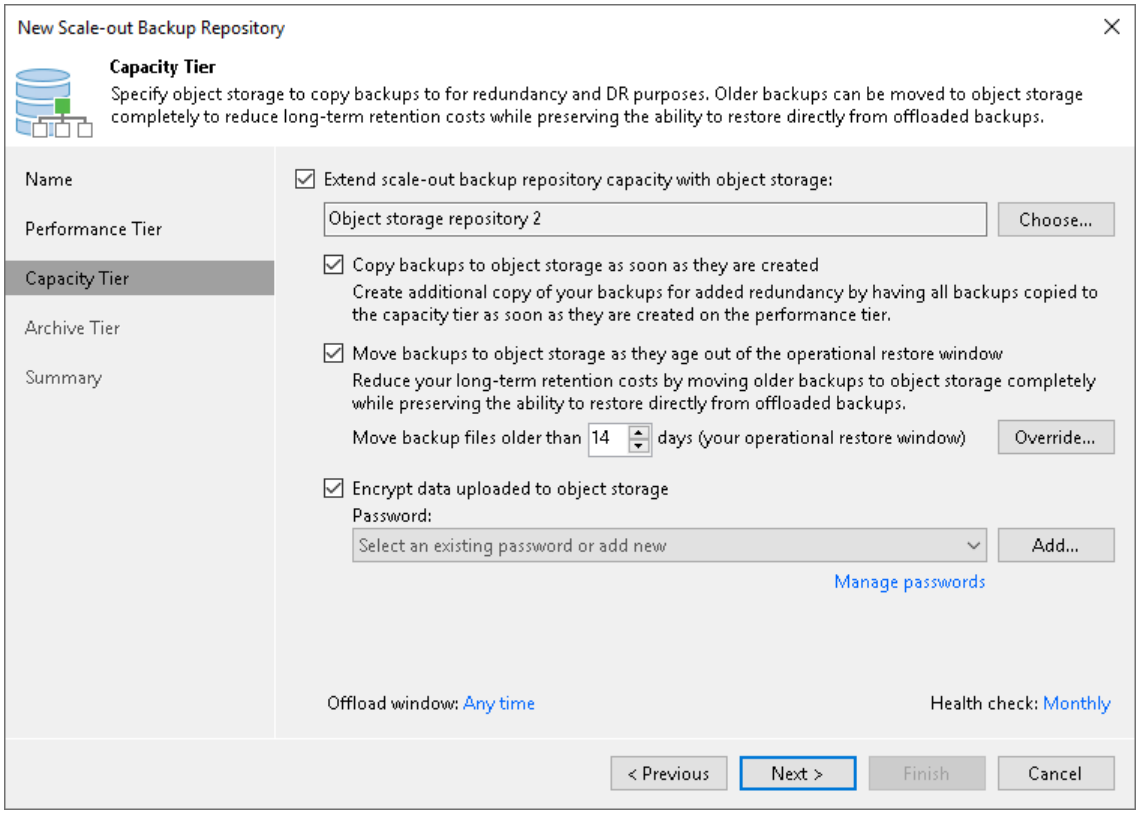


# Synchronizing Capacity Tier Data

When you add as a capacity extent an object storage repository that contains offloaded backup data, you will be prompted to synchronize this data with the performance extents in the scale-out backup repository.

Consider the following:

- An object storage repository can only be added as a capacity extent after existing data (if any) is synchronized.
- During synchronization, Veeam Backup & Replication downloads backup files with metadata located in object storage to the performance extents that are part of the scale-out backup repository that is being added.

  These files are created as described in section Moving Backups to Capacity Tier.
- Extents to which backup data is going to be downloaded (synchronized), will be selected automatically, depending on the available resources.
- The actual data blocks will not be downloaded and will continue to remain in object storage.
- When synchronizing encrypted storage, make sure to provide the same exact password with which the data was encrypted.

After the synchronization is complete, the associated backup files located in object storage will become available as Imported and will be displayed in the **Home** view, under the **Object Storage (Imported)** node in the inventory pane.

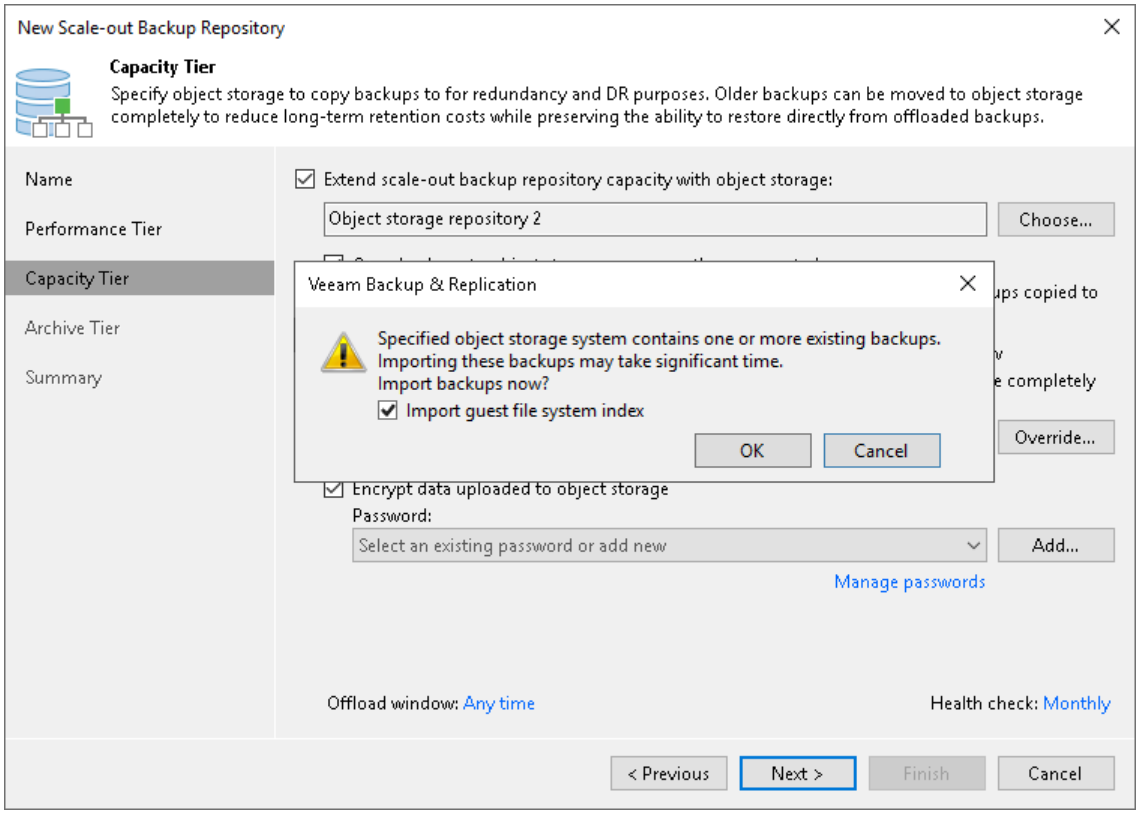

# Step 6. Add Archive Tier

Before you add an archive tier, check the prerequisites.

At the **Archive Tier** step of the wizard, select an object storage repository that you want to add as an archive extent and specify data move settings.

> **TIP**
>
> If you have a compatible type of a repository configured as a capacity extent, you can add an archive extent to the existing scale-out backup repository. To do so, select the scale-out backup repository, click **Edit Scale-out Repository** on the ribbon or right-click the scale-out backup repository and select **Properties**. In the **Edit Scale-out Backup Repository** wizard go to the **Archive Tier** step and proceed with the following steps.

Consider the following:

- The **Archive Tier** step of the wizard will appear only if you have a compatible type of repository configured as a capacity extent or a performance tier. For more information, see Limitations for Archive Tier.
- You can add only one archive extent per scale-out backup repository.

To configure the archive extent, do the following:

1. Select the **Archive GFS full backups to object storage** check box. In the drop-down list, select one of available object storage repositories or click **Add** to add a new one.
2. In the **Archive GFS backups older than N days** field, specify the operational restore window to define a period after which inactive backup chains on your capacity extent will be considered outdated and, therefore, should be moved to the archive extent. Consider that "0" is an acceptable value, which you can specify to archive inactive backup chains on the same day they are created.
3. To encrypt data in the archive extent, select **Encrypt data uploaded to object storage** and provide a strong password. Veeam Backup & Replication will encrypt the entire collection of blocks along with the metadata during the archive job. If you have not created the password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Encryption for Archive Tier.

   > **NOTE**
   >
   > If you have encryption on the capacity tier level, but do not enable encryption on the archive tier level, the backups will not be encrypted in the archive tier.

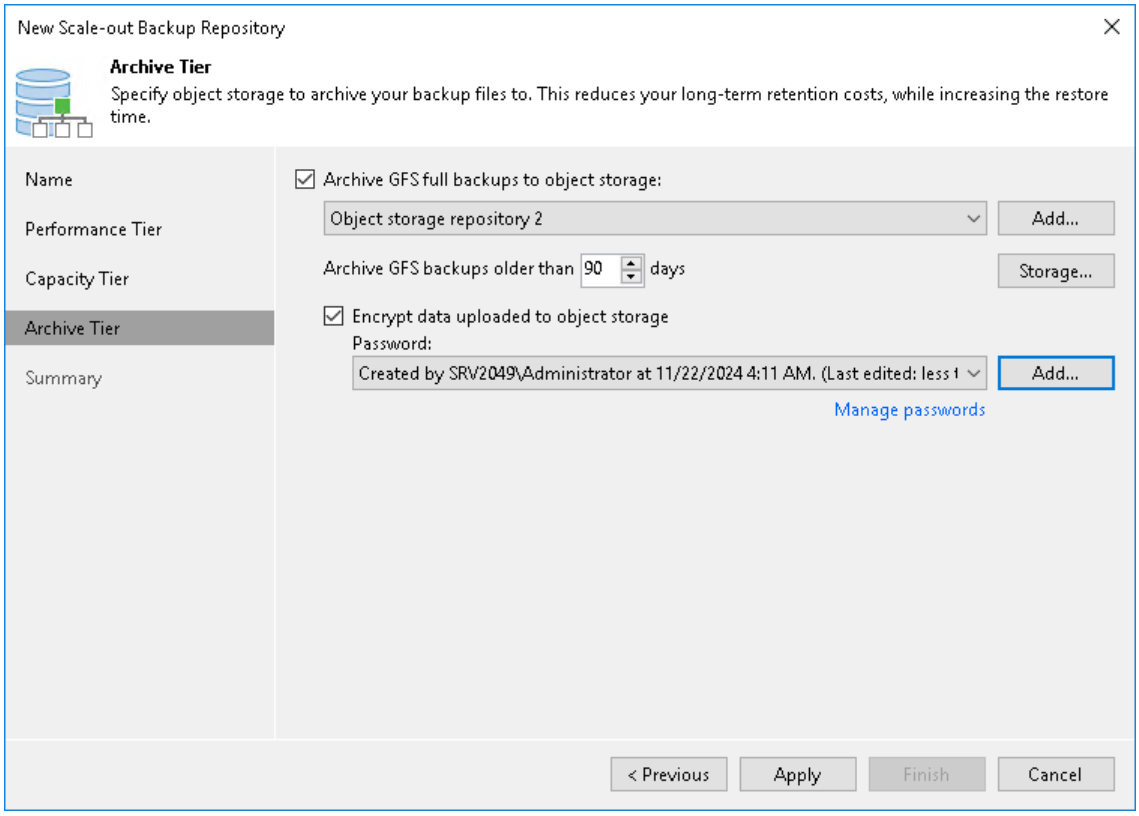


# Specifying Storage Settings

You can use the default storage settings or specify them manually. For that, click **Storage**.

- Select the **Store archived backups as standalone fulls** check box to forbid reuse of the data blocks.
- Select the **Archive backups only if the remaining retention time is above minimal storage period** check box to specify which data blocks can be transported to the archive tier.

When you add as an archive extent an object storage repository that contains archived backup data, you will be prompted to synchronize existing backup chains with data in this scale-out backup repository. After the synchronization is complete, the existing backups will become available as Imported and will be displayed in the **Home** view, under the **Archive Tier (Imported)** node in the inventory pane.

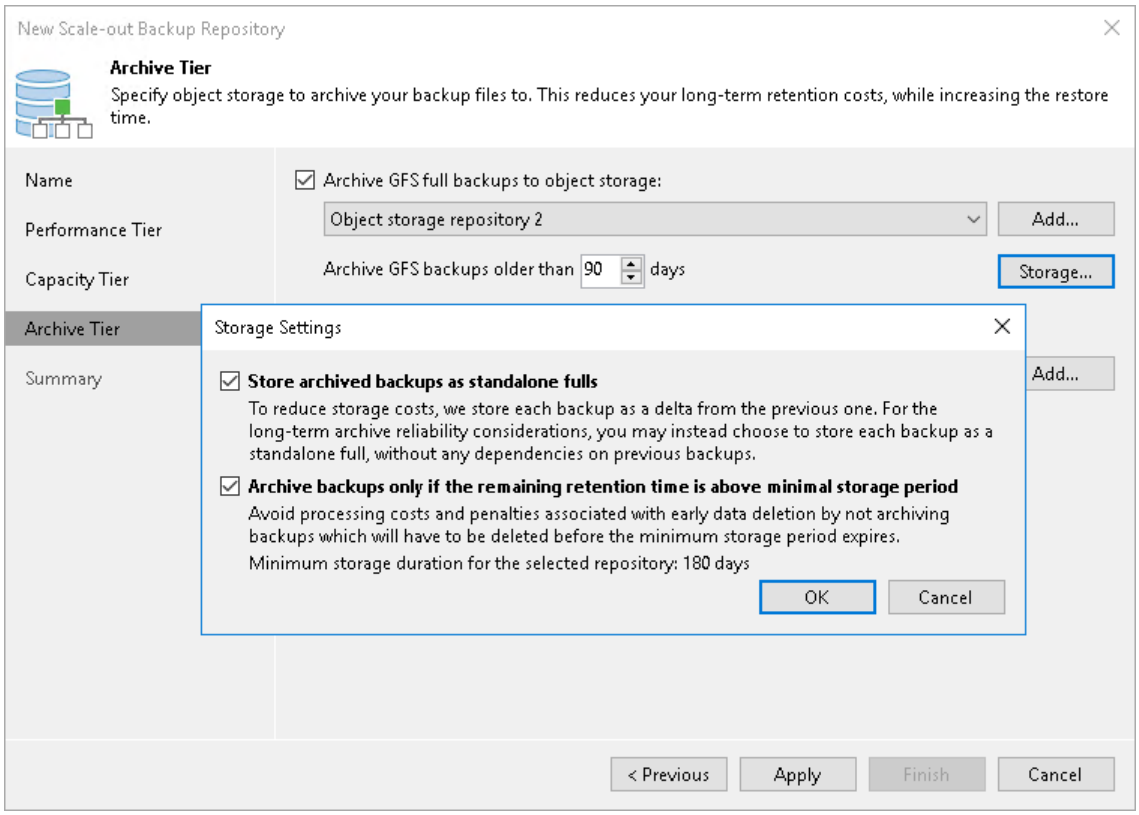

# Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of scale-out backup repository configuration.

Wait for the scale-out backup repository to be added to the backup infrastructure. The process may take some time.

1. Review details of the scale-out backup repository.
2. Click **Finish** to exit the wizard.

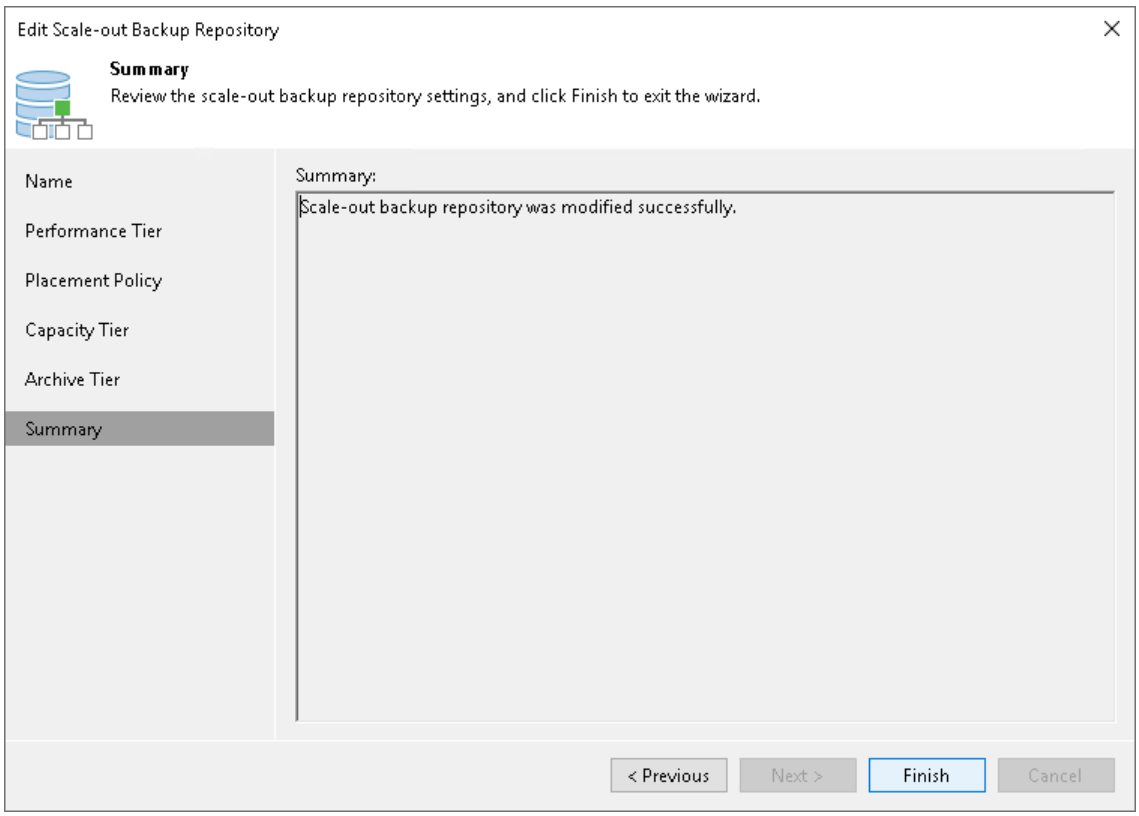

# Managing Scale-Out Backup Repositories

You can manage your scale-out backup repositories and the data there in various ways: edit settings of the scale-out backup repository, rescan scale-out backup repositories automatically or manually, discover on which performance extent of the scale-out backup repository a particular backup file is stored, extend scale-out backup repository or remove certain performance extents from it, perform service actions or remove the scale-out backup repository.

## Editing Settings of Scale-Out Backup Repositories

You can edit settings of the scale-out backup repository, for example, if you want to change the backup file placement policy or specify other advanced settings for the backup repository.

Consider the following:

- If you enable or disable the **Use per-machine backup file** option, Veeam Backup & Replication will apply new settings after a new full backup file is created.
- If you enable or disable the **Perform full backup when required extent is offline** option, Veeam Backup & Replication will apply the new settings starting from the next session of the job targeted at this scale-out backup repository.
- If you change the backup file placement policy settings, Veeam Backup & Replication will apply the new settings starting from the next session of the job targeted at this scale-out backup repository.

To change the scale-out backup repository settings:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Scale-out Repositories**.
3. In the working area, select the scale-out repository and click **Edit Scale-out Repository** on the ribbon or right-click the scale-out backup repository and select **Properties**.
4. Follow the steps of the **Edit Scale-out Backup Repository** wizard and edit settings as required.

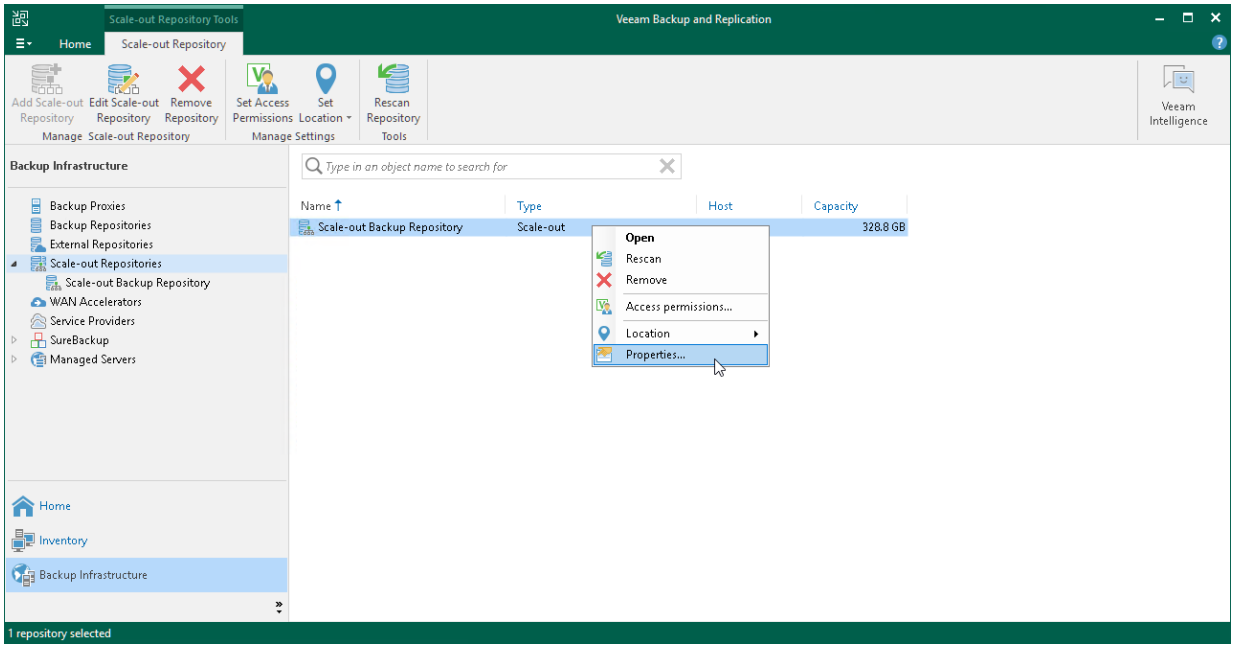

# Rescanning Scale-Out Repositories

Veeam Backup & Replication periodically rescans scale-out backup repositories. During the rescan process, it gets the following information:

- State of every performance extent added to the scale-out backup repository: online or offline.
- Status of Veeam Data Movers on extents: up-to-date or outdated.
- Space available in the scale-out backup repository.

The rescan operation is performed automatically by a rescan process that works permanently in the background. The process is started every 24 hours. It can be also started when a new task session starts, and the Veeam Backup Service requires information about the infrastructure to be refreshed.

In addition to the automated rescan process, you can manually start rescan of the scale-out backup repository. Backup repository rescan may be helpful, for example, if you want to discover backup files that were manually relocated from one performance extent to another one.

Consider the following:

- Information about backup files location is updated only if you perform manual rescan of scale-out backup repositories.
- Veeam Backup & Replication rescans scale-out backup repositories when you perform backup files import.
- Veeam Backup & Replication does not rescan extents that are set into the Maintenance mode.
- To successfully rediscover relocated backups files created by backup copy jobs, make sure to disable these jobs manually prior to rescanning.

  For more information, see Disabling and Deleting Jobs.

To start the rescan process:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane select **Scale-out Repositories**.

3. In the working area, select the scale-out repository and click **Rescan Repository** on the ribbon or right-click the scale-out backup repository and select **Rescan**.

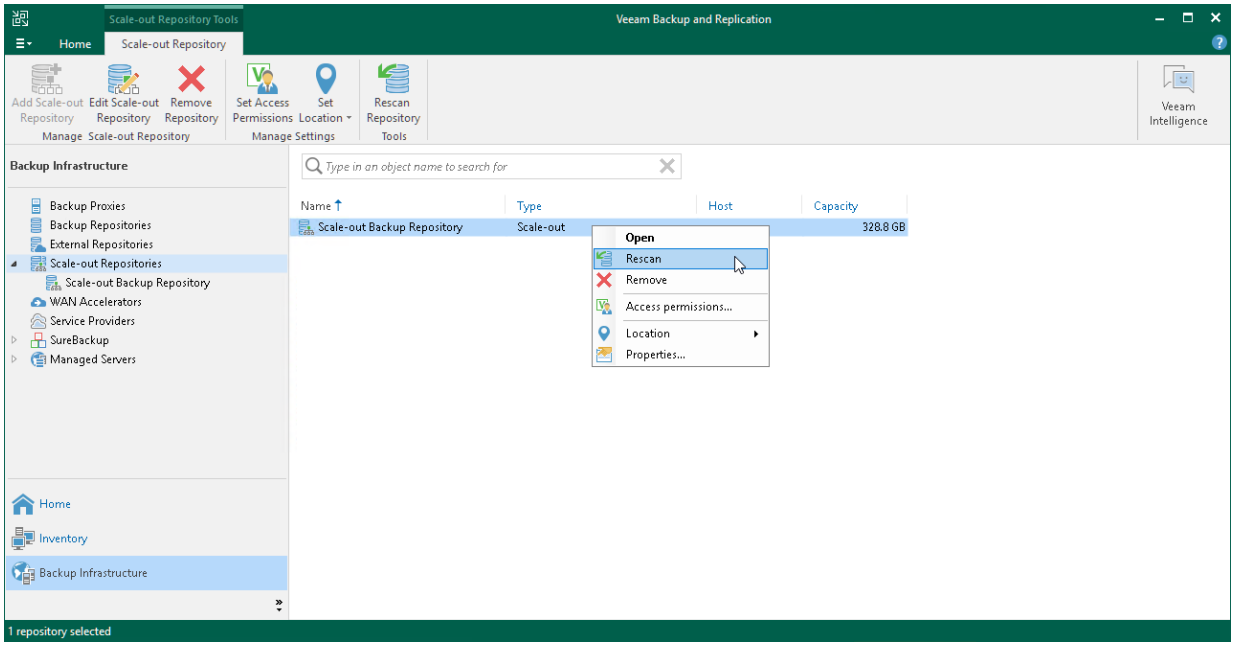


# Checking Rescan Sessions Reports

If necessary, you can view a history of all rescan reports that Veeam Backup & Replication ran for scale-out backup repositories (both automatic and manual).

To do it, perform the following steps:

1. Open the **History** view.
2. In the inventory pane, select **Statistics**.

3. In the working area, select the necessary task session and click **Statistics** on the ribbon. You can also right-click the necessary session and select **Statistics**.

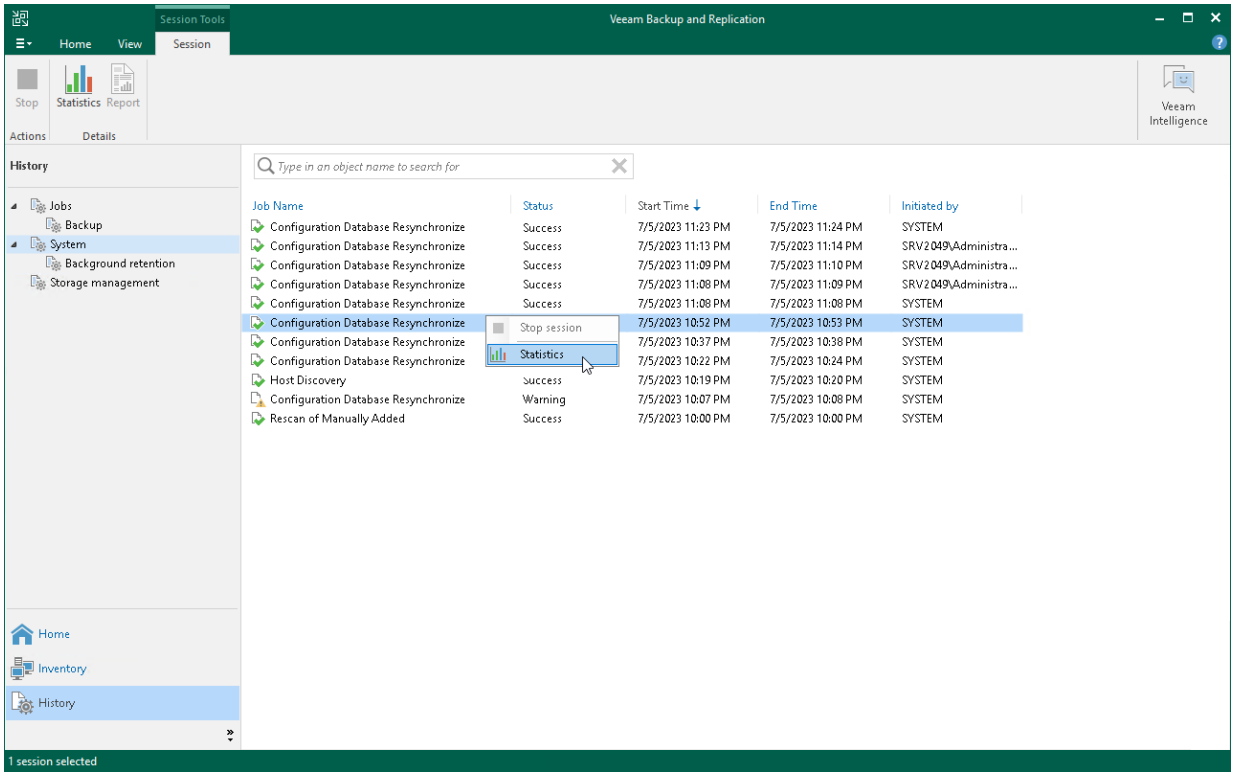


# Discovering Backups in Scale-Out Backup Repositories

To discover on which performance extent of the scale-out backup repository a particular backup file is stored, you can examine the job session statistics or check the backup properties.

To view the job session statistics:

1. Open the **Home** view.
2. In the inventory pane, click **Backup** under **Jobs**.
3. In the working area, right-click the job and select **Statistics**.

4. In the bottom left pane of the window, click the VM name. In the **Action** pane, locate the message: *Using N scale-out repository extent*.

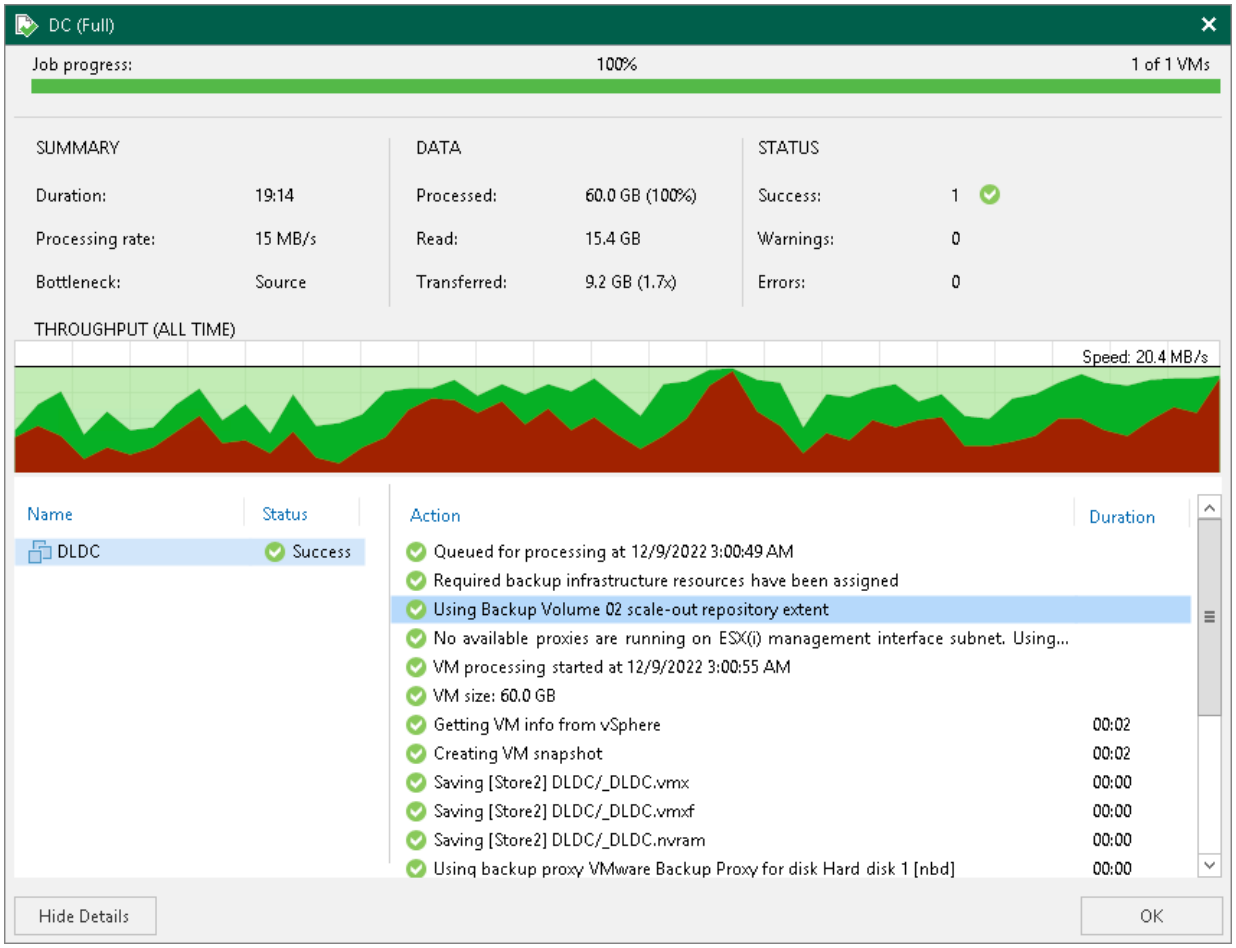


To view the backup properties:

1. Open the **Home** view.
2. In the inventory pane, select **Disk** under **Backups**.
3. In the working area, right-click the backup and select **Properties**.

4. Veeam Backup & Replication will display the extent where the backup file resides in the headline of the **Backup Properties** window. To see the path to the backup file, right-click the job and select **Copy path**.

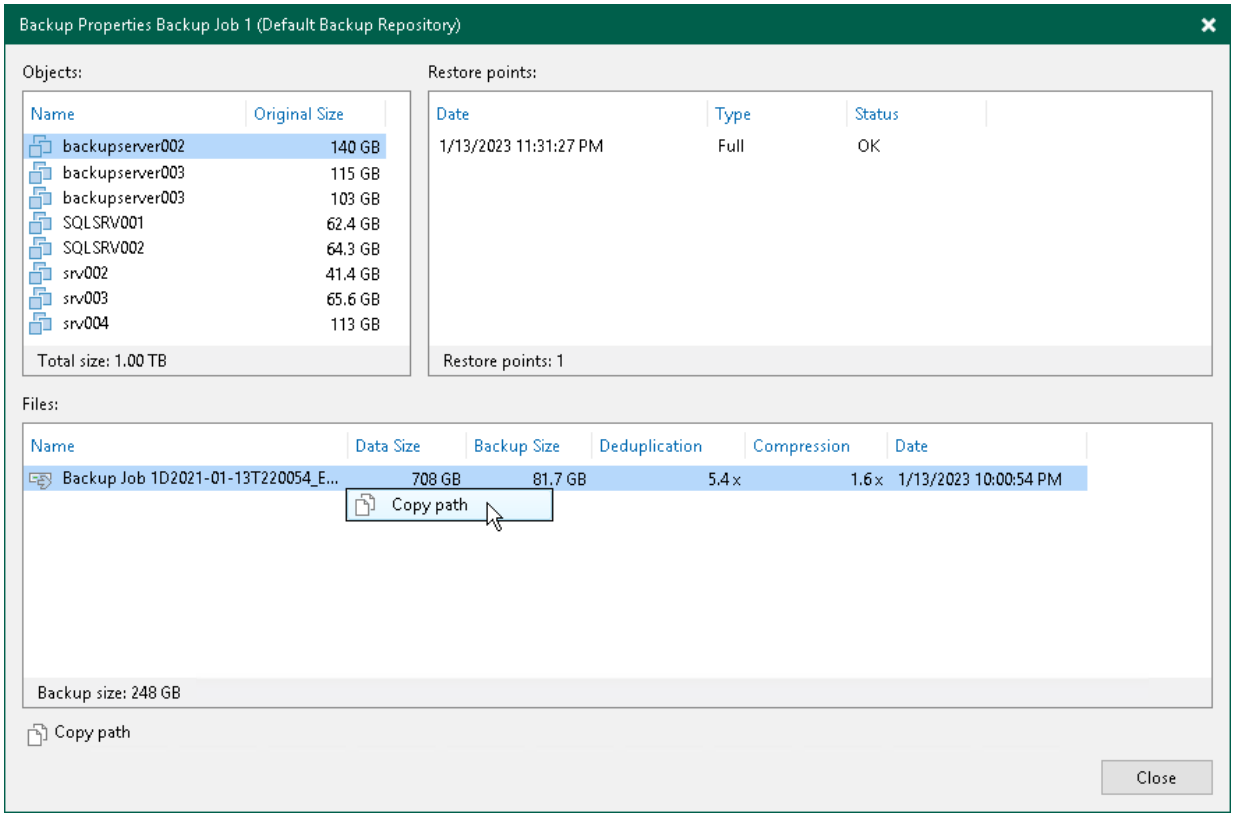


# Backup State Indicators

Restore points icons help you understand current state of a restore point in a scale-out backup repository.

For more information on icons and what they indicate, see Infrastructure Icons.

# Extending Scale-Out Repositories

You can add a backup repository or an object storage repository as a performance extent or a capacity extent to the scale-out backup repository at any time. For example, the scale-out backup repository may run low on space, and you will need to add storage capacity to it.

# Extending Performance Tier

To extend the performance tier, perform the following steps:

1. Check the limitations for performance tier.
2. Open the **Backup Infrastructure** view.
3. In the inventory pane, click **Scale-out Repositories**.
4. In the working area, select the scale-out repository and click **Edit Scale-out Repository** on the ribbon or right-click the backup repository and select **Properties**.
5. Move to the **Performance tier** step of the wizard.

6. Click **Add**.
7. In the **Extents** window, select a check box next to the backup repository that you want to add as a performance extent to the scale-out backup repository.

   If a backup repository that you add as a performance extent is already used by jobs of supported type or there are backups pointing at the backup repository (for example, independent backups created by VeeamZIP), Veeam Backup & Replication will offer you to update a link to the backup repository in the job properties. Click **Yes** to update the link and target the jobs and backups at the scale-out backup repository. If you click **No**, you will not be able to pass to the next steps of the wizard.
8. Pass through the next wizard steps and finish working with the wizard. The new performance extent will be added to the scale-out backup repository.

**NOTE**

After you add a backup repository to the scale-out backup repository as a performance extent, you will not be able to use it as an individual backup repository.

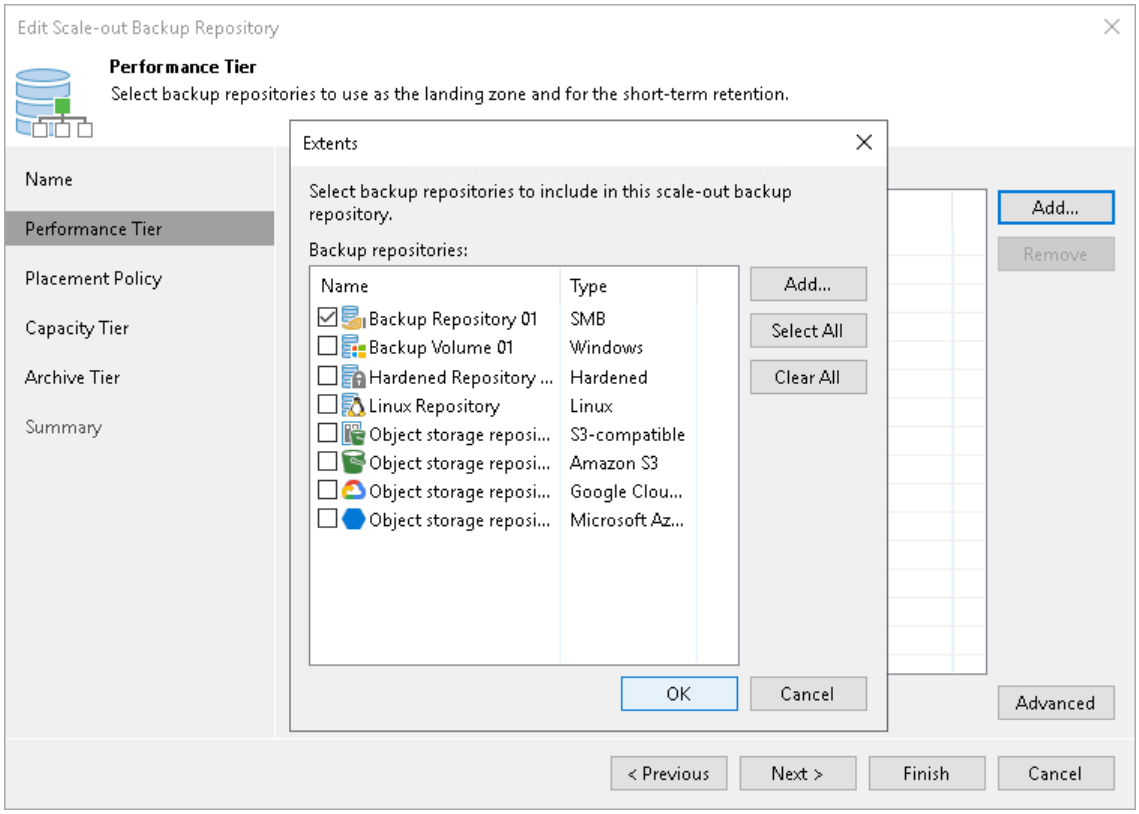


# Extending Capacity Tier

To extend the capacity tier, perform the following steps:

1. Check the limitations for capacity tier.
2. Open the **Backup Infrastructure** view.
3. In the inventory pane, click **Scale-out Repositories**.
4. In the working area, select the scale-out repository and click **Edit Scale-out Repository** on the ribbon or right-click the backup repository and select **Properties**.
5. Move to the **Capacity tier** step of the wizard.

6. Click **Choose**.
7. In the **Capacity Tier Extents** window, select a check box next to the backup repository that you want to add as a capacity extent to the scale-out backup repository.
8. If a backup repository that you add as a capacity extent is already used by jobs of supported type or there are backups pointing at the backup repository (for example, independent backups created by VeeamZIP), Veeam Backup & Replication will offer you to update a link to the backup repository in the job properties. Click **Yes** to update the link and target the jobs and backups at the scale-out backup repository. If you click **No**, you will not be able to pass to the next steps of the wizard.
9. Pass through the next wizard steps and finish working with the wizard. The new capacity extent will be added to the scale-out backup repository.

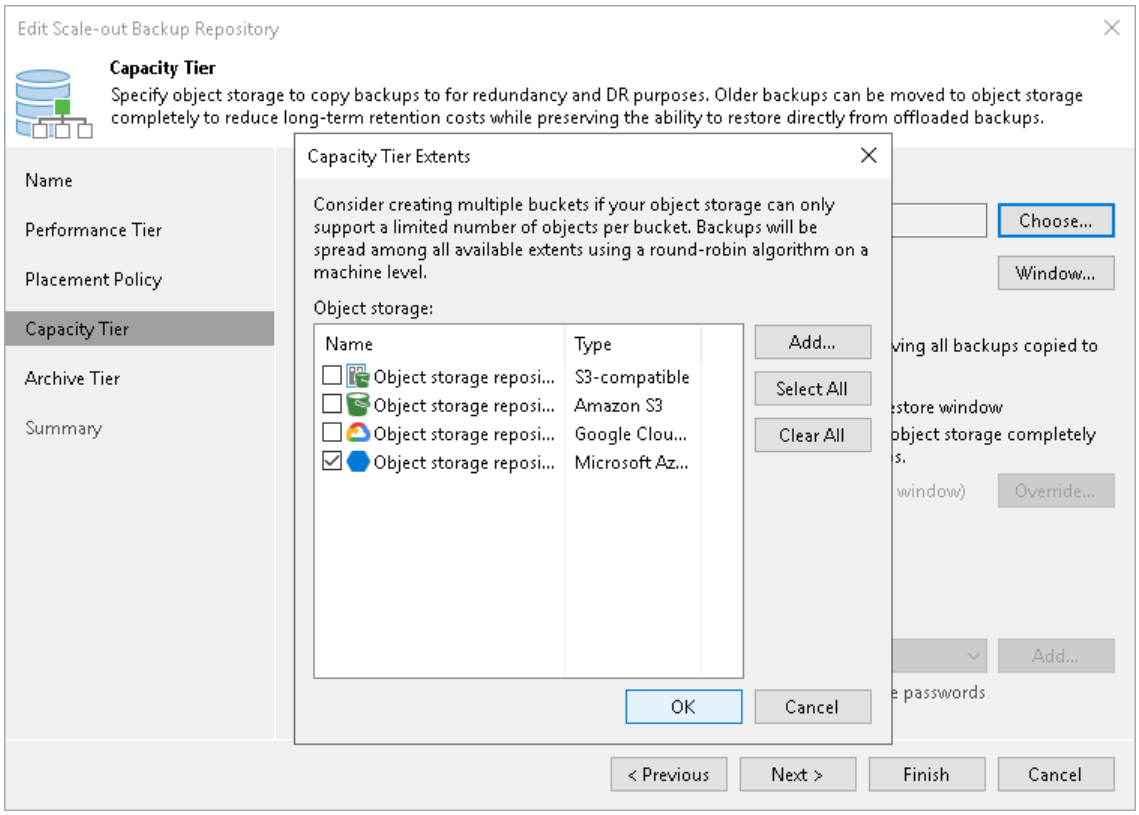


# Service Actions with Scale-Out Backup Repositories

In some cases, you may want to perform service actions with scale-out backup repository extents. For example, you need to upgrade the repository server and add more memory to it. Or you want to replace a storage device backing the extent and need to relocate backup files.

You can perform the following service actions with extents of a scale-out backup repository:

- Switch to Sealed Mode.
- Switch to Maintenance mode.
- Evacuate backups from extents.
- Migrate Backups within Scale-Out Backup Repository.

## Switching to Sealed Mode

Veeam Backup & Replication allows you to put any of the scale-out backup repository extents into the Sealed mode.

Sealing up scale-out backup repository extents allows you to gradually remove data located on these extents by applying a retention policy. You can use this feature to gracefully stop using some of your extents and exclude them from the scale-out backup repository configuration.

After the extent is sealed, no further data is saved to the extent, and only read operations such as restore, merge and remove are allowed.

Backup jobs that are targeted to a scale-out backup repository with the sealed extents that store active backup chains are forced to create a new active full backup on the next run. Veeam Backup & Replication saves the new active full to another available extent in the scale-out backup repository scope and forms a new active backup chain. To select an extent where to keep the new active full, Veeam Backup & Replication compares available resources of all extents and selects the best suitable extent.

> **NOTE**
>
> Consider the following:
>
> - If you create backup jobs with the help of another Veeam backup solution, you will have to trigger active full backup job manually.
> - If you use the per-machine backup chain format for backed-up data, the active full is forced only for machines which backups have active chains on the sealed extents.

# Considerations and Limitations

Consider the following limitations:

- All restore points that exceed the specified retention period will be continuously removed from the sealed extents on each subsequent backup session.
- When you put an extent into the Sealed mode, Veeam Backup & Replication restricts any further data transfer to the extent. For more information on operation restrictions, see Sealed Mode Restrictions.
- All restore points that exceed the specified retention period will be continuously removed from the sealed extent, as described in section Retention Policy.
- An object storage repository can be put into Sealed mode only if it is a member of the scale-out backup repository.

  If an object storage repository was not added as part of any of your scale-out backup repositories, the **Seal Extent** option will not be available.
- An extent can be put into both the Maintenance and the Sealed modes at the same time.

  When both modes are applied, the Maintenance mode overrides Sealed mode.

To put an extent into the Sealed mode:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the scale-out backup repository under the **Scale-out Repositories node**.
3. In the working area, select the extent and click **Seal extent** on the ribbon or right-click the extent and select **Seal extent**.

To remove the extent from the Sealed mode, select the extent and click **Sealed** on the ribbon or right-click the extent and select **Sealed**.

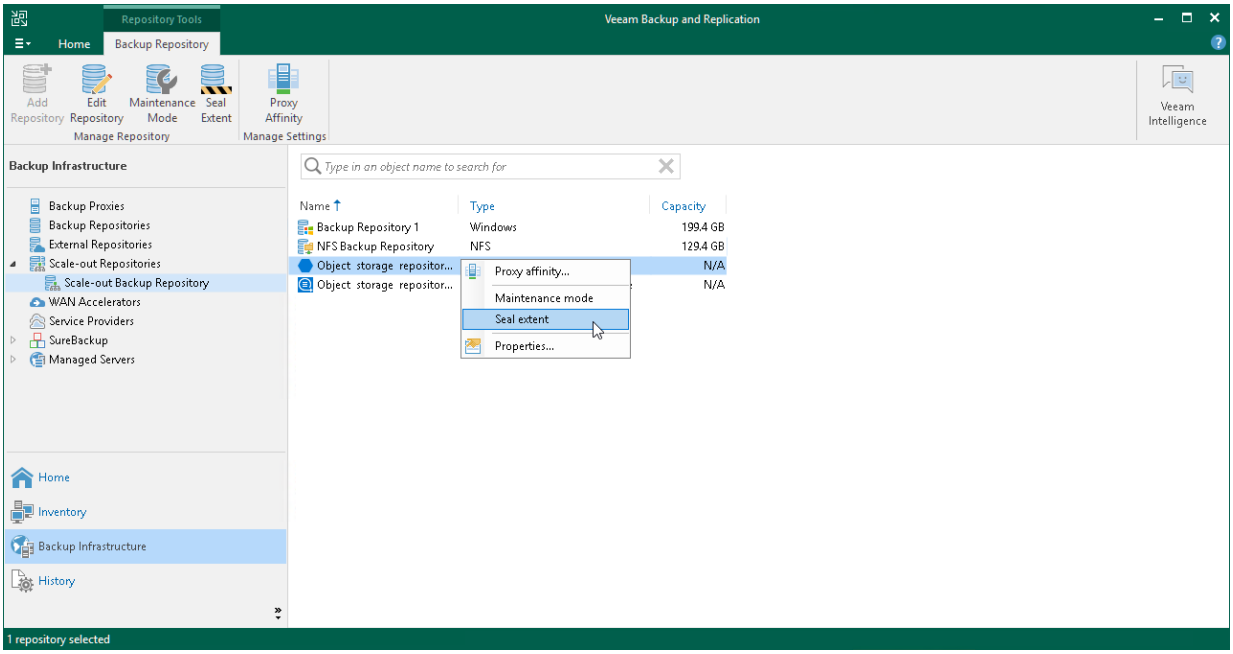


# Sealed Mode Restrictions

The following table lists restrictions that are imposed right after the extent is put into Sealed mode.

| Activity | Restriction Level |
|---|---|
| Moving backups to object storage | Restricted |
| Copying backups to object storage | Restricted |
| Moving backups to capacity tier (manual operation) | Restricted |
| Downloading data from object storage | Allowed |
| Copying backups to performance tier (manual operation) | Allowed |
| Moving backups to performance tier (manual operation) | Allowed |
| Restoring data from backup files residing on object storage | Allowed |
| Exporting as .VBK file from backup files residing on object storage | Allowed |
| Removing backups from configuration | Allowed |

| Activity | Restriction Level |
|---|---|
| Retention policies | Allowed |
| Removing backups or VMs created with the per-machine method | Allowed |
| Removing a VM from a single storage | Allowed |
| Scale-out backup repository rescan | Allowed |
| Evacuating backups | Allowed |

## Switching to Maintenance Mode

# Maintenance Mode

Veeam Backup & Replication allows you to put any of the scale-out backup repository extents into the Maintenance mode. You can use this mode if you need to perform service actions, such as upgrading an extent or installing a patch on it. Putting an extent into the Maintenance mode is mandatory to evacuate backups, as described in section Evacuating Backups from Extents.

When you switch to the Maintenance mode, Veeam Backup & Replication launches the *Repository Maintenance* job. The *Repository Maintenance* job checks the status of jobs and tasks targeted at the extent and puts the extent to one of the following modes:

- If no tasks using the extent are currently running, the job puts the extent to the Maintenance mode immediately.
- If the extent is busy with any task, for example, a backup job, the job puts the extent to the Maintenance pending state and waits for the task to complete. When the task is complete, the extent is put to the Maintenance mode.

# Considerations and Limitations

Consider the following limitations:

- When you put an extent into the Maintenance mode, Veeam Backup & Replication restricts any further data transfer to the extent and allows only operations listed in section Maintenance mode restrictions.
- An extent can be put into both the Maintenance and the Sealed modes at the same time. When both modes are applied, Maintenance mode restrictions override restrictions of Sealed mode.
- You cannot restore VM data from backup files residing on the extent. You also cannot restore VM data from backup files residing on other extents if a part of the backup chain resides on the extent in the Maintenance mode.
- If an extent contains a part of an active backup chain on it, do NOT put this extent into the Maintenance mode. Otherwise, the backup job targeted to the scale-out backup repository will fail. To avoid this, create a full backup manually.

# Putting Extent to Maintenance Mode

To put an extent into the Maintenance mode:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the scale-out backup repository under **Scale-out Repositories**.
3. In the working area, select the extent and click **Maintenance Mode** on the ribbon or right-click the extent and select **Maintenance mode**.

To remove the extent from the Maintenance mode, select the extent and click **Maintenance Mode** on the ribbon or right-click the extent and select **Maintenance mode** once again.

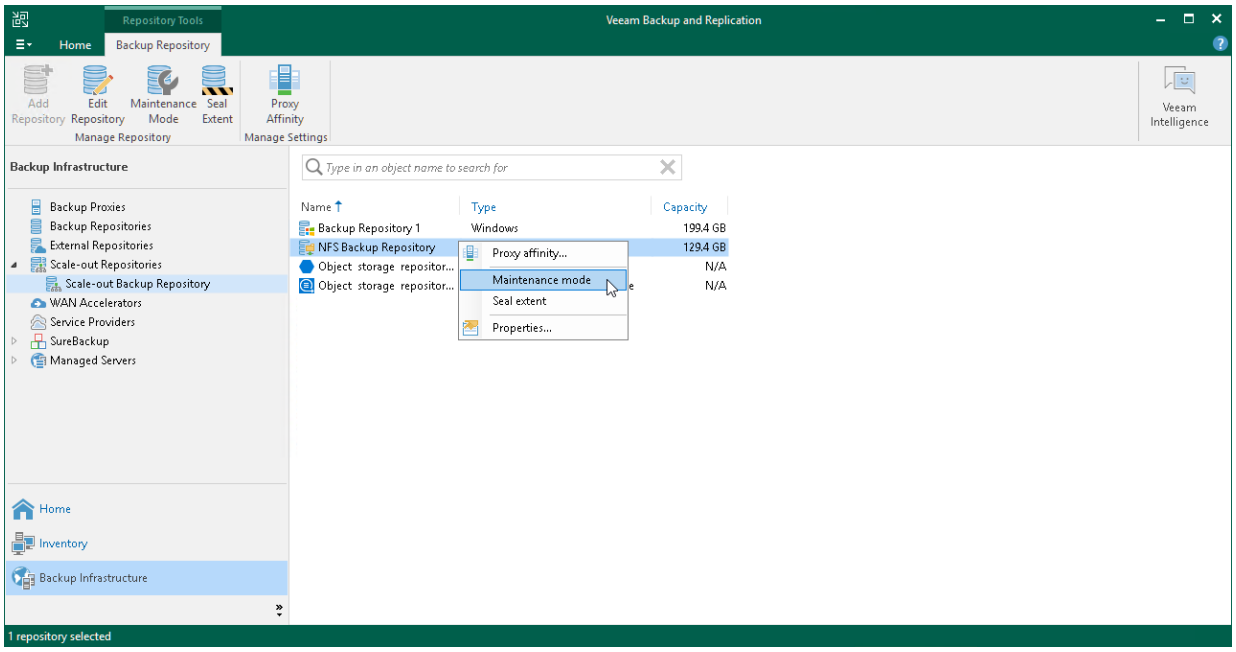


# Maintenance Mode Restrictions

The following table lists restrictions that are imposed right after the necessary extent of a scale-out backup repository is put into Maintenance mode.

| Activity | Restriction Level |
|---|---|
| Moving to object storage. | Restricted |
| Copying to object storage. | Restricted |
| Downloading from object storage. | Restricted |
| Moving to capacity tier (manual operation). | Restricted |
| Copy to performance tier (manual operation). | Restricted |

| Activity | Restriction Level |
|---|---|
| Restore from offloaded backups. | Restricted |
| Export as `.vbk` from offloaded backups. | Restricted |
| Removal of backups or VMs created with the per-machine method. | Restricted |
| Scale-out backup repository rescan. | Allowed |
| Removal of backups from configuration. | Allowed |
| Evacuation of storage along with indexes from on-premises extents. | Allowed |

## Evacuating Backups from Extents

If you want to remove an extent from a scale-out backup repository, you first need to evacuate backups from this extent. When you evacuate backups, Veeam Backup & Replication moves backup files from the necessary extent to other extents within the same scale-out backup repository.

You can evacuate data from the following types of extents:

- Performance extents consisting of backup repositories
- Performance extents consisting of direct object storage repositories
- Capacity extents

# How Evacuating Backups from Extents Works

Depending on the type of the extent, Veeam Backup & Replication applies different scenario to distribute data between extents during evacuation.

### Evacuation from Performance Extents Consisting of Backup Repositories

This scenario applies to data evacuation from performance extents consisting of backup repositories.

When Veeam Backup & Replication selects the target extent for evacuated files, it attempts to follow the backup placement policy specified for remaining extents. For example, you have 3 extents in the scale-out backup repository with the following backup file placement settings:

- On *Extent 1*, full backup files are stored.
- On *Performance Extents 2* and *3*, incremental backup files are stored.

If you evacuate backup files from *Performance Extent 2*, Veeam Backup & Replication will relocate them to *Performance Extent 3*.

### Evacuation from Performance Extents or Capacity Extents

This scenario applies to data evacuation from performance extents or capacity extents.

When Veeam Backup & Replication selects the target extent for evacuated files, it uses multiple conditions to distribute data between extents. Data distribution depends on whether it is the first move or a subsequent move of a backup chain.

To select an extent for backup file placement, Veeam Backup & Replication checks the following conditions:

1. During the first job session, Veeam Backup & Replication checks availability of extents. In case, an extent is set to Sealed or Maintenance mode, these extents will be skipped from backup file placement.
2. After that, Veeam Backup & Replication checks the amount of backup chains in available extents and selects the extent with a minimal amount of backup chains.
3. If some extents have storage space limitations, Veeam Backup & Replication calculates the average storage of the backup chains located in all extents and select the extent that contains a minimum amount of backup chains. If all extents have the same number of backup chains, Veeam Backup & Replication selects the extent that has more free space.
4. During subsequent job sessions, Veeam Backup & Replication moves backup chains to the extent that was specified before.

> **NOTE**
>
> If an extent is set to Sealed mode or Maintenance mode, Veeam Backup & Replication will not move a backup chain to this extent. In this case, a new extent is selected and Veeam Backup & Replication follows the same algorithm as during the first job session. The backup chain located in an unavailable extent will be moved again to another extent. Veeam Backup & Replication will remove backup chains from the unavailable extents according to retention policy.

# Considerations and Limitations

Before you start the backups evacuation, consider the following:

- If performance extents consist of repositories with Fast Clone enabled, Veeam Backup & Replication selects the target extent depending on the whether it contains data blocks that can be reused. In the process of evacuation, Veeam Backup & Replication creates a new backup file in the target extent using these data blocks. Therefore, evacuated backups do not occupy all space in the target performance extent.
- For performance extents that consist of object storage repositories, Veeam Backup & Replication selects the target performance extent as described in section Extent Selection for Object Storage Repositories Added as Performance Extents.
- If immutability is enabled on the source extent, during evacuation the data blocks are copied to the target extent. The data blocks are removed only when immutability expires.

  > **NOTE**
  >
  > The data blocks are not removed after immutability expires, if you evacuated these data blocks from object storage repositories added as performance extents of a scale-out backup repository. When the immutability time period is over, you will need to delete these files manually.

- If you evacuate your backups from a hardened repository that is added as a performance extent, Veeam Backup & Replication will copy these backups instead of moving. For more information, see Hardened Repository as Performance Extent.

- You cannot evacuate data from performance extents that consist of non-object storage backup repositories (for example, direct attached storage, network attached storage or deduplicating storage appliances) to performance extents that consist of direct backup object storage repositories. For this scenario, use the backup move option.

# Evacuating Backups from Extents

Before you evacuate backups from the extent, you must put it into the Maintenance mode. For more information, see Switching to Maintenance Mode.

> **TIP**
>
> If you want to gracefully stop using some of your extents and exclude them from a scale-out backup repository configuration, consider using the Sealed mode instead of evacuating backups.

To evacuate backup files from an extent:

1. [Recommended] Stop and disable jobs targeted at the extent from which you plan to evacuate backups.
2. Open the **Backup Infrastructure** view.
3. In the inventory pane, select the scale-out backup repository under **Scale-out Repositories**.
4. In the working area, select the extent and click **Maintenance Mode** on the ribbon. Alternatively, you can right-click the extent and select **Maintenance mode**.
5. Select the extent and click **Evacuate Backups** on the ribbon. Alternatively, you can right-click the extent and select **Evacuate backups**.
6. If you have disabled jobs, enable them.
7. After you evacuate backups, you can proceed to removing the extent from the scale-out backup repository. For more information, see Removing Performance Extents from Scale-Out Repositories.

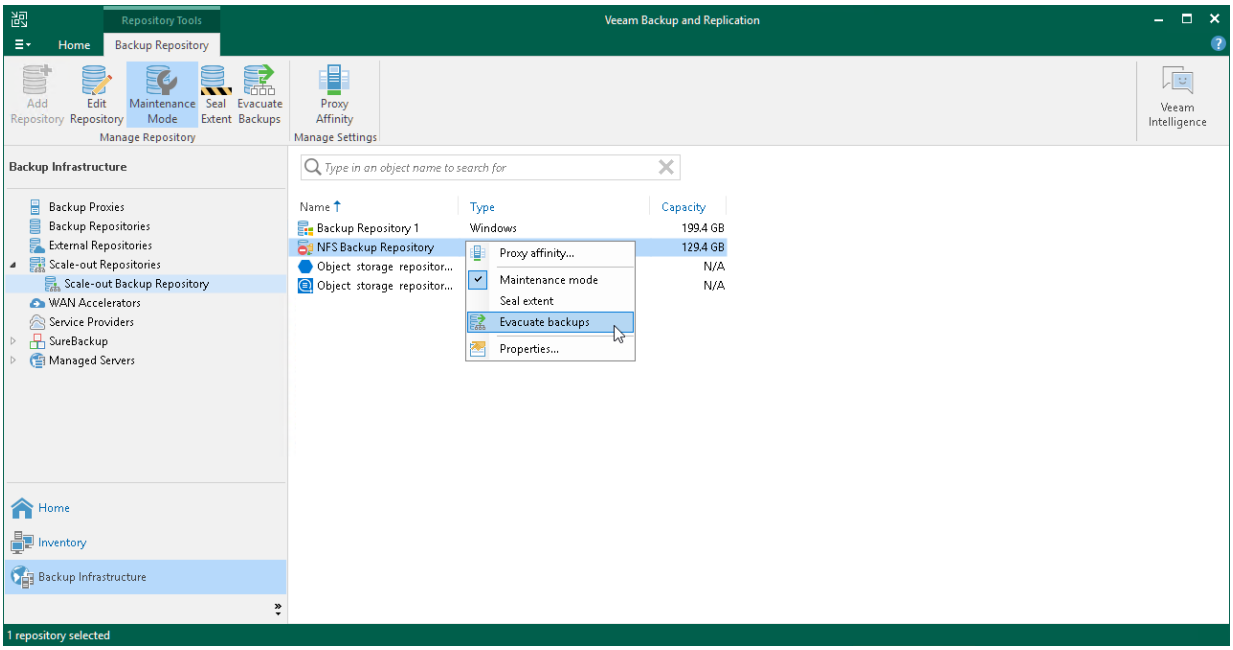

# Monitoring Evacuating Backups

To monitor backups evacuation, do the following:

1. Open the **History** view.
2. In the inventory pane, click **System**.
3. In the working area, select the evacuation session and click **Statistics** on the ribbon or right-click the evacuation session and select **Statistics**.

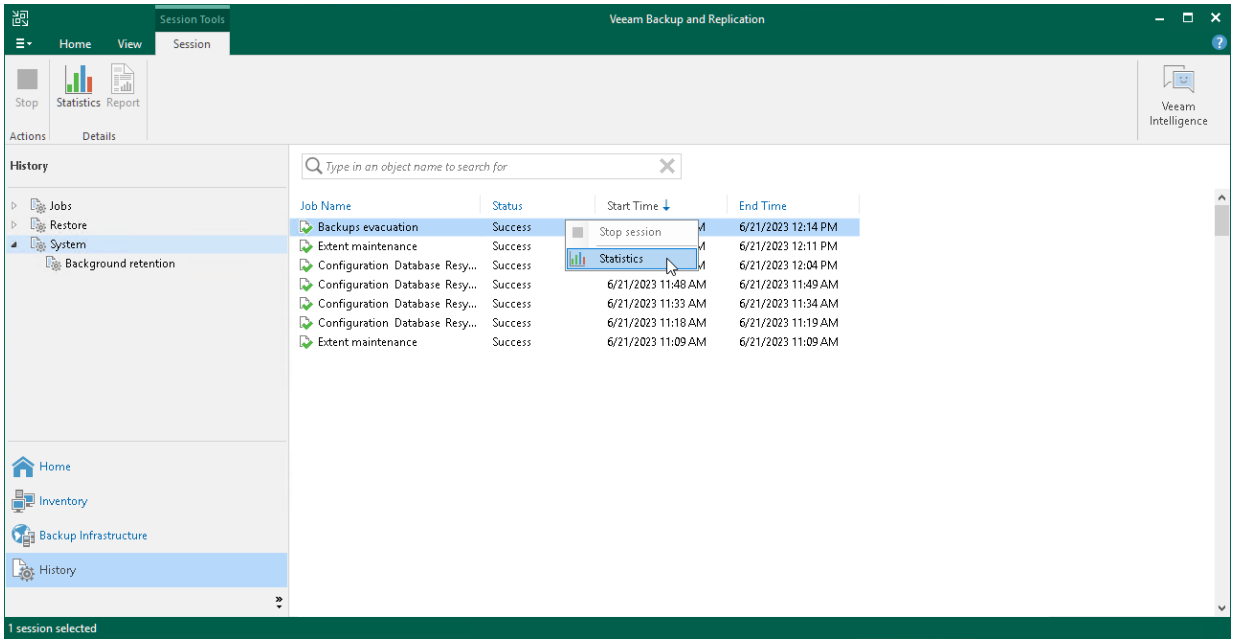


# Stopping Evacuating Backups

You can stop the ongoing evacuation process.

Evacuation cannot be stopped immediately. If you stop evacuation in the middle of moving a large backup file, Veeam Backup & Replication requires additional time to complete the process. Backup files that were not moved to a new extent remain on the current extent.

To stop backups evacuation, do the following:

1. Open the **History** view.
2. In the inventory pane, click **System**.

3. In the working area, select the evacuation session and click **Stop session** on the ribbon or right-click the evacuation session and select **Stop session**.

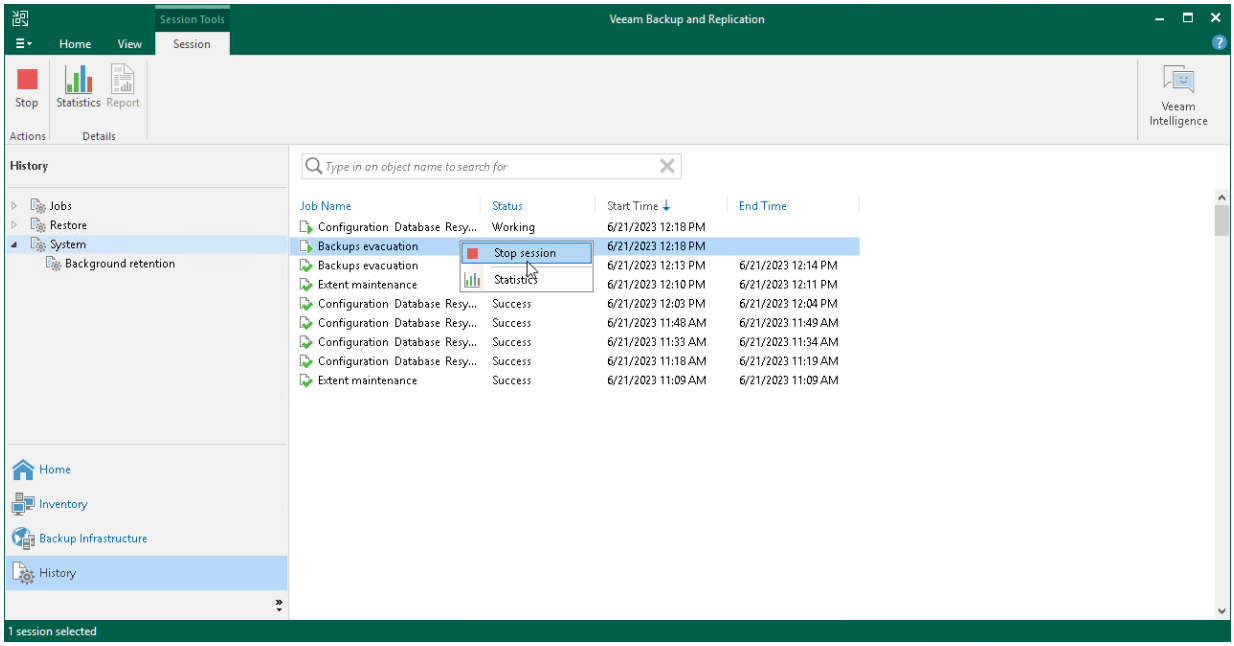


## Migrating Backups within Scale-Out Backup Repository

The migration option allows you to migrate data between performance or capacity extents of different object storage repositories. This option may be helpful if you want to switch from one object storage repository to another or to migrate from the object storage repository with different immutability settings.

> **IMPORTANT**
>
> If you need to migrate your data from a mutable bucket or container to an immutable bucket or container, perform the steps described in the Migrating Data Between Different Buckets or Containers scenario. Do NOT use 3rd party tools for this type of migration.

# Migrating Data Between Different Object Storage Repositories

You can follow this scenario in case you need to migrate data from one type of object storage to another. For example, to migrate data from the Amazon S3 object storage repository to the Azure Blob storage or any other object storage. These object storage repositories can have different or identical immutability settings. If Veeam Backup & Replication migrates data from the mutable object storage repository, it removes data from this repository. If Veeam Backup & Replication migrates data from the immutable object storage repository, this data remains intact after migration and will be removed according to the immutability period settings. Depending on the immutability settings, consider the necessary migration scenario.

> **IMPORTANT**
>
> Before migrating data, configure the connection through the gateway server for both the source and target object storage repositories to ensure data flows entirely within the cloud environment.

### Identical Immutability Settings

To migrate data, do the following:

1. Put an extent from which you want to migrate data to the Sealed mode.
2. Add a new object storage repository to the same (performance or capacity) tier of your existing scale-out backup repository.
3. Switch an extent from which you want to migrate data into the Maintenance mode.
4. Evacuate data from the source extent.
5. Remove the old object storage repository from the backup infrastructure.

### Different Immutability Settings

To migrate data, do the following:

1. Switch the extent from which you want to migrate data to the Sealed and Maintenance mode.
2. Add a new object storage repository to the same (performance or capacity) tier of your existing scale-out backup repository.
3. Evacuate data from the source extent.
4. [If the source repository is immutable] Disable the Maintenance mode for the source object storage repository.

> **NOTE**
>
> Do not turn off the Sealed mode for the source object storage repository. The backups will be removed from it according to its immutability period settings.

## Migrating Data Between Different Buckets or Containers

You can follow this scenario if you need to migrate data between different buckets or containers of the same object storage repository. These buckets or containers can have different or identical immutability settings. If Veeam Backup & Replication migrates data from the mutable bucket or container, it removes data from the source bucket or container. If Veeam Backup & Replication migrates data from the immutable bucket or container, this data remains intact after migration and will be removed according to the immutability period settings.

Depending on the immutability settings, consider the necessary migration scenario.

### Identical Immutability Settings

To migrate data, do the following:

1. Switch the extent from which you want to migrate data to the Maintenance Mode.
2. Add a new object storage repository to the backup infrastructure and specify a new bucket or container. For more information, see the Adding Object Storage Repositories section.
3. Set this new object storage repository as a performance or capacity extent of your existing scale-out backup repository. For more information, see the Extending Scale-Out Repositories section.
4. Evacuate data from the source extent.

5. Synchronize your data. For more information, see the Synchronizing Capacity Tier Data section.

### Different Immutability Settings

To migrate data, do the following:

1. Switch the extent from which you want to migrate data to the Sealed and Maintenance mode.
2. Add a new object storage repository to the backup infrastructure and specify a new bucket or container. For more information, see the Adding Object Storage Repositories section.
3. Set this new object storage repository as a performance or capacity extent of your existing scale-out backup repository. For more information, see the Extending Scale-Out Repositories section.
4. Evacuate data from the source extent.
5. [If the source repository is immutable] Disable the Maintenance mode for the source object storage repository.

   > **NOTE**
   >
   > Do not turn off the Sealed mode for the source object storage repository. The backups will be removed from it according to its immutability period settings.

6. Synchronize your data. For more information, see the Synchronizing Capacity Tier Data section.

## Copying Data Between Different Buckets or Containers with 3rd Party Tool

You can follow this scenario in case you want to use a 3rd party tool to migrate data located in an object storage repository between different buckets or containers of the same same object storage repository, do the following:

1. Use any available 3rd party tool to copy ALL data from an old bucket or container to a new bucket or container.
2. Add a new object storage repository to the backup infrastructure. For more information, see the Adding Object Storage Repositories section.
3. Set a new object storage repository as a capacity extent of the scale-out backup repository which you use.
4. Synchronize your data. For more information, see the Synchronizing Capacity Tier Data section.

## Migrating Data Within One Capacity Tier

You can migrate your data between different extents of the same tier.

To do this, perform the following steps:

1. Switch the necessary extent into the Maintenance mode.
2. Evacuate backups from the extent.

# Receiving Scale-Out Backup Repository Reports

Veeam Backup & Replication is capable of sending reports that contain information about processing results of your scale-out backup repositories data.

Consider the following:

- Reports are sent only after you have enabled and configured email notifications, as described in section Configuring Global Email Notification Settings.
- Reports are sent daily at time specified in global notification settings.
- Reports are sent for all notification types selected in global notification settings, such as *Success*, *Warning* and *Failure*.
- The title of a report is built up of "*Scale-out Backup Repository*" + a repository name. That said, if your scale-out backup repository name is Amazon, then the report title will be *Scale-out Backup Repository Amazon*.

Each report is divided into sections and contains the following information:

- **Performance Tier** (upper-left) section:
  - **Used Space**. Shows the used disk space of your scale-out backup repository.
  - **Capacity**. Shows the total storage capacity of your scale-out backup repository.
- **Capacity Tier** (upper-right) section:
  - **Used Space**. Shows the occupied storage space in your object storage repository.
  - **Space Limit**. Shows the space limit (if any). A space limit is specified when adding a new object storage repository, as described in section Adding Object Storage Repositories.
- **Performance Tier** (middle) section:
  - **Extent**. Shows extents of a scale-out backup repository.
  - **Capacity**. Shows the total storage capacities of your performance extents.
  - **Used Space**. Shows the amount of disk space used on your extents.
  - **Status**. Shows the status of each extent, as described in section Description of Report Statuses.
- **Capacity Tier** (lower) section:
  - **Extent**. Shows the name of the capacity extent.
  - **Space Limit**. Shows the space limit (if any).
  - **Used Space**. Shows the occupied storage space in your capacity extent.
  - **Status**. Shows the status of the capacity extent, as described in section Description of Report Statuses.

    If an automatic offload job session exits with any status other than *Success*, you will see the associated status message in this field. For more information about the offload job, see Moving Backups to Capacity Tier.

# Description of Report Statuses

The following table lists possible combinations of *Warning* and *Error* messages shown under the **Status** column of a report.

If none of the conditions listed in the **Extent state** column is true, then the report status will be shown as *Success*.

| Extent type | Extent state | Status message | Report type |
|---|---|---|---|
| **Performance tier** | Maintenance mode | Maintenance mode | Warning |
| | Threshold limit exceeded. Threshold is specified in the **Backup storage** section, as described in section Specifying Other Notification Settings. | Reaching capacity | Warning |
| | Unavailable | Offline | Error |
| **Capacity tier** | Maintenance mode | Maintenance mode | Warning |
| | Space limit exceeded. Space limit is specified when adding a new object storage repository, as described in section Adding Object Storage Repositories. | Out of capacity | Error |
| | Unavailable | Offline | Error |
| | Threshold limit exceeded. Threshold is specified in the **Backup storage** section, as described in section Specifying Other Notification Settings. | Reaching capacity | Warning |

# Report Examples

## Success Reports

The following figure shows an example of a report consisting of two performance extents (*Backup Volume 01* and *Backup Volume 02*); both share *253.3 GB* of storage capacity, of which *52.4 GB* is occupied.

Both extents have OK status, which means that neither extent was put into the Maintenance mode, nor has any of these performance extents exceeded the allowed threshold limit.

This report also includes the **Capacity Tier** section consisting of a capacity extent with no **Space Limit** applied. This capacity extent stores *29.6 GB* of data and has the *OK* status.

**Scale-out backup repository Amazon** — **Success**
**Remote Object Storage**

Wednesday, December 19, 2022 4:04:12 AM

| Performance Tier | | Capacity Tier | |
|---|---|---|---|
| **Used Space** | 52.4 GB | **Used Space** | 29.6 GB |
| **Capacity** | 253.3 GB (70% free) | **Space Limit** | Not set |

Performance Tier

| Extent | Capacity | Used Space | Status |
|---|---|---|---|
| Backup Volume 01 | 126.7 GB (67% free) | 30.3 GB | OK |
| Backup Volume 02 | 126.7 GB (73% free) | 22.1 GB | OK |

Capacity Tier

| Extent | Space Limit | Used Space | Status |
|---|---|---|---|
| Amazon S3 Object Storage | Not set | 29.6 GB | OK |

## Warning Reports

The following figure demonstrates a report with the *Warning* status.

As per example, the *Backup Volume 01* performance extent has been put into the Maintenance mode, and the *Backup Volume 02* performance extent has exceeded the allowed threshold both of which have caused a report to be generated with the *Warning* status.

**Scale-out backup repository Amazon** — **Warning**
**Remote Object Storage**

Wednesday, December 19, 2022 4:58:35 AM

| Performance Tier | | Capacity Tier | |
|---|---|---|---|
| **Used Space** | 82.0 GB | **Used Space** | 29.6 GB |
| **Capacity** | 253.3 GB (58% free) | **Space Limit** | Not set |

Performance Tier

| Extent | Capacity | Used Space | Status |
|---|---|---|---|
| Backup Volume 01 | 126.7 GB (43% free) | 60.0 GB | Maintenance mode |
| Backup Volume 02 | 126.7 GB (73% free) | 22.1 GB | Reaching capacity |

Capacity Tier

| Extent | Space Limit | Used Space | Status |
|---|---|---|---|
| Amazon S3 Object Storage | Not set | 29.6 GB | OK |

### Error Reports

In the following figure, a report has been generated with the *Error* status caused by the *Amazon S3 Object Storage* performance extent which has exceeded its allowed space limit.

**Scale-out backup repository Amazon** — **Error**
Remote Object Storage

Thursday, December 20, 2022 3:55:02 AM

| Performance Tier | | Capacity Tier | |
|---|---|---|---|
| **Used Space** | 82.0 GB | **Used Space** | 29.6 GB |
| **Capacity** | 253.3 GB (58% free) | **Space Limit** | 2.0 GB (0% free) |

Performance Tier

| Extent | Capacity | Used Space | Status |
|---|---|---|---|
| Backup Volume 01 | 126.7 GB (43% free) | 60.0 GB | OK |
| Backup Volume 02 | 126.7 GB (73% free) | 22.1 GB | OK |

Capacity Tier

| Extent | Space Limit | Used Space | Status |
|---|---|---|---|
| Amazon S3 Object Storage | 2.0 GB (0% free) | 29.6 GB | Out of capacity |

# Removing Backups from Capacity or Archive Tier

To remove moved or copied backups from capacity or archive extent, use the **Delete from disk** feature, as described in section Deleting Backups from Scale-Out Backup Repositories.

Consider the following:

- When removing offloaded backup files from the backup chain that was created with the per-machine method, the associated blocks of data will be removed from the capacity or archive extent altogether.

  For more information about per-machine backups, see Backup Chain Formats.

- When removing offloaded backup files from the backup chain that was created as a single-file backup file, then nothing will be removed until either of the following occurs:
  - All the VMs were removed from the backup.
  - The backup itself was removed.
- Immutable backups cannot be removed.

  For more information, see Immutability for Scale-Out Backup Repositories.

- If the capacity or archive extent has been put into the Maintenance mode, the removal of data from such a repository is not possible until the extent is removed from the Maintenance mode.

  For more information, see Switching to Maintenance Mode.

- During data removal, the entire folder structure starting from the repository folder (`<backup_id>`) will be completely purged.

  For more information on how Veeam Backup & Replication stores data in the capacity extent, see Capacity Extent Structure.

- If backup files with metadata that are located on your extents have been removed locally in any way other than by using the Deleting from Disk feature, Veeam Backup & Replication will not be able to synchronize the backup chain state with that of the capacity or archive extent. Therefore, the offloaded blocks of data will continue to remain in cloud storage. To remove such blocks, use your cloud platform abilities.

# Rebalancing Extents of Scale-Out Backup Repositories

To maintain Backup File Placement policies and distribute backup data between performance extents evenly, you can use a rebalance. It can be useful if one of the performance extents of your scale-out backup repository contains more data than the other extents.

You can use the rebalance to distribute data within extents according to this policy in the following cases:

- To maintain the current placements policy. If data within performance extents is not distributed according to the placement policy and some extent contains more backups than the others, you can rebalance these extents.
- To change the placement policies. If you have switched from one type of policy to another, for example from *Data locality* policy to *Performance* placement policy, the rebalance will help to distribute data according to a new policy.

After you start the rebalance, Veeam Backup & Replication will scan performance extents and backup files located on these extents. In case some backup chain does not meet the backup placement policy, Veeam Backup & Replication will check all available performance extents, analyze their total space and amount of free space and will evaluate the best way to distribute data between extents. After that, Veeam Backup & Replication will move backup chains from the source extent to a new extent.

Depending on the mechanisms that are used to store data on extents, Veeam Backup & Replication chooses one of the following algorithm to perform the rebalance:

- Common rebalance — Veeam Backup & Replication uses this type of the rebalance, if a scale-out backup repository contains at least one extent that does not use the data deduplication or block cloning technologies.
- Deduplication rebalance — Veeam Backup & Replication uses this type of the rebalance, if all extents of a scale-out backup repository use the data deduplication, block cloning technologies and applies the data locality placement policy.

# Considerations and Limitations

Before you start the rebalance, consider the following:

- Before you perform the rebalance, you must stop all activities (for example, backup jobs, restore operations and so on) that are related to the scale-out backup repository.
- The rebalance option is not supported for object storage repositories added as performance extents of the scale-out backup repository.
- If you use the ReFS Microsoft Windows file system or XFS Linux file system as performance extents, Veeam Backup & Replication will utilize Fast Clone technology.
- Veeam Backup & Replication skips from rebalance immutable backups located on immutable extents and hardened repositories added as extents of the scale-out backup repository.
- Veeam Backup & Replication does not rebalance the following types of extents:
  - Capacity tier extents.
  - Archive extents.

- Performance extents that are set to the Maintenance mode.

- You must stop the backup jobs created by Veeam Plug-in for Microsoft SQL Server or Veeam Plug-in for Oracle RMAN.

# How Rebalance of Scale-Out Backup Repositories Works

After you start a rebalance, Veeam Backup & Replication performs the following steps:

1. The rebalance session starts.
2. Veeam Backup & Replication gets a list of available extents.
3. All available extents are set to the Maintenance mode.
4. Veeam Backup & Replication checks extents configuration and selects the best algorithm to perform the rebalance.
5. Veeam Backup & Replication checks if data is distributed across extents according to the placement policy. If at least one backup files does not meet the placement policy, the whole chain on the extents is marked as a violator.
6. Veeam Backup & Replication checks the capacity of every extent, prioritize them according to the placement policy and optimal free space.
7. Veeam Backup & Replication chooses the preferred extent to keep data.
8. Veeam Backup & Replication starts to evacuate data from the current extents and distribute it to the preferred extent.
9. After the data is distributed, extents are removed from the Maintenance mode.

# Rebalancing Extents of Scale-Out Backup Repositories

To start a rebalance session, you must press and hold the [Ctrl] key on your keyboard while you right-click the necessary scale-out backup repository and select **Rebalance**. After that Veeam Backup & Replication will start a session that will show details on how data is moved from one performance extent to another performance extent.

> **TIP**
>
> If you want to exclude a specific performance extent from rebalance, set to the Maintenance mode.

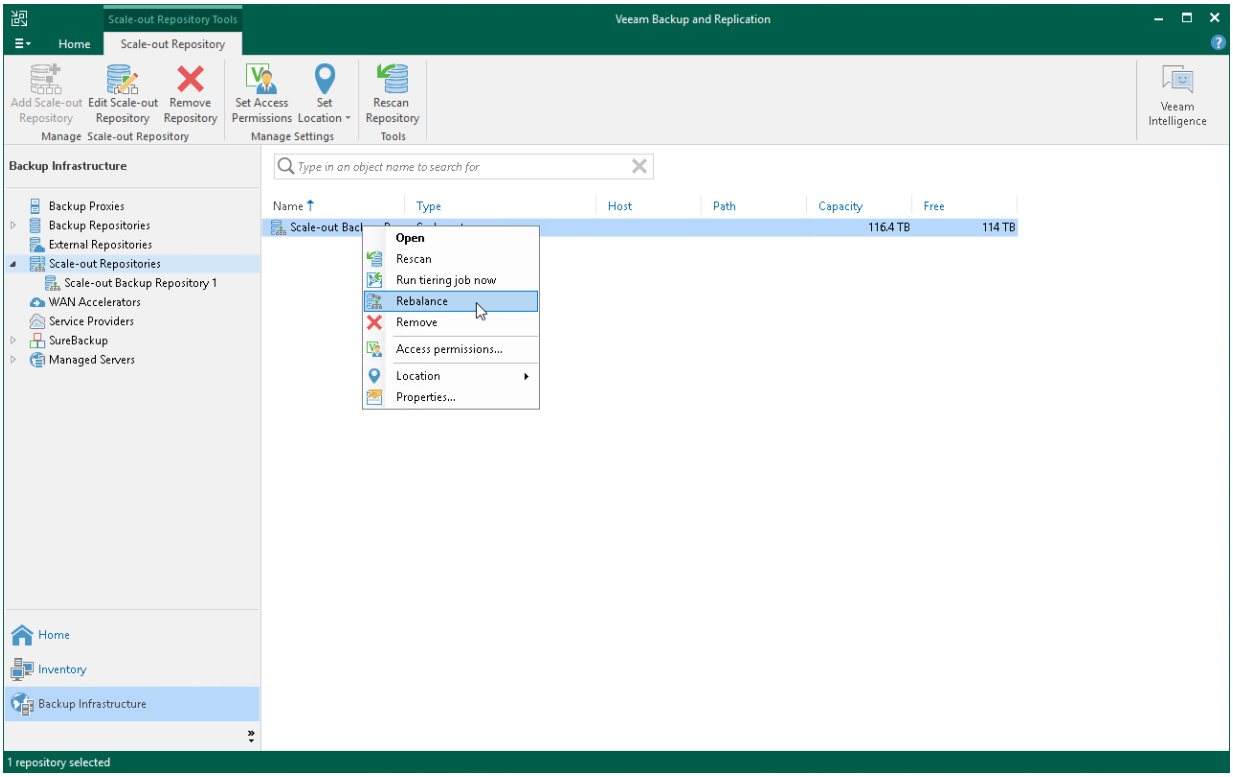


# Removing Scale-Out Backup Repositories

You can remove a scale-out backup repository at any time. When you remove a scale-out backup repository, Veeam Backup & Replication unassigns the extent role from all the backup repositories configured into it, and they become individual backup repositories. Backup files are not removed from the backup repositories — they remain on the disk or an object storage repository.

You cannot remove a scale-out backup repository if at least one job is targeted at it. First, you must move all backup files to the new backup repository and then retarget the jobs. For details, see this Veeam KB article.

To remove a scale-out backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Scale-out Repositories**.

3. In the working area, select the scale-out repository and click **Remove Repository** on the ribbon or right-click the backup repository and select **Remove**.

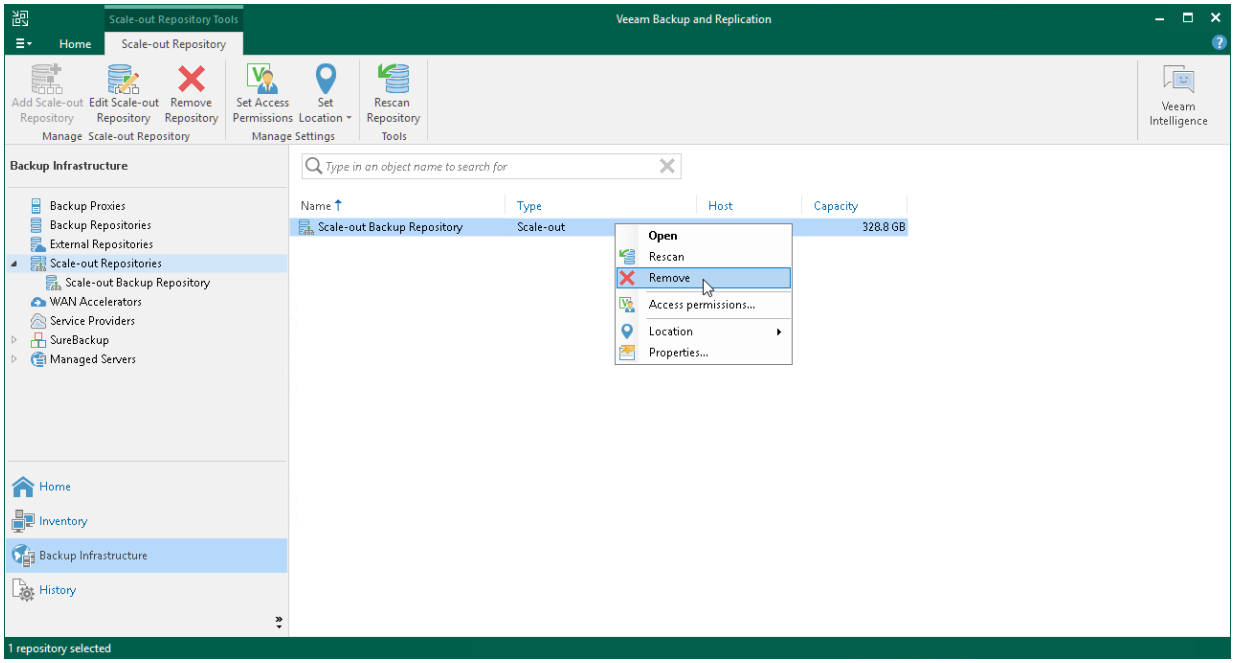

# Guest Interaction Proxies

The guest interaction proxy is a backup infrastructure component that sits between the backup server and processed VM. To interact with the VM guest OS, Veeam Backup & Replication needs either to install non-persistent runtime components or use (if necessary, install) persistent agent components in each VM. The task of deploying these components in a VM is performed by the guest interaction proxy. For more information on the components, see Non-Persistent Runtime Components and Persistent Agent Components.

> **NOTE**
>
> The guest interaction proxy functionality is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.

The guest interaction proxy allows you to communicate with the VM guest OS even if the backup server and processed VM run in different networks.

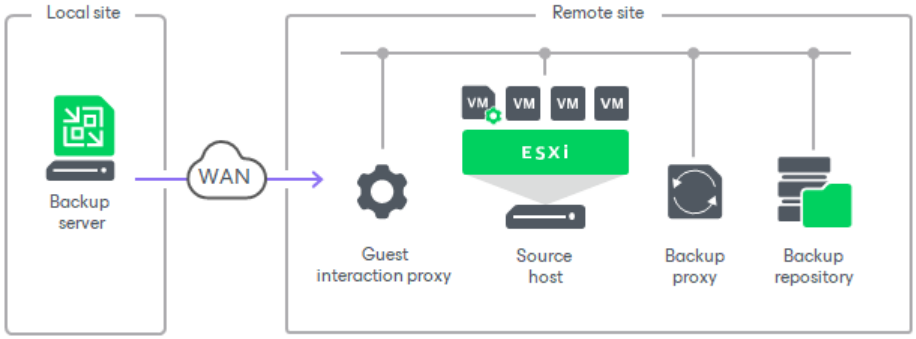


> **IMPORTANT**
>
> The guest interaction proxy deploys the non-persistent runtime components or persistent agent components only in Microsoft Windows VMs and Microsoft Windows-based Oracle, PostgreSQL and Microsoft SQL Server VMs. In VMs with another guest OS, the non-persistent runtime components or persistent agent components are deployed by the backup server.

## Usage Scenarios

This component is needed if the backup or replication jobs perform the following processing of VMs:

- Application-aware processing
- Guest file system indexing
- Transaction logs processing

## Guest Interaction Proxy Deployment

You can use multiple guest interaction proxies to improve performance. Multiple guest interaction proxies will deploy non-persistent runtime components or persistent agent components in VMs faster compared to the same operation performed by one guest interaction proxy.

In a backup infrastructure with multiple remote sites, you can deploy a guest interaction proxy in each site. This can reduce load on the backup server and produce less traffic between the backup server and remote site.

# Requirements for Guest Interaction Proxy

A machine performing the role of a guest interaction proxy must meet the following requirements:

- The role of a guest interaction proxy can be assigned to a Microsoft Windows server (physical or virtual).
- You must add the machine to the Veeam Backup & Replication console as a managed server.
- Guest interaction proxy must have either a LAN or VIX connection to the VM that will be processed. You do not have to set up both connections — only one connection is required. For more information about setting up a connection to the VM, see this Veeam KB article.

The guest interaction proxy role can be performed by any machine that meets the requirements, including VMware backup proxy, backup repository, WAN accelerator or backup server.

# Guest Interaction Proxy Selection

When you add a Microsoft Windows machine to the backup infrastructure, Veeam Backup & Replication deploys Veeam Data Mover on it. Veeam Data Mover includes the components responsible for deployment of non-persistent runtime components or persistent agent components during guest OS interaction.

To assign a guest interaction proxy for the job, you must select a Microsoft Windows machine that will perform the role of the guest interaction proxy at the **Guest Processing** step of the backup or replication job wizard. You can assign the guest interaction proxy manually, or let Veeam Backup & Replication do it automatically. Veeam Backup & Replication uses the following priority rules to select the guest interaction proxy:

1. A machine in the same network as the protected VM that does not perform the backup server role.
2. A machine in the same network as the protected VM that performs the backup server role.
3. A machine in another network that does not perform the backup server role.
4. A machine in another network that performs the backup server role.

If Veeam Backup & Replication finds several available machines of equal priority, it selects the less loaded machine. The load is defined by the number of tasks that the machine already performs.

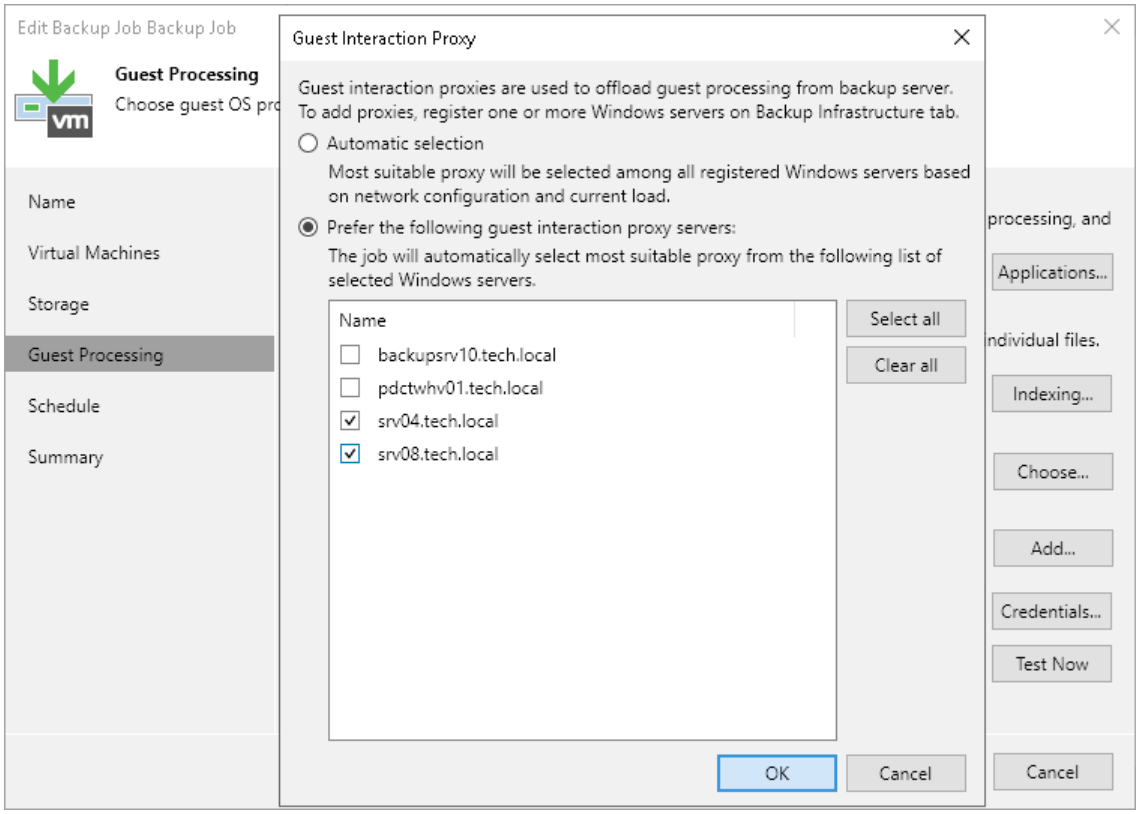


# Failover from Guest Interaction Proxy to Backup Server

If the guest interaction proxy fails to connect to a Microsoft Windows VM, the guest interaction proxy will not be able to access the VM and deploy non-persistent runtime components or persistent agent components in it. In this case, the backup server will take over the role of guest interaction proxy and deploy the non-persistent runtime components or persistent agent components in the VM.

# Gateway Servers

A gateway server is an auxiliary backup infrastructure component that "bridges" the backup server and backup repository. It can also "bridge" a source backup repository and a target backup repository in case of backup copy jobs. The gateway server is required if you deploy the following types of backup repositories in the backup infrastructure:

- Shared folder backup repositories
- Dell Data Domain deduplicating storage appliance
- HPE StoreOnce deduplicating storage appliance
- Object storage repository

Such backup repositories cannot host Veeam Data Movers — Veeam components that establish a connection between a backup proxy and backup repository (in case of backup jobs) or between backup repositories (in case of backup copy jobs). To overcome this limitation, Veeam Backup & Replication uses gateway servers.

In the backup infrastructure, a gateway server hosts the target Veeam Data Mover. Veeam Backup & Replication establishes a connection between the source Veeam Data Mover and target Veeam Data Mover, and transports data from/to backup repositories through gateway servers.

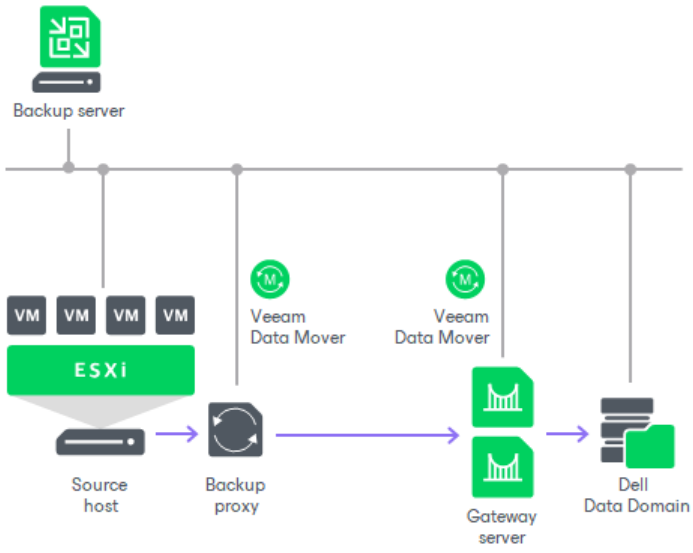


For more information on using gateway servers in backup copy jobs, see Backup Copy Architecture.

# Requirements and Limitations for Gateway Servers

Consider the following:

- The gateway server can run on a Microsoft Windows or Linux machine that is added to the backup infrastructure as a managed server. The machine must meet the system requirements. For more information, see System Requirements.
- The role of a gateway server for Dell Data Domain, HPE StoreOnce storage appliances or SMB repositories must be assigned to a Microsoft Windows machine.
- The role of a gateway server for NFS or object storage repositories can be assigned to a Microsoft Windows or Linux machine.

- The machine must have access to the backup repository — shared folder, Dell Data Domain or HPE StoreOnce.
- For deduplicating storage appliances working over Fibre Channel, you must explicitly select at least one gateway server that will communicate with the appliance over Fibre Channel connection.
- For HPE StoreOnce deduplicating storage appliances, you must assign the role of a gateway server to a 64-bit machine.
- If connection to the gateway server is lost during the job run, the job fails. Veeam Backup & Replication selects a new available gateway server when the job starts next time.

# Gateway Server Deployment

To configure a gateway server, you must first add a machine that you plan to use as a gateway server to the backup infrastructure using the **New Windows Server** or **New Linux Server** wizard. For more information, see Adding Microsoft Windows Servers or Adding Linux Servers.

After that, you must go through the **New Backup Repository** wizard and define gateway server settings. For more information, see Adding Backup Repositories. You can select a gateway server explicitly or instruct Veeam Backup & Replication to select it automatically.

If you plan to select gateway servers explicitly, these servers must be located as close to the backup repository as possible. However, if you use a deduplicating storage appliance with source-side data deduplication, it is reasonable to assign the roles of gateway servers to machines that are located closer to the backup proxy. This will help you reduce the amount of traffic traveling over the network. For more information, see Dell Data Domain and HPE StoreOnce.

# Gateway Selection

Whenever possible, Veeam Backup & Replication distributes the backup workload between multiple available gateway servers. This helps optimize performance of multiple concurrent tasks. Veeam Backup & Replication assigns a separate gateway server for each task, based on gateway server connectivity and their current load. When Veeam Backup & Replication selects gateway servers, the following applies:

- If the number of tasks is greater than the number of available gateway servers, Veeam Backup & Replication uses one gateway server for multiple tasks.
- Veeam Backup & Replication uses one gateway server for the whole job if the **Use per-machine backup files** option is disabled for the repository to which the job is targeted at. For more information on this option, see Backup Chain Formats.
- For per-machine VM backups, Veeam Backup & Replication uses one gateway server to process all disks of a VM.

You can select gateway servers explicitly or instruct Veeam Backup & Replication to select them automatically. For more information on which backup infrastructure components Veeam Backup & Replication uses as gateway servers during automatic selection, see the Automatic Selection section.

## Manual Gateway Selection

If you select gateway servers explicitly, Veeam Backup & Replication uses only the selected servers and performs all operations on them. Veeam Backup & Replication analyzes the gateway server connectivity and their current task load, and picks the most suitable gateway server for the next task.

**IMPORTANT**

If you select only one gateway server explicitly and it is not accessible, the job will fail.

## Automatic Gateway Selection

If you instruct Veeam Backup & Replication to select gateway servers automatically, Veeam Backup & Replication uses the backup infrastructure components described in the following table. Note that principles described in the Gateway Selection section also apply, and multiple gateway servers may be used for concurrent tasks. If the primary selection gateway server is not accessible, Veeam Backup & Replication fails over to the next available option.

| Type of job | Component used as gateway server | Component used as gateway server for synthetic operations |
|---|---|---|
| **Backup job / File backup job** | Backup proxy that was assigned first to process workload data / file share for a backup job*.<br><br>For backup and file backup jobs to an object storage repository with the direct connection mode, Veeam Backup & Replication does not use only one proxy as the gateway server for the entire job duration. Instead, it assigns the role of a gateway server to the least loaded available proxy before each operation. | Synthetic operations are performed on the mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server. |

| Type of job | Component used as gateway server | Component used as gateway server for synthetic operations |
| --- | --- | --- |
| **Backup copy job / File backup copy job** | For backup copy and file copy jobs, the selected gateway depends on the type of the source backup repository:<br>• Direct attached repository. The mount server associated with the backup repository is used as the gateway server. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server.<br>• Backup repository with the gateway server connection. The gateway of the source backup repository is used.<br>For backup copy and file copy jobs to an object storage repository with the direct connection mode, the source backup repository is used as the gateway server.<br>For backup copy jobs that work over WAN accelerators, the role of a gateway server is assigned to source or target WAN accelerator (depending on the shared folder backup repository location). File backup copy job does not support WAN accelerators. | Synthetic operations are performed on the mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server.<br>These rules are applied to the direct data path and processing over WAN accelerators. File backup copy job does not support WAN accelerators. |
| **Tape job** | If there is a direct connection between a backup repository and tape device, the role of a gateway server is assigned to the tape server.<br>Otherwise, the role of a gateway server is assigned to the backup server. | Synthetic operations are performed on the mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server. |
| **Veeam Agent backup job** | Mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server. | Synthetic operations are performed on the mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server. |

| Type of job | Component used as gateway server | Component used as gateway server for synthetic operations |
| --- | --- | --- |
| **Backup job created by Veeam Plug-in for Oracle RMAN/SAP HANA/SAP on Oracle/Microsoft SQL Server/DB2** | Backup server. | — |
| **Restore operations** | Backup proxy used for a restore operation*. | — |
| **Replication from backup** | Target backup proxy assigned for a replication operation*. | — |
| **Repository rescan** | Mount server associated with the backup repository. | — |
| **Offload job** | If you offload backups to an object storage repository with the direct connection mode, the selected gateway depends on the type of the source backup repository:<br>• Direct attached repository. The source backup repository is used as the gateway server. For hardened repository, the role of the gateway server is assigned to the hardened repository.<br>• Non-direct attached repository. The gateway of the source backup repository is used.<br>• Object storage repository with the direct connection mode. The mount server associated with the source backup repository is used as the gateway server. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server.<br>• Object storage repository with the gateway server connection. The gateway of the source backup repository is used. | — |

| Type of job | Component used as gateway server | Component used as gateway server for synthetic operations |
|---|---|---|
| **Move and copy operations** | If you copy or move backups to an object storage repository with the direct connection mode, the selected gateway depends on the type of the source backup repository:<br>• Direct attached repository. The source backup repository is used as the gateway server. For hardened repository, the role of the gateway server is assigned to the hardened repository.<br>• Non-direct attached repository. The gateway of the source backup repository is used.<br>• Object storage repository with the direct connection mode. The mount server associated with the source backup repository is used as the gateway server. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server.<br>• Object storage repository with the gateway server connection. The gateway of the source backup repository is used.<br>If you copy or move backups to other non-direct attached backup repository, the role of the gateway server is assigned to mount server. If mount server is not accessible,<br>Veeam Backup & Replication fails over to the backup server. | – |
| **Health check operation** | Mount server associated with the backup repository. | – |

* [For repositories for which Linux gateway servers cannot be used – Dell Data Domain, HPE StoreOnce and SMB]: If the backup proxy is a Linux machine, the role of the gateway server is assigned to the mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server.

# Related Topics

- Scale-Out Backup Repositories

- Specifying Server or Shared Folder Settings

# Mount Servers

The mount server is a server required for restores that work with guest OS files and application items. To access files or items stored in a backup file, Veeam Backup & Replication mounts the content of the backup to the mount server. Only after the content is mounted, Veeam Backup & Replication can get files and copy them to the restore destination.

The mount server is required if you perform the following operations:

- Guest OS file restore
- Application items restore
- Secure restore
- Instant recovery
- Instant disk recovery
- Instant file share recovery
- Restore of specific objects from file backup and object storage backup

To reduce the load on the network and speed up the restore process, the mount server must be located in the same site as the backup repository where backup files are stored. In this case, you will be able to keep the traffic in one site. If the mount server is located in some other site, the data will travel across the network between the sites.

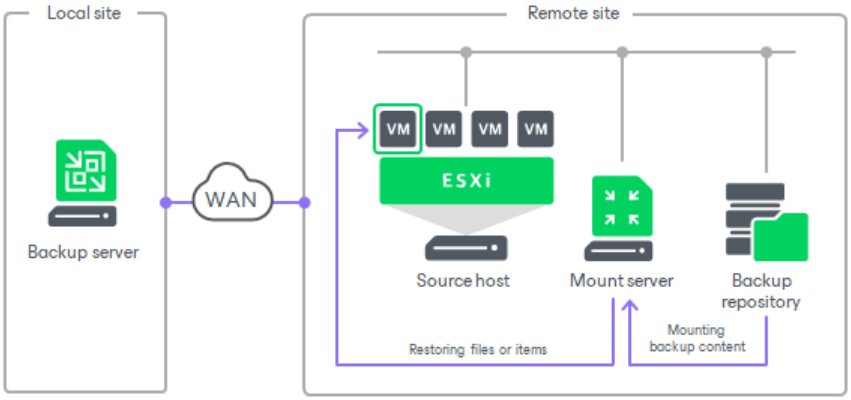


**NOTE**

In some scenarios, Veeam Backup & Replication can mount content of backups to machines other than mount servers. For more information, see Mount Points and Restore Scenarios.

## Mount Server Deployment

The mount server is created for every backup repository and is associated with it. When you configure a backup repository, you specify to which server you want to assign the role of the mount server.

You can assign the mount server role to any 64-bit Microsoft Windows machine added to the backup infrastructure. This machine and the backup repository must be located as close to each other as possible. If you have several sites, we recommend you configure at least one mount server in each site.

It is recommended to enable the vPower NFS Service on the mount server. In this case, the mount server acts as the vPower NFS server.

By default, Veeam Backup & Replication suggests assigning the mount server role to the following infrastructure components depending on the OS of the backup repository:

- For Microsoft Windows backup repositories, Veeam Backup & Replication suggests the backup repository itself.
- For Linux, shared folder backup repositories and deduplicating storage appliances, Veeam Backup & Replication suggests the backup server.

> **NOTE**
>
> Consider the following:
>
> - For scale-out backup repositories, you must specify the mount server for every extent.
> - For cloud repositories and hosts that store replicas or backups from storage snapshots, the mount server role is assigned to the backup server. For such repositories, you cannot assign the mount server role to a different machine.

# Mount Server Services and Components

Mount servers run light-weight services that take a few seconds to deploy. Deployment is fully automated. Veeam Backup & Replication installs the following services:

- **Veeam Mount Service** mounts backups and replicas for file-level access, browsing the guest file system and restoring guest OS files and application items.
- **Veeam Data Mover** handles traffic.
- **Veeam vPower NFS Service** (if you enable it when configuring the mount server).
- **Veeam Threat Hunter Service** runs a Veeam Threat Hunter scan.

# Requirements for Mount Servers

A machine that performs the role of a mount server must meet the following requirements:

- The machine must meet the system requirements. For more information, see System Requirements.
- The mount server must have access to the backup repository with which it is associated and to the original VM (the VM to which you restore files or application items). For restore from storage snapshots, the mount server must also have access to the ESXi host on which the temporary VM is registered.

# Veeam Data Mover Service

Veeam Data Mover performs data processing tasks on behalf of Veeam Backup & Replication, such as retrieving source machine data, performing data deduplication and compression, and storing backed-up data on the target storage.

Veeam Data Mover can be persistent or non-persistent:

- For Microsoft Windows servers, Veeam Data Movers are persistent, that is, Veeam Data Mover is uploaded and installed on a server only once.

  Veeam Backup & Replication automatically installs Veeam Data Mover when you add a Microsoft Windows server to the backup infrastructure.

- For Linux servers, Veeam Data Movers can be persistent or non-persistent:
  - Non-persistent Veeam Data Mover is uploaded and removed through the SSH connection each time Veeam Backup & Replication addresses a server. If you specify an account with non-root permissions when adding a Linux server, Veeam Data Movers will be non-persistent.
  - For Veeam Data Mover to be persistent, you must specify an account with root or equivalent to root permissions when adding a Linux server. After the Linux server is added to the backup infrastructure, root or equivalent to root permissions can be revoked. The permissions are required only when adding, editing or removing a server from the backup infrastructure.

    Persistent Veeam Data Movers are required for hardened repositories and backup proxies. For other backup infrastructure components based on Linux servers, Veeam Data Movers can be persistent or non-persistent.

    > **NOTE**
    >
    > Persistent Veeam Data Movers are required for Linux servers with the backup proxy role assigned. When you upgrade to the latest version, Veeam Backup & Replication automatically updates all Linux backup proxies and installs Veeam Data Movers on them.
    > Veeam Backup & Replication does not automatically update Linux servers that have other roles. If you want to install Veeam Data Movers on such servers, open the Edit Linux Server wizard for the necessary server and click **Finish**.

## Requirements and Limitations for Veeam Data Movers

Before you use Veeam Data Movers, consider the following requirements and limitations:

- For Microsoft Hyper-V and Microsoft Windows servers:
  - File and printer sharing must be enabled in network connection settings of the added server.
  - Make sure the user account that you use to add a Microsoft Windows server is in the local administrators group on the server being added.
- For Linux servers:
  - Linux server version must be 64-bit. For more information, see System Requirements. Note that Perl is required only for non-persistent Veeam Data Movers. Check the full list of required Perl modules in this Veeam KB article.
  - Veeam Backup & Replication does not install Veeam Data Movers on deduplicating storage appliances based on Linux.

- Make sure the user account that you use to add a Linux server has required permissions:
  - If you want to deploy persistent Veeam Data Movers, the user account must have root or elevated to root permission at the moment of deployment. Otherwise, Veeam Data Movers will be non-persistent, that is, the Linux server will not host Veeam Data Movers permanently.
  - For hardened repository, when you add a Linux server with single-use credentials, the user account must have elevated to root permission in order for Veeam Data Movers to be persistent. For more information, see Step 3. Specify Linux Server.

> **TIP**
>
> If Veeam Data Mover on a Linux server fails for some reason, you can re-install it manually. For more information, see this Veeam KB article.

# Veeam vPower NFS Service

The vPower technology enables the following features:

- SureBackup
- SureReplica
- Instant Recovery
- Instant Disk Recovery
- Staged restore
- Universal Application-Item Recovery (U-AIR)
- Multi-OS guest OS file restore

The key construct of the vPower technology is the vPower NFS Service. The vPower NFS Service is a Microsoft Windows service that runs on a Microsoft Windows machine and enables this machine to act as an NFS server.

On the vPower NFS server, Veeam Backup & Replication creates a special directory — the vPower NFS datastore. When you start a VM or a VM disk from a backup, Veeam Backup & Replication "publishes" VMDK files of the VM from the backup on the vPower NFS datastore. Technically, Veeam Backup & Replication emulates the presence of VMDK files on the vPower NFS datastore — the VMDK files themselves are still located in the backup file in the backup repository.

The vPower NFS datastore is then mounted to the ESXi host. As a result, the ESXi host can "see" backed-up VM images with the help of the vPower NFS datastore and work with them as with regular VMDK files. The emulated VMDK files function as pointers to the real VMDK files in the backup repository.

The vPower NFS datastore stays mounted to the ESXi host. Each time you start an operation from the list, Veeam Backup & Replication checks whether vPower NFS datastore is mounted to the host. If you manually unmounted the datastore or the datastore was unmounted by other causes, Veeam Backup & Replication remounts the datastore.

> **IMPORTANT**
>
> Veeam vPower NFS datastores are service datastores that can be used for vPower operations only. You cannot use them as regular VMware vSphere datastores — for example, you cannot place files of replicated VMs on such datastores.

## vPower NFS Server Location

If you store backups on a Microsoft Windows backup repository, it is strongly recommended that you enable the vPower NFS server on this backup repository. In this case, Veeam Backup & Replication will be able to set up a direct connection between the backup repository and ESXi to which the vPower NFS datastore is mounted.

The Veeam vPower NFS Service can also run on any Microsoft Windows server in the backup infrastructure, including the backup server itself. However, in this case the recovery verification performance may decrease. The connection between the ESXi host and backup repository will be split into two parts:

1. From ESXi host to the vPower NFS server
2. From the vPower NFS server to the backup repository

# vPower-Specific Settings

To establish a connection between the ESXi host and vPower NFS server, you must make sure that the ESXi host has a proper network interface configuration and can access the vPower NFS server.

> **NOTE**
>
> Veeam Backup & Replication uses IP address authorization to restrict access to the vPower NFS server. By default, the vPower NFS server can be accessed only by the ESXi host that provisioned the vPower NFS datastore. However, you can disable this option with a registry value. For more information, contact Veeam Customer Support.

When connecting to the vPower NFS server, the ESXi host uses a VMkernel interface. For this reason, the ESXi host must have a VMkernel interface. Otherwise, Veeam Backup & Replication will fail to mount the vPower NFS datastore on the ESXi host.

# WAN Accelerators

WAN accelerators are dedicated components that Veeam Backup & Replication uses for WAN acceleration. WAN accelerators are responsible for global data caching and data deduplication.

Technically, WAN accelerators add a new layer in the backup infrastructure — between Veeam Data Movers on the source side and the Veeam Data Mover on the target side.

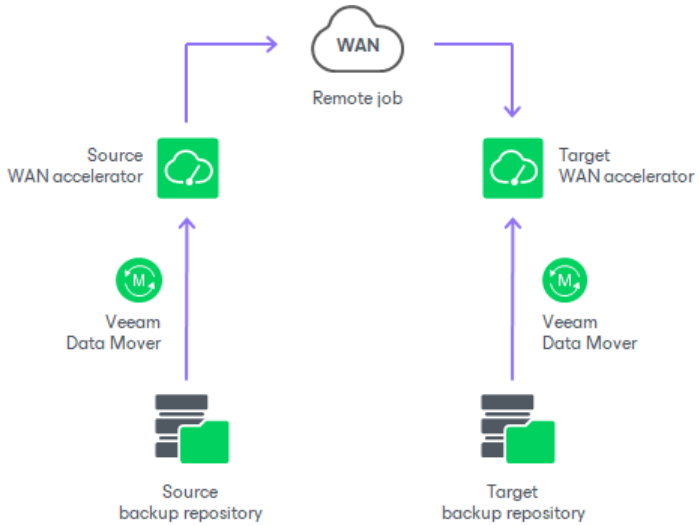


## WAN Accelerators Deployment

To enable WAN acceleration and data deduplication technologies, you must deploy a pair of WAN accelerators in your backup infrastructure.

- One WAN accelerator is deployed on the source site, closer to the source backup repository or source host.
- The other WAN accelerator is deployed on the target site, closer to the target backup repository or target host.

On each WAN accelerator Veeam Backup & Replication creates the `VeeamWAN` folder containing the following data:

- The `VeeamWAN` folder on the source WAN accelerator stores files with digests required for deduplication. For more information, see How WAN Acceleration Works.
- The `VeeamWAN` folder on the target WAN accelerator stores global cache data.

> **NOTE**
>
> Global cache is not used if both WAN accelerators in the pair (the source one and the target one) operate in the **High bandwidth mode**.

To learn how to add a WAN accelerator to the Veeam Backup & Replication infrastructure, see Adding WAN Accelerators.

# Recommendations for WAN Accelerators

You should not assign one source WAN accelerator to several jobs to remote locations that you plan to run simultaneously. The source WAN accelerator requires a lot of CPU and RAM resources, and does not process multiple tasks in parallel. As an alternative, you can create one job to the remote location for all VMs you plan to process over one source WAN accelerator.

The target WAN accelerator, however, can be assigned to several jobs to remote locations.

# WAN Global Cache

From the technical point of view, the global cache is a folder on the target WAN accelerator. By default, global cache data is stored in the `VeeamWAN` folder on the disk with the most amount of space available. However, you can define any folder of your choice when you configure the target WAN accelerator.

> **NOTE**
>
> Global cache is not used if both WAN accelerators in the pair (the source one and the target one) operate in the **High bandwidth mode**.

By default, the size of the global cache is 100 GB. You can increase the size or decrease it if necessary. The more space you allocate, the more repeating data blocks will be written to the global cache and the more efficient WAN acceleration will be. It is recommended that you allocate at least 40 GB to the global cache storage.

The global cache size is specified per source WAN accelerator. That is, if you plan to use one target WAN accelerator with several source WAN accelerators, the specified amount of space will be allocated for every source WAN accelerator that will be working with the target WAN accelerator and the size of the global cache will increase proportionally. For more information, see WAN Accelerator Sizing.

The WAN global cache is a "library" that holds data blocks repeatedly going from the source side to the target side. The global cache is populated at the first cycle of a job to the remote location. The priority is given to data blocks of Windows-based OSes, other OSes like Linux/Unix, and Microsoft Exchange Server.

Veeam Backup & Replication constantly maintains the global cache in the actual state. To do that, it continuously monitors data blocks going over WAN and data blocks in the global cache.

- If some new data block is constantly sent over WAN, it is added to the global cache.
- If some data block in the global cache is not sent over WAN and are not re-used for some period of time, it is removed from the global cache to make room for new data blocks.

Veeam Backup & Replication also performs periodic consistency checks. If some data block in the global cache gets corrupted, Veeam Backup & Replication removes it from the global cache.

The efficiency of the WAN acceleration increases with every new backup copy interval in the backup copy job. During the first backup copy interval in the backup copy job, the WAN acceleration level is minimal. Veeam Backup & Replication populates the global cache. With every new job cycle, Veeam Backup & Replication updates the global cache to include the most "popular" data blocks and the WAN acceleration efficiency increases.

> **NOTE**
>
> You can populate the global cache before you run the job to the remote location for the first time. In this case, Veeam Backup & Replication will use the global cache starting from the first session of the job to the remote location, and the WAN traffic will be minimal. For more information, see Manual Population of Global Cache.

# Many to One WAN Acceleration

The WAN global cache can be used by several source WAN accelerators simultaneously. For example, if you have several remote/branch offices, you can configure several source WAN accelerators in remote sites and one target WAN accelerator in the head office.

In this case, the global cache will hold cache data for separate source WAN accelerators. The cache data for every source WAN accelerator will be stored in a dedicated subfolder in the global cache folder.

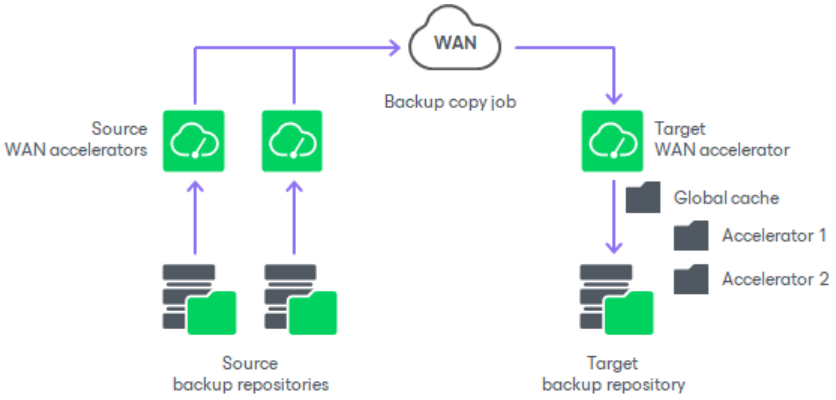


When one target WAN accelerator is used by several source WAN accelerators, Veeam Backup & Replication can copy data blocks between global cache of these WAN accelerators. This mechanism works if there are no matching backups of VMs in the target backup repository, but matching data is available in cache of other WAN accelerators.

For example, you have two backup copy jobs: *Job 1* and *Job 2*. The *Job 1* uses the source WAN accelerator *Source 1* and the target WAN accelerator *Target 3*. The *Job 2* uses the source WAN accelerator *Source 2* and the same target WAN accelerator *Target 3*. In the global cache folder, Veeam Backup & Replication will store data for 2 WAN accelerators: *Source 1* and *Source 2*.

- *Job 1* processes a VM running Microsoft Windows Server 2008 R2, and it has been running for some time. In the global cache, there is already data for this type of OS.
- *Job 2* also processes a VM running Microsoft Windows Server 2008 R2. When you start *Job 2* for the first time, there is no data for this type of OS in the global cache for *Source 2* WAN accelerator. In such situation, Veeam Backup & Replication will copy the necessary data block from the *Source 1* cache to the *Source 2* cache and will not transport this data block over WAN.

> **NOTE**
>
> Beside using global cache of other WAN accelerator, Veeam Backup & Replication also utilizes backup files residing in the backup repository. For example, if the backup repository contains a backup file created by a backup job and the backup copy job starts copying a backup of a VM of the same type, Veeam Backup & Replication will populate global cache on the WAN accelerator from the backup file not to transfer redundant data over WAN.

# Manual Population of Global Cache

You can manually pre-populate the global cache to avoid the situation where the cache remains empty. As a result, by the time a job to the remote location starts, the global cache will contain data blocks that can be used for data deduplication.

Manual population of the global cache can be helpful in the following scenarios:

- First run of a job to the remote location. When you run a first session of a job to the remote location, the global cache is empty, and the whole amount of VM data needs to be transferred over WAN. It is recommended that you manually populate the global cache before you start a job to the remote location for the first time.

- Global cache corruption. If the global cache gets corrupted for some reason, Veeam Backup & Replication needs to perform at least one session of the job to the remote location to replace corrupted data blocks with valid data blocks. In this situation, you can clean the global cache and manually populate it with valid data before the job to the remote location begins.

> **IMPORTANT**
>
> Veeam Backup & Replication does not use encrypted backups for manual global cache population.

# Limitations for Manual Population of Global Cache

The manually performed global cache population task has the following limitations:

- Veeam Backup & Replication does not use encrypted backups for global cache population.
- Veeam Backup & Replication writes only data blocks for Windows-based OSes to the default cache. Data blocks for other OSes like Linux/Unix and application data blocks are not written to the cache.
- You can start the global cache population task for the target WAN accelerator that is not currently used by any job to the remote location.
- If the global cache population task is currently running, the target WAN accelerator is locked. You cannot start any job to the remote location by using this target WAN accelerator.
- [For global cache corruption scenario] You must clean the global cache before you populate it with valid data.
- [Veeam Cloud Connect] Veeam Backup & Replication does not use tenant backups to populate global cache on the service provider side.

# How Manual Population of Global Cache Works

Global cache population is a manual operation performed by the user. When you run the global cache population task, Veeam Backup & Replication creates a 'default cache' on the target WAN accelerator. The default cache is used as a basic, universal cache for every new job to the remote location.

To populate the default cache, Veeam Backup & Replication uses backup files stored in backup repositories as a source of data.

The procedure of global cache population includes the following steps:

1. The user manually starts the global cache population tasks and selects backup repositories from which data blocks should be retrieved.
2. Veeam Backup & Replication scans backup repositories and makes up a list of OSes whose data blocks are available in backup files.
3. Veeam Backup & Replication copies data blocks from backup repositories and populates the default cache with these data blocks.

When a job to the remote location starts, Veeam Backup & Replication renames the default cache to a folder for the source WAN accelerator engaged in the job. As a result, Veeam Backup & Replication can use data blocks in this folder for deduplication starting from the very first session of the job to the remote location.

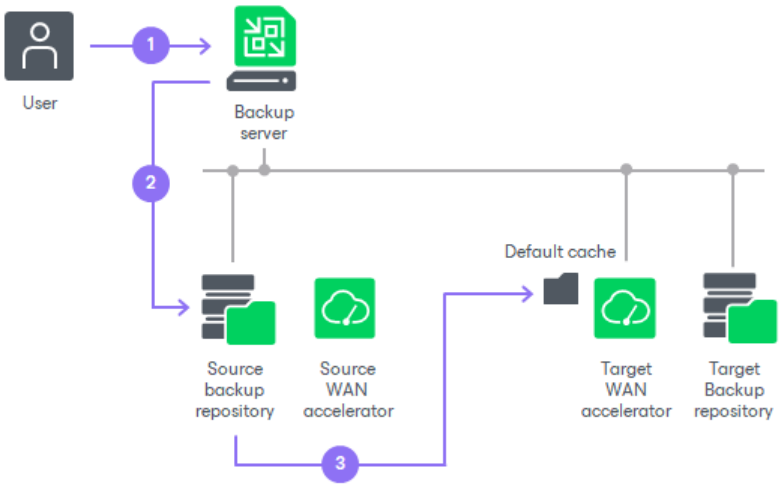


In many-to-one WAN accelerator deployment scenarios, the global cache may have folders for other source WAN accelerators, and these folders may contain data blocks for some OSes. If the global cache already contains some data, the procedure of global cache population includes the following steps:

1. The user manually starts the global cache population tasks and selects backup repositories from which data blocks should be retrieved.
2. Veeam Backup & Replication scans backup repositories and makes up a list of OSes whose data blocks are available in backup files.
3. Veeam Backup & Replication scans folders for other source WAN accelerators in the global cache and makes up a list of OSes whose data blocks are available there.
4. The list of OSes in the global cache is compared to the list of OSes in backup repositories. This way, Veeam Backup & Replication detects data blocks for which OSes are missing in the global cache.
5. In the global cache, Veeam Backup & Replication detects a folder with the maximum amount of data. This folder is used as a basis for the default cache.

6. Veeam Backup & Replication copies data blocks only for missing OSes from backup repositories and populates the default cache with these blocks. Data blocks for OSes available in folders for other source WAN accelerators are not copied to the default cache during the population task.
Veeam Backup & Replication copies these data blocks on the fly, when a job to the remote location runs. For more information, see Many to One WAN Acceleration.

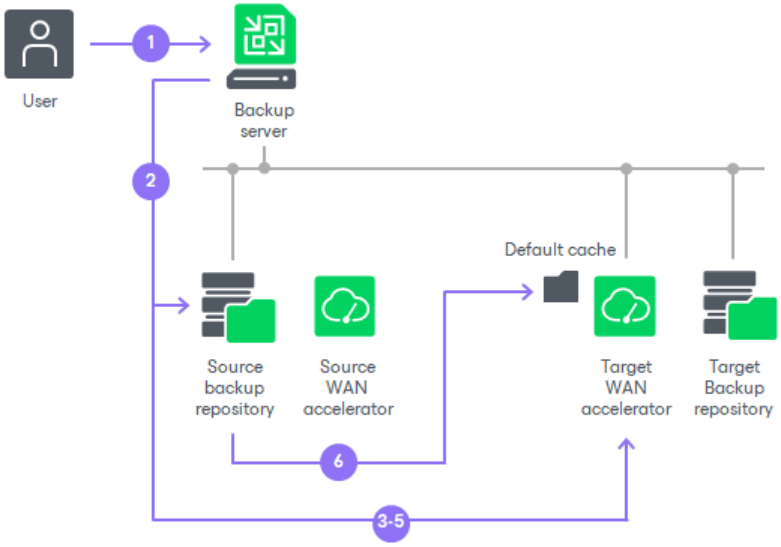


## Manually Populating Global Cache

To manually populate the global cache:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **WAN Accelerators** node.
3. In the working area, right-click the target WAN accelerator and select **Populate cache**.

   If the selected WAN accelerator is not assigned as a target WAN accelerator to any job to the remote location, Veeam Backup & Replication will display a warning.
4. In the **Source Backup Repositories** window, select backup repositories from which OS data blocks must be retrieved.

   It is strongly recommended that you select backup repositories on the same site where the target WAN accelerator is located. In the opposite case, the traffic will travel between sites, which will increase load on the network.

5. Click **OK**.

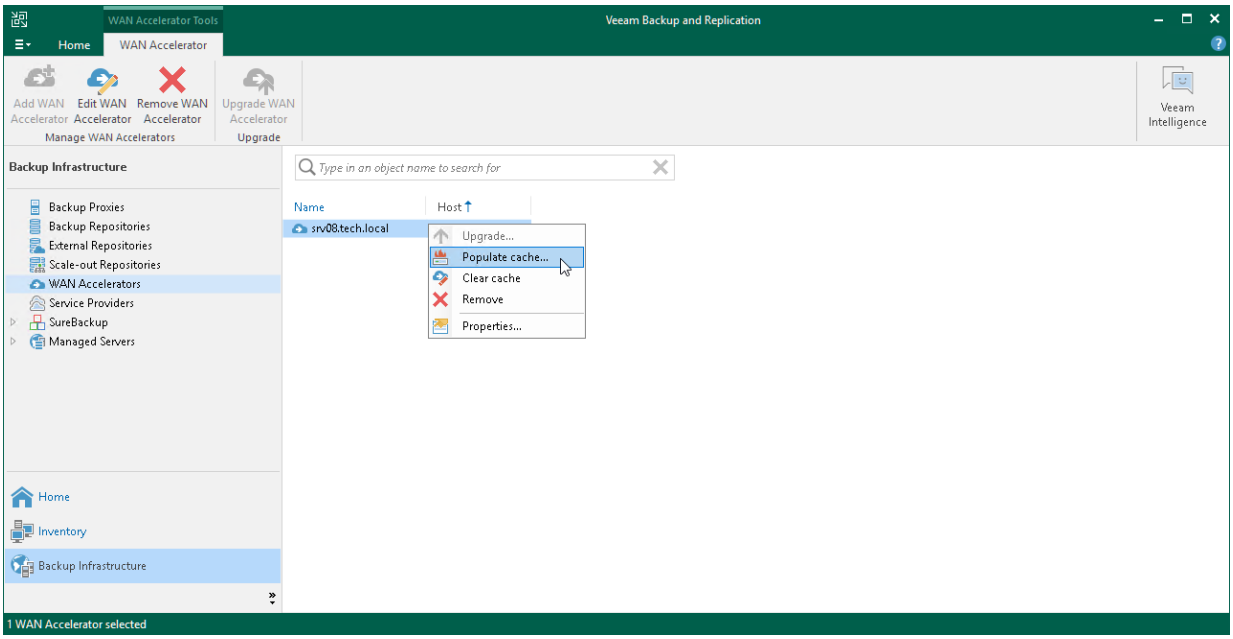


# Clearing Global Cache

You can clear the global cache on the target WAN accelerator. It is recommended that you clear the global cache in the following situations:

- Global cache is corrupted.
- Global cache contains data that is no longer needed. This situation may occur, for example, if you have decided to fully switch to the **High bandwidth mode** that does not use the global cache data.

In such cases, it is recommended that you clear the global cache and populate it anew before you start jobs to remote locations processing new types of VMs.

To clear the global cache:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **WAN Accelerators**.
3. In the working area, right-click the target WAN accelerator and select **Clear cache**.

> **IMPORTANT**
>
> Consider the following:
>
> - Before you clear the global cache, make sure that you do not have any running jobs that use this target WAN accelerator. When the global cache is cleared, Veeam Backup & Replication will restart the Veeam WAN Accelerator Service, and running jobs will complete with the Failed status.
> - When you clear the global cache, you also clear all digest data stored at this WAN accelerator.

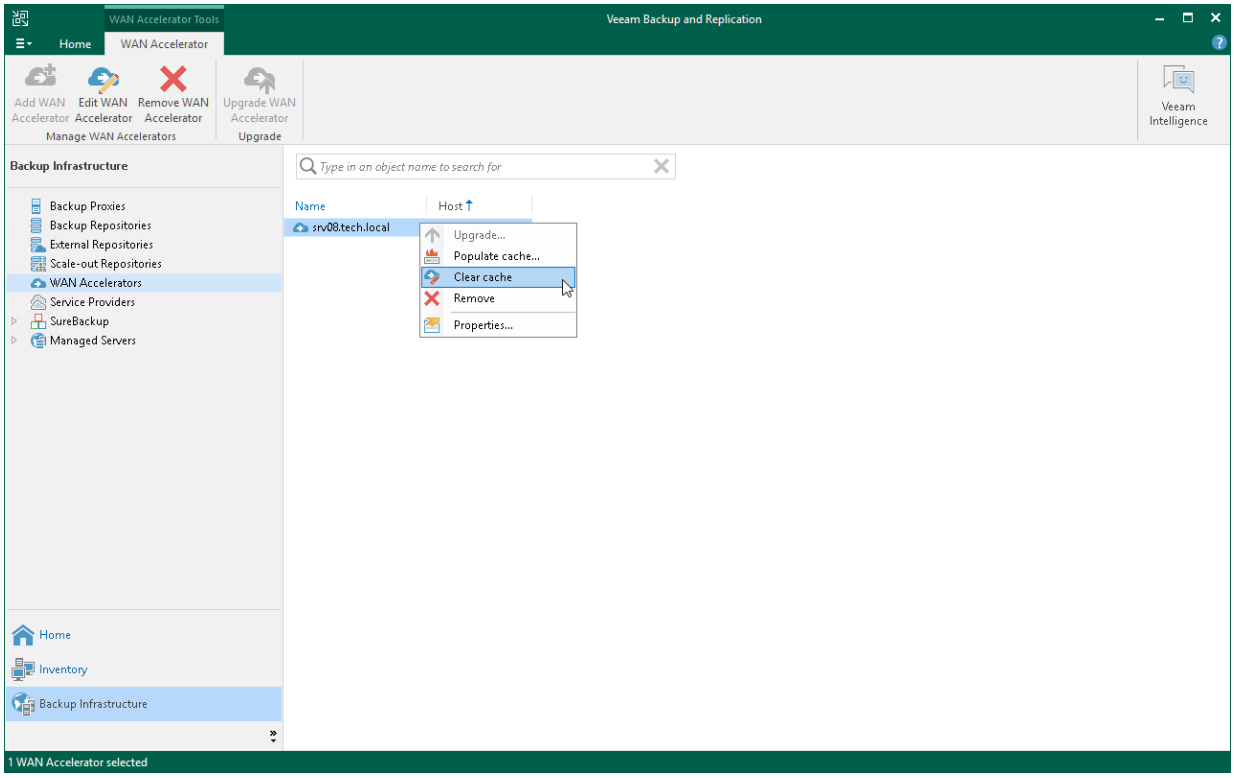
WAN Accelerator Tools
Veeam Backup and Replication
Home
WAN Accelerator
Add WAN Accelerator
Edit WAN Accelerator
Remove WAN Accelerator
Manage WAN Accelerators
Upgrade WAN Accelerator
Upgrade
Veeam Intelligence
Backup Infrastructure
Backup Proxies
Backup Repositories
External Repositories
Scale-out Repositories
WAN Accelerators
Service Providers
SureBackup
Managed Servers
Type in an object name to search for
Name
Host
srv08.tech.local
Upgrade...
Populate cache...
Clear cache
Remove
Properties...
Home
Inventory
Backup Infrastructure
1 WAN Accelerator selected

# WAN Accelerator Sizing

To ensure correct work of remote jobs over WAN accelerators, you must provide enough free space for service data on source and target WAN accelerators.

When configuring WAN accelerators, note that there can be situations when WAN acceleration switches from the **High bandwidth mode** to the **Low bandwidth mode**: for example, the link to the remote location changes and you decide to disable the **High bandwidth mode** for one of the accelerators in the pair. If you disable the **High bandwidth mode** and start a job which utilizes this WAN accelerator, Veeam Backup & Replication deletes digest data that was used in the **High bandwidth mode** and recreates it for the **Low bandwidth mode**. Besides, Veeam Backup & Replication will also use the global cache at the target WAN accelerator.

To avoid problems caused by the lack of free space when switching from the **High bandwidth mode** to the **Low bandwidth mode**, we recommend that you configure WAN accelerators as if you planned to use them in the **Low bandwidth mode**.

## Source WAN Accelerator

When you run a remote job over WAN accelerators, Veeam Backup & Replication analyses data blocks going to target and calculates digests for these data blocks. Digests data is stored on the source WAN accelerator, in the `VeeamWAN` folder on the disk that you select when you configure the WAN accelerator.

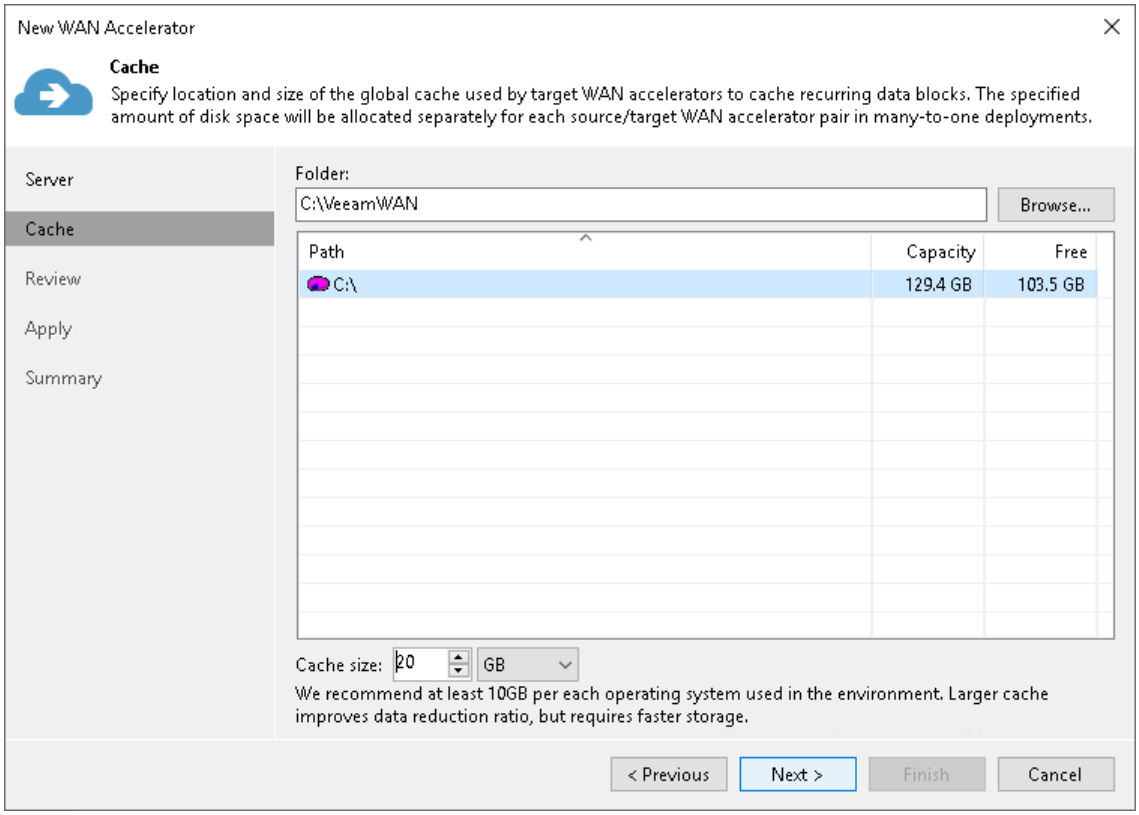


You must make sure that there is enough disk space on the source WAN accelerator to store digest data.

The amount of disk space required for a source WAN accelerator operating in the **Low bandwidth mode** is calculated by the following formula:

```
Digest Size = 2% of Provisioned VM Size
```

For example, if you plan to process 10 VMs whose provisioned size is 2 TB, you must allocate 40 GB of disk space for digest data on the source WAN accelerator operating in the **Low bandwidth mode**.

The amount of disk space required for a source WAN accelerator operating in the **High bandwidth mode** is calculated by the following formula:

```
Digest Size = 1% of Provisioned VM Size
```

For example, if you plan to process 10 VMs whose provisioned size is 2 TB, you must allocate 20 GB of disk space for digest data on the source WAN accelerator operating in the **High bandwidth mode**.

You can increase the throughput by adjusting the number of upload streams. To learn how to configure them on the source WAN accelerator, see the Choose Server step of the **New WAN Accelerator** wizard.

# Target WAN Accelerator

You must make sure that you provide enough free space for the following data on the target WAN accelerator:

- Global cache data
- Digest data

> **NOTE**
>
> For the target WAN accelerator operating in the **High bandwidth mode** only, you must provide enough free space to generate digest data. Global cache data is not used in the **High bandwidth mode**.
>
> When you enable the **High bandwidth mode** for an existing WAN accelerator, Veeam Backup & Replication does not automatically remove the global cache that was previously used for acceleration. If you are planning to use the **High bandwidth mode** for WAN acceleration and you do not need the global cache anymore, you can free the disk space by manually removing the cache. If you are planning to use WAN acceleration in the **Low bandwidth mode** in the future, we recommend that you keep the global cache. You can disable the **High bandwidth mode** and switch back to the **Low bandwidth mode** at any time.

## Global Cache Data

Global cache is stored on the target WAN accelerator, in the `VeeamWAN` folder on the disk that you select when you configure the WAN accelerator. The size of global cache is specified in the properties of the target WAN accelerator.

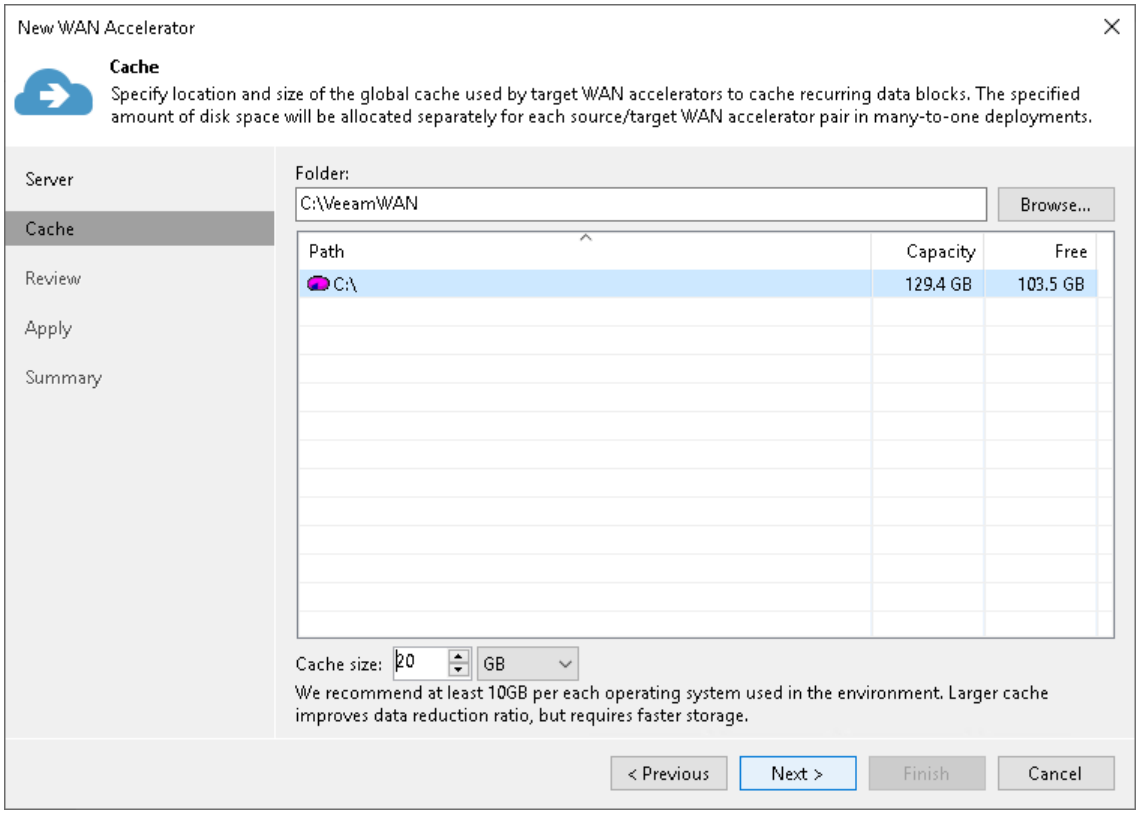


You must provide enough free space for global cache data. It is recommended that you provide 10 GB per every type of OS on VMs that you plan to process. By default, Veeam Backup & Replication allocates 100 GB for the global cache size.

For example, you want to process the following VMs:

- 1 VM that runs Microsoft Windows 7
- 3 VMs that run Microsoft Windows Server 2008 R2
- 2 VMs that run Microsoft Windows Server 2012 R2

There are 3 types of OSes so you must allocate 30 GB for the global cache on the target WAN accelerator.

> **NOTE**
>
> Global cache is stored only on the target WAN accelerator. You do not have to provide space for global cache on the source WAN accelerator.

## Digest Data

In some cases, Veeam Backup & Replication may require more space on the target WAN accelerator than specified in the WAN accelerator properties. This can happen if digest data on the source WAN accelerator is missing or cannot be used. For example:

- You have performed the **Clear Cache** operation on the source WAN accelerator and it no longer contains digest data. For more information, see Clearing Global Cache.

- Veeam Backup & Replication has attempted to resume operation of backup data transfer, but the backup file was not prepared for the operation in a proper way. The digest data must be recalculated.

In such situations, the target WAN accelerator calculates digest data on its own, which requires additional space. After the digest data is calculated, the target WAN accelerator transfers it to the source WAN accelerator. After the transfer, the copy of the digest data is removed from the target WAN accelerator.

For safety reasons, it is recommended that you provide the following amount of space for digest data on the target WAN accelerator:

The amount of disk space required for digest data at a target WAN accelerator operating in the **Low bandwidth mode** is calculated by the following formula:

```
Digest Size = 2% of Provisioned VM Size
```

The amount of disk space required for digest data at a target WAN accelerator operating in the **High bandwidth mode** is calculated by the following formula:

```
Digest Size = 1% of Provisioned VM Size
```

This amount of space is required for digest data recalculation. If you do not provide this amount of space and a situation where Veeam Backup & Replication needs to recalculate digest data occurs, the job to the remote location will work in the limited mode. Veeam Backup & Replication will not deduplicate data against the previous restore points copied to target. For more information, see Global Data Deduplication.

> **IMPORTANT**
>
> When you specify the global cache size for a target WAN accelerator, you do not allocate any space for storing digest data. To let Veeam Backup & Replication recalculate digest data, you must make sure that necessary amount of free space is available on the target WAN accelerator (in addition to the space allocated for the global cache).

For example:

- You have allocated 100 GB for global cache on the target WAN accelerator operating in the **Low bandwidth mode**.
- Provisioned size of VMs to be processed is 2 TB.

In this case, the needed amount of free disk space for the global cache on the target WAN accelerator is:

```
100 GB + 40 GB = 140 GB
```

# Many-to-One WAN Acceleration Scenario

Global cache size is calculated per 1 source WAN accelerator working with the target WAN accelerator. If you plan to use several source WAN accelerators with 1 target WAN accelerator, you must increase the size of the global cache proportionally. The cache data for every source WAN accelerator will be stored in a dedicated subfolder in the global cache folder of the target WAN accelerator. The global cache size is calculated by the following formula:

```
Total Global Cache Size = (# of Source WAN Accelerators) * (Size of Global Cach
e Configured in Target WAN Accelerator Properties)
Total Free Disk Space to Provide = Total Global Cache Size + Digest Size
```

For example:

- You have 4 source WAN accelerators in the source side working with 1 target WAN accelerator in the disaster recovery (DR) site.
- The global cache size configured in properties of the target WAN accelerator is 100 GB.
- The size of VMs to be processed is 2 TB.

In this case, the needed amount of free disk space for the global cache and digests on the target WAN accelerator is:

```
4*100 GB + 40 GB = 440 GB
```

**NOTE**

For more information and recommendations on WAN accelerator cache sizing, see this Veeam KB article.

# Adding WAN Accelerators

To add a WAN accelerator, you must assign the WAN accelerator role to a Microsoft Windows server added to the backup infrastructure.

You must deploy a pair of WAN accelerators: one WAN accelerator on each side of the WAN link.

Before you add a WAN accelerator, check prerequisites. Then use the **New WAN Accelerator** wizard to add a WAN accelerator.

## Before You Begin

Before you add a WAN accelerator, check the following prerequisites:

- The machine that will operate as a WAN accelerator must meet the system requirements. For more information, see System Requirements.
- You must assign the WAN accelerator role to a Microsoft Windows server (physical or virtual). The WAN accelerator role can be assigned to backup proxies and Microsoft Windows backup repositories already configured in the backup infrastructure.
- You must use 64-bit Microsoft Windows machines as WAN accelerators. Veeam Backup & Replication does not support 32-bit versions of Microsoft Windows used as WAN accelerators.
- WAN acceleration operations are resource-consuming. When assigning the WAN accelerator role, consider available CPU and memory resources of the Microsoft Windows server. It is recommended that you assign the WAN accelerator role to servers with 8 GB RAM and more.
- The machine must have enough free disk space to store digests or global cache data. For more information, see WAN Accelerator Sizing.
- You must add the machine to the Veeam Backup & Replication console as a managed server before adding it as a WAN accelerator.

# Step 1. Launch New WAN Accelerator Wizard

To launch the **New WAN Accelerator** wizard, do either of the following:

- Open the **Backup Infrastructure** view, in the inventory pane select **WAN Accelerators** and click **Add WAN Accelerator** on the ribbon.
- Open the **Backup Infrastructure** view, in the inventory pane right-click **WAN Accelerators** and select **Add WAN Accelerator**.

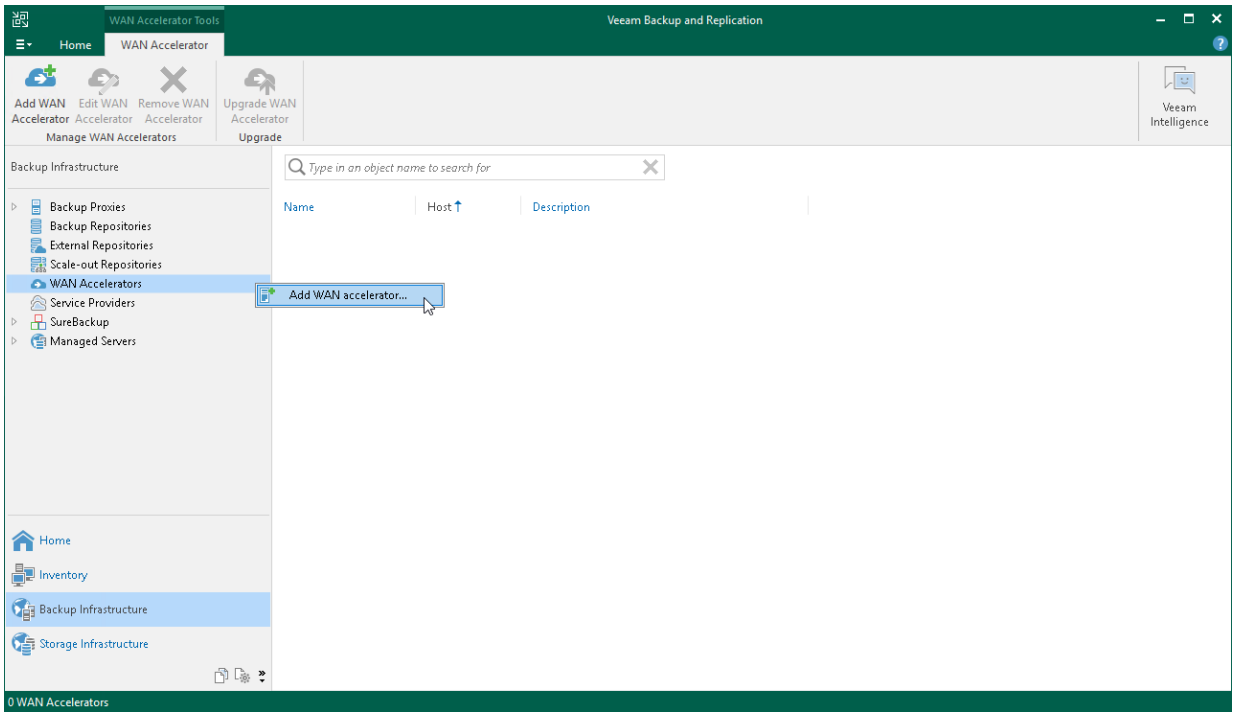

# Step 2. Choose Server

At the **Server** step of the wizard, select a Microsoft Windows server that you plan to use as a WAN accelerator and define port and connection settings for this server.

1. From the **Choose server** list, select a Microsoft Windows server added to the backup infrastructure. If the server is not added to the backup infrastructure yet, you can click **Add New** to open the **New Windows Server** wizard. For more information, see Adding Microsoft Windows Servers.
2. In the **Description** field, provide a description for future reference.

   It is recommended that you describe the added WAN accelerator as the source or target one. When you create a job to the remote location, this hint will be displayed in brackets next to the WAN accelerator name, which will help you choose the necessary WAN accelerator to be used on the source or target side.
3. In the **Traffic port** field, specify the number of the port over which WAN accelerators must communicate with each other. By default, port 6165 is used.
4. In the **Streams** field, specify the number of connections that must be used to transmit data between WAN accelerators. By default, 5 connections are used.

   This setting applies only to the source WAN accelerator. The greater is the number of streams, the more bandwidth resources Veeam Backup & Replication will use. A great number of streams engage more CPU and memory resources of the source WAN accelerator.

   If the link has low latency and high bandwidth, the default setting (5 streams) may be enough to fully saturate it. If the link is still not saturated, the number of streams may be increased. Tests show that with high latency links, link speed x1.5 is a good best practice for estimating the number of streams required. The following example shows benchmark on a 10 Mbit/s WAN link with 100 milliseconds of latency.

| Link (Mbit/s) | Latency (ms) | Packet Loss | Streams | Throughput (Mbps) |
|---|---|---|---|---|
| 10 | 100 | 0 | 3 | 3.5 |
| 10 | 100 | 0 | 10 | 7.5 |
| **10** | **100** | **0** | **15** | **10** |
| 10 | 100 | 0 | 20 | 10 |

   Increasing the number of streams to more than required for fully saturating the link will slow down data transfers, as the data transfer will wait for all streams to initialize and stabilize before starting to transfer any data.

> **TIP**
>
> To test different scenarios in the lab before deploying WAN acceleration, you can use a WAN emulator (such as WANem).

5. If your network bandwidth is more than 100 Mbps, we recommend that you select the **High bandwidth mode** check box. The High bandwidth mode provides significant bandwidth savings comparable to the direct mode on WAN links under 1 Gbps.

   To use the High bandwidth mode, enable the option for WAN accelerators at both sites of the data transfer: the source one and the target one. If the High bandwidth mode is enabled for the target WAN accelerator, different source accelerators can parallelly interact with it in different modes, depending on the mode selected for each source WAN accelerator.

   When you enable the High bandwidth mode for an existing pair of WAN accelerators, Veeam Backup & Replication does not automatically remove the global cache that was previously used for acceleration. If you are planning to use only the High bandwidth mode for WAN acceleration and you do not need the global cache anymore, you can free the disk space by manually removing the cache. If you are planning to use WAN acceleration in the Low bandwidth mode in the future, we recommend that you keep the global cache. You can disable the High bandwidth mode at any time.

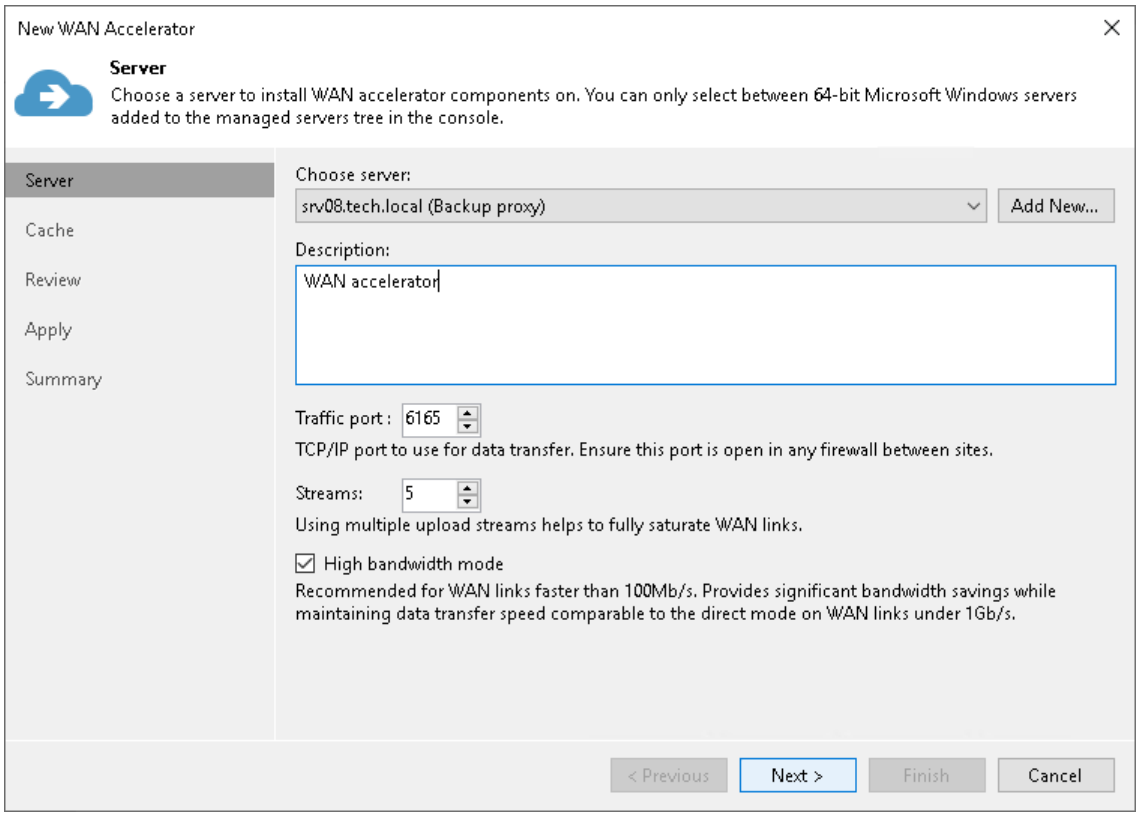

# Step 3. Define Cache Location and Size

At the **Cache** step of the wizard, define settings for the folder where service files and global cache data will be stored.

> **NOTE**
>
> If both WAN accelerators (the source one and the target one) work in the High bandwidth mode, WAN acceleration does not use the global cache. But note that you can disable the High bandwidth mode and switch back to the Low bandwidth mode at any time, which will require to use the global cache. To provide correct operation of WAN accelerators, we recommend that you allocate enough disk space for the global cache folder when adding a new WAN accelerator, no matter if the High bandwidth mode is enabled for it or not.

1. In the **Folder** field, specify a path to the folder in which service files (for source and target WAN accelerators) and global cache data (for target WAN accelerator) must be stored. When selecting a folder on the target WAN accelerator, make sure that there is enough space for storing global cache data.
2. [For target WAN accelerator] In the **Cache size** field, specify the size for the global cache. The global cache size is specified per source WAN accelerator. If you plan to use one target WAN accelerator with several source WAN accelerators, the specified amount of space will be allocated to every source WAN accelerator and the size of the global cache will increase proportionally. For more information, see WAN Accelerator Sizing.

> **IMPORTANT**
>
> Do not nest the global cache folder deep in the file tree. During WAN acceleration operations, Veeam Backup & Replication generates service files with long file names. Placing such files to a folder of significant depth may cause problems on the NTFS file system.

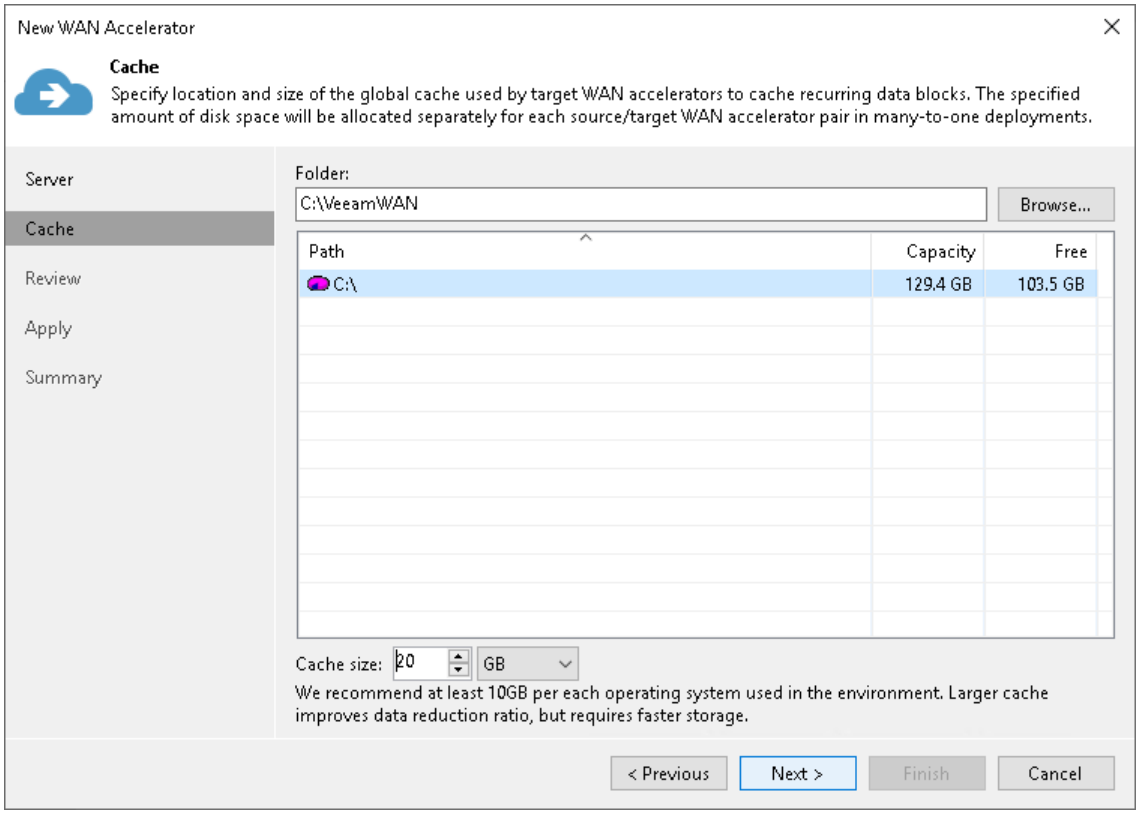

# Step 4. Review Components

Veeam Backup & Replication installs the following components on the server used as a WAN accelerator:

- Veeam Data Mover
- Veeam WAN Accelerator Service

At the **Review** step of the wizard, review what components are already installed on the server and what components will be installed.

1. Review the components.
2. Click **Next** to install the components on the server.

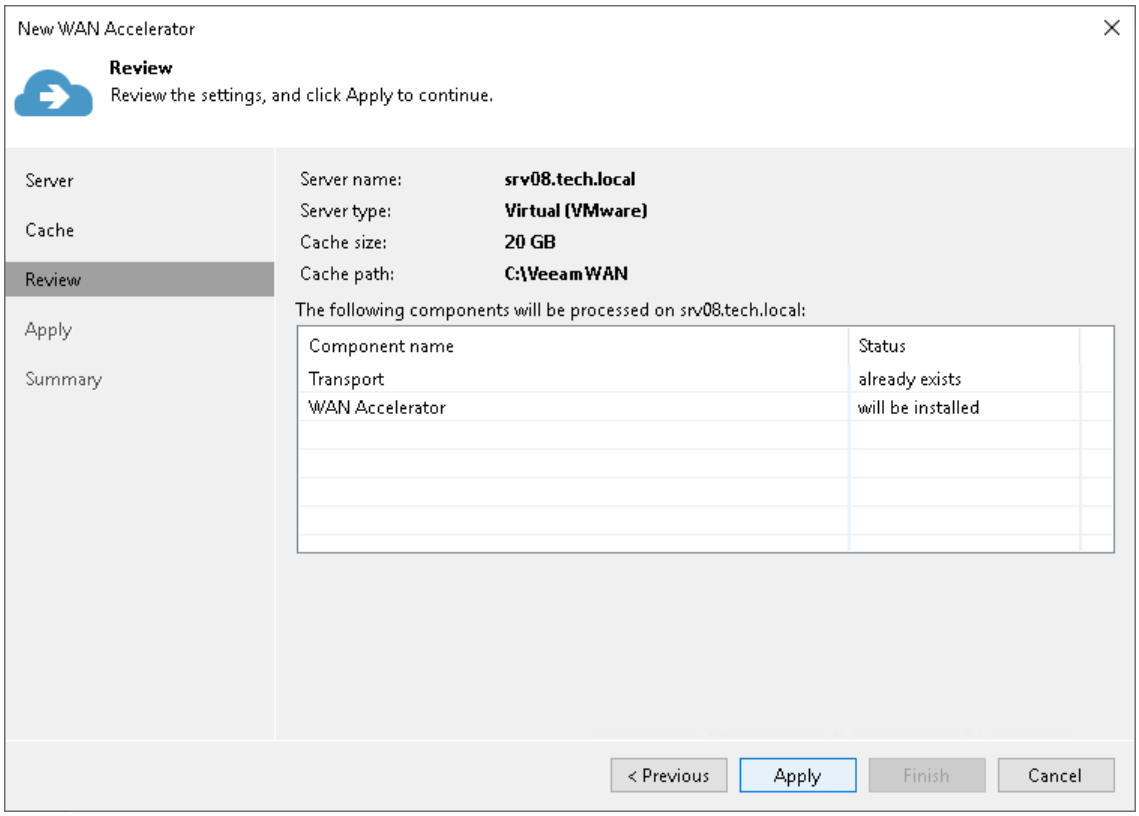

# Step 5. Apply WAN Accelerator Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components. Click **Next** to complete the procedure of creating the WAN accelerator and adding it to the backup infrastructure.

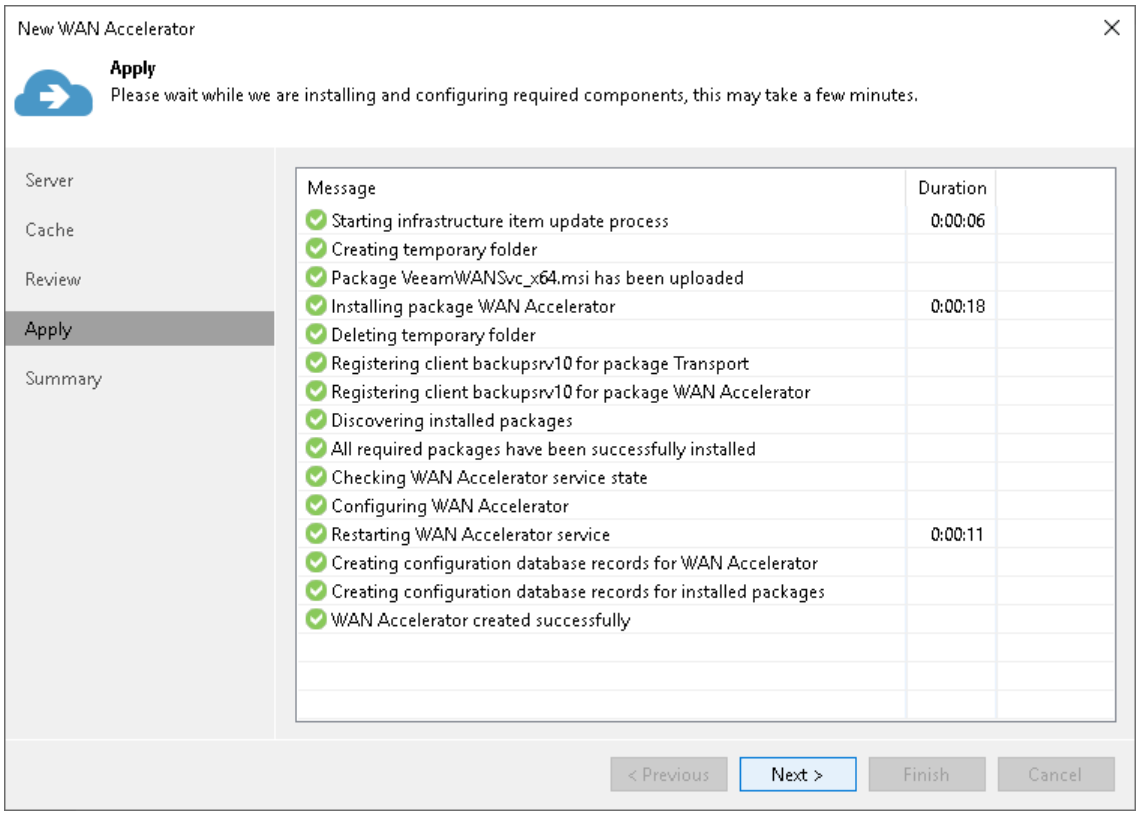

# Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added WAN accelerator and click **Finish** to exit the wizard.

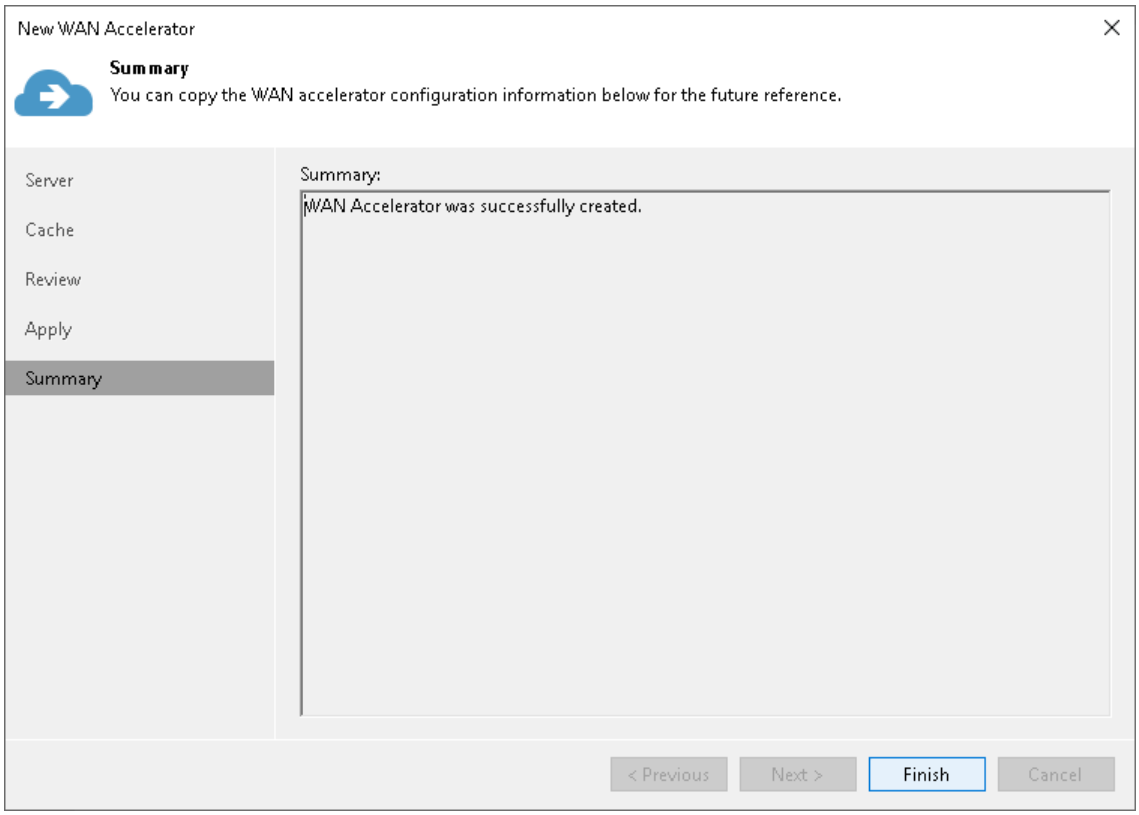

# Removing WAN Accelerators

You can permanently remove a WAN accelerator from the backup infrastructure. When you remove a WAN accelerator, Veeam Backup & Replication unassigns the WAN accelerator role from the server, and this server is no longer used as a WAN accelerator. The server itself remains in the backup infrastructure.

To remove a WAN accelerator:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **WAN accelerators**.
3. In the working area, select the WAN accelerator and click **Remove WAN Accelerator** on the ribbon or right-click the WAN accelerator and select **Remove**.

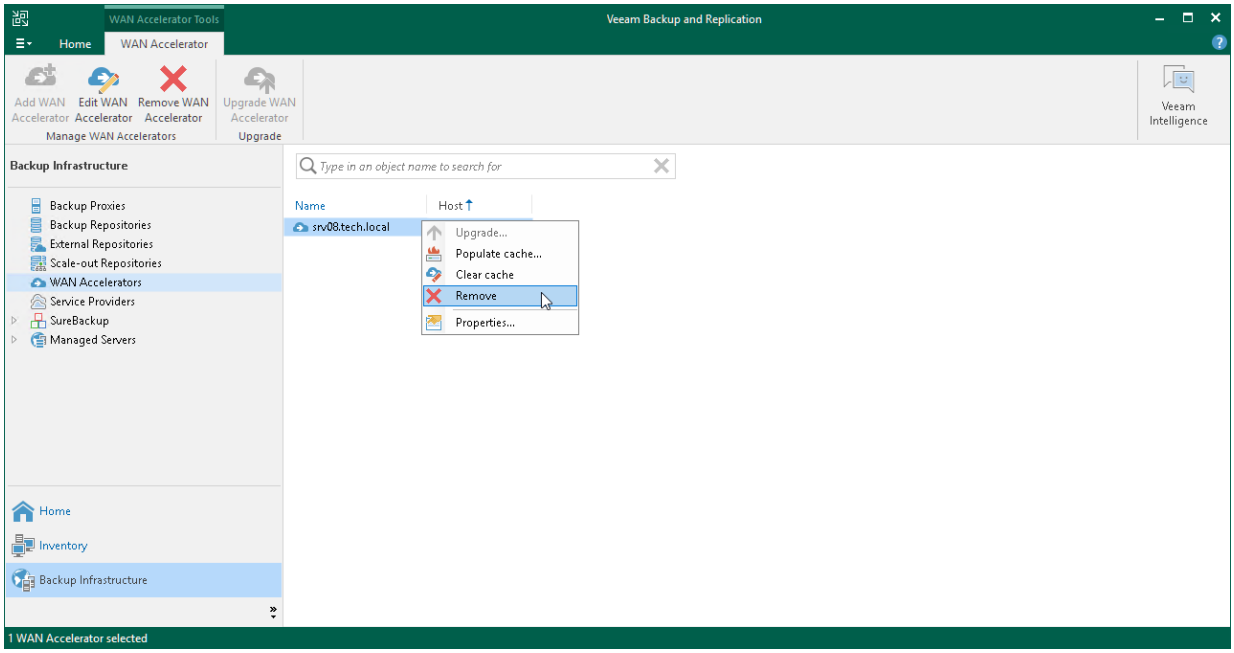

# WAN Acceleration

WAN acceleration is a Veeam technology that optimizes data transfer to remote locations. It is specific for off-site backup copy jobs and replication jobs.

To enable WAN acceleration and data deduplication technologies, you must deploy a pair of WAN accelerators in your backup infrastructure. For more information, see WAN Accelerators.

> **NOTE**
>
> WAN acceleration is included in the Veeam Universal License. When using a legacy socket-based license, the Enterprise or Enterprise Plus editions of Veeam Backup & Replication are required.

Off-site backup copy and replication always involve moving large volumes of data between remote sites. The most common problems that backup administrators encounter during off-site backup and replication are insufficient network bandwidth to support VM data traffic and transmission of redundant data. To solve these problems, Veeam Backup & Replication offers the WAN acceleration technology that combines:

- Network traffic compression
- Multistreaming upload
- Global data deduplication
- Variable block size deduplication

These technologies help optimize the data transfer and reduce the amount of data going over WAN.

## High Bandwidth Mode

We recommend using the **High bandwidth mode** for WAN connections faster than 100 Mbps. If compared with the **Low bandwidth mode**, it does not leverage the global cache, but utilizes a faster compression method, optimized digests processing and an alternative deduplication mechanism. As a result, the new mode provides better performance and higher speed of data transfer.

Note that to use the **High bandwidth mode**, you must enable this option for WAN accelerators at both sites of the data transfer: the source and the target. If the target WAN accelerator has the **High bandwidth mode** enabled, different source accelerators can parallelly interact with it in different modes, depending on the mode selected for each source WAN accelerator.

## How WAN Acceleration Works

When you create a job to the remote location, you can select to use WAN acceleration in its properties.

The procedure of data copying with WAN acceleration enabled is performed in the following way:

1. [For backup copy job] Veeam Backup & Replication decompresses the backup file to analyze its content.
2. The Veeam WAN Accelerator Service on the source WAN accelerator analyzes data blocks of the file to be transferred and creates a file with digests for these data blocks. The created file with digests is stored to the `VeeamWAN` folder on the source WAN accelerator.
3. Veeam Backup & Replication compresses the backup file data (for backup copy) or VM data (for replication) and copies it to the target side.

   At this point, Veeam Backup & Replication can perform deduplication within the VM itself — that is, deduplicate identical data blocks in every VM disk.

4. During the data transfer process, the Veeam WAN Accelerator Service on the target WAN accelerator populates the global cache storage with data blocks from the copied file.
5. During the next job cycle, the Veeam WAN Accelerator Service on the source WAN accelerator analyzes data blocks in the file that must be transferred this time and creates digests for these data blocks.
6. The Veeam WAN Accelerator Service compares the created digests with the digests that have been previously stored to the `VeeamWAN` folder on the source WAN accelerator. If duplicate data blocks are found, the actual data block in the backup file is not copied over WAN. Instead, it is taken from the global cache and written to the restore point in the backup copy folder or on the target data volume.
7. Additionally, Veeam Backup & Replication analyzes restore points that have been previously copied to the target side. If duplicates are found, Veeam Backup & Replication does not copy such blocks over WAN but takes them from the previously copied restore points.

As a result, Veeam Backup & Replication copies only new data blocks to the target side and uses data blocks that are already stored in the global cache or in restore points in the target backup repository.

If the target WAN accelerator is used by several jobs, the target backup repository may already contain data blocks of the necessary VM type. In this situation, Veeam Backup & Replication will copy the required data blocks to the global cache before the copying process starts and use these data blocks further on. For more information, see Many to One WAN Acceleration.

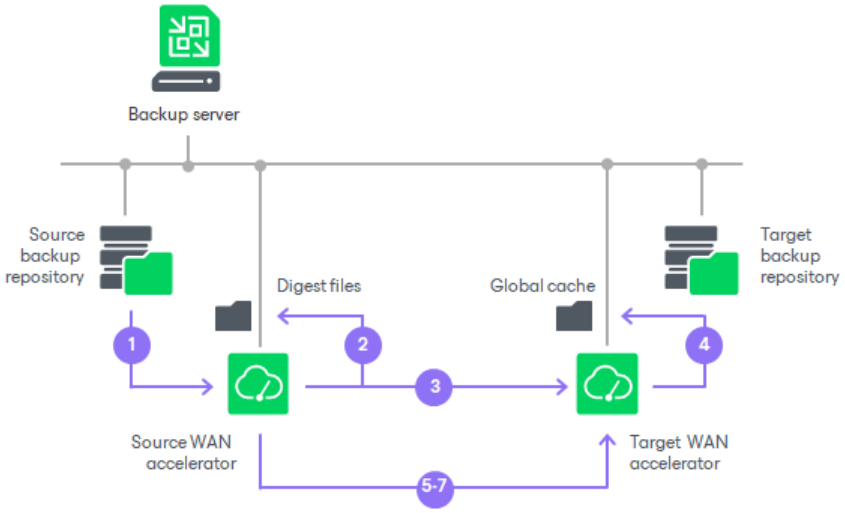


If WAN acceleration is performed in the **High bandwidth mode**, the procedure of data transfer with WAN acceleration has the following peculiarities:

- Global cache is not used. Thus, the target WAN accelerator does not need extra disk space to store the global cache folder.
- Deduplication is performed only by using previous restore points for the processed VM on the target repository. Therefore, Veeam Backup & Replication performs less deduplication operations and saves resources and time for data processing.
- The data chunk size used by the Changed Block Tracking mechanism during deduplication is smaller if compared with the **Low bandwidth mode**. This reduces the size of redundant data to transfer.

# Global Data Deduplication

The goal of WAN acceleration is to send less data over the network. To reduce the amount of data going over WAN, Veeam Backup & Replication uses the global data deduplication mechanism.

1. When you first run a job to the remote location, Veeam Backup & Replication analyzes data blocks going over WAN.

2. With every new cycle of a job to the remote location, Veeam Backup & Replication uses the data redundancy algorithm to find duplicate data blocks in copied files. Veeam Backup & Replication analyzes data blocks in files on the source side and compares them with those that have been previously transferred over WAN. If an identical data block is found, Veeam Backup & Replication deduplicates it.

As a result, only unique data blocks go over WAN. Data blocks that have already been sent are not sent. This way, Veeam Backup & Replication eliminates transfer of redundant data over WAN.

Veeam Backup & Replication uses three sources for data deduplication:

- VM disks. Veeam Backup & Replication analyses data blocks within the same VM disk. If identical blocks are found, duplicates are eliminated.
  For example, in case of a virtualized Microsoft Exchange server, the same email is typically stored in sender's Outbox folder and recipient's Inbox folder, which results in duplicate data blocks. When a job to the remote location runs, Veeam Backup & Replication detects such VM data blocks and performs deduplication.
- Previous restore points for the processed VM on the target repository. Veeam Backup & Replication analyses data in the restore point that is about to be copied and the restore points that are already stored on the target side. If an identical block is found on the target side, Veeam Backup & Replication eliminates the redundant data block in the copied restore point.
- Global cache. Veeam Backup & Replication creates a global cache holding data blocks that repeatedly go over WAN. In a new job session, Veeam Backup & Replication analyzes data blocks to be sent and compares them with data blocks stored in the global cache. If an identical data block is already available in the global cache, its duplicate on the source side is eliminated and not sent over WAN.

**NOTE**

Consider the following:

- Veeam Backup & Replication deduplicates data blocks within one VM disk and in restore points for one VM only. Deduplication between VM disks and restore points of different VMs is performed indirectly, using the global cache. For more information, see WAN Global Cache.
- Global data deduplication and deduplication within the same VM disk are not used if both WAN accelerators in the pair (the source one and the target one) operate in the **High bandwidth mode**.

# Data Block Verification

During the VM copy process, Veeam Backup & Replication performs a CRC check for the VM traffic going between the source and target WAN accelerators. The CRC check helps ensure that the correct VM data goes to the target side and no corrupted data blocks are written to the global cache or to backup files in the target backup repository.

The check is performed in the following way:

1. Before sending a data block to the target side, Veeam Backup & Replication calculates a checksum for the copied data block.
2. Once the data block is copied over WAN and before it is written to the global cache or to the target backup repository, Veeam Backup & Replication recalculates the checksum for this data block on the target side.
3. The source and target checksums are compared. If the checksums do not coincide, the target WAN accelerator sends a request to the source WAN accelerator for the correct data block. The source WAN accelerator re-sends the necessary data blocks to the target WAN accelerator as is and the re-sent data block is written to the global cache or to the backup file in the target backup repository on the fly.

# Data Transport on WAN Disconnect

If you replicate VMs over WAN accelerators, and a WAN connection drops for short periods of time (less than 30 minutes), Veeam Backup & Replication transparently handles disconnect issues. It automatically resumes the data transport process from the point when the connection was lost. The resume on disconnect capability improves the reliability of off-site replication, reduces the backup window and minimizes the load on the WAN link.

If a WAN connection is lost for more than 30 minutes, Veeam Backup & Replication still does not finish the job with a failed status. After a WAN connection is resumed, Veeam Backup & Replication starts a new data transfer cycle. Data transported with every new transport cycle is written to a new working snapshot of a VM replica. As the WAN connection may drop several times, Veeam Backup & Replication can create a number of working snapshots.

Not to keep long snapshot chains, Veeam Backup & Replication merges earlier snapshots and maintains only two working snapshots for the VM replica. When all VM data is transferred to the target host, the two working snapshots are also merged to create one fully functional VM restore point.

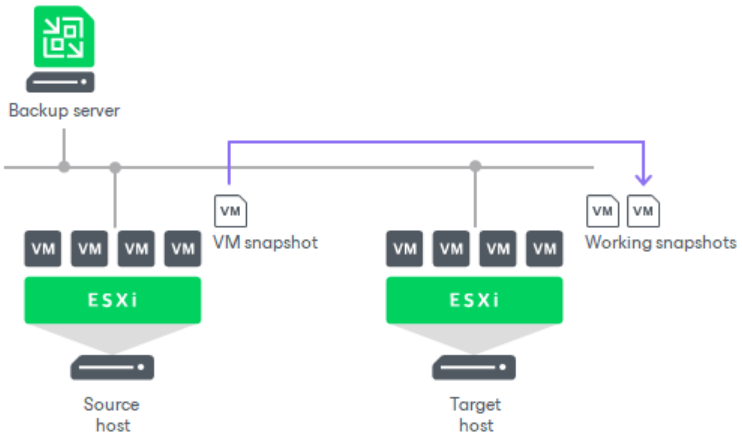


If the WAN link is weak and drops constantly, Veeam Backup & Replication may fail to transport VM data by the time a new replication job session starts. In this case, during a new replication job session Veeam Backup & Replication attempts to transfer VM data that have changed since the last replication job session and VM data that were not transferred during the previous replication job session.

# Log Shipping Servers

Log shipping servers are dedicated components that Veeam Backup & Replication uses for backup of Microsoft SQL Server transaction logs, PostgreSQL WAL files and Oracle archive logs. For more information, see Microsoft SQL Server Log Backup, PostgreSQL WAL Files Backup and Oracle Log Backup.

# Tape Servers

Tape servers are dedicated components responsible for transferring data between data source and tape device. For more information, see the Veeam Backup & Replication User Guide.

# NDMP Servers

If your NAS device supports the NDMP protocol, you can back up data from it to tape. To do this, you need to add the NAS device as an NDMP server. For more information, see the Veeam Backup & Replication User Guide.

# Veeam Backup Enterprise Manager

Veeam Backup Enterprise Manager is an optional component intended for distributed enterprise environments with multiple backup servers. Veeam Backup Enterprise Manager federates backup servers and offers a consolidated view of these servers through a web browser interface. You can centrally control and manage all jobs through a single "pane of glass", edit and clone jobs, monitor job state and get reporting data across all backup servers. Veeam Backup Enterprise Manager also enables you to search for VM guest OS files in all current and archived backups across your backup infrastructure, and restore these files in one click.

## Veeam Backup Enterprise Manager Deployment

Veeam Backup Enterprise Manager can be installed on a physical or virtual machine. You can deploy it on the backup server or use a dedicated machine.

## Veeam Backup Enterprise Manager Services and Components

Veeam Backup Enterprise Manager uses the following services and components:

- **Veeam Backup Enterprise ManagerService** coordinates all operations of Veeam Backup Enterprise Manager, aggregates data from multiple backup servers and provides control over these servers.
- **Veeam Backup Enterprise Manager Configuration Database** is used by Veeam Backup Enterprise Manager for storing data. The database instance can be located on PostgreSQL or Microsoft SQL Server installed either locally (on the same machine as Veeam Backup Enterprise Manager Server) or remotely.
- **Veeam Guest Catalog Service** replicates and consolidates VM guest OS file system indexing data from backup servers added to Veeam Backup Enterprise Manager. Index data is stored in Veeam Backup Enterprise Manager Catalog (a folder on the Veeam Backup Enterprise Manager Server) and is used to search for VM guest OS files in backups created by Veeam Backup & Replication.

# Backup

Veeam Backup & Replication produces image-level backups of VMs. It treats VMs as objects, not as a set of files. When you back up VMs, Veeam Backup & Replication copies a VM image as a whole at a block level. Image-level backups can be used for different types of restore, including Instant Recovery, entire VM restore, VM file recovery, file-level recovery, and so on.

The backup technology is typically used for VMs with lower RTOs. When the primary VM fails, you need some time to restore VM data from a compressed and deduplicated backup file.

# About Backup

Veeam Backup & Replication is built for virtual environments. It operates at the virtualization layer and uses an image-based approach for VM backup.

Veeam Backup & Replication does not install agent software inside the VM guest OS to retrieve VM data. To back up VMs, it leverages VMware vSphere snapshot capabilities. When you back up a VM, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. You can think of the VM snapshot as a point-in-time copy of a VM that includes virtual disks, system state, configuration, and so on. Veeam Backup & Replication uses this point-in-time copy as a data source for backup.

Veeam Backup & Replication copies VM data from the source datastore at a block level. It retrieves VM data, compresses and deduplicates it, and stores in backup files in the backup repository in Veeam proprietary format.

In Veeam Backup & Replication, backup is a job-driven process. To perform backup, you need to configure backup jobs. A backup job is a configuration unit of the backup activity. The backup job defines when, what, how and where to back up. One backup job can be used to process one or several VMs. You can instruct Veeam Backup & Replication to run jobs automatically by schedule or start them manually.

The first backup job session always produces a full backup of the VM image. Subsequent backup job sessions are incremental — Veeam Backup & Replication copies only those data blocks that have changed since the last backup job session. To keep track of changed data blocks, Veeam Backup & Replication uses different approaches. For more information, see Changed Block Tracking.

# How Backup Works

Veeam Backup & Replication performs VM backup in the following way:

1. When a new backup job session starts, Veeam Backup & Replication starts the Veeam Backup Manager process on the backup server. Veeam Backup Manager reads job settings from the configuration database and creates a list of VM tasks to process. For every disk of VMs added to the job, Veeam Backup & Replication creates a new task.
2. Veeam Backup Manager connects to the Veeam Backup Service. The Veeam Backup Service includes a resource scheduling component that manages all tasks and resources in the backup infrastructure. The resource scheduler checks what backup infrastructure resources are available and assigns backup proxies and backup repositories to process job tasks.
3. Veeam Backup Manager establishes a connection with Veeam Data Movers on the target backup repository and backup proxy and sets a number of rules for data transfer, such as network traffic throttling rules, and so on.
4. Veeam Data Movers on the backup proxy and backup repository establish a connection with each other for data transfer.
5. Veeam Backup Manager queries information about VMs and virtualization hosts from the Veeam Broker Service.
6. If application-aware image processing is enabled for the job, Veeam Backup & Replication connects to VM guest OSes, deploys non-persistent runtime components or, if necessary, persistent agent components on VM guest OSes and performs in-guest processing tasks.
7. Veeam Backup & Replication requests vCenter Server or ESXi host to create a VM snapshot. VM disks are put into the read-only state, and every virtual disk receives a delta file. All changes the user makes to the VM during backup are written to delta files.
8. The source Veeam Data Mover reads the VM data from the read-only VM disk and transfers the data to the backup repository in one of the transport modes. During incremental job sessions, the source Veeam Data Mover uses CBT to retrieve only those data blocks that have changed since the previous job session. If CBT is not available, the source Veeam Data Mover interacts with the target Veeam Data Mover on the backup repository to obtain backup metadata and uses this metadata to detect blocks that have changed since the previous job session.

   While transporting VM data, the source Veeam Data Mover performs additional processing. It filters out zero data blocks, blocks of swap files and blocks of excluded VM guest OS files. The source Veeam Data Mover compresses VM data and transports it to the target Veeam Data Mover.

9. After the backup proxy finishes reading VM data, Veeam Backup & Replication requests the vCenter Server or ESXi host to commit the VM snapshot.

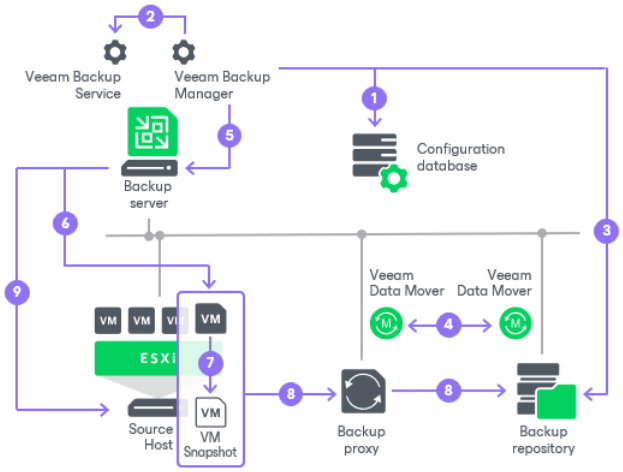

# Backup Infrastructure for Backup

Veeam Backup & Replication uses the following components for the backup process:

- One or more source hosts with associated datastores
- One or more backup proxies
- Backup repository
- [Optional] One or more guest interaction proxies
- [For shared folder backup repository] Gateway server

All backup infrastructure components engaged in the job make up a data pipe. The source host and backup repository produce two terminal points for the data flow. Veeam Backup & Replication processes VM data in multiple cycles, moving VM data over the data pipe block by block.

Veeam Backup & Replication collects VM data, transforms and transports it to the target with the help of Veeam Data Movers. Veeam Backup & Replication uses two-service architecture — one Veeam Data Mover controls interaction with the source host and the other controls interaction with the backup repository. Veeam Data Movers communicate with each other and maintain a stable connection.

When a new backup session starts, Veeam Backup & Replication performs the following actions:

1. Veeam Backup & Replication deploys non-persistent runtime components or, if necessary, persistent agent components on VM guest OSes using the guest interaction proxy (for Microsoft Windows VMs) or backup server (for VMs with other OSes).
2. The target-side Veeam Data Mover obtains job instructions and communicates with the source-side Veeam Data Mover to begin data collection.
3. The source-side Veeam Data Mover copies VM data from the source storage in one of the transport modes. During incremental job runs, the source-side Veeam Data Mover retrieves only those data blocks that have changed since the previous job session.

   While copying, the source-side Veeam Data Mover performs additional data processing. It filters out zero data blocks, blocks of swap files and blocks of excluded VM guest OS files; it compresses and deduplicates VM data blocks and moves them to the target-side Veeam Data Mover.
4. The target-side Veeam Data Mover deduplicates similar blocks of data on the target side and writes the result to the backup file in the backup repository.

# On-Site Backup

To back up to a Microsoft Windows or Linux backup repository in the local site, you need to deploy a backup proxy on a machine with access to the source datastore and point the backup job to this backup proxy. In this scenario, the source-side Veeam Data Mover is started on the backup proxy, and the target-side Veeam Data Mover is started on the Microsoft Windows or Linux repository. VM data is sent from the backup proxy to the backup repository over the LAN.

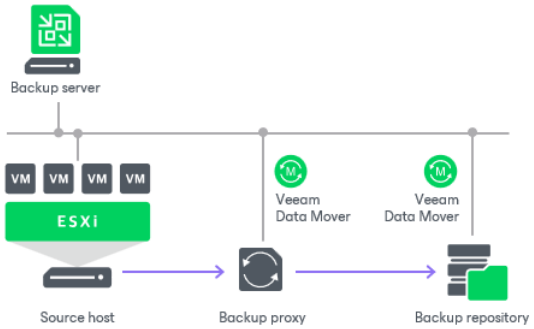


To back up to a shared folder in the local site, you need to deploy a gateway server with access to the shared folder backup repository. You can assign the role of a gateway server to the backup server itself or any Microsoft Windows machine added to the backup infrastructure.

You can use the same Microsoft Windows machine as the backup proxy and gateway server for SMB. In this scenario, Veeam Backup & Replication starts the source-side and target-side Veeam Data Movers on the same machine and sends VM data from the backup proxy to the shared folder backup repository over the LAN.

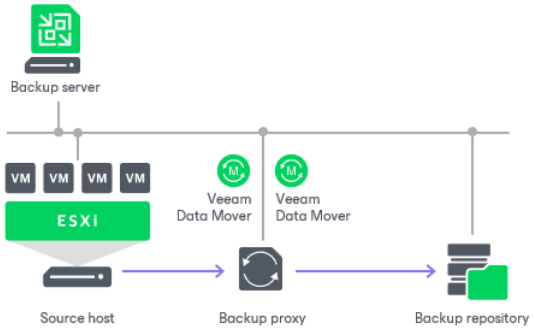


# Off-Site Backup

The common requirement for off-site backup is that one Veeam Data Mover runs in the production site (closer to the source datastore), and the other Veeam Data Mover runs in the remote site, closer to the backup repository. During backup, Veeam Data Movers maintain a stable connection, which allows for uninterrupted operation over the WAN or slow links.

To back up to a Microsoft Windows or Linux repository in the remote site, you need to deploy a backup proxy in the production site, closer to the source datastore. In this scenario, the source-side Veeam Data Mover is started on the backup proxy, and the target-side Veeam Data Mover is started on the Microsoft Windows or Linux repository. VM data is sent from the backup proxy to the backup repository over the WAN.

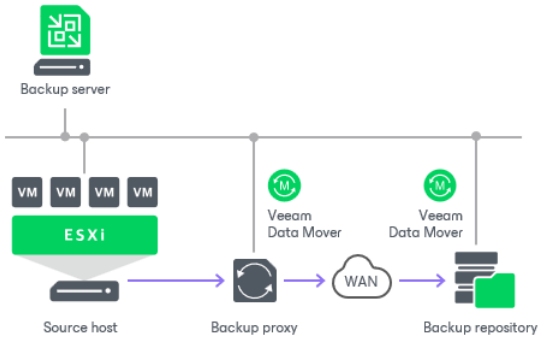


To back up VMs to a shared folder backup repository in the remote site, you must deploy a backup proxy in the source site and a gateway server in the remote site. The shared folder backup repository must be pointed at the target-side gateway server. During backup, the source-side Veeam Data Mover is started on the source backup proxy in the production site, and the target-side Veeam Data Mover is started on the target gateway server in the remote site. VM data is transferred between the backup proxy and gateway server over the WAN.

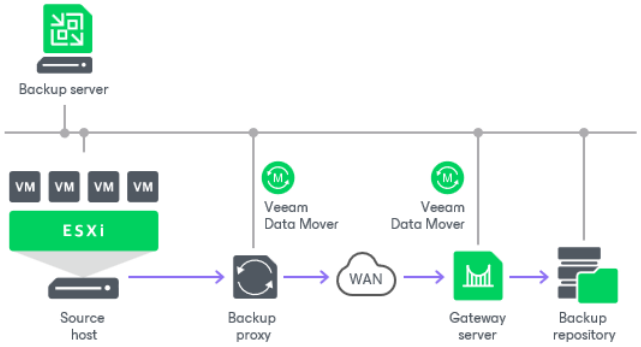

# Backup Chain

A backup chain is a sequence of backup files created by jobs. The backup chain provides the ability to recover data.

The backup chain consists of the first full backup file, incremental backup files, metadata files and some additional files. Full and incremental backup files correspond to restore points of the backed-up VM. You can think of restore points as "snapshots" of VM data at specific points in time. Restore points let you roll back VMs to the necessary state.

The type of backup files and how Veeam Backup & Replication orders them in the backup chain depend on the chosen backup method. For more information, see Backup Methods.

The amount of backup chains for backed-up VMs depends on the chosen backup chain format. For more information, see Backup Chain Formats.

## Backup Files

Veeam Backup & Replication creates and maintains the following types of backup files:

- VBK — full backup files that store copies of full VM images.
- VIB or VRB — incremental backup files that store incremental changes of VM images.
- VBM — backup metadata files that store information about the backup job, VMs processed by the backup job, number and structure of backup files, restore points, and so on. Metadata files facilitate the import of backups, backup mapping and other operations.

In addition to these file types, Veeam Backup & Replication can create the following files in the backup repository:

- VSB — virtual synthetic backup files used to generate virtual full backups on tapes. For more information, see the Virtual Full Backup section in the Veeam Backup & Replication User Guide.
- VLB, VSM and VLM — files that store Microsoft SQL Server transaction log data. For more information, see Microsoft SQL Server Log Backup.
- VLB, VOM and VLM — files that store Oracle archived log data. For more information, see Oracle Log Backup.
- VLB, VPM and VLM — files that store PostgreSQL WAL files. For more information, see PostgreSQL WAL Files Backup.

All backup files created by the backup job reside in a dedicated job folder in the backup repository. For example, if you create a backup job with the *DC Backup* name, Veeam Backup & Replication will create the `DC Backup` folder on the target backup repository and store all backup files produced with this job in this folder.

## Rolling Back VMs

To roll back a VM to a specific point in time, you need a chain of backup files: a full backup file plus a set of incremental backup files dependent on this full backup file. If some file in the backup chain is missing, you will not be able to roll back to the necessary state. For this reason, you must not manually delete separate backup files from the backup repository. Instead, you must specify retention policy settings that will let you maintain the desired number of backup files in the backup repository.

# Backup Methods

Veeam Backup & Replication provides the following methods for creating backup chains:

- **Forever forward incremental (FFI)**

  When the forever forward incremental (FFI) backup method is used, Veeam Backup & Replication creates a backup chain that consists of the first full backup file (VBK) and a set of forward incremental backup files (VIBs) following it.

  This backup method helps you save space on the backup storage because Veeam Backup & Replication stores only one full backup file and removes incremental backup files once the retention period is exceeded. To meet retention policy settings, Veeam Backup & Replication injects data of an incremental file into the full backup file before deleting the increment. Such transformations can lead to the fragmentation of the full backup file, and you have to schedule the compact of full backup file operation. This operation produces an additional I/O load on the backup storage. Overall, the FFI method produces a moderate I/O impact on the backup storage compared to other backup methods.

  Restore to the earliest restore point from backup files created using the FFI method is the fastest compared to other methods because the first available restore point is always a full backup. Restore to other restore points can be compared by speed with the FI method.

  For more information on the FFI backup method and how it works, see Forever Forward Incremental Backup.

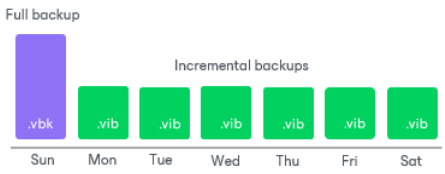


- **Forward incremental (FI)**

  When the forward incremental (FI) backup method is used, Veeam Backup & Replication creates a backup chain that consists of multiple full backup files (VBKs) and sets of forward incremental backup files (VIBs) following each full backup file. Full backups can be created using synthetic full and active full methods. By doing regular full backups, the backup chain is split into shorter series. This lowers the chances of losing the backup chain completely and makes this backup method the most reliable.

  This backup method requires more storage space than other methods because the backup chain contains multiple full backup files, and sometimes Veeam Backup & Replication stores more restore points than specified in the retention policy settings due to the specifics of the FI retention policy. For more information on how backups are retained, see Forward Incremental Backup Retention Policy.

  The FI backup method produces the lowest I/O impact on the backup storage. However, the impact on the backup storage increases on days when synthetic full backups are scheduled; the impact on the production storage increases on days when active full backups are scheduled.

  Restore from backup files created using the FI method is the most optimal in time compared to other methods (in cases when you do not restore to the earliest or latest restore point). That is because the backup chain is usually divided into short series of full backup files and incremental files, and the aggregation of the desired restore point does not take a long time.

  For more information on the FI backup method and how it works, see Forward Incremental Backup.

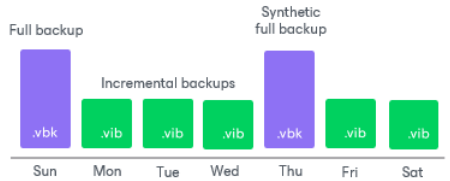

- **Reverse incremental (RI)**

  When the reverse incremental (RI) backup method is used, Veeam Backup & Replication creates a backup chain that consists of the full backup file (VBK) and a set of reverse incremental backup files (VRBs) preceding it.

  This backup method helps you save space on the backup storage because Veeam Backup & Replication stores only one full backup file (if you do not schedule periodic full backups) and removes incremental backup files once the retention period is exceeded. For more information on how backups are retained, see Reverse Incremental Backup.

  The RI method produces the heaviest I/O impact on the backup storage compared to other backup methods. That is because, during backup, Veeam Backup & Replication injects changed data blocks into the full backup file and also creates reverse incremental backup files. Such transformations can lead to the fragmentation of the full backup file, and you have to schedule the compact of full backup file operation. This operation produces an additional I/O load on the backup storage.

  Restore to the latest restore point from backup files created using the RI method is the fastest compared to other methods because the most recent restore point is always a full backup and gets updated after every backup cycle. Restore to earlier restore points is slower than for other backup methods.

  For more information on the RI backup method and how it works, see Reverse Incremental Backup.

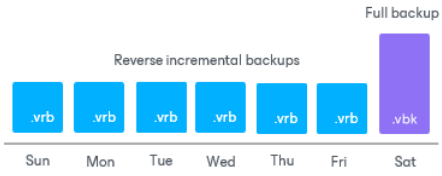


## Forever Forward Incremental Backup

The forever forward incremental backup method produces a backup chain that consists of the first full backup file (VBK) and a set of forward incremental backup files (VIB) following it.

Veeam Backup & Replication creates a forever forward incremental backup chain in the following way:

1. During the first session of a backup job, Veeam Backup & Replication creates a full backup file in the backup repository.
2. During subsequent backup job sessions, Veeam Backup & Replication copies only VM data blocks that have changed since the last backup job session (full or incremental) and saves these blocks as an incremental backup file in the backup chain.
3. After adding a new restore point to the backup chain, Veeam Backup & Replication checks the retention policy for the job. If Veeam Backup & Replication detects an outdated restore point, it transforms the backup chain to make room for the most recent restore point. For more information, see Forever Forward Incremental Backup Retention Policy.

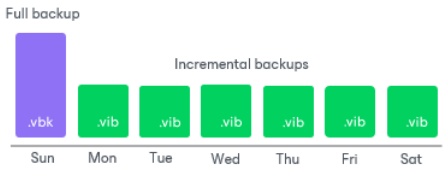


To use the forever forward incremental backup method, you must select the following options in the backup job settings:

1. Select the **Incremental** backup mode.

2. Do not enable synthetic full backups and active full backups. If you enable synthetic and active full backups, Veeam Backup & Replication will produce a forward incremental backup chain.

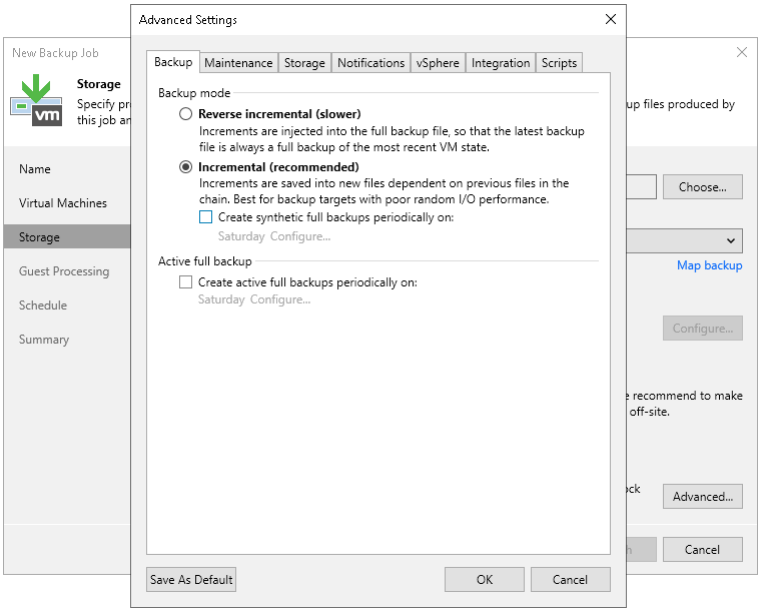


# Forward Incremental Backup

The forward incremental backup method produces a backup chain that consists of the first full backup file (VBK) and a set of forward incremental backup files (VIB) following it. Additionally, the forward incremental backup chain contains synthetic full and active full backup files that "split" the backup chain into shorter series.

Veeam Backup & Replication creates a forward incremental backup chain in the following way:

1. During the first backup job session, Veeam Backup & Replication creates a full backup file in the backup repository.
2. During subsequent backup job sessions, Veeam Backup & Replication copies only VM data blocks that have changed since the last backup job session (full or incremental) and saves these blocks as an incremental backup file in the backup chain.
3. On a day when the synthetic full or active full backup is scheduled, Veeam Backup & Replication creates a full backup file and adds it to the backup chain. Incremental restore points produced after this full backup file use it as a new starting point.
4. After adding a new restore point to the backup chain, Veeam Backup & Replication checks the retention policy set for the job. If Veeam Backup & Replication detects an outdated restore point, it attempts to remove this point from the backup chain. For more information, see Forward Incremental Backup Retention Policy.

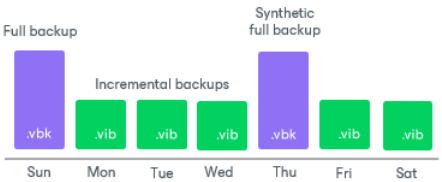


The forward incremental backup with synthetic full backup enabled is a default method for backup chain creation. To use the forward incremental backup method, you can leave the default settings or select the following options in the backup job settings:

1. Select the **Incremental** backup mode.

2. Enable synthetic full backups or active full backups. If the synthetic full backup and active full backups are not enabled, Veeam Backup & Replication will produce a forever forward incremental backup chain.

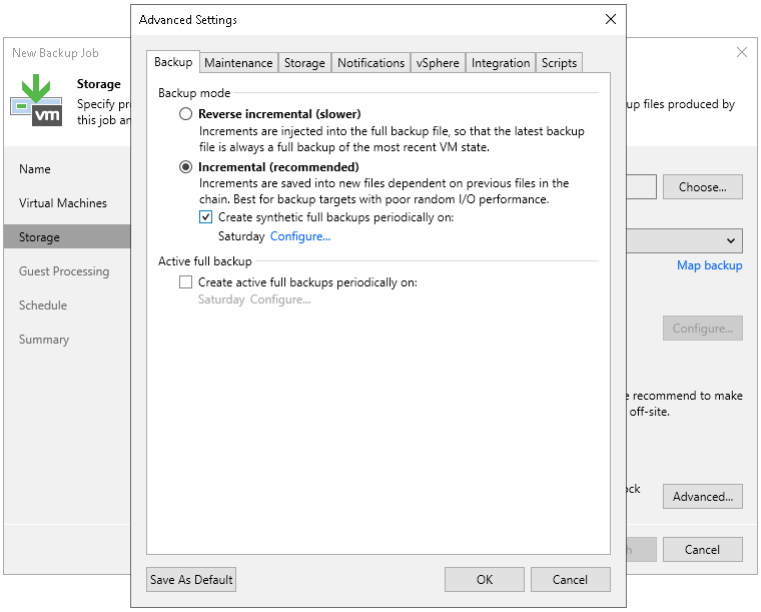


## Reverse Incremental Backup

The reverse incremental backup method produces a backup chain that consists of the last full backup file (VBK) and a set of reverse incremental backup files (VRB) preceding it.

Veeam Backup & Replication creates a reverse incremental backup chain in the following way:

1. During the first backup job session, Veeam Backup & Replication creates a full backup file in the backup repository.
2. During subsequent backup job sessions, Veeam Backup & Replication copies only VM data blocks that have changed since the last backup job session. Veeam Backup & Replication "injects" copied data blocks into the full backup file to rebuild it to the most recent state of the VM. Additionally, Veeam Backup & Replication creates a reverse incremental backup file containing data blocks that are replaced when the full backup file is rebuilt and adds this reverse incremental backup file before the full backup file in the backup chain.
3. After adding a new restore point to the backup chain, Veeam Backup & Replication checks the retention policy set for the job. If Veeam Backup & Replication detects an outdated restore point, it removes this point from the backup chain. For more information, see Reverse Incremental Backup Retention Policy.

As a result, the most recent restore point in the backup chain is always a full backup, and it gets updated after every successful backup job session.

The reverse incremental backup method lets you immediately restore a VM to the most recent state without extra processing because the most recent restore point is a full backup file. If you need to restore a VM to a particular point in time, Veeam Backup & Replication applies the required VRB files to the VBK file to get to the required restore point.

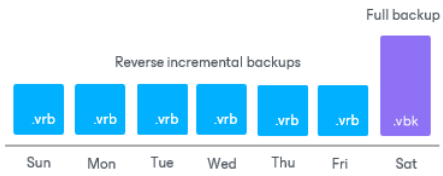

To use the reverse incremental backup method, you must select the **Reverse incremental** option in the backup job settings.

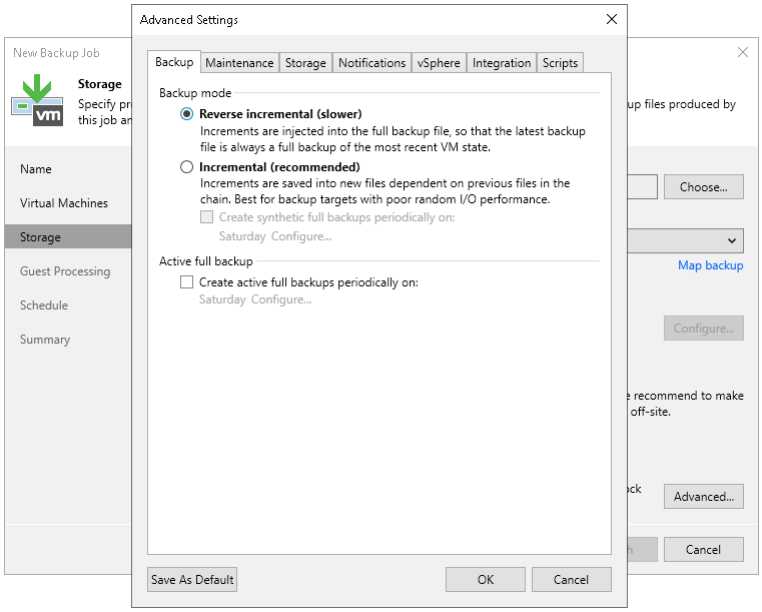


## Switching Between Backup Methods

You can easily switch between backup methods. Veeam Backup & Replication does not transform the previously created chain. Instead, it creates a new backup chain next to the existing one in the following manner:

- If you switch from the reverse incremental method to the forever forward incremental or forward incremental method, Veeam Backup & Replication creates a set of incremental backup files next to the reverse incremental chain. The full backup file in the reverse incremental chain is used as a starting point for incremental backup files.
- If you switch from the forever forward incremental or forward incremental method to the reverse incremental method, Veeam Backup & Replication first creates a full backup file next to incremental backup files. During every new job session, Veeam Backup & Replication transforms this full backup file and adds reverse incremental backup files to the backup chain.
- If you switch from the forever forward incremental method to the forward incremental method, Veeam Backup & Replication creates synthetic full backups according to the specified schedule. The old backup chain is deleted when the number of restore points in the new chain reaches the retention limit.
- If you switch from the forward incremental method to the forever forward incremental method, synthetic full backups are no longer created. When the number of restore points created since the last full backup reaches the retention limit, the old backup chain is deleted. Thereafter, with each restore point creation the earliest increment file will merge with the full backup file.

For more information on which settings are required for each method, see the following sections: Forever Forward Incremental Backup, Forward Incremental Backup and Reverse Incremental Backup.

# Backup Chain Formats

When you add repositories to the backup infrastructure, you can configure how many backup files Veeam Backup & Replication creates for the protected workloads. The type of backup files created depends on the selected backup method.

Veeam Backup & Replication provides the following backup chain formats:

- **Per-machine backup with separate metadata files (default format)**

  Veeam Backup & Replication creates one metadata file for each workload in a job. During a job session, Veeam Backup & Replication creates one backup file for each workload. Compared to other formats, this format allows managing individual workloads, for example, moving backups to other jobs or creating full backups.

- **Per-machine backup with single metadata file (legacy format)**

  Veeam Backup & Replication creates one metadata file for all workloads in a job. During a job session, Veeam Backup & Replication creates one backup file for each workload.

- **Single-file backup**

  Veeam Backup & Replication creates one metadata file for all workloads in a job. During a job session, Veeam Backup & Replication creates one backup file for all workloads.

The backup chain format depends on the **Use per-machine backup files** option in the backup repository settings and also the version of Veeam Backup & Replication when the repository was added. You can find the details further in this section.

**IMPORTANT**

For backup copy jobs, consider the following:

- New backup copy jobs always create per-machine backup files with separate metadata files.
- Backup copy jobs created in the previous versions of Veeam Backup & Replication continue to create backups in the specified formats: single-file backups if the **Use per-machine backup files** option was disabled or per-machine backups with single metadata file if the option was enabled. You can upgrade the backup chain format as described in section Upgrading Backup Chain Formats.

**NOTE**

In the Veeam Backup & Replication guide, we use "per-machine backups" when information relates to both per-machine backups, with single and separate metadata files. If information relates to one backup chain format only, we name it explicitly — per-machine backup with a single metadata file or per-machine with separate metadata files.

# Per-Machine Backup with Separate Metadata Files

The per-machine backup with separate metadata files format is the default option for new repositories since Veeam Backup & Replication version 12. When you create a repository, check that the **Use per-machine backup files** option is selected in the repository settings.

When per-machine backup with separate metadata files format is used, a backup job writes data in a separate stream for each workload. Veeam Backup & Replication saves data of each workload into a separate backup file and also creates a separate metadata file (.VBM) for each workload. Veeam Backup & Replication perceives each backup created during one job session as one restore point. As a result, each workload has its own independent backup chain.

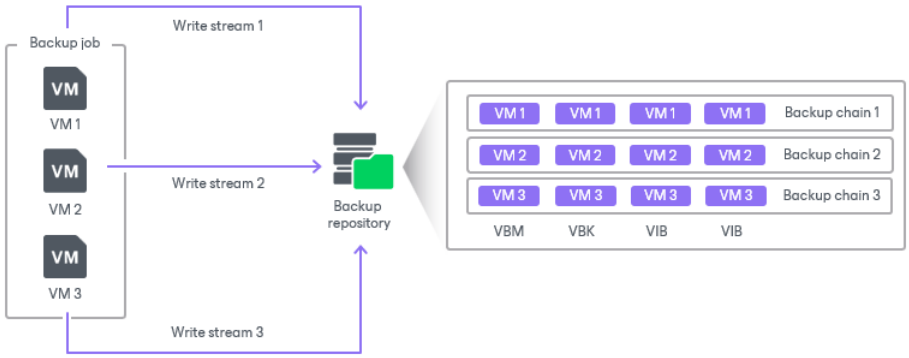


Such a way of storing backups is recommended and makes environment management more flexible. Veeam Backup & Replication can perform operations with individual workloads. It can, for example, move a workload and its backups from one job to another or launch active full backups for individual workloads.

When Veeam Backup & Replication needs to remove restore points by retention, it analyzes the backup chain of an individual workload, not all workloads. For more information on the restore point removal, see Removal of Restore Points.

# Per-Machine Backup with Single Metadata File

The per-machine backup with single metadata file format is legacy since Veeam Backup & Replication version 12. However, backup repositories created in Veeam Backup & Replication prior to version 12, for which the **Use per-machine backup files** option was enabled, still create backups in the per-machine backup with single metadata file format.

When per-machine backup with single metadata file format is used, a backup job writes data in a separate stream for each workload. Veeam Backup & Replication saves data of each workload into a separate backup file. However, Veeam Backup & Replication creates a single metadata file (.VBM) for all workloads and perceives all backups created during one job session as one restore point.

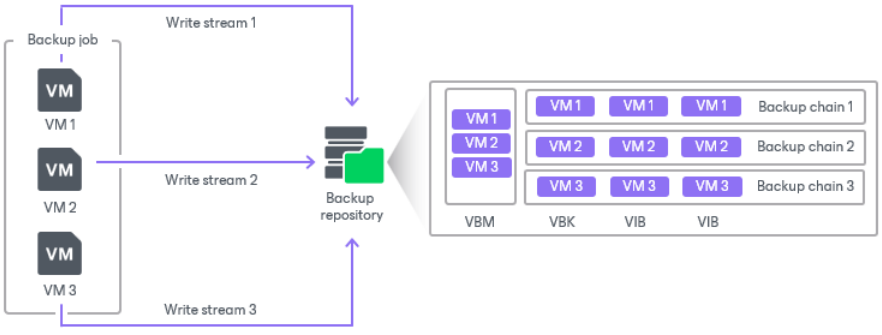


Such a way of storing backups is more efficient than a single-file backup format. Compared to per-machine backup with separate metadata files format, Veeam Backup & Replication cannot perform operations with individual workloads. These operations include moving backups, launching active full for individual workloads, and so on.

When Veeam Backup & Replication needs to remove earlier restore points by retention policy, it removes backup files of all workloads created during one job session. Veeam Backup & Replication does not remove data for separate workloads. In some situations, a certain workload can have fewer restore points than it is specified in retention policy settings. For more information on the restore point removal, see Removal of Restore Points.

# Single-File Backup

This is the default way of storing backup files for Veeam Backup & Replication prior to version 12. The single-file backups are created in backup repositories for which the **Use per-machine backup** files option is not selected.

When single-file backup format is used, a backup job writes workload data to a repository in one write stream. Veeam Backup & Replication saves data of all workloads into the same backup file and creates one metadata file (.VBM). Such behavior is not optimal if the target storage device is able to write data in multiple streams simultaneously. In this situation, the backup repository may become the bottleneck for the data transfer, even though its resources will not be fully utilized.

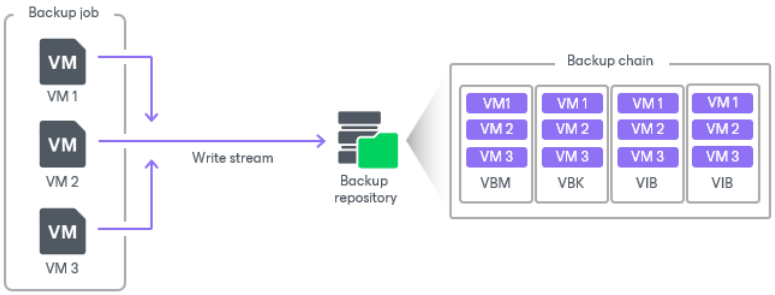


When Veeam Backup & Replication needs to remove earlier restore points by retention policy, it removes backup files of all workloads created during one job session. Veeam Backup & Replication does not remove data for separate workloads. In some situations, a certain workload can have fewer restore points than it is specified in retention policy settings. For more information on the restore point removal, see Removal of Restore Points.

# Limitations and Considerations for Per-Machine Backup Files

When planning to use per-machine backup files, consider the following limitations:

- We recommend that you use the **Use per-machine backup files** option, especially for deduplicating storage appliances used as backup repositories. Veeam Backup & Replication will write machine data to the backup repository in several streams, which will improve the backup job performance. However, deduplication may be less effective compared to deduplication for single-file backups. Deduplication works within one backup file. For single-file backups, one backup file stores data for all machines, which is why data is deduplicated for all machines. Per-machine backups store data for one machine, which is why data is deduplicated for one machine only.
- [For Veeam Agent backup jobs] Veeam Backup & Replication ignores the **Use per-machine backup files** option. The way Veeam Backup & Replication creates backup files depends on the objects included in the backup job:
  - If several Veeam Agent machines are included in the backup job, Veeam Backup & Replication creates a separate backup file for each machine.
  - If failover clusters are included in the backup job, Veeam Backup & Replication creates a separate backup file for each cluster.

# Upgrading Backup Chain Formats

Changing the **Use per-machine backup files** option for repositories does not affect the way backups are created. To change the backup chain format, use the following instructions:

- Upgrading Backup Chain Formats for backup jobs.
- Upgrading Backup Chain Formats for backup copy jobs.

# Full Backup Methods

Veeam Backup & Replication provides the following methods for creation of full backup files:

- Active Full Backup

  When you perform an active full backup, Veeam Backup & Replication retrieves VM data from the source datastore where the VM resides, compresses and deduplicates it and writes it to the VBK file in the backup repository.

- Synthetic Full Backup

  When you perform a synthetic full backup, Veeam Backup & Replication does not retrieve VM data from the source datastore. Instead, it synthesizes a full backup from data you already have in the backup repository. Veeam Backup & Replication accesses the previous full backup file and a chain of subsequent incremental backup files on the backup repository, consolidates VM data from these files and writes consolidated data into a new full backup file. As a result, the created synthetic full backup file contains the same data you will have if you create an active full backup.

> **TIP**
>
> You can perform both active and synthetic full backups. For more information on how to do that, see Backup Settings.

## Active Full Backup

In some cases, you need to create a full backup regularly. For example, your corporate backup policy may require you to create a full backup on weekends and run incremental backup on work days. To let you conform to these requirements, Veeam Backup & Replication lets you create active full backups.

The active full backup produces a full backup of a VM, just as if you run the backup job for the first time. Veeam Backup & Replication retrieves data for the whole VM from the source, compresses and deduplicates it and stores it in the full backup file — VBK.

The active full backup resets a backup chain. All incremental backup files use the latest active full backup file as a new starting point. A previously used full backup file remains on disk until it is automatically deleted according to the retention policy.

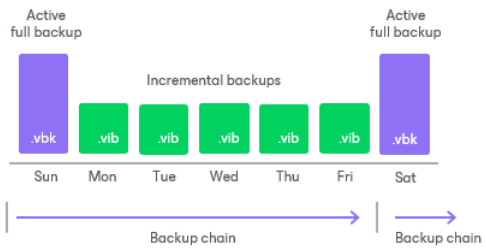

You can create active full backups manually or schedule a backup job to create active full backups periodically.

- To create an active full backup manually, follow the steps described in section Performing Active Full Backup.
- To schedule active full backups, specify scheduling settings in the **Advanced** section of a backup job. You can schedule active full backups to run weekly, for example, every Saturday. You can also schedule them monthly, for example, every fourth Sunday of a month.

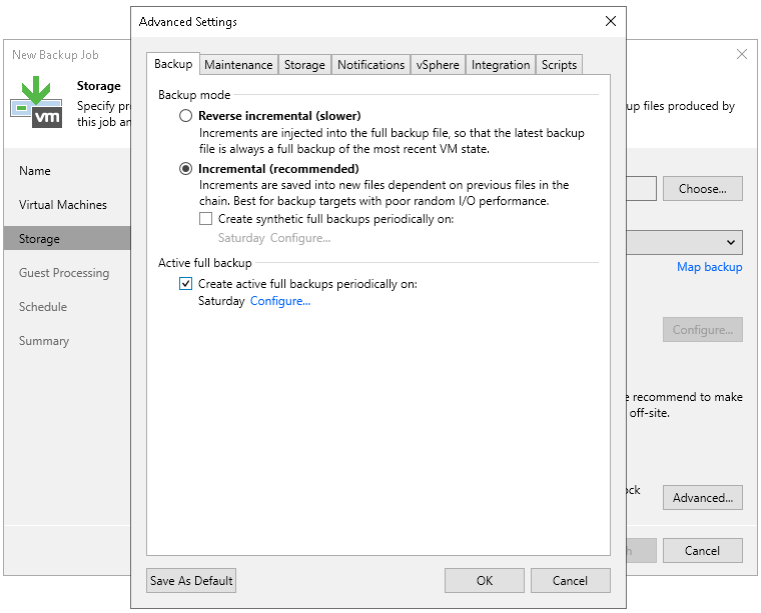


# Active Full Backup Schedule

Veeam Backup & Replication automatically triggers a backup job to create an active full backup, even if a regular backup job session is not scheduled on this day. The job session is started at the same time the parent backup job is scheduled. For example, if you schedule the parent backup job at 12:00 AM Sunday through Friday, and schedule active full backup on Saturday, Veeam Backup & Replication will start a backup job session that will produce an active full backup at 12:00 AM on Saturday.

If the parent backup job is not scheduled to run automatically or is disabled, Veeam Backup & Replication will not perform active full backup.

If a regular backup job is scheduled together with an active full backup, Veeam Backup & Replication will produce only an active full backup that will contain the latest state of the source VM. An incremental backup file that should have been created by the backup job schedule will not be added to the backup chain.

Veeam Backup & Replication creates an active full backup only once a day on which active full backup is scheduled (unless you create a full backup manually). If you run the backup job again on the same day, Veeam Backup & Replication will perform incremental backup in a regular manner.

> **IMPORTANT**
>
> If you schedule a job to start after another job (initial job), but the initial job does not run on days when the active full backup is scheduled for the chained job, Veeam Backup & Replication will not create active full backups.

# Synthetic Full Backup

In some situations, periodically running active full backups may not be an option. Active full backups are resource-intensive and consume a considerable amount of network bandwidth. As an alternative, you can create synthetic full backups.

In terms of data, the synthetic full backup is identical to a regular full backup. A synthetic full backup produces a VBK file that contains data of the whole VM. The difference between active and synthetic full backup lies in the way VM data is retrieved:

- When you perform an active full backup, Veeam Backup & Replication retrieves VM data from the source datastore where the VM resides, compresses and deduplicates it and writes it to the VBK file in the backup repository.
- When you perform a synthetic full backup, Veeam Backup & Replication does not retrieve VM data from the source datastore. Instead, it synthesizes a full backup from data you already have in the backup repository. Veeam Backup & Replication accesses the previous full backup file and a chain of subsequent incremental backup files on the backup repository, consolidates VM data from these files and writes consolidated data into a new full backup file. As a result, the created synthetic full backup file contains the same data you have if you create an active full backup.

The synthetic full backup has several advantages:

- The synthetic full backup does not use network resources; it is created from backup files you already have on disk.
- The synthetic full backup produces less load on the production environment: it is synthesized right on the backup repository.

Veeam Backup & Replication treats synthetic full backups as regular full backups. As well as any other full backup file, the synthetic full backup file resets the backup chain. All subsequent incremental backup files use the synthetic full backup file as a new starting point. A previously used full backup file remains on disk until it is automatically deleted according to the retention policy.

> **IMPORTANT**
>
> Consider the following:
>
> - If you enable both synthetic and active full backups and schedule their creation on the same day, the synthetic full backup is not created.
> - If you schedule a job to start after another job (initial job), but the initial job does not run on days when the synthetic full backup is scheduled for the chained job, Veeam Backup & Replication will not create synthetic full backups.
> - Synthetic full backups cannot be created independently for backup jobs targeted at object storage. To include synthetic full backups in the backup process, you must enable the GFS policy.

To create synthetic full backups, you must enable the **Create synthetic full backups periodically** option and schedule the creation of synthetic full backups on specific days.

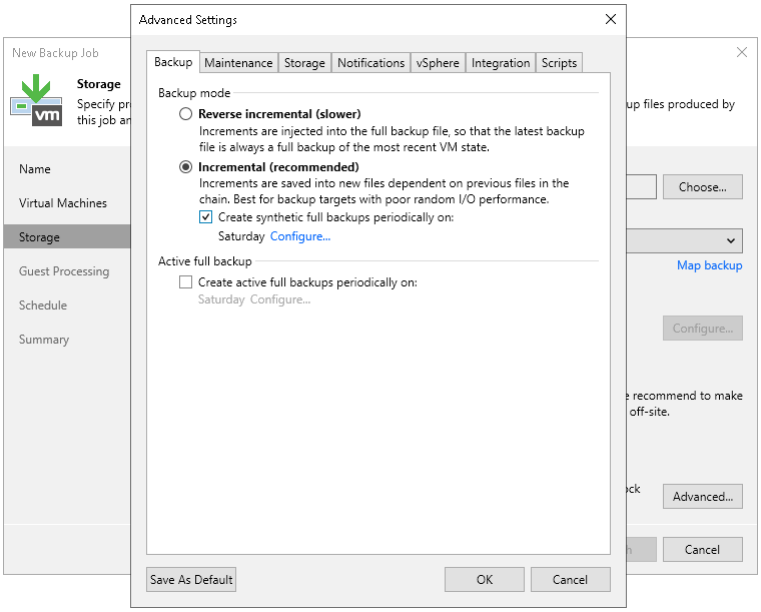


## How Synthetic Full Backup Works

To understand how synthetic full backups are created, consider the following example. The backup job starts on Sunday, with synthetic full backups scheduled on Friday. Active full backups are not configured.

Veeam Backup & Replication creates a synthetic full backup in the following way:

1. On a day when a synthetic full backup is scheduled, Veeam Backup & Replication triggers a new backup job session. During this session, Veeam Backup & Replication first performs incremental backup in a regular manner and adds a new incremental backup file to the backup chain.

   Veeam Backup & Replication retrieves VM data for this incremental backup file from the production storage. Incremental backup helps Veeam Backup & Replication ensure that the synthetic full backup includes the latest changes of the source VM in the production environment.

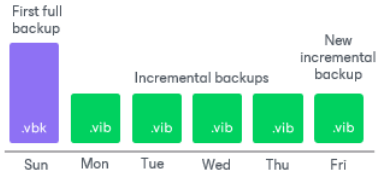


2. At the end of the backup job session, Veeam Data Mover on the backup repository builds a new synthetic full backup using backup files that are already available in the backup chain, including the newly created incremental backup file.

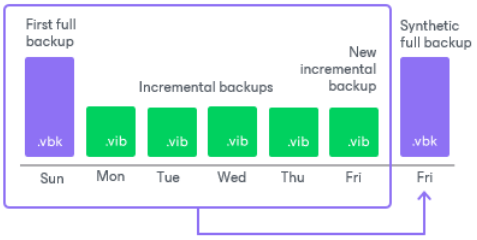

3. When the synthetic full backup is created, Veeam Data Mover on the backup repository deletes the incremental backup file created at the beginning of the job session. As a result, you have a backup chain that consists of a full backup file, a set of incremental backup files and a synthetic full backup file.

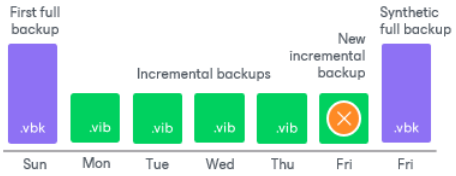


4. Every next job session creates a new incremental restore point, starting from the synthetic full backup until the day on which the synthetic full backup is scheduled. On that day, Veeam Backup & Replication creates a new synthetic full backup.

# Synthetic Full Backup Schedule

Veeam Backup & Replication automatically triggers a backup job session to create a synthetic full backup, even if a regular backup job session is not scheduled on this day. The job session is started at the same time the parent backup job is scheduled. For example, if you schedule the parent backup job at 12:00 AM Sunday through Friday and schedule a synthetic full backup on Saturday, Veeam Backup & Replication will start a backup job session that will produce a synthetic full backup at 12:00 AM on Saturday.

If a regular backup job is scheduled together with a synthetic full backup, Veeam Backup & Replication will produce only one backup file — a synthetic full backup that will contain the latest state of the source VM. An incremental backup file that should have been created by the backup job schedule will not be added to the backup chain.

If an active full backup is scheduled together with a synthetic full backup, Veeam Backup & Replication will create only the active full backup.

Veeam Backup & Replication creates a synthetic full backup only once a day on which synthetic full backup is scheduled. If you run the backup job again on the same day, Veeam Backup & Replication will perform incremental backup in a regular manner.

If an active or synthetic full backups are scheduled for the chained job, but the initial job does not run on these days, the active and synthetic full backups will not be created for the chained job. For more information, see Chained Jobs.

# Short-Term Retention Policy

Every successful backup job session creates a new restore point that lets you roll back VM data to an earlier point in time. To control the number of restore points in the backup chain, you must specify retention policy settings. The retention policy defines how many restore points you want to retain on disk and, thus, how 'far' you can roll back. After the allowed number of restore points is exceeded, Veeam Backup & Replication applies the retention policy — it removes the earliest restore point from the backup chain.

To define the retention policy for a backup job, you must specify the necessary number of restore points or days in the **Retention policy** field in the backup job settings. By default, Veeam Backup & Replication keeps the restore points for the last 7 days. In the **Storage** step of the **New Backup Job** wizard, you can select the following units of retention policy.

- **Restore points**: Veeam Backup & Replication keeps the last *N* restore points, where *N* is the specified number of restore points.

- **Days**: Veeam Backup & Replication keeps the restore points created during the last *N* days, where *N* is the specified number of days.

  The daily retention policy can be helpful if you periodically create off-schedule backups. For example, you create a daily backup job and you want to store retention points for 14 calendar days. If you select 14 restore points and manually create off-schedule backups, there will be several restore points in one day. So, you will have retention points for less than 14 days. In this case, you can use the **days** option.

  Consider the following for the daily retention policy:

  - The minimum number of retained restore points is 3. This number does not depend on the number of days set in the retention policy. For example, the retention policy is set to 5 days. You launch the job after it was stopped for 10 days. Normally, Veeam Backup & Replication will delete all previous restore points. However, due to the minimum number of retained restore points, you will still have at least 3 restore points: the newly created restore point and the two previous ones.

    You can change the minimum number of retained restore points with a registry value. For more information, contact Veeam Customer Support.

  - If the backup job starts at the end of the day and finishes the next day, Veeam Backup & Replication assumes that the restore point is created at the moment when the backup job started. However, Veeam Backup & Replication starts counting retention policy days only after the backup job finishes processing VMs.

  - When determining whether the number of allowed days is exceeded, Veeam Backup & Replication ignores the day when the daily retention policy is applied.

    In fact, Veeam Backup & Replication keeps the restore points for the *N + 1* days, where *N* is the number of days you specify in the job settings. In this case, Veeam Backup & Replication applies a retention policy after the *N+1* days have passed. For example, if you set a retention policy to keep the restore points for 6 days, Veeam Backup & Replication will keep the restore points for 7 days and apply the retention policy on the 8th day. Note that the retention period may be longer depending on the specified backup method.

  - When determining whether the number of allowed days is exceeded, Veeam Backup & Replication also counts days when the backup job did not create any backups.

When the specified number is exceeded, the earliest restore points will be removed from the backup chain or merged with the next closest restore point. Veeam Backup & Replication handles restore points in different ways for forever forward incremental, forward incremental and reverse incremental backup chains:

- Forever Forward Incremental Backup Retention Policy
- Forward Incremental Backup Retention Policy
- Reverse Incremental Backup Retention Policy

> **NOTE**
>
> When the specified number of restore points or days is exceeded, Veeam Backup & Replication deletes the whole backup file, not separate VMs from it. For more information, see Removal of Restore Points.

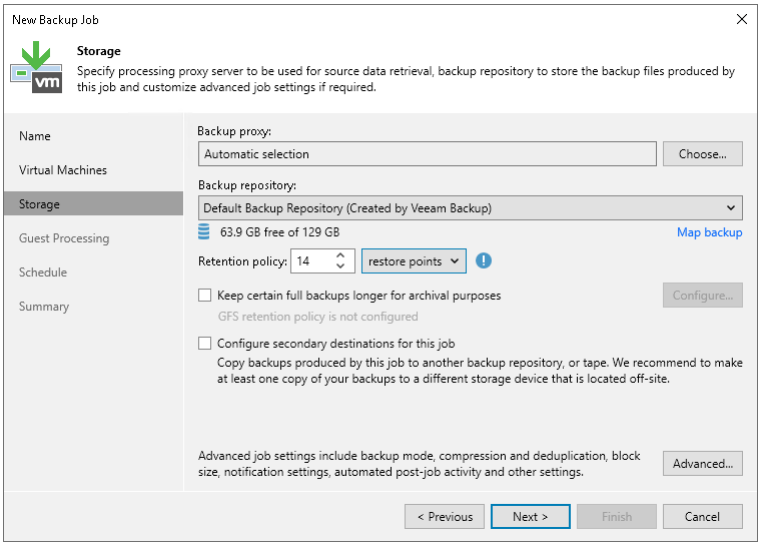


# Forever Forward Incremental Backup Retention Policy

If the number of days or restore points in forever forward incremental backup chains exceeds retention policy settings, Veeam Backup & Replication transforms the backup chain to make room for the most recent restore point. The transformation process is performed in the following way:

1. Veeam Backup & Replication adds a new incremental backup file to the backup chain and detects that the number of allowed restore points or days is exceeded.
2. Veeam Backup & Replication reuses empty data blocks in the full backup file to include changes of the incremental backup file that follows the full backup. To do that, Veeam Backup & Replication injects data blocks from the first incremental backup file in the chain into the full backup file. As a result, the full backup file 'moves' one step forward in the backup chain.

> **NOTE**
>
> If the forever forward incremental backup chain resides on a deduplicating storage appliance, Veeam Backup & Replication does not reuse empty data blocks of the full backup file. Instead, Veeam Backup & Replication appends data from the first incremental backup file in the chain to the full backup file. As a result, the backup chain may consume more disk space on the appliance.

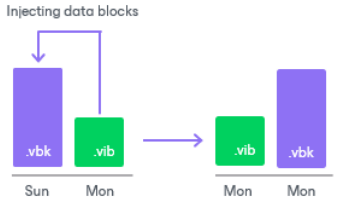


3. The first incremental backup file is removed from the backup chain as redundant. Its data has already been injected into the full backup file, and the full backup file contains the same data as this incremental backup file.

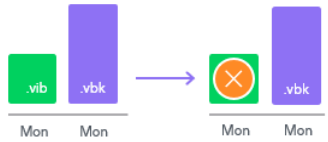

For example, you want to keep 7 restore points in the backup chain. The backup job starts on Sunday and runs daily. In this case, Veeam Backup & Replication will create the backup chain in the following way:

1. During the first backup job session on Sunday, Veeam Backup & Replication creates the first restore point — a full backup file.
2. Monday through Saturday Veeam Backup & Replication adds six incremental backup files to the backup chain.

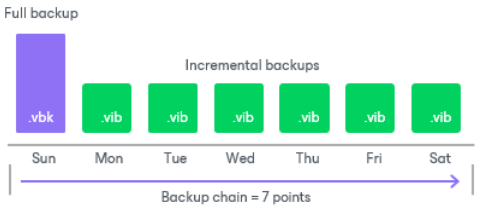


3. The next Sunday, Veeam Backup & Replication adds a new incremental backup file to the backup chain.
4. Veeam Backup & Replication detects that the number of allowed restore points is exceeded and starts the transformation process:
   a. Veeam Backup & Replication merges data blocks from the incremental backup file created on Monday into the full backup file created on Sunday. This way, the full backup file 'moves' one step forward — from Sunday to Monday.

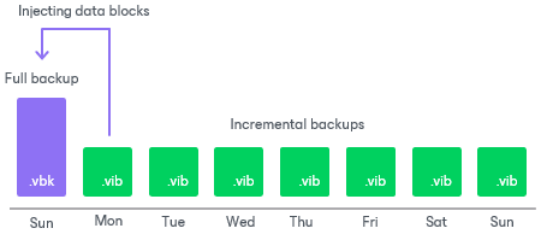


   b. The incremental backup created on Monday becomes redundant and is removed from the backup chain.

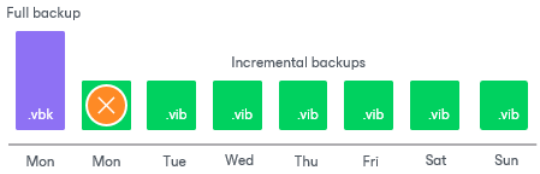


As a result, you have a chain of a full backup file as of Monday and six incremental backup files Tuesday through Sunday.

## Forward Incremental Backup Retention Policy

To be able to restore from a forward incremental backup, you need to have a full backup file and a chain of subsequent incremental backup files on disk. If you delete a full backup file, the whole chain of incremental backup files will become useless. Similarly, if you delete any incremental backup file before the point to which you want to roll back, you won't be able to restore VM data (since later incremental backup files depend on earlier incremental backup files).

For this reason, if you select the forward incremental backup method, on some days, there will be more restore points on disk than specified by retention policy settings. Veeam Backup & Replication will remove the full backup chain only after the last incremental backup file in the chain becomes outdated.

For example, the retention policy is set to 3 restore points. The first full backup file is created on Sunday, incremental backup files are created Monday through Saturday, and the second full backup is created on Thursday. Although the retention policy is already reached on Wednesday, the first full backup is not deleted. Without the full backup, the backup chain will be useless, leaving you without any restore point at all. Veeam Backup & Replication will wait for the second full backup file and 2 incremental backup files to be created, and only then will delete the whole previous chain, which will happen on Saturday.

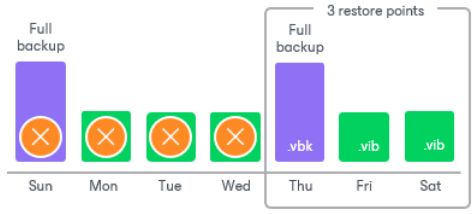


## Reverse Incremental Backup Retention Policy

In case of reverse incremental backup, Veeam Backup & Replication immediately deletes the earliest reverse incremental backup file as soon as it becomes outdated.

For example, you configure a backup job in the following way:

- The backup job starts on Sunday.
- The backup method is reverse incremental.
- Retention policy is set to 6 restore points.

Veeam Backup & Replication will start the backup job on Sunday. Monday through Friday, it will add new restore points to the backup chain and rebuild the full backup file. On Saturday, Veeam Backup & Replication will add a new restore point and remove the earliest reverse incremental backup file (VRB) from the backup chain.

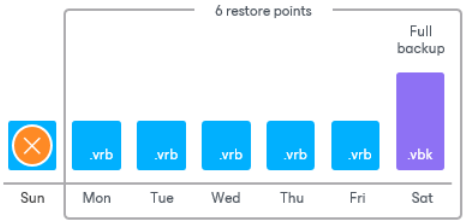


## Retention Policy for Deleted Items

In some situations, after you configure and run backup jobs in Veeam Backup & Replication, you may want to change something in the virtual infrastructure or in the backup strategy. For example, you may remove some machines from the virtual infrastructure or move them to another location. You may also exclude some machines from jobs that have already run for some time.

The retention policy for deleted items functions differently depending on the **Use per-machine backup files** option. For details, see Backup Chain Formats.

- [If per-machine is enabled] When you enable the retention policy for deleted items, Veeam Backup & Replication will remove data for machines that are no longer processed by the backup job from the backup repository.

- [If per-machine is disabled] When you enable the retention policy for deleted items, Veeam Backup & Replication will remove the data about deleted items from the backup job and Veeam Backup & Replication database. The stored blocks of deleted machines will remain in the repository. The stored blocks of deleted machines will be removed only when the restore point retention limit is reached or by the compact full backup file option.

The retention policy for deleted items data is set at the level of the backup job. You must enable the **Remove deleted items data after** option in backup job settings and specify the period of time for which data for deleted items must be retained in the backup repository.

Consider the following:

- You must use the retention policy for deleted items data carefully. We strongly recommend that you set the retention policy to 7 days or more to prevent unwanted data loss.
- The **Remove deleted items data after** option lets you control data of deleted or excluded items. In addition to it, Veeam Backup & Replication applies general retention policy rules to maintain the necessary number of restore points in the backup chain. For more information, see Short-Term Retention Policy.
- If the **Remove deleted items data after** option is not enabled, data for the VMs that are no longer being processed will remain in the backup repository forever.

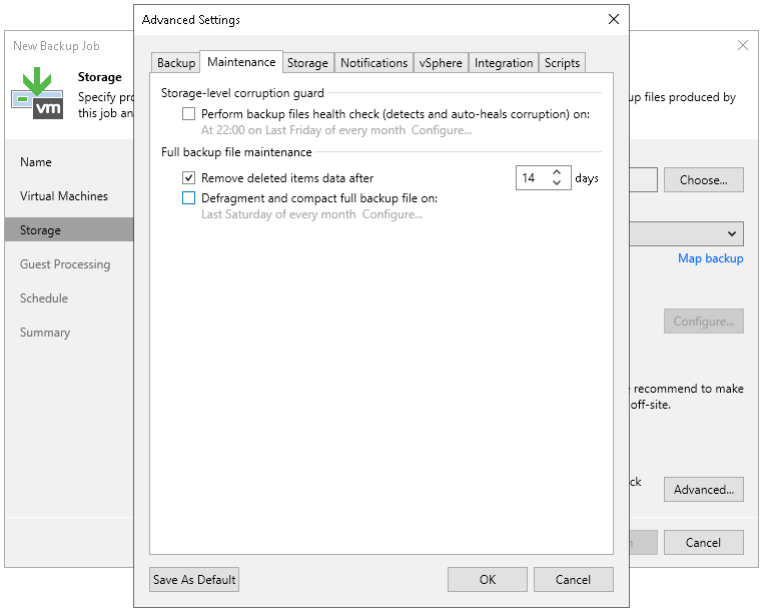


# How Retention Policy for Deleted Items Works

If you enable a retention policy for deleted items data in backup job settings, Veeam Backup & Replication performs the following actions:

1. If all machines in the job are processed with the *Success* status, at the end of the backup job session Veeam Backup & Replication gets a list of machines in the backup.
2. For every machine in the backup, Veeam Backup & Replication checks the configuration database and gets the date of the latest backup job session completed with the *Success* status.
3. Veeam Backup & Replication checks if any machine in the backup meets the following conditions:
   - There are no successful backups for the machine for the last N days.
   - There are no corrupted backups for the machine for the last N days.

   Where N is the number of days specified in the **Remove deleted items data after N days** setting.

4. If both conditions are true for some machines, Veeam Backup & Replication removes data for this machine from the backup. Note that if per-machine is disabled, it does not free up space in the backup repository. It marks the space as available to be overwritten, and this space is removed during subsequent job sessions or the full backup file compact operation.

### Example 1

You create a backup job for 2 VMs and set the retention policy for deleted items to 5 days. The backup job runs once a day for 7 times and processes VMs in the following way:

- *VM 1* is successfully processed during all job sessions.
- *VM 2* is successfully processed during the 1st and 2nd backup job sessions. Before the 3rd job session, *VM 2* is excluded from the job and is not processed by subsequent job sessions.

During the 8th job session, Veeam Backup & Replication will remove data for *VM 2* from backups in the backup repository since there are no successful and corrupted backups for *VM 2* for the last 5 days.

### Example 2

You create a backup job for 2 machines and set the retention policy for deleted machines to 5 days. The backup job runs once a day for 7 times and processes machines in the following way:

- *VM 1* is successfully processed during all job sessions.
- *VM 2* is successfully processed during the 1st and 2nd backup job sessions. Starting from the 3rd job session, *VM 2* fails to be processed, for example, due to power loss while machine data is transported.

During the 8th job session, Veeam Backup & Replication will not remove data for *VM 2* from backups in the backup repository. Even though there are no successfully created backups for *VM 2* for the last 5 days, Veeam Backup & Replication will detect that the configuration database contains information about corrupted backups for *VM 2* for the last 5 days.

# Limitations for Retention Policy for Deleted Items

- [Per-machine is disabled] Retention policy for deleted items does not function if you enable synthetic full backups and active full backup.

  [Per-machine is enabled] Retention policy for deleted items functions without limitations.
- Retention policy for deleted items does not apply to full backup files to which GFS flags are assigned.
- [For VMware Cloud Director backup jobs] To apply retention policy for deleted items, Veeam Backup & Replication checks backups created for the vApp itself, not for a machine in this vApp. Thus, the retention policy is applied only if the job stops creating backups for the entire vApp.

## Removal of Restore Points

Retention works in different ways for per-machine backup with separate metadata files, single-file backup and per-machine with single metadata file formats. For more information on backup chain formats, see Backup Chain Formats.

# Retention for Per-Machine Backups with Separate Metadata Files

When you use the per-machine backup with separate metadata files format, Veeam Backup & Replication creates a separate backup file for each workload during one session. Veeam Backup & Replication perceives each backup file created during one job session as one restore point. As a result, each workload has its own backup chain. When Veeam Backup & Replication needs to remove earlier restore points by retention, it analyzes the backup chain of an individual workload, not all workloads.

## Removal of Restore Points from Forward Incremental Chains

In the case of a forward incremental backup chain, Veeam Backup & Replication does not remove a restore point immediately. Instead, Veeam Backup & Replication waits for a new full backup (synthetic or active) to be created and a new backup chain to be started. As soon as the last incremental restore point in the "old" backup chain is marked as redundant, Veeam Backup & Replication removes the whole "old" backup chain from the backup repository. For more information, see Forward Incremental Backup Retention Policy.

For example, a backup job processes 2 VMs: *VM 1* and *VM 2*. According to the retention policy settings, the backup chain must contain 3 restore points. The backup job ran 3 times and VMs were processed in the following way:

- *VM 1* was successfully backed up 3 times and has 3 restore points.
- *VM 2* was not backed up in 2 job sessions and has 1 valid restore point.

During the 4th job run, Veeam Backup & Replication will create new restore points for *VM 1* and *VM 2* — full backups. However, Veeam Backup & Replication will not remove the earliest restore point.
Veeam Backup & Replication will wait until a full backup file and 2 incremental backup files are present in the backup chains of each VM. After that, Veeam Backup & Replication will remove the whole outdated backup chain from the backup repository. During the 5th job run, Veeam Backup & Replication will create a new restore point for *VM 1* and *VM 2*.

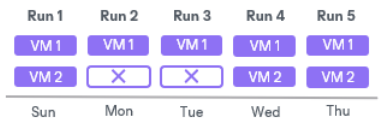


During the 6th job run, the processing of *VM 2* will fail. Veeam Backup & Replication will remove the earliest restore points only of *VM 1* because the number of restore points for *VM 1* in the active backup chain equals 3 — a full backup file and 2 incremental backup files. The number of restore points in the active backup chain of *VM 2* is not sufficient to remove restore points.

During the 7th job run, Veeam Backup & Replication will create restore points for *VM 1* and *VM 2*. Veeam Backup & Replication will remove the earliest restore points only for *VM 2* because a new backup chain is not started for *VM 1* and the number of restore points for *VM 2* in the active backup chain equals 3 — a full backup file and 2 incremental backup files.

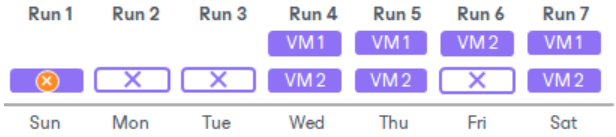


## Removal of Restore Points from Reverse Incremental Chains

In the case of a reverse incremental backup chain, Veeam Backup & Replication immediately removes a redundant restore point when the allowed number of restore points is exceeded. For more information, see Reverse Incremental Backup Retention Policy.

For example, a backup job processes two VMs: *VM 1* and *VM 2*. According to the retention policy settings, the backup chain must contain 5 restore points. The backup job ran 5 times and VMs were processed in the following way:

- *VM 1* was successfully backed up 5 times and has 5 valid restore points.
- *VM 2* was not backed up in 1 job session and has 4 valid restore points.

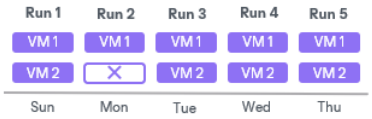


During the 6th job run, Veeam Backup & Replication will create restore points for *VM 1* and *VM 2*. Veeam Backup & Replication will remove the earliest restore point of *VM 1* because the number of restore points in the backup chain exceeds 5.

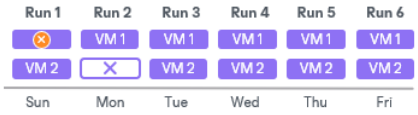


During the 7th job run, Veeam Backup & Replication will create restore points for *VM 1* and *VM 2*. Veeam Backup & Replication will remove the earliest restore points of *VM 1* and *VM 2* because the number of restore points in the backup chain of each VM exceeds 5.

# Retention for Single-File Backups and Per-Machine Backups with Single Metadata File

When you use the single-file backup chain format, Veeam Backup & Replication creates one backup file for all workloads during one session. When you use the per-machine backup with a single metadata file format, Veeam Backup & Replication creates a separate backup file for each workload during one session. However, Veeam Backup & Replication perceives all backup files created during one job session as one restore point.

When Veeam Backup & Replication needs to remove earlier restore points by retention policy, it removes backup files of all workloads created during one job session. Veeam Backup & Replication does not remove data for separate VMs. In some situations, a certain VM may have fewer restore points than it is specified in retention policy settings. It can happen if a backup job processes multiple workloads, and some workloads fail to be processed during some job sessions.

## Removal of Restore Points from Forward Incremental Chains

In the case of a forward incremental backup chain, Veeam Backup & Replication does not remove a restore point immediately. Instead, Veeam Backup & Replication waits for a new full backup (synthetic or active) to be created and a new backup chain to be started. As soon as the last incremental restore point in the "old" backup chain is marked as redundant, Veeam Backup & Replication removes the whole "old" backup chain from the backup repository. For more information, see Forward Incremental Backup Retention Policy.

For example, a backup job processes 2 VMs: *VM 1* and *VM 2*. According to the retention policy settings, the backup chain must contain 3 restore points. The backup job ran 3 times and VMs were processed in the following way:

- *VM 1* was successfully backed up 3 times and has 3 restore points.
- *VM 2* was not backed up in 2 job sessions and has 1 valid restore point.

When Veeam Backup & Replication adds a new restore point to the backup chain, it will not remove the earliest restore point. Veeam Backup & Replication will wait until a new full backup file and 2 incremental backup files are added to the backup chain. After that, it will remove the whole outdated backup chain from the backup repository. Restore points in the new backup chain, at the same time, may contain data for both VMs or for one VM only: Veeam Backup & Replication regards backup files as restore points, not separate VMs in these files.

## Removal of Restore Points from Reverse Incremental Chains

In the case of a reverse incremental backup chain, Veeam Backup & Replication immediately removes a redundant restore point when the allowed number of restore points is exceeded. For more information, see Reverse Incremental Backup Retention Policy.

For example, a backup job processes two VMs: *VM 1* and *VM 2*. According to the retention policy settings, the backup chain must contain 5 restore points. The backup job ran 5 times and VMs were processed in the following way:

- *VM 1* was successfully backed up 5 times and has 5 valid restore points.
- *VM 2* was not backed up in 2 job sessions and has 3 valid restore points.

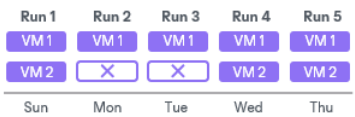

After that, Veeam Backup & Replication will run a new backup job session in which *VM 1* and *VM 2* will be successfully processed. When a new restore point is added to the chain, Veeam Backup & Replication will remove the earliest restore point because the number of restore points in the backup chain exceeds 5. As a result, you will have 5 restore points for *VM 1* and 3 restore points for *VM 2*.

# Long-Term Retention Policy (GFS)

The long-term or Grandfather-Father-Son (GFS) retention policy allows you to store backup files for long periods of time — for weeks, months and even years. For this purpose, Veeam Backup & Replication does not create any special new backup files — it uses backup files created while the backup job runs and marks these backups with specific GFS flags.

To mark a backup file for long-term retention, Veeam Backup & Replication can assign to the file the following types of GFS flags: weekly (W), monthly (M) and yearly (Y). The types of GFS flags that Veeam Backup & Replication assigns depend on the configured GFS retention policy settings. Depending on which flag is assigned to a backup, it will be stored for a specified number of weeks, months or years.

> **IMPORTANT**
>
> Consider the following:
>
> - GFS flags can be assigned only to full backup files created during the time period specified in GFS policy settings.
> - For backups located in object storage repositories, GFS flags are assigned to active full backup files that match the established GFS policy. If you enable the GFS retention policy but do not schedule active full backups, or the active full backup does not match the GFS policy, Veeam Backup & Replication will generate a synthetic full backup and assign a GFS flag to it.

As soon as Veeam Backup & Replication assigns a GFS flag to a full backup file, this backup file can no longer be deleted or modified. Also, Veeam Backup & Replication does not apply short-term retention policy settings to the full backup file — that is, Veeam Backup & Replication ignores the backup file when determining whether the number of allowed backup files is exceeded. For more information on how GFS flags are assigned, see Assignment of GFS Flags.

When the specified retention period ends, Veeam Backup & Replication unassigns the GFS flag from the full backup file. If the backup file does not have any other GFS flags assigned, it can be modified and deleted according to the short-term retention policy. For more information on when Veeam Backup & Replication removes GFS flags, see Removal of GFS Flags.

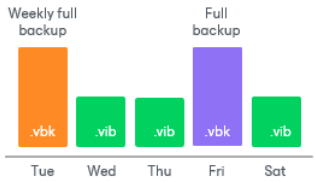


# Limitations

When planning to use GFS retention policy, consider the following limitations:

- GFS retention policy does not apply to reverse incremental backup chains.

- GFS retention policy applies to forever forward incremental backup chain only if you periodically create full backups manually or using scheduled scripts. For more information on cmdlets that you can use in scripts, see the Veeam PowerShell Reference.
- If you store backups with GFS flags in the capacity tier with immutability enabled, the GFS retention policy is not considered. Immutability period depends on the object storage repository settings.
- As Veeam Backup & Replication does not create new full backup files while applying the GFS retention policy, you must configure your backup jobs in a way you do not lose any essential data due to an insufficient number of full backup files. For example, if you configure monthly GFS retention, you need at least one full backup file per month.
- If a GFS flag is assigned to a full backup file in an active backup chain, Veeam Backup & Replication is not able to merge data from incremental backup files into the full backup file. For forever forward incremental backup chain, this means, that the short-term retention policy does not apply.
- Veeam Backup & Replication assigns GFS flags only after you save GFS retention policy settings. This means that GFS flags are assigned only to those backup files created after the configuration, while backup files created earlier are not affected and previously assigned flags are not modified.
- You cannot store full backups to which GFS flags are assigned in backup repositories with rotated drives.
- Retention policy for deleted items does not apply to full backup files to which GFS flags are assigned.

## Assignment of GFS Flags

When configuring GFS retention policy settings, you can choose a number of GFS flag types that Veeam Backup & Replication will use to mark backup files for long-term retention. Depending on this number, Veeam Backup & Replication will apply slightly different algorithms when assigning GFS flags:

- Algorithm for One Flag Type
- Algorithm for Multiple Flag Types

> **IMPORTANT**
>
> When you reconfigure the GFS retention policy and save it, the following applies:
>
> - If you change the amount of time for which backups with GFS flags must be kept, the previously created GFS backups will also be kept according to the new settings.
> - If you change the time period when GFS flags must be assigned, the previously assigned GFS flags are not considered when determining whether new GFS flags must be assigned.

### Algorithm for One Flag Type

If you select only one type of GFS flag when configuring retention policy settings for a backup job, Veeam Backup & Replication assigns the flags depending on the scheduled period. The scheduled period is a period when a new GFS flag must be assigned and is set according to the following schema: for weekly GFS — the day selected in the GFS retention policy settings; for monthly GFS — the selected week; for yearly GFS — the selected month.

Depending on the scheduled period settings, Veeam Backup & Replication performs the following steps as soon as the job finishes:

- If the job finishes within the period when a new GFS flag must be assigned (the scheduled period), Veeam Backup & Replication checks whether the backup job has created a full backup file.
  - Yes, the full backup has been created: If the GFS flag has already been assigned to another backup file during the scheduled period, Veeam Backup & Replication does not assign another GFS flag to the backup file. If the GFS flag has not been assigned, Veeam Backup & Replication assigns it and stops waiting for the full backup file (removes the waiting flag).
  - No, the full backup has not been created: Veeam Backup & Replication waits for a full backup file to be created (sets the waiting flag).
- If the job finishes outside the scheduled period, Veeam Backup & Replication checks whether the full backup file is being waited for (check whether the waiting flag is set).
  - Yes, the waiting flag is set: Veeam Backup & Replication checks whether the backup job has created a full backup file.
    - Yes: Veeam Backup & Replication assigns the GFS flag and stops waiting for the full backup file (removes the waiting flag).
    - No: Veeam Backup & Replication does not assign a new GFS flag.
  - No, the waiting flag is not set: Veeam Backup & Replication does not assign a new GFS flag.

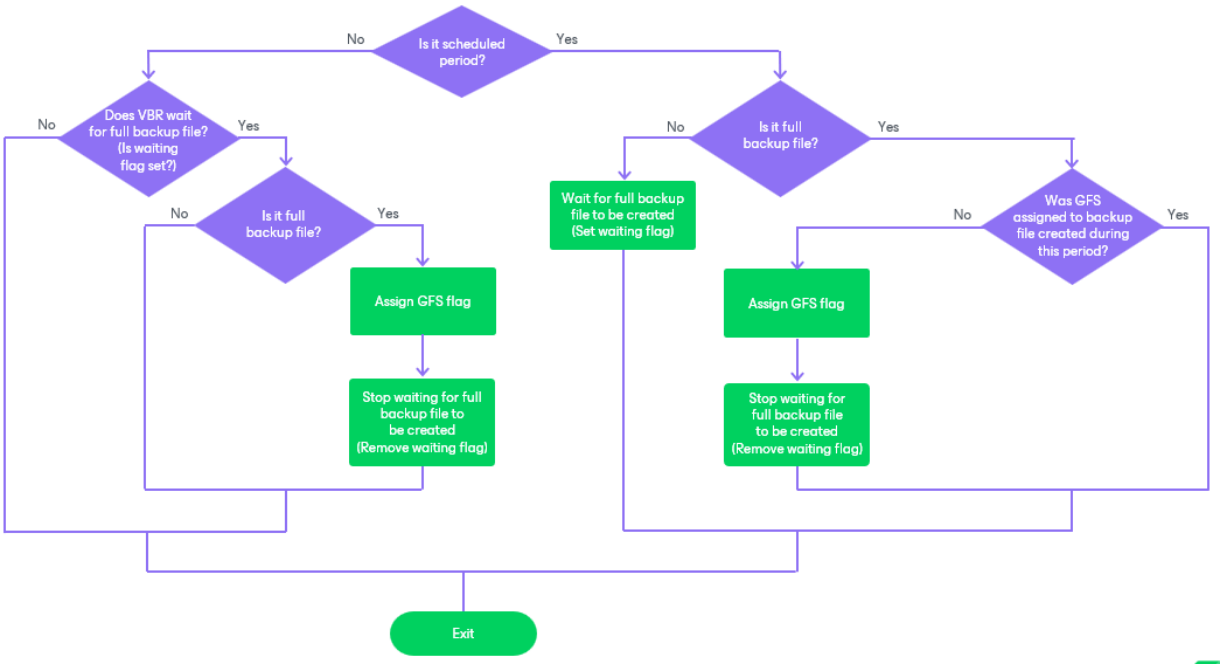


Consider the following example. On Monday, you configure the GFS policy settings of a backup job in a way weekly GFS flags must be assigned every Wednesday. In this example, Veeam Backup & Replication will take the following steps.

### Example Description

1. Until Wednesday, Veeam Backup & Replication will not assign any GFS flags because the scheduled period has not started yet.
2. On Wednesday, the backup job will produce an incremental backup file. Veeam Backup & Replication will start waiting for a full backup file to be created since the scheduled period is now started.

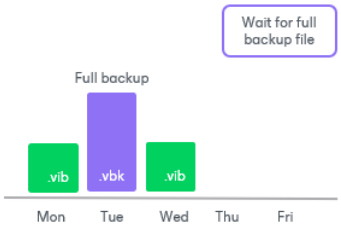


3. On Thursday, the backup job will produce another incremental backup file, while Veeam Backup & Replication will still be waiting for a full backup file.

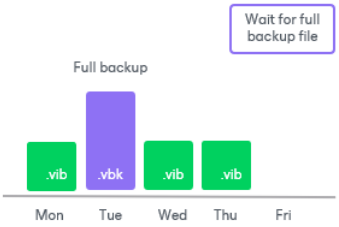


4. On Friday, the backup job will produce a full backup file, and Veeam Backup & Replication will immediately assign the weekly GFS flag to the backup file.

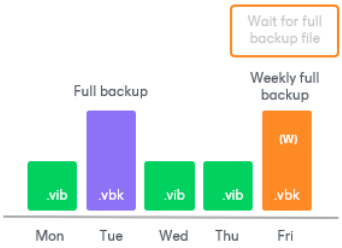


## Algorithm for Multiple Flag Types

If you select multiple types of GFS flags when configuring retention policy settings for a backup job, GFS flags depend on each other. Yearly flags (high level) depend on monthly flags, monthly flags depend on weekly flags, and weekly flags (low level) do not depend on any flags. This means that Veeam Backup & Replication can only assign GFS flags of a higher level to backup files with GFS flags of a lower level — this mechanism helps you save space in the backup repository.

> **IMPORTANT**
>
> The described mechanism does not apply to a situation where you select only two types of GFS flags: yearly and weekly. If you select these two types and do not select the monthly type, Veeam Backup & Replication treats the case as if you select only one flag type. It means that flags of each type are assigned according to the algorithm for one flag type.

Veeam Backup & Replication performs the following steps as soon as the job finishes and performs them for GFS flags of each level. Note that GFS flags of lower levels are processed before flags of higher levels. In the described algorithm, the "scheduled period" is the following: for weekly GFS — the day selected in the GFS retention policy settings; for monthly GFS — the selected week; for yearly GFS — the selected month.

- If the job finishes within the period when a new GFS flag must be assigned, Veeam Backup & Replication checks whether the flag of a lower level is configured in the job settings.
    - Yes, the flag of a lower level is configured: If the GFS flag of a lower level has already been assigned to the backup file created by the job, Veeam Backup & Replication assigns the GFS flag of the current level and stops waiting for the flag of the current level to be assigned (removes the current-level waiting flag). If the GFS flag of a lower level has not been assigned, Veeam Backup & Replication checks whether the GFS flag of the current level has already been assigned to another backup file during the scheduled period:
        - Yes, the GFS flag of the current level has already been assigned: Veeam Backup & Replication does not assign any GFS flags.
        - No, the flag has not been assigned: Veeam Backup & Replication waits for the flag of the current level to be assigned (sets the current-level waiting flag).
    - No, the flag of a lower level is not configured: Veeam Backup & Replication checks whether the backup job has created a full backup file.
        - Yes: If the GFS flag of the current level has already been assigned to another backup file during the scheduled period, Veeam Backup & Replication does not assign the GFS flag of the current level to the backup file. If the GFS flag has not been assigned, Veeam Backup & Replication assigns it and stops waiting for the GFS flag of the current level to be assigned (removes the current-level waiting flag).
        - No: Veeam Backup & Replication waits for a full backup file to be created (sets the current-level waiting flag).
- If the job finishes outside the scheduled period, Veeam Backup & Replication checks whether the GFS flag of the current level is being waited for (checks whether the current-level waiting flag is set).
    - Yes, the flag of the current level has already been assigned: Veeam Backup & Replication does not assign any GFS flags.
    - No, the flag has not been assigned: Veeam Backup & Replication checks whether the backup job has created a full backup file.
        - Yes: Veeam Backup & Replication assigns the flag of the current level and stops waiting for the GFS flag of the current level to be assigned (removes the current-level waiting flag).

- No: Veeam Backup & Replication does not assign any GFS flags.

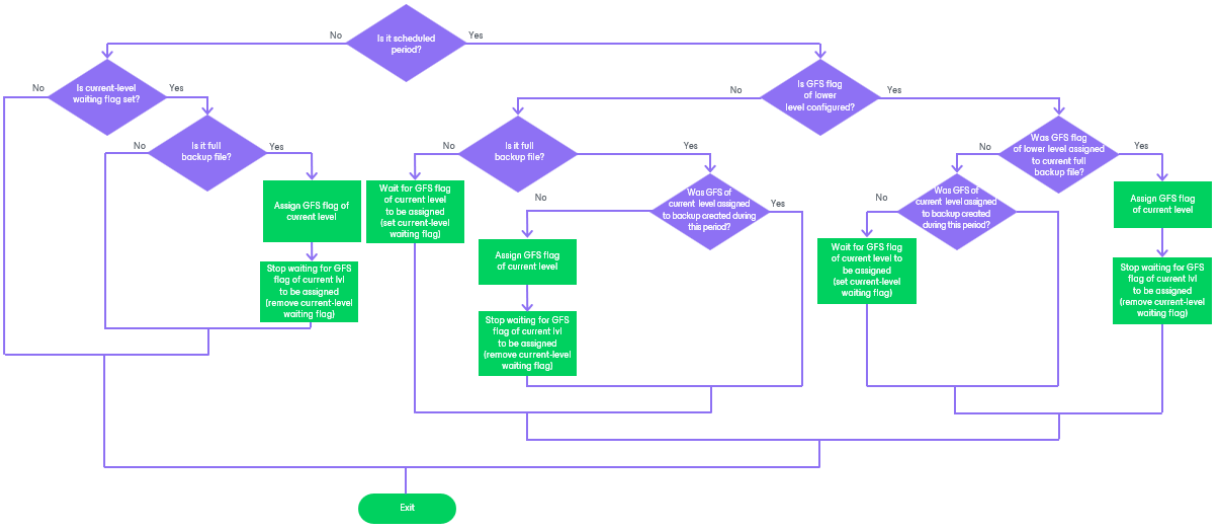


Consider the following example. On Monday, you configure GFS policy settings of a backup job in a way weekly GFS flags must be assigned every Wednesday and monthly GFS flags must be assigned every first week of a month. In this example, Veeam Backup & Replication will take the following steps.

### Example Description

1. On Monday, the backup job will produce an incremental backup file. Veeam Backup & Replication will start waiting for a full backup file since the monthly scheduled period is now started.
2. On Tuesday, the backup job will produce a full backup file. Although Veeam Backup & Replication will still be waiting for a full backup file, the full backup file with weekly GFS flag is required.

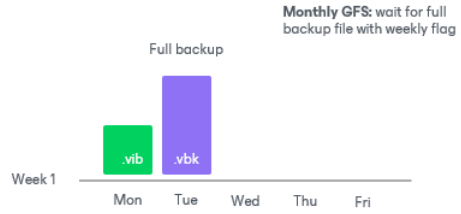


3. On Wednesday, the backup job will produce an incremental backup file.
4. On Thursday, the backup job will produce an incremental backup file. Veeam Backup & Replication will start waiting for a full backup file to assign the weekly GFS flag to it since the weekly scheduled period is now started.

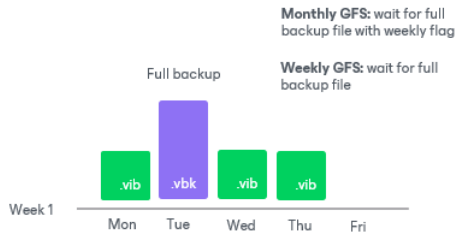


5. On Friday, the backup job will produce a full backup file. Veeam Backup & Replication will immediately assign the weekly GFS flag to the backup file. As the weekly GFS flag will have been assigned, Veeam Backup & Replication will also assign the monthly GFS flag to the backup file.

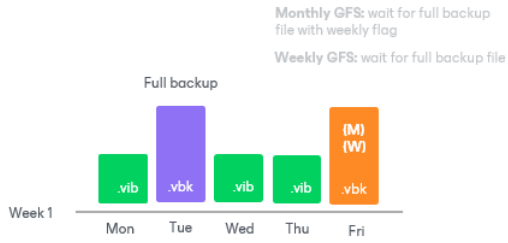


## Removal of GFS Flags

When configuring GFS retention policy settings, you can specify the retention period for each type of GFS flag. After the specified retention period exceeds, Veeam Backup & Replication removes GFS flags.

The date when Veeam Backup & Replication can remove the GFS flag is calculated by the following formulas:

- Weekly: `date of GFS flag assignment + N * 7 days`
- Monthly: `date of GFS flag assignment + N months + 1 day`

  When calculating the `date of GFS flag assignment + N months`, Veeam Backup & Replication increases the month ordinal number by `N`. If the calculated date does not exist, Veeam Backup & Replication uses the last day of the calculated month.
- Yearly: `date of GFS flag assignment + N years + 1 day`

Where N is the value specified in the **Keep weekly/monthly/yearly full backups for** field. For more information, see Configure Long-Term Retention.

> **NOTE**
>
> Consider the following:
>
> - Veeam Backup & Replication removes GFS flags during running backup job sessions. This means that if the backup job does not run on the calculated date, Veeam Backup & Replication will remove the GFS flag later during the next job session.
> - GFS flags can also be removed during the background GFS retention process. It detects backup files that have assigned GFS flags and removes flags with an expired retention period. For more information, see Background GFS Retention.
> - After you change GFS retention policy settings, the date of GFS flag removal is recalculated for already created restore points.

Consider the following example. At the beginning of January, you create a backup job whose GFS retention policy settings are configured to assign monthly GFS flags. You want to keep backup files with monthly flags for 1 month and set the value of the **Keep monthly full backups for** field to 1. Veeam Backup & Replication will perform the following steps to assign and remove the flags.

1. Veeam Backup & Replication will assign the monthly GFS flag on `1/31/2019`.
2. To calculate the date when the monthly flag must be removed, the following formula is used: `date of GFS flag assignment + 1 month`. This means that the flag must be removed on `2/31/2019`. However, this date does not exist since the last date of February is `2/28/2019`. That is why Veeam Backup & Replication will remove the GFS flag on `3/1/2019` (which is `2/28/2019 + 1 day`).

# Background Retention

In addition to applying a retention policy (short-term retention and long-term retention) within a job session, Veeam Backup & Replication performs background retention for backups. The background retention aims mostly at backups that are no longer processed by jobs (orphaned backups shown in the node with the **(Orphaned)** postfix). However, this retention can also be helpful for standard backups, in case backups are created by jobs without a schedule, the job retention has not been applied yet or failed for some reason, and so on.

The background retention starts automatically every 24 hours at 00:30, runs in the background and consists of the following activities:

- Background basic retention
- Background GFS retention
- Background log retention
- Backup cleanup
- Deleted Agent retention

# Limitations

The background retention does not apply to the following backups:

- Backups stored on backup repositories with rotated drives.
- Backups stored in the archive tiers of scale-out backup repositories.
- Backups stored in the capacity tier. For details, see section Considerations.

- Imported backups.
- Unstructured Data Backup.
- Replicas (including CDP replicas).
- VeeamZIP backups.
- Exported backups.
- Copied backups.
- Backups exported by Kasten policies.
- Backups created by Veeam Plug-ins for Enterprise Applications and Veeam Cloud Plug-ins (Veeam Backup for AWS, Veeam Backup for Google Cloud, Veeam Backup for Microsoft Azure).

# Considerations

Consider the following:

- Like the job retention, the background retention cannot delete immutable backup files. The background retention waits until the immutability and retention periods end for these files.
- [For backups linked to jobs] If the retention is set in days, Veeam Backup & Replication leaves at least 3 backup files in a backup chain regardless of the set retention. If the retention is set in restore points, Veeam Backup & Replication leaves backup files in a backup chain. The minimum number of backup files left equals the current retention period. You can delete these backup files manually as described in section Deleting Backups from Disk.
- [For orphaned backups] If the retention is set in days, Veeam Backup & Replication can remove all outdated backup files in a backup chain. If the retention is set in restore points, Veeam Backup & Replication leaves backup files in a backup chain. The minimum number of backup files left equals the current retention period. You can delete these backup files manually as described in section Deleting Backups from Disk.
- [For backups stored in the capacity tier] Background retention job does not delete capacity tier copies of backup data directly. However, if background retention removes local copies of backups, they may also be marked for removal on the capacity tier. In such a case, cleanup during the next SOBR offloading session will remove them from the capacity tier.
- [For backups created with Veeam Agent for Windows operating in standalone or managed by Veeam Agent mode] For backups linked to jobs basic retention does not apply.
- Unlike the job retention, the background retention does not merge data from one backup file to another; it just deletes files. In the case of forever forward incremental and forward incremental backup chains, the background retention deletes incremental files only after the last increment in the related part of the backup chain becomes outdated.
- The background retention does not delete backup files if they are locked by other processes. The retention waits until the backup file is unlocked.
- You can launch the background retention manually as described in section Launching Background Retention.

# Background Basic Retention

Background basic retention process analyzes the retention period set for backup files. In the case of orphaned backups, Veeam Backup & Replication analyzes the retention period that was set in the job that created this backup. If the retention period has expired, Veeam Backup & Replication removes the backup files.

# Background GFS Retention

Background GFS retention process detects backup files with GFS flags and analyzes their retention period. If the retention period for GFS flags of such backup files has expired, Veeam Backup & Replication removes GFS flags. Then Veeam Backup & Replication deletes these backup files according to the short-term retention policy if the following conditions are met:

- The backup file does not have any other GFS flags assigned.
- The backup file does not have any dependent files.

Background retention job applies to backup files created by all types of jobs that use GFS. For more information about GFS retention policy, see Long-Term Retention Policy (GFS).

# Background Log Retention

Background log retention detects logs whose related image-level backup was deleted by background retention. Then Veeam Backup & Replication deletes these log files.

# Backup Cleanup

Backup cleanup applies only to orphaned backups. After the whole orphaned backup chain is deleted, Veeam Backup & Replication also deletes the .VBM file and the folder where the backup chain was stored.

# Deleted Agent Retention

Deleted agent retention is a process that detects and deletes outdated backup files according to Veeam Agent retention policy for outdated backups.

An outdated backup file is a backup file for which no new restore points were created and no new backup job sessions were started within the last N days, where N is the retention period specified in the Veeam Agent backup job settings.

For more information, see Retention Policy for Outdated Backups article in the Veeam Agent for Microsoft Windows User Guide and Maintenance Settings article in the Veeam Agent Management Guide.

# Cloud Connect Background Retention

In the Veeam Cloud Connect infrastructure, independent retention is applied to tenant backups on the service provider side. The background retention job runs automatically in Veeam Backup & Replication on the service provider backup server according to the same rules as in the regular Veeam backup infrastructure. The service provider can also launch background retention manually as described in section Launching Background Retention.

On the tenant side, Veeam Cloud Connect repositories are skipped from processing by the background retention job launched automatically and manually. For more information, see the Background Retention for Tenant Backups section in the Veeam Cloud Connect Guide.

# Changed Block Tracking

To perform incremental backup, Veeam Backup & Replication needs to know what data blocks have changed since the previous job session.

For VMware VMs with hardware version 7 and later, Veeam Backup & Replication employs a native VMware vSphere feature — VMware vSphere Changed Block Tracking (CBT). Instead of scanning VMware VMFS (Virtual Machine File System), Veeam Backup & Replication queries CBT through VMware VADP (VMware vStorage API for Data Protection) and gets a list of blocks that have changed since the last job session. The use of CBT increases the speed and efficiency of block-level incremental backups.

Veeam Backup & Replication uses CBT for the following operations:

- Backup
- Replication
- Entire VM restore
- VM disk restore

Veeam Backup & Replication enables CBT by default. If necessary, you can disable it in job settings.

> **NOTE**
>
> If you back up a VM to which the backup proxy role is assigned and that uses Virtual appliance (HotAdd) transport mode, CBT for this VM is disabled and cannot be enabled.
>
> Also, ensure the virtual machine has no snapshots before enabling VMware vSphere CBT. Once you enable CBT, snapshots can be created on the virtual machine and backed up. For more information, see this KB article.

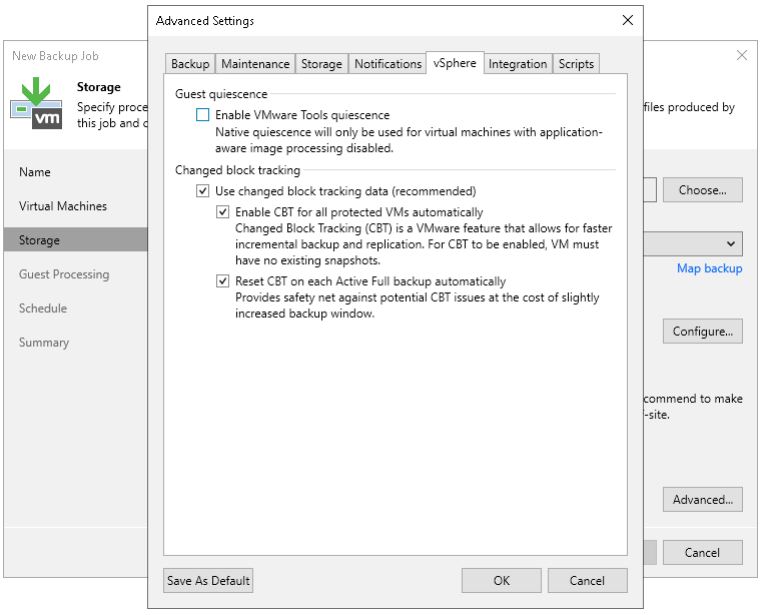

> **NOTE**
>
> For VMs with virtual disks in thin format, Veeam Backup & Replication also uses CBT during active full backup sessions to detect unallocated regions of virtual disks and skip them. For VMs with virtual disks on an NFS datastore, Veeam Backup & Replication uses CBT as well but cannot leverage CBT on the first full run (for more information, see this KB article).

In some situations, Veeam Backup & Replication cannot leverage VMware vSphere CBT, for example, if VMs run an earlier version of virtual hardware. If Veeam Backup & Replication cannot leverage VMware vSphere CBT, it fails over to the Veeam proprietary filtering mechanism. Instead of tracking changed blocks of data, Veeam Backup & Replication filters out unchanged data blocks.

During VM processing, Veeam Backup & Replication consolidates virtual disk content, scans through the VM image and calculates a checksum for every data block. Checksums are stored as metadata to backup files next to VM data. When incremental backup is run, Veeam Backup & Replication opens all backup files in the chain of previous full and incremental backups, reads metadata from these files and compares it with checksums calculated for a VM in its current state. If a match is found (which means the block already exists in the backup), this block is filtered out.

# Data Compression and Deduplication

Veeam Backup & Replication provides mechanisms of data compression and deduplication. Data compression and deduplication let you decrease traffic going over the network and disk space required to store backup files and VM replica files.

## Data Compression

Data compression decreases the size of created files but affects the duration of the backup or replication procedure. Veeam Backup & Replication allows you to select one of the following compression levels:

- **None** compression level is recommended if you plan to store backup files and VM replica files on storage devices supporting hardware compression and deduplication. Disabled compression can reduce performance due to the increased amount of data to be transferred.
- **Dedupe-friendly** compression level is recommended for some deduplicating storage appliances and when external WAN accelerators are used.
- **Optimal** compression level is the recommended compression level. It provides the best ratio between compression and performance, the lowest backup proxy CPU usage and the fastest restore.
- **High** compression level provides up to 60% of additional data reduction over the **Optimal** compression level at the cost of 2x higher CPU usage and 2x slower restore.
- **Extreme** compression level provides up to 33% of additional data reduction over the **High** compression level at the cost of 5x higher CPU usage.
- [For backup copy job] **Auto** is the recommended compression level for backup copy jobs. Select this level to use the compression settings of the copied backup files.

> **NOTE**
>
> If encryption is enabled for a job and the **Decompress backup data blocks before storing** check box is selected in the settings of the target backup repository, Veeam Backup & Replication does not compress VM data. Therefore, in the job statistics, you may observe a higher amount of transferred data (the **Transferred** counter) as compared to a job for which encryption is disabled. For more information on job statistics, see Viewing Real-Time Statistics.
>
> In the backup properties created with encryption, you may also see that the backup size (the **Backup Size** column) is larger than the original VM size (the **Original Size** column). For more information on backup properties, see Viewing Backup Properties.

The decreased size of created files is not guaranteed, as the effectiveness of compression and deduplication depends on various factors, such as the type of source data and the target file system.

### Changing Data Compression Settings

You can change data compression settings for existing jobs. After you change the settings, you will not need to create new full backups to use the new settings. Veeam Backup & Replication will automatically apply the new compression level to the newly created backup files after you save the settings. Previously created backup files will not be affected.

However, if you use the reverse incremental backup method, the newly created backup files will contain a mixture of data blocks compressed at different levels. For example, you have a backup job that uses the reverse incremental backup method and the Optimal level of compression. After several job sessions, you change the compression level to High. In the reverse incremental backup chains, the full backup file is rebuilt with every job session to include new data blocks. As a result, the full backup file will contain a mixture of data blocks: data blocks compressed at the Optimal level and data blocks compressed at the High level. The same behavior applies to synthetic full backups: synthetic full backups created after the compression level change will contain a mixture of data blocks compressed at different levels.

If you want the newly created backup file to contain data blocks compressed at one level, you can create an active full backup. Veeam Backup & Replication will retrieve data for the whole VM image from the production infrastructure and compress it at the new compression level. All subsequent backup files in the backup chain will also use the new compression level.

# Deduplication

Data deduplication decreases the size of files. With data deduplication enabled, Veeam Backup & Replication does not store the resulting file identical data blocks and space that have been pre-allocated but not used.

We recommend enabling data deduplication for your backup jobs if they contain several VMs with a lot of free space on their logical disks or VMs with similar data blocks — for example, VMs created from the same template. Data deduplication optimizes storage usage within each individual backup job. However, note that it may decrease job performance.

Veeam Backup & Replication uses Veeam Data Movers to deduplicate VM data:

- Veeam Data Mover in the source side deduplicates VM data at the level of VM disks. Before the source-side Veeam Data Mover starts processing a VM disk, it obtains digests for the previous restore point in the backup chain from Veeam Data Mover in the target side. The source-side Veeam Data Mover consolidates this information with CBT information from the hypervisor and filters VM disk data based on it. If some data block exists in the previous restore point for this VM, the source-side Veeam Data Mover does not transport this data block to the target. In addition, in the case of thin disks, the source-side Veeam Data Mover skips unallocated space.
- Veeam Data Mover in the target side deduplicates VM data at the level of the backup file. It processes data for all VM disks of all VMs in the job. The target-side Veeam Data Mover uses digests to detect identical data blocks in transported data and stores only unique data blocks in the resulting backup file.

You can change data deduplication settings for existing jobs. After you change the settings, you will not need to create new full backups to enable or disable the deduplication. Veeam Backup & Replication will automatically apply the change to the newly created backup files after you save the settings. Previously created backup files will not be affected.

# Storage Optimization

To optimize job performance and storage usage, Veeam Backup & Replication allows you to choose the minimum data block size to process VMs. The optimal data block size depends on the type of storage you select as a target and the size of your files.

When selecting the data block size, consider the following aspects:

- When reading the VM image, Veeam Backup & Replication "splits" the VM image into blocks of the selected size. The more data blocks there are, the more time is required to process the VM image.
- [For replication and VMware Cloud Director replication] Veeam Backup & Replication writes information about every data block to the VM replica metadata stored in the backup repository. The more data blocks there are, the more metadata is written to the backup repository.

- [For changed block tracking enabled] During incremental job runs, Veeam Backup & Replication uses CBT to define changed data blocks in the VM. The larger the size of the found changed data block, the greater the amount of data transferred to the target site.

The incorrectly chosen data block size may decrease the performance. For example, when you deduplicate a large backup file to small data blocks, Veeam Backup & Replication produces a very large deduplication metadata table, which can potentially overgrow the memory and CPU resources of your backup repository. For large backup files, it is better to use large data blocks.

Veeam Backup & Replication provides several storage optimization options with different block sizes used. The following table will help you to choose the optimal option according to the size of your backup files and the storage type.

| Storage optimization option | Description |
|---|---|
| 4 MB (former **Local target (large blocks)**) | Recommended for backup files that are larger than 16 TB.<br>This option provides the lowest deduplication ratio and the largest size of incremental files. |
| 1 MB (former **Local target**) | Recommended for backup and replication to SAN, DAS or local storage.<br>This option provides the fastest job performance but reduces the deduplication ratio because, with larger data blocks, it is less likely to find identical blocks. |
| 512 KB (former **LAN target**) | Recommended for backup and replication to NAS, and on-site backup and replication.<br>This option provides a better deduplication ratio and reduces the size of a backup file because of reduced data block sizes. |
| 256 KB (former **WAN target**) | Recommended if you are planning to use WAN for off-site backup and replication.<br>This option provides the maximum deduplication ratio and the smallest size of backup files that allow you to reduce the amount of traffic over WAN. However, note that various factors can affect the overall size of the backup files. As a result, the smallest file size may not always be achieved. |

## Changing Storage Optimization Settings

You can change storage optimization settings for existing jobs. New settings will not have any effect on previously created files in the chain. They will be applied to new files created after the settings were changed.

For Veeam Backup & Replication to apply the new settings, use the following instructions.

### Backup Jobs

To apply new storage optimization settings in backup jobs, you must create an active full backup after you change storage optimization settings. Veeam Backup & Replication will use the new block size for the active full backup and subsequent backup files in the backup chain.

### Backup Copy Jobs

To change data block size for backup copy jobs, you must perform the following actions:

1. Change the data block size in the initial backup job settings.
2. Create an active full backup with the initial backup job.
3. Create an active full backup with the backup copy job.

# Data Exclusion

When you configure a backup or replication job, you can define what data you want to back up and replicate and exclude data you do not need. Data exclusion helps reduce the size of the VM backup or replica and decrease the load on the network.

You can exclude data at the VM level and at the VM guest OS level.

At the VM level:

- VMs added as part of the container
- VM disks
- VM templates (only for backup)

At the VM guest OS level:

- Swap files on the VM guest OS
- Deleted file blocks on the VM guest OS (BitLooker)
- Files and folders on the VM guest OS

> **NOTE**
>
> To reduce the size of the backup file, Veeam Backup & Replication automatically excludes VM log files from processing.

## VMs and VM Disks

When you configure a backup or replication job, you can exclude the following objects from processing:

- VMs added as a part of a VM container
- Individual VM disks
- [For backup jobs] VM templates

## VMs as Part of Container

If you want to back up or replicate a VM container that holds several VMs but want to skip some VMs, you can exclude specific VMs from the job processing. This option will help you reduce the size of the resulting backup or replica and increase the job performance.

You can define which VMs you want to skip at the **Virtual Machines** step of the backup or replication job wizard.

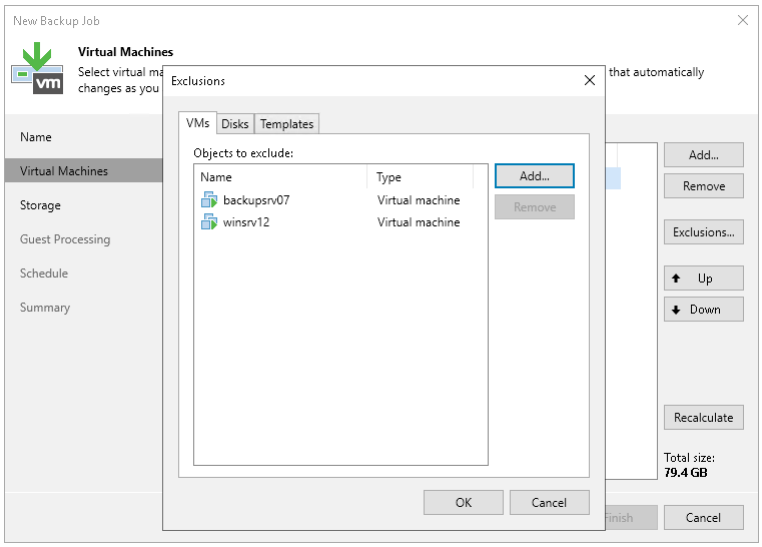


# Individual VM Disks

You can choose what VM disks you want to back up or replicate:

- All VM disks
- 0:0 disks (which are commonly the VM system disks)
- Specific IDE, SCSI or SATA disks

For example, you may want to back up or replicate only the system disk instead of creating a backup or replica of a full VM. VM disks exclusion reduces the size of the backup or replica.

You can define which VM disks you want to back up or replicate at the **Virtual Machines** step of the backup or replication job wizard. You can specify disk processing settings granularly for every VM in the job or for the whole VM container. In the latter case, Veeam Backup & Replication will apply the configured rule to all VMs in this container.

You can additionally instruct Veeam Backup & Replication to modify the configuration file of the VM. When you start a VM from the backup or failover to the VM replica, you will be able to use such a VM immediately. You will not have to edit its configuration file and remove excluded disks from it.

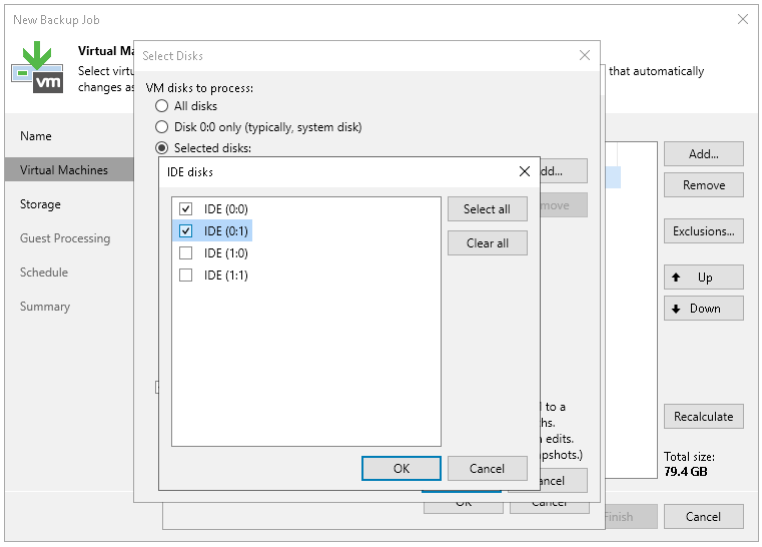

# VM Templates

You can include VM templates in the backup. Backing up VM templates provides additional safety of your production environment but requires additional space in the backup repository.

Veeam Backup & Replication allows you to include a VM template only in the full backup and omit it in incremental backups. Note that Veeam Backup & Replication cannot use Direct SAN transport mode to back up VM template data. If you have configured failing over to the Network mode, the Network mode will be used.

You can define how Veeam Backup & Replication must process VM templates at the **Virtual Machines** steps of the wizard.

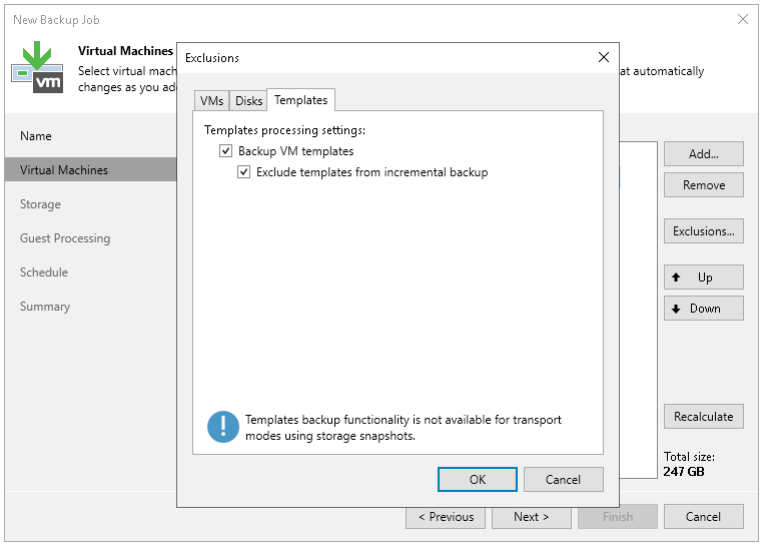


# Deleted File Blocks (BitLooker)

Deleted file block exclusion is a technology that reduces the backup size and time needed to create a backup. As a rule, file systems do not zero out blocks of permanently deleted files; they only remove information about such files from the file allocation table. It means the file content still exists on the disk image and enlarges the backup size. Deleted file block exclusion helps avoid copying this unnecessary data or, in Veeam terms, "dirty" blocks.

By default, deleted file block exclusion is enabled. If you do not want to exclude deleted file blocks from backups or replicas, you can disable the **Exclude deleted file blocks** option in the backup or replication job settings.

> **NOTE**
>
> If you enable or disable the **Exclude deleted file blocks** setting for the existing job, Veeam Backup & Replication will apply the new setting from the next job session.

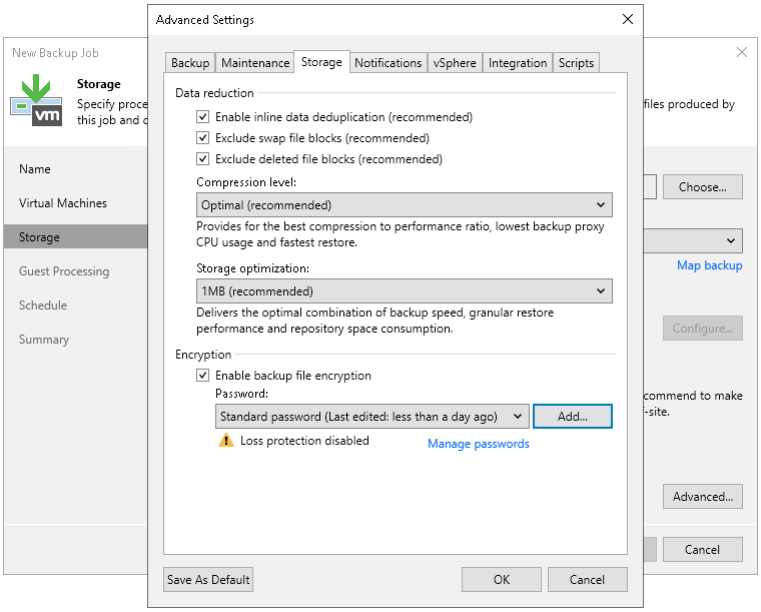


# How Deleted File Block Exclusion Works

With this option enabled, Veeam Backup & Replication performs the following operations during the job session:

1. Veeam Backup & Replication accesses the MFT file on the VM guest OS to identify deleted file blocks and zeros out these blocks.
2. Veeam Backup & Replication processes and transports data blocks of the VM image in the following manner:
   - If a data block of the VM image contains only the deleted file blocks, Veeam Backup & Replication does not read this data block from the source datastore.
   - If a data block of the VM image contains zeroed-out blocks and other data, Veeam Backup & Replication copies this block to the target. Due to data compression, data blocks that are marked as deleted are compressed, and the size of the resulting backup or replica file reduces.

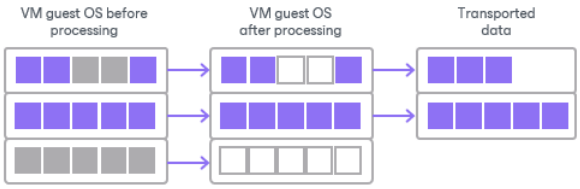


# Limitations for Deleted File Blocks Exclusion

Deleted file blocks exclusion has the following limitations:

- Veeam Backup & Replication can exclude deleted file blocks only on the VM guest OS with Microsoft NTFS.
- Veeam Backup & Replication supports both basic and dynamic disks. For the dynamic disks, simple type of volumes is supported. Spanned, mirrored and striped volumes are not supported.

# Swap Files

You can instruct Veeam Backup & Replication to exclude `pagefile.sys` and `hiberfil.sys` files from backups or replicas of Microsoft Windows VMs. As a result, the size of incremental backups and replicas will be smaller.

- `hiberfil.sys` is a system file created by the OS for the correct work of the hibernate mode.
- `pagefile.sys` is a swap file. Swap files are dynamic and can change intensively between job sessions, even if a VM itself does not change much.

# How Swap File Exclusion Works

When you exclude `pagefile.sys` and `hiberfil.sys` files, Veeam Backup & Replication performs the following operations during the job session:

1. Veeam Backup & Replication accesses the MFT file on the VM guest OS to identify data blocks of `pagefile.sys` and `hiberfil.sys` files and zeros them out.
2. Veeam Backup & Replication processes and transports data blocks of the VM image in the following manner:
   - If a data block of the VM image contains only blocks of these files, Veeam Backup & Replication does not copy this data block to the target.
   - If a data block of the VM image contains blocks of these files and other data, Veeam Backup & Replication copies this block to the target.

# Limitations for Swap File Exclusion

Veeam Backup & Replication can exclude blocks of `pagefile.sys` and `hiberfil.sys` files only on the VM guest OS with Microsoft Windows NTFS.

# VM Guest OS Files

If you do not want to back up or replicate some files and folders on the VM guest OS, you can exclude them from the backup or replica. File exclusion reduces the size of the backup or replica but may affect job performance.

You can specify file exclusion settings granularly for every VM in the job or for the whole VM container. In the latter case, Veeam Backup & Replication will apply the configured rule to all VMs in this container.

To define which VM guest OS files must and must not be processed, you can use the following options:

- Disable file exclusion.

  Veeam Backup & Replication will back up or replicate the whole content of the VM guest file system.

- Exclude specific files and folders from the backup or replica.

  Veeam Backup & Replication will back up or replicate all files and folders except the specified ones.

- Include only specific files and folders in the backup or replica.

  Veeam Backup & Replication will back up or replicate only the specified files and folders.

# Defining File Exclusions and Inclusions

To define a list of exclusions or inclusions, you can use the following methods:

- Specify a full path to a folder on the VM guest OS, for example, `C:\Documents\`.
- Specify a full path to a file on the VM guest OS, for example, `C:\Documents\MyReport.docx`.

  If a path is not full, Veeam Backup & Replication will expand it to the root directory on the computer volume and attempt to detect such files on all computer volumes. For example, you have `C`, `D` and `E` disks on the VM. In the list of exclusions, you specify `Document.docx`. Veeam Backup & Replication will scan the whole file system and exclude the following files (if any): `C:\Document.docx`, `D:\Document.docx`, `E:\Document.docx`. If there is a `C:\MyDocuments\Document.docx` file, it will not be excluded — this file is not located in the root directory.

- Use environmental variables, for example, `%TEMP%`, `%windir%`.

  Environment variables must be defined for the user account that you use to connect to the VM guest OS and under which the non-persistent runtime components or persistent agent components are started. For example, you connect to the VM guest OS under the *Administrator* account. If you want to use the `%windir%` variable in the list of exclusions or inclusions, you must ensure that the `%windir%` variable is added to the list of user variables for *Administrator* on the VM guest OS.

- Use file masks. You can use the following characters for masks:
  - (*) — a substitution for one or more characters in the file name or path. It can be used for any sequence of characters (including no characters). For example, `*.pdf`.
  - (?) — a substitution of one character in the file name or path. For example, `repor?.pdf`.
  - (;) — mask separator, for example, `report.*;reports.*`.

  In the following table, the mask stands for any sequence of characters.

| Mask format | Affects paths/files |
|---|---|
| `*mask*` | All paths that contain the given sequence. |
| `mask*` | If the asterisk character (*) is not specified at the beginning of the mask, the mask will be applied to all volumes on the VM guest OS, and Veeam Backup & Replication will include/exclude files and folders in the root folder on the volume: `A:\mask*, B:\mask*, …, Z:\mask*`. |
| `<drive_letter>:\*mask*` | All paths on the specified volume that contain the given sequence. |

| Mask format | Affects paths/files |
| --- | --- |
| `*mask1*;*mask2*;*mask3*` | All paths that contain at least one of the given character sequences: `*mask1*` or `*mask2*` or `*mask3*`. |

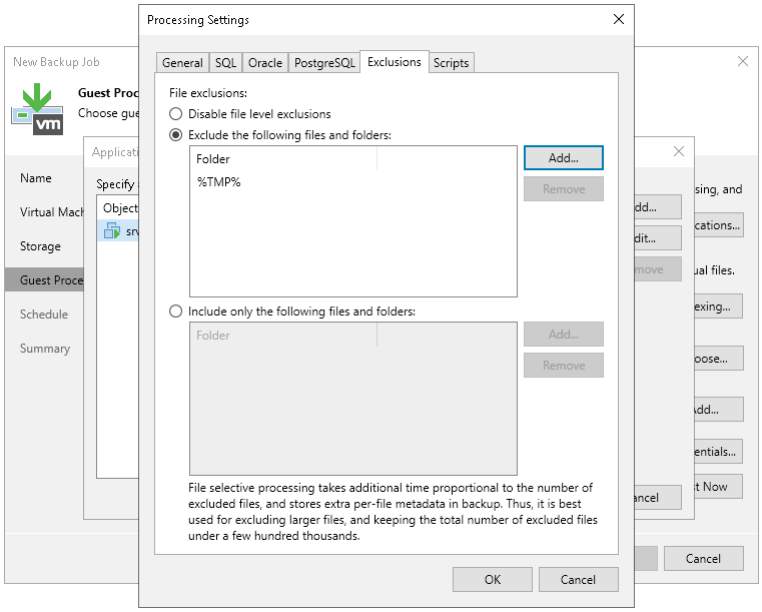


# Requirements and Limitations for VM Guest OS File Exclusion

VM guest OS files exclusion has the following limitations:

- File exclusion works only on Microsoft Windows NTFS.
- File exclusion is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.
- To exclude VM guest OS files, Veeam Backup & Replication must be able to deploy the non-persistent runtime components or use (in necessary, deploy) persistent agent components inside the VM. For this reason, the VM must be running and accessible by an IP address or through the Installer Service on VM, and credentials for application-aware processing must be valid.
- Veeam Backup & Replication supports both basic and dynamic disks. For the dynamic disks, simple type of volumes is supported. Spanned, mirrored and striped volumes are not supported.
- It is not recommended that you use VM guest files exclusion in Microsoft Windows for volumes with enabled Data Deduplication. If you decide to use VM guest files exclusion for such volumes and set up a list of inclusions, you must add the System Volume Information folder to the list of inclusions.

Consider the following:

- Be careful when using masks with double wildcard characters. If you specify masks of such type, Veeam Backup & Replication will exclude all files and paths that contain the given mask. For example, if you specify the `*.doc*` mask, Veeam Backup & Replication will exclude files like `MyReport.docx`, `Report.doc.txt` and so on.
- If you use file masks for file exclusion, Veeam Backup & Replication will scan the VM guest file system, and thus the time of VM disk processing will increase.

- The number of entries in the list of exclusions or inclusions must not exceed a few hundreds. The number of entries in the list influences the job performance — the more files are included or excluded from the backup or replica, the more time Veeam Backup & Replication requires to process these files.
- It is recommended that you do not exclude system files without the necessity.
  Veeam Backup & Replication does not perform any checks to verify the VM image integrity.
- Exclusion of small files (less than 2 KB in size) is ineffective and will not reduce the size of the backup or replica significantly.

# Exclusion Rules for VMs with Several Volumes

The VM guest file exclusion and inclusion functionality works at the volume level. Consider the following situations:

## Data exclusion

A VM has several volumes: `C:\`, `D:\` and `E:\`. You want to exclude from the backup the `Archive` files and folders that are present at the root folder on all volumes of the VM. If you add the `C:\Archive` folder to the list of exclusions, Veeam Backup & Replication will back up the following data:

- Whole content of the `C:\` volume except the `Archive` folder
- Whole content of `D:\` and `E:\` volumes

To exclude the `Archive` files and folders from the root folder from all volumes of the VM, you must `Archive*` to the list of exclusions.

## Data inclusion

A VM has several volumes: `C:\`, `D:\` and `E:\`. You want to include to the backup only the `D:\Documents` folder. If you add the `D:\Documents` folder to the list of inclusions, Veeam Backup & Replication will back up the following data:

- `D:\Documents` folder
- Whole content of `C:\` and `E:\` volumes

To include only the `D:\Documents` folder to the backup, you must add the `D:\Documents` folder to the list of inclusions and, additionally, exclude unnecessary disks (that contain `C:\` and `E:\` volumes) at the **Virtual Machines** step of the wizard. For more information, see Exclude Objects from Backup Job .

# How VM Guest OS File Exclusion Works

When you exclude VM guest OS files from the backup or replica, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication checks the job settings to identify what VM guest OS files must be excluded.
2. Veeam Backup & Replication opens the MFT file from the VM guest file system in the memory cache on the backup proxy, and marks data blocks of excluded files as deleted.

3. When Veeam Backup & Replication copies VM data to the target, it reads data both from the VM snapshot and memory cache on the backup proxy. On the target, Veeam Backup & Replication creates a "merged" version of VM disks that do not contain excluded VM guest OS files. Due to data compression, data blocks that are marked as deleted are compressed, and the size of the resulting backup or replica file reduces.

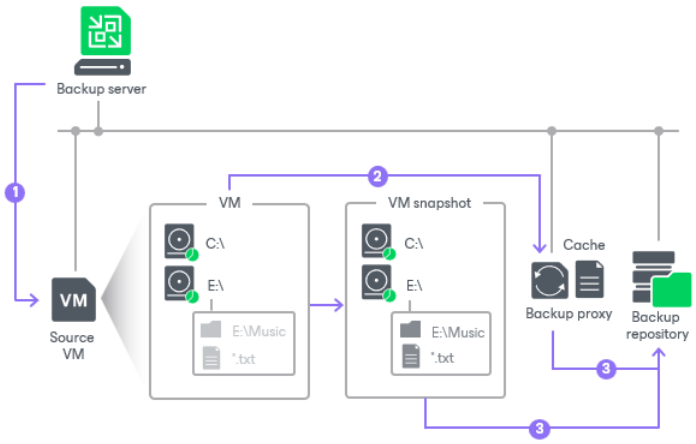


During the job session with file exclude, Veeam Backup & Replication makes changes to processed VM disks at the NTFS level using the cache on the backup proxy. However, these changes are not visible to the CBT mechanism. For this reason, Veeam Backup & Replication saves information about excluded data blocks in the backup file and replica metadata. During the next job session with use of CBT, Veeam Backup & Replication retrieves a list of data blocks that were excluded during the previous job session from the backup file or replica metadata and analyzes what data needs to be processed during the current job session. To do this, Veeam Backup & Replication regards the following data:

- Data blocks that are marked as new with CBT
- Data blocks that were excluded during the previous job session
- Data blocks that must be excluded during the current job session

# VMware Tools Quiescence

When you back up or replicate a running VM, you need to quiesce or 'freeze' the VM to bring its file system and application data to a consistent state. Data quiescence is crucial for highly-transactional applications. It helps create transactionally consistent backups or replicas and guarantees safety of application data.

To create consistent backups and replicas for VMs that do not support Persistent VSS Snaphots (for example, Linux VMs), Veeam Backup & Replication uses the VMware Tools to freeze the file system and application data on VMs before it creates a backup or a replica. VMware Tools also allow you to create backups and replicas for Microsoft Windows-based VMs that support Microsoft VSS. For this, VMware Tools use VMware VSS component. For details on the supported OSes and quiescence features, see the VMware documentation.

To create transactionally consistent backups or replicas of VMs that do not support Microsoft VSS, you must define the following settings at the job level:

1. Enable VMware Tools quiescence. For more information on how to enable VMware Tools quiescence, see vSphere Settings.
2. Specify scripts that will bring applications to a transactionally consistent state before freezing and to the initial state after freezing VMs. For more information, see Pre-Freeze and Post-Thaw Scripts.

**NOTE**

By default, Veeam Backup & Replication uses application-aware processing when both VMware Tools quiescence and application-aware processing are enabled. If some VMs cannot be quiesced with application-aware processing or it is disabled for specific VMs in the job (the **Disable application processing** is set for VMs in the job settings), Veeam Backup & Replication uses VMware Tools quiescence to prepare these VMs for backup or replication.

# Guest Processing

If you back up or replicate running VMs, you can enable guest processing options. Guest processing options are advanced tasks that help Veeam Backup & Replication to communicate with the VM guest OS. Using guest processing options, Veeam Backup & Replication creates transactionally consistent backups, replicas, and catalogs of files and folders on the VM guest OS.

To coordinate guest processing activities, Veeam Backup & Replication deploys non-persistent runtime components or uses (if necessary, deploys) persistent agent components on the VM guest OS. For more information, see the Non-Persistent Runtime Components and Persistent Agent Components section. The non-persistent runtime components run only during guest processing and are stopped immediately after the processing is finished (depending on the selected option, during the backup job session or after the backup job finishes).

Veeam Backup & Replication offers the following guest processing options:

## Application-Aware Processing

By default, Veeam Backup & Replication does not process application logs and creates a crash-consistent backup of VMs with applications that use transaction logs for operations. You can create a transactionally consistent backup - in this case Veeam Backup & Replication will process application logs. In case a disaster strikes, Veeam Backup & Replication will use backups of logs to perform recovery operations. You can create transactionally consistent backups and replicas of VMs that run Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL Server, Oracle Database or PostgreSQL instances. Transactionally consistent backups guarantee proper recovery of these applications without data loss. For information on supported applications, see Supported Platforms and Applications.

When you define the job settings, you can set up the following application-aware processing settings:

- Transaction log backup for Microsoft SQL Server Log Backup, Oracle Log Backup and PostgreSQL WAL Files Backup — to back up transaction logs from Microsoft SQL Server, Oracle server and PostgreSQL server.
- Transaction log truncation — to truncate transaction logs on the VM guest OS after the VM is successfully processed.
- Pre-freeze and post-thaw scripts — to run pre-freeze and post-thaw scripts to quiesce VMs running applications that do not support Microsoft VSS.
- VM guest OS files exclusion — to exclude or include individual files and folders from or to backup or replica.

## VM Guest File System Indexing

You can set up the backup job to create a catalog of files and folders on the VM guest OS. The catalog lets you search for VM guest OS files and perform 1-click restore in Veeam Backup Enterprise Manager.

If you do not enable this option in the backup job settings, you will still be able to perform 1-click restore from backups created by the backup job. For more information, see the Preparing for File Browsing and Searching section in the Enterprise Manager User Guide.

> **NOTE**
>
> If you use Kerberos authentication, consider requirements and limitations described in section Kerberos Authentication.

# Requirements and Limitations

Consider the following requirements and limitations for guest processing:

- Check that accounts that you plan to use for guest processing have permissions described in section Permissions.
- Veeam Backup & Replication excludes from application-aware processing Microsoft SQL databases that are mounted to the Microsoft SQL Server using a remote UNC path. If at least one file of the database is located on a network shared folder, this database will be backed up in the crash-consistent state. Other databases on this server will be backed up in the transactionally consistent state.
- Veeam Backup & Replication excludes the master database from guest processing and does not process transaction logs for it.

  If you want to exclude other databases from the transaction log processing workflow, see this Veeam KB article. Consider that exclusion configured this way will be treated as a global setting.
- Transaction log backups cannot be copied to capacity tier.
- [For Oracle databases except Oracle RAC] Check that databases that you want to back up are listed in the `/etc/oratab` file.
- [For Oracle databases on Microsoft Windows VMs] Make sure that the PATH system variable on the machine with your Oracle database includes the path to the Oracle database software (`%ORACLE_HOME%\bin`). For example, `C:\app\Administrator\product\19.0.0\client_1\bin` or `d:\oracle[...]\bin`.
- To back up Microsoft SQL transaction logs with Veeam Backup & Replication, you must make sure that the recovery model is set to *Full* or *Bulk-logged* recovery model for required databases on Microsoft SQL Server VMs. If the recovery model is set to *Simple*, Veeam Backup & Replication will not detect and process transaction logs on Microsoft SQL Server VMs.
- To back up Oracle transaction logs with Veeam Backup & Replication, you must make sure that ARCHIVELOG is turned on for required databases on Oracle VMs. If ARCHIVELOG is turned off, Veeam Backup & Replication will not detect and process transaction logs on Oracle VMs.
- Due to Microsoft limitations, you cannot use Microsoft Entra ID (formerly Azure Active Directory) credentials to perform guest processing on VMs running Microsoft Windows 10 (or later).
- Veeam Backup & Replication supports backup of Microsoft Exchange, SharePoint and SQL Server databases existing in mount point volumes.

# Non-Persistent Runtime Components and Persistent Agent Components

To perform guest processing tasks, Veeam Backup & Replication must use special components when protecting VMs. These components can be either non-persistent runtime components or persistent agent components:

- Non-Persistent Runtime Components
- Persistent Agent Components

# Non-Persistent Runtime Components

Non-persistent runtime components help to avoid agent-related drawbacks such as pre-installing, troubleshooting and updating. These components are deployed on every VM added to the job when the job starts. As soon as the job finishes, the components are removed. This method is used for guest processing by default.

To use non-persistent runtime components, do not select the **Use persistent guest agent** check box when specifying application-aware processing settings as described in section Application-Aware Processing for VM backup jobs. By default, Veeam Backup & Replication performs processing of the backup jobs with enabled guest processing according to the following algorithm.

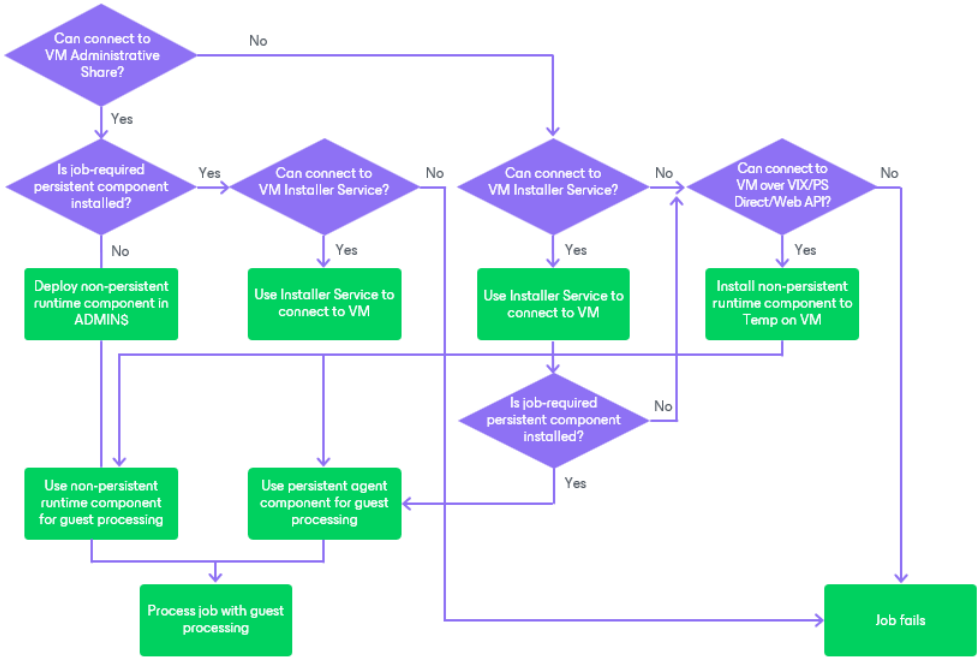


> **NOTE**
>
> If the **Use persistent guest agent** option is disabled but persistent agent components were previously installed on the VM, Veeam Backup & Replication may use them for guest processing. To use non-persistent runtime components, uninstall persistent agent components from protected VMs.

Veeam Backup & Replication deploys the non-persistent runtime components on VMs in two ways:

- For VMs running Microsoft Windows, non-persistent runtime components are deployed using guest interaction proxies. For more information, see Guest Interaction Proxies. If there are no guest interaction proxies or guest interaction proxies fail for some reason, Veeam Backup & Replication will deploy non-persistent runtime components on Microsoft Windows VMs from the backup server.
- For VMs running Linux or Unix operating systems, non-persistent runtime components are deployed from the backup server.

When you start a job with guest processing tasks enabled, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication defines the machines that will perform the guest interaction proxy role.
2. Veeam Backup & Replication deploys the non-persistent runtime components on VMs:
    - [For Microsoft Windows VMs] The guest interaction proxy connects to VMs and deploys the non-persistent runtime components on them.

- [For VMs running Linux or Unix] The backup server connects to VMs and deploys the non-persistent runtime components on them.

3. The job session proceeds as usual.
4. When the job session is complete, Veeam Backup & Replication deletes the non-persistent runtime components on VMs.

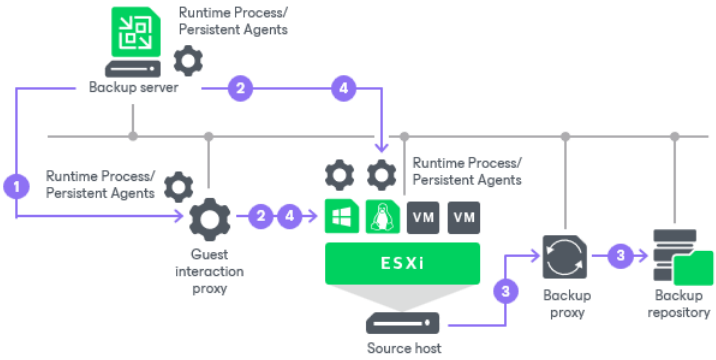


If a network connection breaks during the job session, Veeam Backup & Replication makes attempts to re-establish the connection:

- If a network connection between the backup server/guest interaction proxy and VM guest OS breaks, Veeam Backup & Replication makes one attempt to reconnect.
- If a network connection between the backup server and guest interaction proxy breaks, Veeam Backup & Replication makes 10 attempts to reconnect.

If attempts are unsuccessful, guest processing tasks fail. The job proceeds with the scenario defined in the job settings. For example, if you have instructed a backup job to try application processing but ignore failures, Veeam Backup & Replication will not perform guest processing tasks but will proceed with the VM backup.

## Persistent Agent Components

If you want to perform guest processing in a highly secure way, you can use persistent agent components (Guest Helper, Log Shipping Service) on protected VMs. Persistent agent components require very limited and clearly defined ports to communicate with Veeam Backup & Replication. For more information about ports, see Guest Processing Components.

# Prerequisites

To use persistent agent components, you must ensure the Veeam Installer Service is installed on the guest VM. For more information, see Veeam Installer Service.

# Installing Persistent Agent Components on Microsoft Windows VMs

For Microsoft Windows VMs, Veeam Backup & Replication deploys persistent agent components using guest interaction proxies. For more information, see Guest Interaction Proxies. If there are no guest interaction proxies or guest interaction proxies fail for some reason, Veeam Backup & Replication will deploy persistent agent components on Microsoft Windows VMs from the backup server.

To use persistent agent components, select the **Use persistent guest agent** check box when specifying application-aware processing settings as described in section Application-Aware Processing for VM backup jobs. The Veeam Installer Service will install persistent guest agents during the first run of the backup job.

When you select the **Use persistent guest agent** option, Veeam Backup & Replication performs processing of the backup jobs with enabled guest processing according to the following algorithm. Persistent agent components provide closing access to the VM administrative share (ADMIN$) and access to the VM over VIX API/vSphere Web Services.

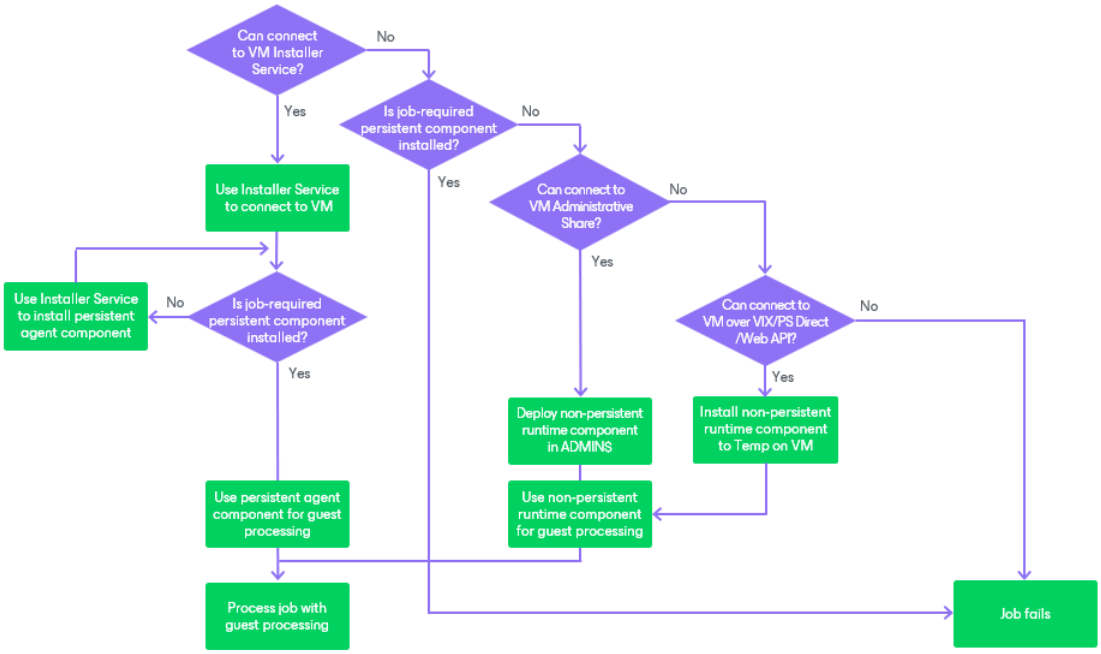


When you start a job with guest processing tasks enabled, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication defines the machines that will perform the guest interaction proxy role.
2. The guest interaction proxy connects to VMs and deploys persistent agent components on them.

   > **NOTE**
   >
   > If the account used for guest processing is a local (non-domain) account, the remote UAC may block the connection. It is a default security policy in Windows that prevents local user accounts from being used to remotely connect to a server. In this case, either use the local administrator account or disable the remote UAC to connect to the persistent agent. For more information, see this Veeam KB article.

3. The job session proceeds as usual.

Veeam Backup & Replication performs operation 4 visible in the following schema only when non-persistent runtime components are used.

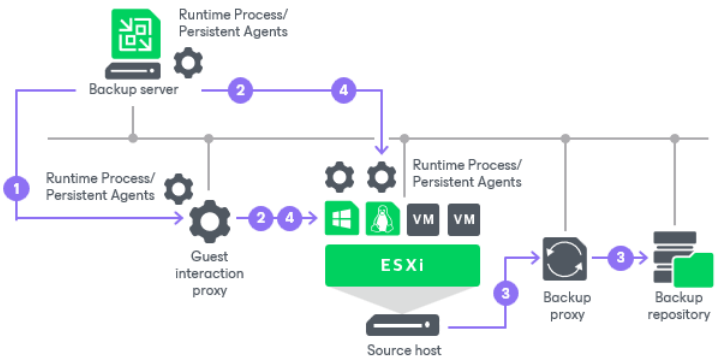

If a network connection breaks during the job session, Veeam Backup & Replication makes attempts to re-establish the connection:

- If a network connection between the backup server/guest interaction proxy and VM guest OS breaks, Veeam Backup & Replication makes one attempt to reconnect.
- If a network connection between the backup server and guest interaction proxy breaks, Veeam Backup & Replication makes 10 attempts to reconnect.

If attempts are unsuccessful, guest processing tasks fail. The job proceeds with the scenario defined in the job settings. For example, if you have instructed a backup job to try application processing but ignore failures, Veeam Backup & Replication will not perform guest processing tasks but will proceed with the VM backup.

# Installing Persistent Agent Components on Linux VMs

For VMs running Linux operating systems, persistent agent components are deployed manually from the backup server using Management Agent. In this case, SSH credentials are used only once to deploy the Veeam Installer Service that installs the Veeam Data Mover Service. These credentials are not stored in the backup infrastructure. The Veeam Data Mover Service will be used to perform guest processing tasks without the SSH connection.

> **NOTE**
>
> For more information about Linux operating systems supported by guest processing, see Supported Platforms and Applications.

To install Management Agent, perform the following steps:

1. Open the **Inventory** view. Select the VM and click **Install Management Agent** on the ribbon or right-click the VM and select **Install Management Agent**.

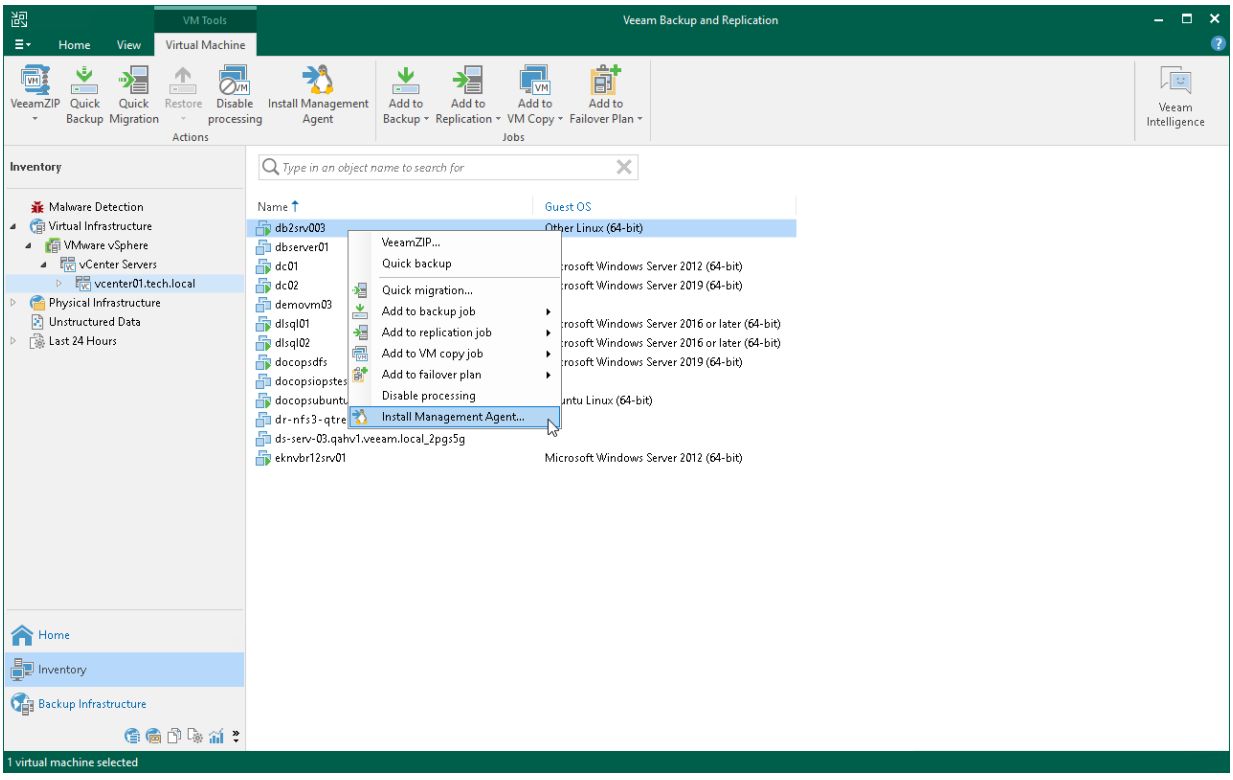

2. At the **Credentials** step of the wizard, specify credentials to connect to the VM and deploy Veeam Data Mover. Specify a user account with root privileges or a non-root user with selected the **Elevate account privileges automatically** check box. If you did not add the user account to the `sudoers` file, select the **Use "su" if "sudo" fails** check box and enter the password for the root account. For more information on these check boxes, see SSH Credentials.

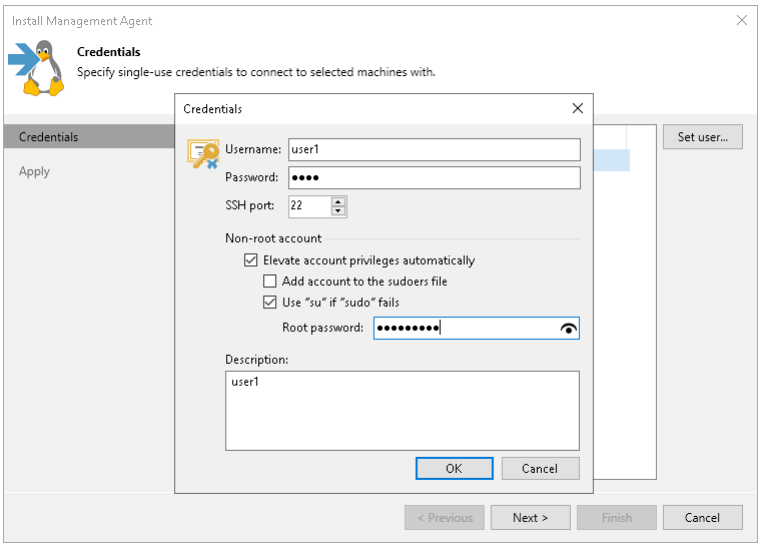


3. Apply changes.

When you start a job with guest processing tasks enabled, Veeam Backup & Replication performs the following operations:

1. The backup server connects to VMs using the credentials specified in the backup job settings. For more information, see Specify Guest Processing Settings.
2. The job session proceeds as usual.

If the transport service connection between the backup server and the VM guest OS fails, Veeam Backup & Replication tries to use the SSH connection with credentials specified in the backup job settings. If the SSH connection also fails, Veeam Backup & Replication uses networkless guest processing over VIX API/vSphere Web Services.

If all attempts are unsuccessful, guest processing tasks fail. The job proceeds with the scenario defined in the job settings. For example, if you have instructed a backup job to try application processing but ignore failures, Veeam Backup & Replication will not perform guest processing tasks but will proceed with the VM backup.

# Application-Aware Processing

By default, Veeam Backup & Replication does not process application logs and creates a crash-consistent backup of VMs with applications that use transaction logs for operations. You can create a transactionally consistent backup — in this case, Veeam Backup & Replication will process application logs. In case a disaster strikes, Veeam Backup & Replication will use backups of logs to perform recovery operations.

To create transactionally consistent backups or replicas of VMs that run the following applications, you must enable application-aware processing in job settings:

- Microsoft Active Directory
- Microsoft SQL Server
- Microsoft SharePoint

- Microsoft Exchange
- Oracle
- PostgreSQL

Application-aware processing is the Veeam technology that allows Veeam Backup & Replication to prepare applications running on the VM and create a consistent view of application data on the VM guest OS. Once the application is ready, Veeam Backup & Replication triggers the VM snapshot and starts to copy VM data to the target. Depending on the VM OS, Veeam Backup & Replication utilizes the following technologies to create transactionally consistent backups:

- For Windows-based VMs, Veeam Backup & Replication uses Microsoft Volume Shadow Copy Service (Microsoft VSS). Microsoft VSS ensures that there are no unfinished database transactions or incomplete application files. For more information, see Microsoft Docs.
- For Linux-based VMs, Veeam Backup & Replication uses the agent that connects to the VM guest OS and prepares databases and instances for a consistent backup.

# Requirements and Limitations

Application-aware processing is supported for Linux-based VMs and Microsoft Windows client versions starting from Windows Vista and for server versions starting from Windows Server 2008. To use application-aware processing, you must have VMware Tools and the latest updates installed on the VM guest OS. For more information on supported guest OS versions, see Supported Platforms and Applications.

> **IMPORTANT**
>
> If a VM runs an application that does not support Microsoft VSS (there is no VSS writer for this particular type of application, for example, MySQL), Veeam Backup & Replication will not be able to utilize Microsoft VSS and application-aware processing for this VM. To process such VMs, you can use VMware Tools quiescence with pre-freeze and post-thaw scripts. For more information, see VMware Tools Quiescence and Pre-Freeze and Post-Thaw Scripts.

# How Application-Aware Processing Works for Windows-Based Machines

If you enable application-aware processing in job settings, Veeam Backup & Replication performs the following operations as a part of the backup or replication process:

1. Veeam Backup & Replication deploys the non-persistent runtime components or, if necessary, persistent agent components on the VM and detects if the VM runs any of the supported applications.
2. Veeam Backup & Replication collects information about applications installed on VMs — this information is required for VSS-aware restore.

   VSS-aware restore is performed when the VM is started after you restore it from the backup or fail over to a VM replica.
3. Veeam Backup & Replication prepares applications for VSS-aware restore.
4. Microsoft VSS communicates with applications and freezes I/O activities at a specific point in time.
5. Veeam Backup & Replication acts as a VSS requestor and triggers a VM VSS snapshot.
6. Veeam Backup & Replication triggers a VMware vSphere snapshot of the VM.

7. Microsoft VSS resumes frozen I/O activities on the VM guest OS.
8. The job session proceeds as usual.
9. If you have instructed Veeam Backup & Replication to truncate transaction logs, Veeam Backup & Replication truncates transaction logs on the VM guest OS after the backup or replica is successfully created.

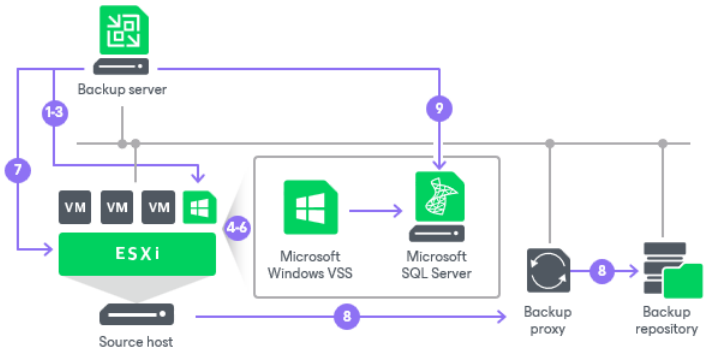


# How Application-Aware Processing Works for PostgreSQL

If you enable application-aware processing in job settings, Veeam Backup & Replication performs the following operations as a part of the backup or replication process

1. Veeam Backup & Replication installs either non-persistent components or persistent agent components to the VM guest OS and detects if the VM runs any of the supported applications.

   > **NOTE**
   >
   > By default, Veeam Backup & Replication installs non-persistent components to the VM guest OS and uninstalls them after the job completes. You can also install a Linux management agent to the VM guest OS — in this case, the agent will remain installed on the VM and Veeam Backup & Replication will use it to access the VM guest OS instead of SSH. For more information, see Persistent Agent Components.

2. The pgsqlagent agent looks for PostgreSQL instance configuration files. It uses these files to get information on the PostgreSQL instance settings.

> **NOTE**
>
> By default, configuration files are located in the following directories:
>
> - [For Ubuntu, Debian] — `/etc/`
> - [For RHEL, SLES] — `/var/lib/`
>
> If you keep the configuration file in the custom directories or if you want to exclude some directories from the scan, you can manually create the `/etc/veeam/VeeamPostgreSQLAgent.xml` file. In this case, the pgsqlagent agent will use commands from this file. To explicitly include or exclude specific configuration files from rescan, add the following commands to the `/etc/veeam/VeeamPostgreSQLAgent.xml` file:
>
> - `ExcludeConfigDirs` — use this command to exclude configuration files.
> - `AddConfigDirs` — use this command to include configuration files.
>
> Note that you must embed the commands into the `<config />` tag. To specify several directories, separate them by a comma.
>
> For example: `<config ExcludeConfigDirs="/etc/postgresql/13/cl4/,/etc/postgresql/13/cl6/" AddConfigDirs="/home/user/pgconfdir/"/>`.
>
> This command skips 2 directories and includes 1 custom directory.

3. The pgsqlagent agent returns a list of PostgreSQL configuration files to Veeam Backup & Replication.
4. Veeam Backup & Replication accesses the PostgreSQL VM guest OS over SSH or over management agent.
5. Veeam Backup & Replication connects to the PostgreSQL instance, gets a list of databases added to the instance and information that is necessary for data recovery operations.
6. The pgsqlagent agent sets the PostgreSQL instance to the ready for a backup state.
7. Agent freezes the VM guest OS and creates a snapshot of the PostgreSQL instance.
8. Veeam Backup & Replication completes a backup of the PostgreSQL instance and resumes stopped activities on the VM guest OS.
9. Veeam Backup & Replication saves a backup of a machine with a PostgreSQL instance to a backup repository.

# Pre-Freeze and Post-Thaw Scripts

If you back up or replicate VMs running applications that do not support Microsoft VSS, you can instruct Veeam Backup & Replication to run custom scripts for VMs. For example, the pre-freeze script may quiesce the file system and application data on the VM guest OS to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script may bring the VM and applications to their initial state.

You can use pre-freeze and post-thaw scripts for the following types of jobs:

- Backup job
- Replication job
- VM copy job

You must create scripts beforehand and specify paths to them in the job settings. Script execution settings can be configured per VM or per container, depending on the objects included in the job.

When the job starts, Veeam Backup & Replication uploads scripts to the VM guest OS and executes them under the account specified in the **Guest OS credentials** section of the job settings.

- Scripts for Microsoft Windows VMs are uploaded to `\\<vmname>\admin$` over the network or VIX API/vSphere Web Services if Veeam Backup & Replication fails to connect to the VM guest OS over the network. Scripts are executed from the `C:\Windows` directory.
- Scripts for Linux VMs are uploaded over SSH or VIX API/vSphere Web Services if the SSH connection fails. Scripts are executed from the `/tmp` directory. If you use Renci or Rebex SSH library, and the Elevate account privileges automatically check box is selected for the user that you have specified in **Guest OS credentials**, the scripts will first be uploaded to the `/home/<username>` and then moved to `/tmp`.

The script is considered to be executed successfully if "0" is returned.

The default time period for script execution is 10 minutes. If the script fails to execute before the timeout expires, Veeam Backup & Replication displays an error message in the job session and error or warning messages issued during script execution.

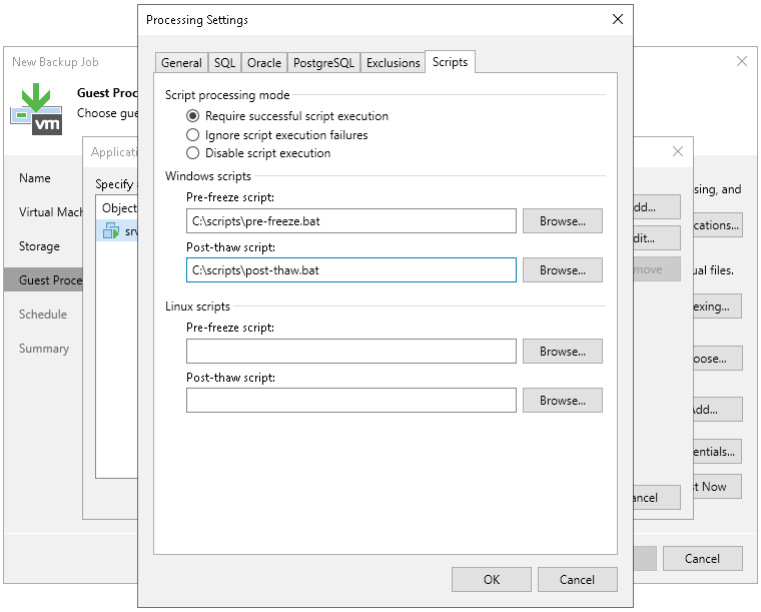


# Supported Script Formats

Pre-freeze and post-thaw scripts can be used for Microsoft Windows and Linux VMs.

- For Microsoft Windows VMs, Veeam Backup & Replication supports scripts in the EXE, BAT, CMD, WSF, JS, VBS and PS1 file format.
- For Linux VMs, Veeam Backup & Replication supports scripts in the SH file format.

# Limitations for Pre-Freeze and Post-Thaw Scripts

Veeam Backup & Replication has the following limitations for pre-freeze and post-thaw scripts:

- You cannot stop a job when the pre-freeze or post-thaw script is executed. If the script hangs up, Veeam Backup & Replication waits for 10 minutes and terminates the job.

- If you want to run several scripts that depend on each other, you must upload them to the VM guest OS manually. For example, you have *script1.bat* that sequentially starts *script2.bat*, *script3.bat* and *script4.bat*. In this case, you must specify a path to *script1.bat* in the job properties and upload *script2.bat*, *script3.bat* and *script4.bat* to the VM guest OS.
- You must not use the standard error (STDERR) stream for error output in Linux scripts. Scripts with STDERR cause failures in Veeam Backup & Replication.

# Transaction Log Truncation

If you back up or replicate virtualized database systems that use transaction logs, for example, Microsoft Exchange or Microsoft SQL Server, you can instruct Veeam Backup & Replication to truncate transaction logs so that logs do not overflow the storage space on the VM. Veeam Backup & Replication provides the following options of transaction logs handling:

- Truncate logs
- Do not truncate logs
- Back up logs periodically

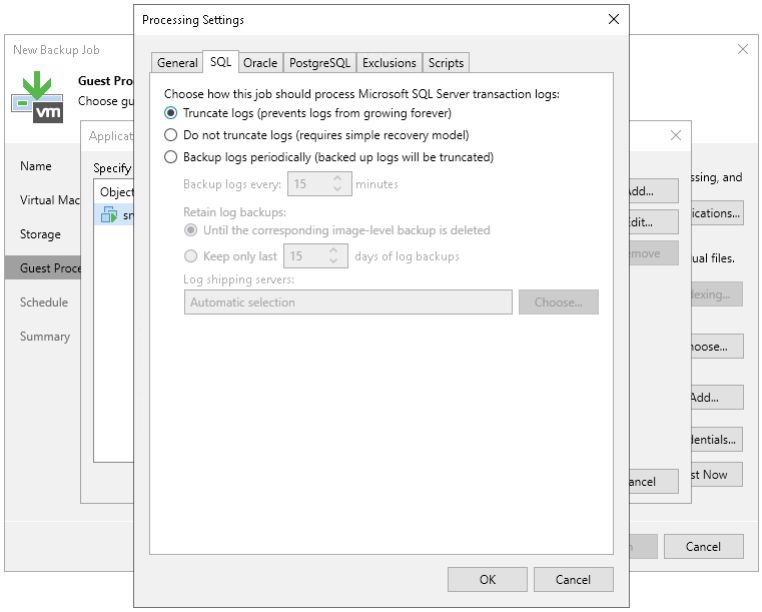


# Truncate Logs

You can instruct Veeam Backup & Replication to truncate logs after a backup or VM replica is successfully created. With this option selected, Veeam Backup & Replication behaves in the following way:

- If the job completes successfully, Veeam Backup & Replication produces a backup file or VM replica and truncates transaction logs on the original VM. As a result, you have the backup file or replica that contains a VM image at a specific point in time.

  In this scenario, you can recover a database to the point in time when the backup file or replica was created. As transaction logs on the VM are truncated, you cannot use them to get the restored database to some point in time between job sessions.
- If the backup or replication job fails, Veeam Backup & Replication does not truncate transaction logs on the VM. In this scenario, you can restore a VM from the most recent backup or replica restore point. Then, you can use database system tools to apply transaction logs and get the database system to the necessary point in time after the restore point.

# Do not Truncate Logs

You can choose not to truncate transaction logs on the VM. This option is recommended if, together with Veeam Backup & Replication, you use another backup tool.

For example, you can use Veeam Backup & Replication to create a VM image backup and instruct the native Microsoft SQL Server log backup job to back up transaction logs. If you truncate transaction logs with Veeam Backup & Replication, the chain of transaction logs will be broken, and the Microsoft SQL Server log backup job will not be able to produce a consistent log backup.

With this option selected, Veeam Backup & Replication produces a backup file or VM replica and does not trigger transaction log truncation. As a result, you have a backup file or VM replica that contains a VM image captured at a specific point in time and transaction logs on the VM. You can use transaction logs to restore the VM to any point in time between job sessions. To do this, you must recover the VM from the backup file or perform a replica failover and use database system tools to apply transaction logs and get the database system to the necessary point in time.

# Back Up Logs Periodically

This option can be used if you back up Microsoft SQL Server VMs and Oracle VMs.

You can choose to back up logs with Veeam Backup & Replication. For more information, see Microsoft SQL Server Log Backup and Oracle Log Backup.

# Support for Database Availability Groups (DAG)

Veeam Backup & Replication supports any configuration of DAGs, in particular, with all databases active on one node, or with active databases on every node. Transaction logs will be truncated on all DAG members, regardless of whether Veeam Backup & Replication backs up an active or passive database.

For more information and recommendations on Microsoft Exchange Server backup, you can also refer to this Veeam KB article.

## Copy-Only Backup

Some organizations prefer to back up Microsoft SQL Server databases and transaction logs with native Microsoft SQL Server tools or 3rd party backup tools. To restore database systems properly, database administrators must ensure they have database backups and a sequence of transaction log backups associated with these backups at hand.

If you use native Microsoft SQL Server tools or 3rd party backup tools and also want to back up Microsoft SQL Server VMs with Veeam Backup & Replication, you must enable the **Perform copy only** option in the job settings.

The **Perform copy only** option indicates that a chain of database backups is created by native Microsoft SQL Server means or by a 3rd party tool and instructs Veeam to preserve this chain (backup history).
Veeam Backup & Replication backs up the Microsoft SQL Server VM using the *VSS_BS_COPY* method for snapshot creation. The *VSS_BT_COPY* method produces a copy-only backup — the backup that is independent of the existing chain of database backups and does not contain transaction logs data. As a result, the copy-only backup does not change the log sequence number and transaction log backup time.

> **IMPORTANT**
>
> Veeam Backup & Replication does not truncate transaction logs after copy-only backup. For this reason, if you instruct the backup job to perform copy-only backup, you cannot specify transaction log handling settings for this job.

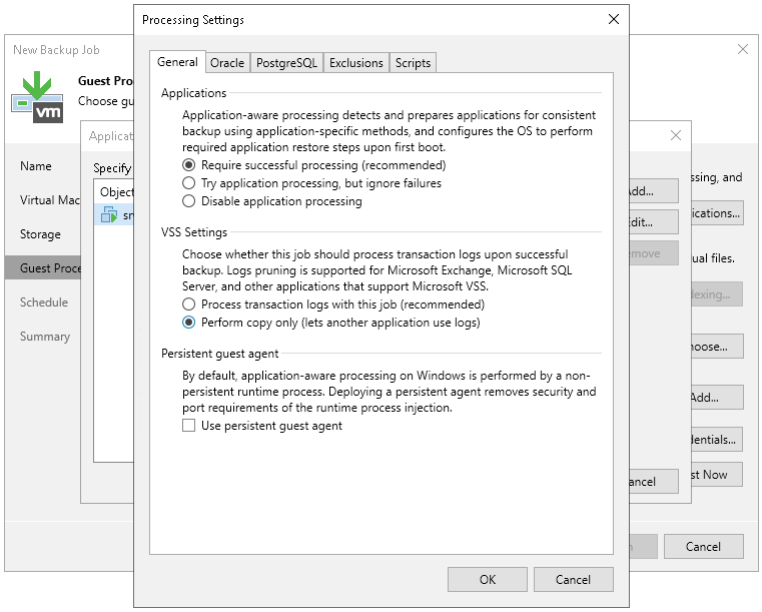


# VM Guest OS File Indexing

You can instruct Veeam Backup & Replication to create an index of files and folders on the VM guest OS during backup. VM guest OS file indexing allows you to search for VM guest OS files inside VM backups and perform a 1-click restore in Veeam Backup Enterprise Manager.

VM guest OS file indexing is enabled at the job level. You can specify granular indexing settings for every VM in the job.

> **NOTE**
>
> VM guest OS file indexing is optional. If you do not enable this option in the backup job settings, you will still be able to perform 1-click restore from the backup created by such backup job. For more information, see the Preparing for File Browsing and Searching section in the Enterprise Manager User Guide.
>
> Note, however, that if you do not enable indexing in the backup job, during a 1-click restore from Linux and other OS backups, Veeam Backup Enterprise Manager will not display symlinks to folders in the file system browser.

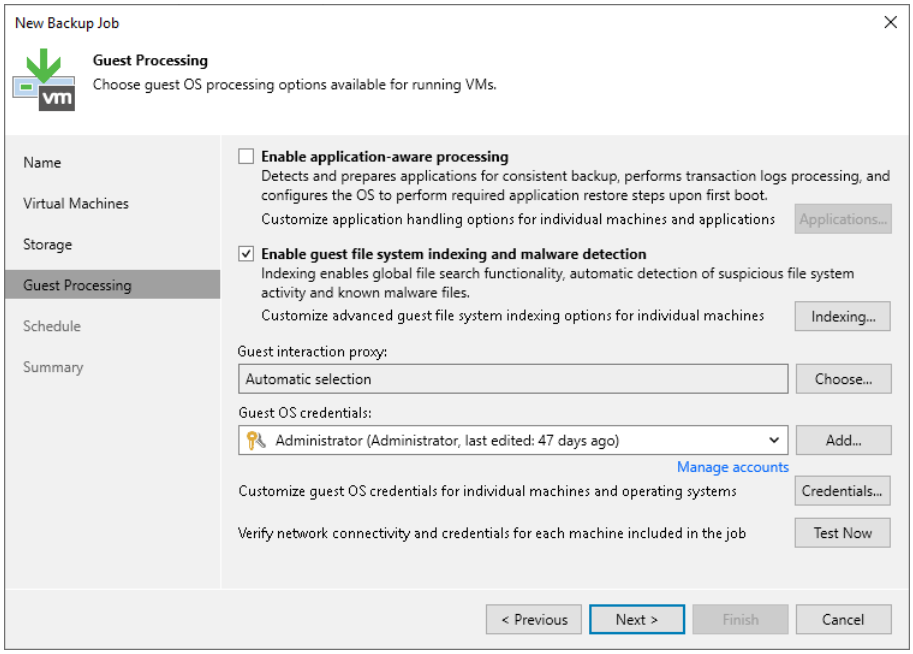


# Requirements and Limitations for VM Guest OS File Indexing

Consider the following requirements and limitations for VM guest OS file indexing:

- Veeam Backup & Replication supports file indexing for VMs running Microsoft Windows and Linux OS.
- [For Linux VMs] To perform guest OS file indexing, Veeam Backup & Replication requires several utilities to be installed on the Linux VM: openssh, gzip and tar. If these utilities are not found, Veeam Backup & Replication will prompt you to deploy them on the VM guest OS.
- To store indexing data, you need enough disk space on the backup server. Indexing data comprises uncompressed index files (IFD) and one compressed *GuestIndexData.zip*. This ZIP file contains the compressed copies of IFD files. The file sizes are calculated in the following way:
  - Each IFD file needs 50 MB of disk space per 1 million files and folders, multiplied by 2. The number of IFD files equals the number of disks on the machine.
  - The size of the *GuestIndexData.zip* file depends on whether Veeam Backup & Replication and Veeam Backup Enterprise Manager are installed on the same machine.
    - If they are on the same machine, the *GuestIndexData.zip* file requires 12.5 MB per 1 million files and folders, multiplied by the longest retention period — the job retention or the retention of the Enterprise Manager guest file system catalog.
    - If they are on different machines, the *GuestIndexData.zip* file requires 12.5 MB per 1 million files and folders, multiplied by the number of days set in the job retention policy.

  Keep in mind that the size of the files is approximate and may vary depending on the length of the file names on the guest OS and their paths.

## Veeam Backup Catalog

For VM guest OS file indexing, Veeam Backup & Replication uses the Veeam Guest Catalog Service. In the backup infrastructure, the Veeam Guest Catalog Service is installed on the Veeam backup server and Veeam Backup Enterprise Manager server.

- The Veeam Guest Catalog Service on the Veeam backup server works as a local catalog service. It collects indexing data for backup jobs and stores it in the Veeam Backup Catalog folder.

  By default, the indexing data is stored in the `VBRCatalog` folder on the backup server.
  Veeam Backup & Replication creates the folder on a volume with the maximum amount of free space, for example, `C:\VBRCatalog`.

- The Veeam Guest Catalog Service on Veeam Backup Enterprise Manager works as a global, federal catalog service. It communicates with Veeam Guest Catalog Services on backup servers connected to Veeam Backup Enterprise Manager and performs the following tasks:

  - Replicates indexing data from backup servers to create a global catalog for the whole backup infrastructure.

    On the Veeam Backup Enterprise Manager server, the default folder for storing indexing data (the `VBRCatalog` folder) is located on a volume with the maximum amount of free space.

  - Maintains indexing data retention.

  - Allows you to search for VM guest OS files in current and archived backup files.

## How VM Guest OS File Indexing Works

When you run a backup job with the file indexing option enabled, Veeam Backup & Replication performs the following operations:

1. When the backup job starts, Veeam Backup & Replication connects to the VM whose file system must be indexed and deploys non-persistent runtime components or connects to (if necessary, deploys) persistent agent components inside this VM. These components are responsible for coordinating indexing activities inside the VM.

2. The non-persistent runtime components or persistent agent components start indexing the VM file system. The indexing procedure is carried out in parallel with the backup procedure. If indexing takes too long, Veeam Backup & Replication will not wait for the indexing procedure to complete. It will start copying VM data and continue file indexing inside the VM. If you have enabled application-aware processing for the VM, Veeam Backup & Replication performs indexing using the VSS snapshot, not the VM guest OS itself. As a result, the created file index exactly reflects the state of the backed-up VM.

3. When file indexing is complete, the non-persistent runtime components or persistent agent components collect indexing data and write it to the `GuestIndexData.zip` file. The `GuestIndexData.zip` file is stored in a temporary folder on the backup server.

4. When the backup job is complete, Veeam Backup & Replication notifies the local Veeam Guest Catalog Service, and the service saves indexing data in the Veeam Catalog folder on the backup server.

5. During the next catalog replication session, the global Veeam Guest Catalog Service replicates data from the backup server to the Veeam Catalog folder on the Veeam Backup Enterprise Manager server.

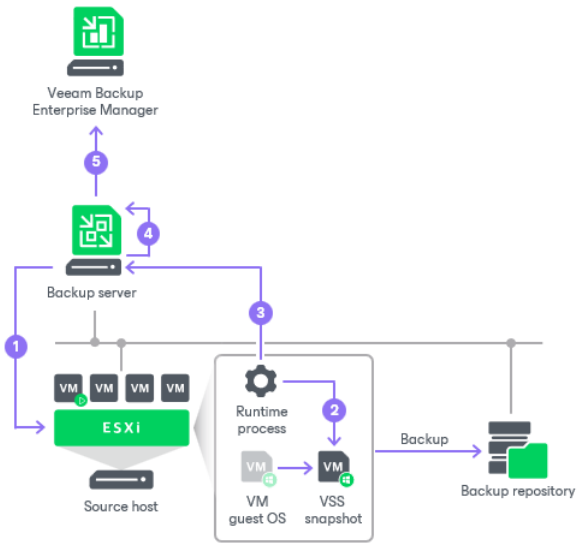


# Persistent VSS Snapshots

During application-aware processing, Veeam Backup & Replication uses a VSS writer for a required application to freeze application data and bring it to a consistent state.

According to Microsoft limitations, applications cannot be kept frozen longer than 60 seconds (20 seconds for Microsoft Exchange). If the Microsoft VSS writer keeps application data frozen longer than this period, a VSS processing timeout occurs, and Veeam Backup & Replication fails to create a transactionally consistent backup of the VM. The VSS processing timeout is a common problem for highly transactional applications such as Microsoft Exchange.

To overcome this limitation, Veeam Backup & Replication uses the persistent VSS snapshots technology for the backup of Microsoft Exchange VMs. If Microsoft Exchange has to be kept frozen for a longer period of time than the allowed one, Veeam Backup & Replication automatically fails over to the persistent VSS snapshot mechanism.

Backup of Microsoft Exchange VMs is performed in the following way:

1. Veeam Backup & Replication triggers the Microsoft VSS framework to prepare Microsoft Exchange on the VM for backup.
2. The Microsoft VSS writer attempts to quiesce Microsoft Exchange.
3. If the Microsoft VSS writer fails to quiesce Microsoft Exchange within the allowed period of time, the control is passed to the native Veeam VSS writer. The Veeam VSS writer holds the freeze operation for the necessary amount of time.
4. After Microsoft Exchange data is brought to a consistent state, the control is passed to the Microsoft VSS provider. The Microsoft VSS framework creates a persistent VSS snapshot for VM disks except the system VM disk.
5. The job session proceeds as usual.
6. After the backup operation is complete, Veeam Backup & Replication triggers Microsoft VSS to remove the persistent VSS snapshot on the production VM. The persistent VSS snapshot holding consistent application data inside the created VM backup remains.

During the entire VM restore, Veeam Backup & Replication recovers data from the backup and reverts VM disks to the persistent VSS snapshot inside the backup. As a result, the Microsoft Exchange VM is restored from the backup in a consistent state without data loss.

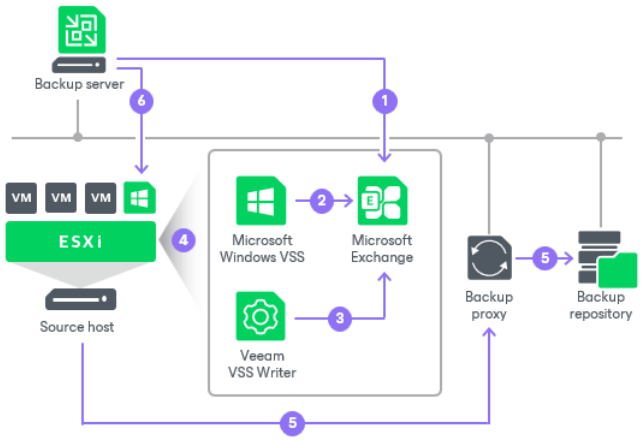


# Limitations for Persistent VSS Snapshot Technology

Veeam Backup & Replication uses the persistent VSS snapshot technology if the VM meets the following requirements:

- The version of Microsoft Exchange installed on the VM is listed in the Supported Applications section of Supported Platforms and Applications.
- The VM does not perform the role of a domain controller.
- Microsoft Exchange databases and log files are located on a non-system disk of the VM. During backup, Veeam Backup & Replication does not trigger a persistent VSS snapshot for system VM disks. As a result, system disks are restored in a crash-consistent, not transactionally consistent state.

# Microsoft SQL Server Log Backup

To protect Microsoft SQL Server VMs, you can instruct the backup job to create image-level VM backups and periodically back up database transaction logs. If Microsoft SQL Server fails, you can restore the Microsoft SQL Server VM from the necessary restore point of the image-level backup. Afterward, you can use Veeam Explorer for Microsoft SQL Server to apply transaction logs and get databases on the Microsoft SQL Server to the necessary state between backups.

## Transaction Log Backup Jobs

To back up transaction logs, you must create a backup job, add Microsoft SQL Server VMs to it and specify advanced settings for transaction log backup in the job settings. The resulting job will comprise two jobs:

- Parent backup job — the backup job that creates an image-level backup of the Microsoft SQL Server VM. The parent backup job is named *<job_name>*, for example, *DB Backup*. You can configure the parent job in the Veeam Backup & Replication console just like any other backup job.
- Child job — a transaction log backup job. To form a name of the child job, Veeam Backup & Replication adds a suffix to the name of the parent backup job: *<parent_job_name> + SQL Server Transaction Log Backup*, for example, *DB Backup SQL Server Transaction Log Backup*. Veeam Backup & Replication automatically creates the child job if it detects a backup job that is scheduled to back up at least one Microsoft SQL Server VM, and transaction log backup is enabled for this job. Session data of the transaction log backup job is stored in the configuration database and displayed in the Veeam Backup & Replication console.

The parent job runs in a regular manner — it starts by schedule or is started manually by the user. The transaction log backup job is triggered by the parent backup job. This sequence ensures that the VM (and the database) restore point is present when it comes to transaction log replay.

## Sessions of Transaction Log Backup Jobs

The transaction log backup job runs permanently in the background, shipping transaction logs to the backup repository at a specific time interval (by default, every 15 minutes). A sequence of time intervals between sessions of the parent backup job makes up a session of the transaction log backup job.

The transaction log backup session starts and stops in the following way:

- The initial session starts when the parent backup job schedule is enabled. After that, the session starts with every new session of the parent backup job.
- The session ends before the next session of the parent backup job or when this parent backup job is disabled.
- When the session ends, Veeam Backup & Replication stops the non-persistent runtime components and uninstalls them from the VM guest OS. When a new session starts, the runtime components are deployed again.

## How Microsoft SQL Server Log Backup Works

To perform transaction log backup, Veeam Backup & Replication installs the *Veeam Guest SQL Log Shipper* runtime component on the VM guest OS.

This component runs during the transaction log backup job session and checks the following details:

- Collects information about databases that require transaction log backup.

- Verifies that it is possible to ship logs directly to the backup repository. If it is not possible, Veeam Backup & Replication uses the log shipping server.

When the transaction log backup job session ends, the component is stopped and removed from the VM guest OS. When a new session starts, the component is installed on the VM guest OS again.

The transaction log backup is performed in the following way:

1. Veeam Backup & Replication launches the parent backup job by schedule.
2. The parent backup job creates an image-level backup of a Microsoft SQL Server VM and stores it in a backup repository.
3. A new session of the transaction log backup starts. Veeam Backup & Replication accesses the VM (directly or through the guest interaction proxy) and installs the runtime components for guest processing, database information collection and transaction log handling on the VM guest OS.
4. Veeam Backup & Replication detects what databases currently exist on the Microsoft SQL Server and maps this data with the information kept in the configuration database. This periodic mapping reveals the databases for which Veeam Backup & Replication must process transaction logs during this time interval.

   The runtime component backs up transaction log files and stores them as a *.bak file in a temporary folder on the VM guest file system. For the information on the temporary folder location, see this Veeam KB article.
5. Veeam Backup & Replication transports transaction log backup copies from the temporary folder on the Microsoft SQL Server VM to the backup repository, either directly or through the log shipping server, and saves them as VLB files. As soon as copies of transaction log backups are saved to the backup repository, transaction log backups in the temporary folder on the Microsoft SQL Server VM are removed.

The session of the transaction log backup job remains working until the next start of the parent backup job. When a new session of the parent job starts, the transaction log backup job stops the current session and then starts a new session, performing steps 1-5.

Transaction logs that, for some reason, were not processed during the log backup interval remain in the temporary folder and are processed during the next log backup interval. To detect these remaining logs, Veeam Backup & Replication enumerates log files in the temporary folder.

> **NOTE**
>
> Backup of Windows Server Failover clusters is not supported; use Veeam Agent managed by Veeam Backup & Replication instead. For more information, see Failover Cluster Support.

> **NOTE**
>
> If a new session of the transaction log backup starts and the parent backup job has not created a new restore point yet, the transaction log backup job will remain in the idle state, waiting for a new restore point to be created.

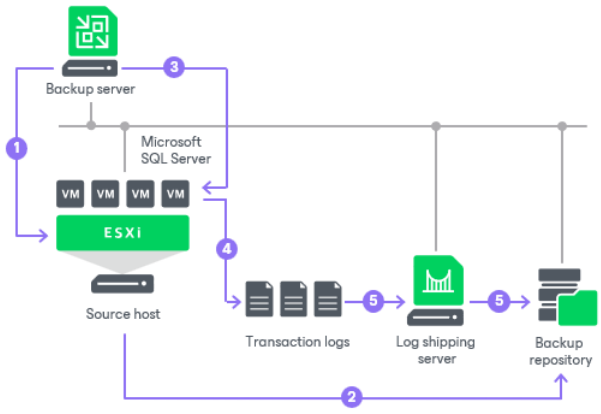


# Retention for Transaction Log Backups

Transaction log backups are stored in files of the proprietary Veeam format — VLB. Veeam Backup & Replication keeps transaction log backups together with the VM image-level backup. The target location of VLB files depends on the type of the backup repository:

- If you store the VM image-level backup in a backup repository, Veeam Backup & Replication writes transaction log backups to the same folder where files of the image-level backup reside.
- If you store the VM image-level backup in a scale-out backup repository, Veeam Backup & Replication writes transaction log backups to the extent where the latest incremental backup file of the VM image-level backup is stored.

Veeam Backup & Replication removes transaction log backups by retention. You can choose one of the following retention methods:

- Retain logs according to the image-level backup
- Retain logs for the specified number of days

# Retain Logs with Image-Level Backup

By default, Veeam Backup & Replication retains transaction log backups together with the image-level backup of the Microsoft SQL Server VM. Veeam Backup & Replication retains VM backup and log backups according to the short-term retention configured for VM backups. When Veeam Backup & Replication removes a restore point of the image-level backup from the backup chain, it also removes a chain of transaction logs relating to this image-level backup. Note that even if long-term retention is configured for the VM backup, Veeam Backup & Replication retains transaction log backups according to the short-term retention policy and deletes them after the short-term retention is exceeded.

This method allows you to have both the image-level backup and necessary transaction log backups at hand. If you need to recover a database to some state, you can restore the Microsoft SQL Database from the necessary restore point and perform transaction log replay to bring the database to the desired state.

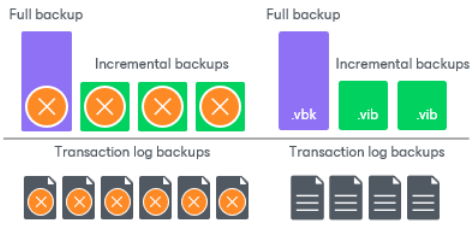

# Retain Logs for a Number of Days

You can instruct Veeam Backup & Replication to keep transaction logs only for a specific period. This retention setting can be used, for example, if you want to save on storage space and plan to retain transaction log backups for the last few days. In this case, you will be able to restore the database only to one of the most recent states.

If you select this retention method, you must make sure that retention policies for the image-level backup and transaction log backup are consistent. The restore point of the image-level backup must always be preserved. If a backup of the database itself is missing, you will not be able to perform transaction log replay.

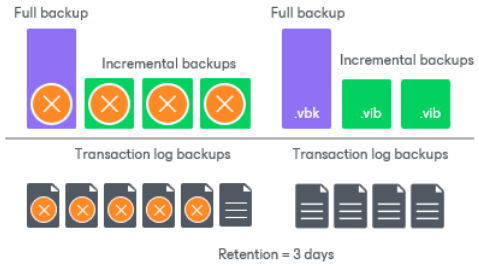


# Related Topics

Microsoft SQL Server Transaction Log Settings

# Log Shipping Servers

For every Microsoft SQL Server VM whose transaction logs you want to back up, Veeam Backup & Replication defines how to ship logs to the backup repository. Transaction logs can be shipped in the following ways:

- **Direct connection**. If it is possible to establish a direct connection between the VM guest OS and backup repository, log files will be shipped directly from the VM guest OS to the backup repository. This is the optimal method, as it does not involve additional resources and puts less load on the VM guest OS.
- **Over log shipping server**. If a direct connection is not possible, files will be shipped through log shipping servers. You can instruct Veeam Backup & Replication to choose a log shipping server automatically from the list of available ones or to use a specific server.

  Note that if a direct connection is possible, files will always be transferred from VM guest to repository directly (regardless of the configured log shipping server, as this server will not be involved). This approach helps to optimize performance at file transfer.

A log shipping server is a Microsoft Windows server added to the backup infrastructure. You can explicitly define what servers you want to use for log shipping or instruct Veeam Backup & Replication to automatically choose an optimal log shipping server. Veeam Backup & Replication chooses the log shipping server based on two criteria: possible data transfer methods and location of the Microsoft SQL Server VMs and log shipping server. For more information, see Location of Log Shipping Server and VMs.

# Data Transfer Methods

Log shipping servers can transport data in two ways:

- **Over the network**. In this scenario, Veeam Backup & Replication obtains files from the VM guest OS and transfers them over the network.

  To offload the VM guest OS, logs are created one by one (not simultaneously). One log creation request is issued for every DB.

- **Over VIX API/vSphere Web Services**. In this scenario, Veeam Backup & Replication obtains transaction logs from the VM guest OS over the VIX API/vSphere Web Services, bypassing the network. For each Microsoft SQL Server instance one log creation request is created for all DBs (grouped by instance).

The default method is log shipping over the network.

# Location of Log Shipping Server and VMs

When choosing a log shipping server, Veeam Backup & Replication considers the location of the Microsoft SQL Server VM and log shipping server. Veeam Backup & Replication uses the following priority rules to select the log shipping server:

1. Log shipping server is located on the same ESXi host as the Microsoft SQL Server VM.
2. Log shipping server and Microsoft SQL Server VM are located in the same network.
3. Log shipping server and Microsoft SQL Server VM are located in different networks (the production infrastructure is isolated from the backup infrastructure).

That is, when choosing a log shipping server, Veeam Backup & Replication will give the top priority to a Microsoft Windows VM that is located on the same ESXi host as the Microsoft SQL Server VM and that has a network connection to the Microsoft SQL Server VM.

Log shipping servers are assigned per job session. When a new job session starts, Veeam Backup & Replication detects log shipping servers anew. Veeam Backup & Replication can also re-detect available servers during the job session. If a log shipping server becomes unavailable for some reason, Veeam Backup & Replication will fail over to another log shipping server.

> **IMPORTANT**
>
> If you do not want to use some servers for transaction logs transport, you can manually define what server Veeam Backup & Replication must use as a log shipping server in the job settings. It is recommended that you assign the log shipping server role to a number of servers for availability purposes.

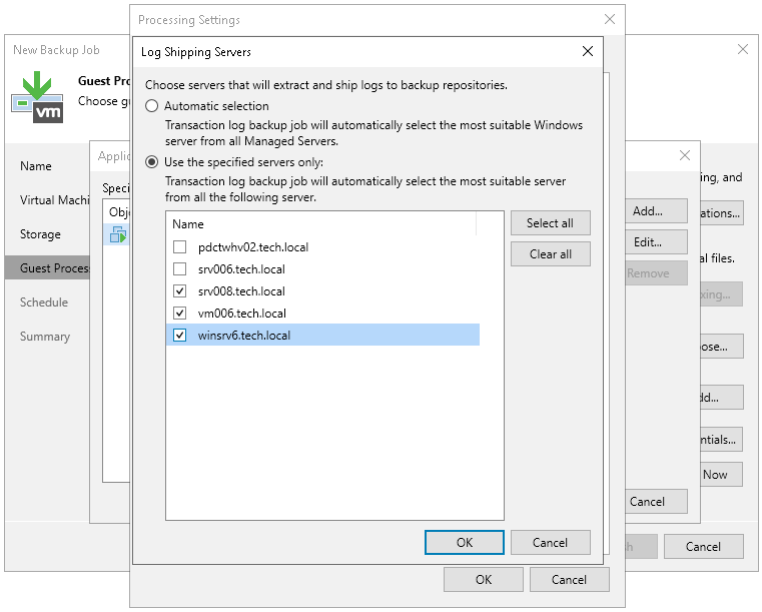

# Transaction Log Backup Statistics

You can view the statistics of the transaction log backup job in the **History** view or in the **Home** view in Veeam Backup & Replication.

In the statistics window, you can examine the overall statistics for the transaction log backup job, as well as view per-VM information.

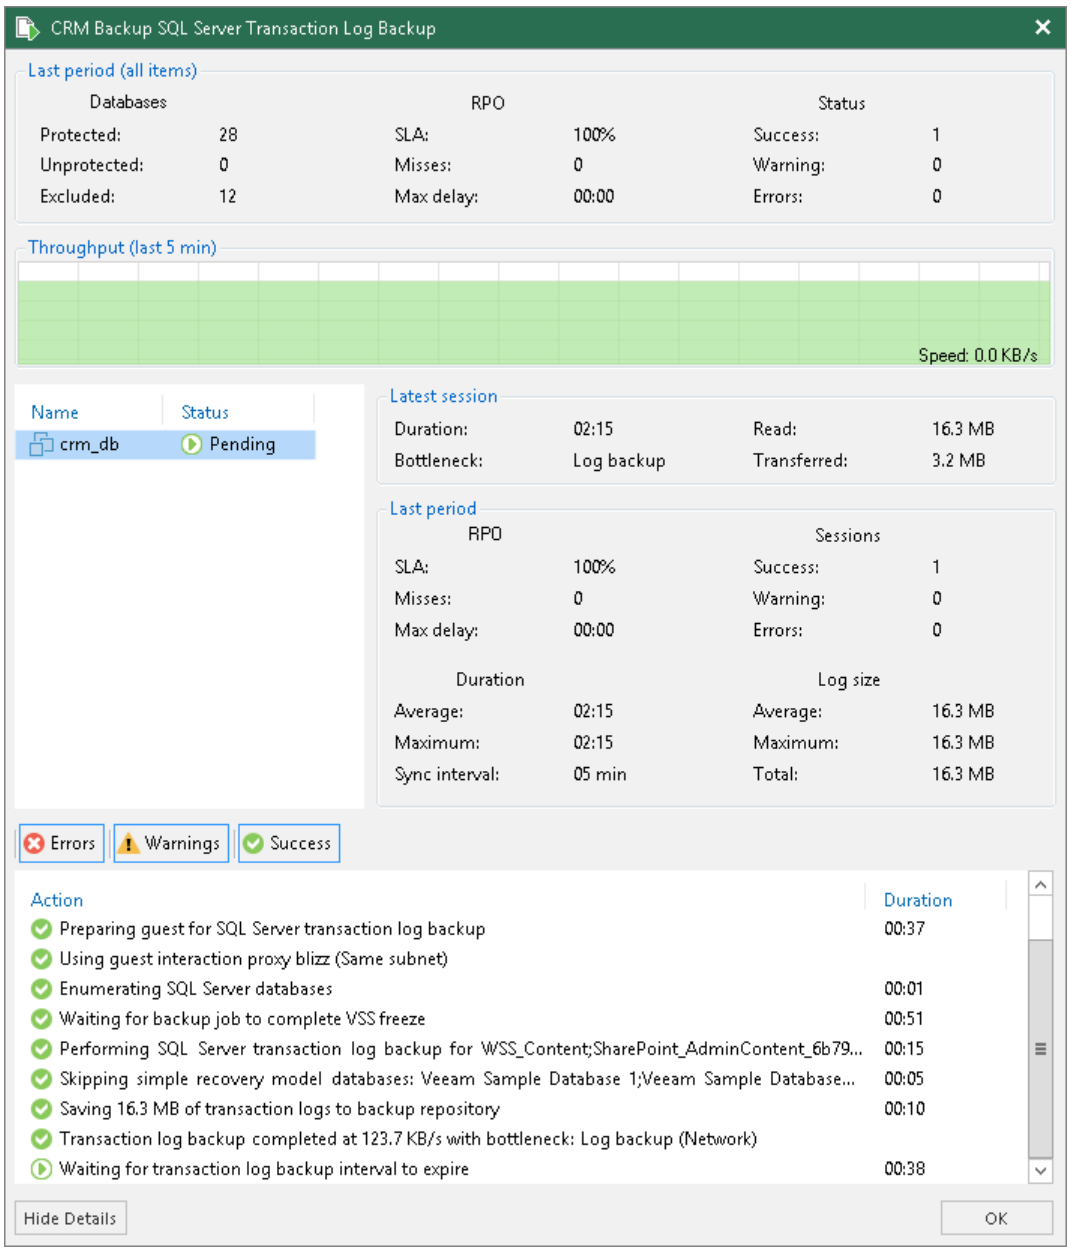


In the upper part of the statistics window, Veeam Backup & Replication displays information about the transaction log backup job for all VMs included in the parent backup job.

The **Last period (all items)** section contains statistics data for the selected session of the backup job.

In the **Databases** column, you can view the following information:

- *Protected* — number of databases that were backed up at least once during the last session.
- *Unprotected* — number of databases that failed to be backed up during the last session.

- *Excluded* — databases excluded from processing. Databases may be excluded for the following reasons:
  - The database status is *Offline.*
  - The database recovery model is set to *Simple.*
  - The database is read-only.
  - The database was deleted after the latest full backup.
  - The AutoClose property is enabled for the database.
  - The database was excluded from application-aware processing. For details, see this Veeam KB article.
  - The database was excluded from SQL log backup processing. For details, see this Veeam KB article.
  - The database belongs to vCenter Server. For details, see this Veeam KB article.

> **NOTE**
>
> Unprotected databases do not comprise *Excluded* databases, as they have different reasons for being non-processed.

In the **RPO** column, you can view the following information:

- *SLA value* — how many log backup intervals were completed in time with successful log backup (calculated as a percentage of a total number of intervals).
- *Misses* — how many intervals were missed (number of intervals).
- *Max delay* — difference between the configured log backup interval and the time actually required for log backup. If exceeded, a warning is issued.

In the **Status** column, the following information is displayed (per job): number of VMs processed successfully, with warnings or with errors.

The **Latest session** section displays the following information for the latest log processing interval for the selected VM:

- *Duration* — duration of log shipment from the VM guest OS to the backup repository since the current log processing interval has started.
- *Bottleneck* — operation with the greatest duration in the last completed interval. The operation may have the following bottlenecks:

| Display Name | Slowing-down Operation |
|---|---|
| **Log backup** | Saving BAK files to a temporary location on VM guest OS |
| **Network** | Uploading log files to the log shipping server |
| **Target** | Saving files to the target repository |

- *Read* — amount of data read from the temporary folder on VM guest OS.
- *Transferred* — amount of data transferred to the target repository.

The **Last period** section displays the following statistics of log backups per VM for the latest session of the transaction log backup job:

- The **RPO** column displays statistics on the log processing interval.
- The **Sessions** column includes statistics of log backups per VM, calculated based on their status:
  - *Success* — number of intervals when all database logs were backed up successfully.
  - *Warning* — number of sequential intervals with failed log processing (if not more than 4 intervals in a sequence).
  - *Errors* — number of sequential intervals with failed log processing (more than 4 intervals in a sequence).
- The **Duration** column includes the following information:
  - *Average* — average duration of log data transfer (through all intervals in the session).
  - *Max* — maximal duration of log data transfer (through all intervals in the session).
  - *Sync interval* — duration of periodic intervals specified for log backup in the parent job settings (default is 15 min).
- The **Log size** column displays the following information:
  - *Average* — average amount of data read from the VM guest OS through all intervals.
  - *Max* — maximal amount of data read from the VM guest OS over all 15-min intervals.
  - *Total* — total amount of data written to the backup repository.

**NOTE**

- Statistics on transaction log processing are updated periodically, simultaneously for the parent backup job and transaction log backup job.
- For Always On Availability groups, Veeam Backup & Replication collects logs only from one node. Thus, in reports, the status of database replicas will be the same for all nodes (Protected or Excluded).

# Log Files

At each start of the SQL Server backup job (parent), a new .VLB is created to store log backups in the repository:

- If the **Use per-machine backup files** option is selected for the repository, then Veeam Backup & Replication will create a separate .VLB for each server processed by the job.
- If this option is cleared, then a single .VLB will be created for all servers processed by the job.

For example, if a job processes only one SQL Server, the repository will contain a number of .VLB files for it (a so-called chain).

As described in the section above, during database log backup (child) job session, transaction log backup is performed by native means of the SQL Server and stored as .BAK file to a temporary folder in the SQL Server VM guest file system. Then, Veeam Backup & Replication copies .BAK file to the current .VLB in the repository. When the new parent job session starts, another .VLB is created, and the .BAK files that appear after that will be stored there during the child job session. The resulting chain of .VLBs will look like the following one, depicted for a single SQL Server VM1:

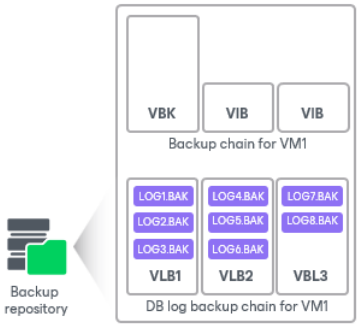


Total number of all `LOG<N>.BAK` files stored at the moment in all VLBs is reported as **a number of restore points for the child job** that backs up database logs. So, in the example above, the log backup job for SQL Server VM1 has created 8 restore points by the moment.

In the Veeam Backup & Replication console, this number of restore points for the log backup job can be seen in the **Restore Points** column of the preview pane.

# Support for Always On Availability Groups

Always On Availability Groups allow you to increase fault tolerance between active and hot-standby databases without involving shared physical disks, which is quite important for the virtualization of Microsoft SQL Servers. Veeam Backup & Replication supports Always On Availability Groups for virtualized Microsoft SQL Server 2012 or later.

Veeam Backup & Replication supports the following Always On Availability Groups:

- Always On Availability Groups based on the Windows Server Failover Cluster
- Always On Clusterless Availability Groups

Veeam Backup & Replication does not support Always On Availability Groups based on SQL Server Failover Cluster Instances.

# Image-level Backup of Microsoft SQL Server VMs

During an image-level backup of a Microsoft SQL Server VM, Veeam Backup & Replication requests and analyzes information about databases that are included in the Always On Availability Groups. Depending on the retrieved information, Veeam Backup & Replication creates a VSS snapshot with or without *COPY_ONLY* flag. The *VSS_BS_COPY* flag for VSS snapshot is triggered if the VM represents a secondary node for at least one Always On Availability Group.

Veeam Backup & Replication also detects to what cluster the database belongs. If the backup job does not include all VMs from the cluster, an information message will be issued.

Retrieved information is saved for further log identification.

# Transaction Log Backup

Transaction log backup can be performed only for those databases that were successfully backed up, either on the primary or on the secondary node of Always On Availability Group.

The transaction logs processing interval may be the same or may differ through VMs included in Always On Availability Group. If the interval is different, Veeam Backup & Replication will use a minimal value (by default, 15 minutes).

At each log processing interval, Veeam Backup & Replication chooses the Always On Availability Group node for which transaction logs will be backed up.

Logs are backed up from one node of the Always On Availability Group. To become a subject for a log backup, the node must meet the following criteria:

- The necessary Veeam Backup & Replication components must be installed on this node and the VM included in Always On Availability Group must be running. For more information on the necessary components, see How Microsoft SQL Server Log Backup Works.
- The database backup preferences settings must allow a backup of the node you want to process. For example, if you want to back up the primary node, you must not exclude this node from a backup, or select the **Secondary only** option in the database backup preferences settings.
- Databases in the Always On Availability Groups for this node were successfully backed up for the last two processing intervals.
- Veeam Backup & Replication can establish a network connection to the node or VIX connection if a connection over the network cannot be established.

> **NOTE**
>
> When you configure a backup job to process Distributed Availability Groups transaction logs, select either primary or secondary distributed availability group. Otherwise, the log chain of the distributed group databases may become inconsistent.
>
> When you configure a backup job to back up transaction logs for other Distributed Availability Groups, use the Perform copy only mode. See Application-Aware Processing to learn more about the copy only mode. You can also use the exclude feature to prevent Guest-OS database from being processed. See Exclude Objects from Backup Job to learn more on excluding objects. To read about distributed availability group limitations, see Configure distributed availability group.

# Oracle Log Backup

Veeam Backup & Replication supports backup of Oracle database archived logs and restore of Oracle databases.

Database archived logs are created by the Oracle system. The Oracle database can run in one of the following logging modes:

- ARCHIVELOG — logs are saved and can be used for recovery purposes.
- NOARCHIVELOG — no logs are saved. This mode is not recommended as it does not provide proper disaster recovery.

With ARCHIVELOG mode enabled, the Oracle system stores database archived logs in a certain location on the VM guest OS, as specified by the database administrator. Veeam Backup & Replication allows you to set up the following ways of log handling:

- Instruct the backup job to collect log files from the Oracle VM and ship them to the backup repository, where they are stored next to image-level backups of the Oracle VM.
- Skip log processing — log files remain untouched on the Oracle VM and are preserved within the image-level backup.

If you enable application-aware processing for an Oracle VM, during the job session Veeam Backup & Replication installs non-persistent runtime components or uses (if necessary, installs) persistent agent components on this VM to collect information about the database and process archived logs according to job settings. Application-specific settings are configured at the **Guest Processing** step of the backup job wizard — you can specify how logs should be backed up and deleted for Oracle databases.

## Requirements for Archived Log Backup

- Veeam Backup & Replication supports archived log backup and restore for Oracle Database version 11 and later. Oracle Database may run on a Microsoft Windows VM or Linux VM.
- Automatic Storage Management (ASM) is supported for Oracle Database 11 and later.
- Oracle Database Express Edition (XE) is supported if running on Microsoft Windows machines only.

## Archived Log Backup Jobs

To back up archived logs, you must create a backup job, add Oracle VMs to it and specify advanced settings for archived logs backup in the job settings. The resulting job will comprise two jobs:

- Parent backup job — the backup job that creates an image-level backup of the Oracle VM. The parent backup job is named <job_name>, for example, Daily Job. You can configure the parent job in the Veeam Backup & Replication console just like any other backup job.
- Child job — an archived log backup job. To form a name of the child job, Veeam Backup & Replication adds a suffix to the name of the parent backup job: *<parent_job_name>* + *Oracle Redo Log Backup*, for example, *Daily Job Oracle Redo Log Backup*. Veeam Backup & Replication automatically creates the child job if it detects a backup job that is scheduled to back up at least one Oracle VM, and archived log backup is enabled for this job. Session data of the archived log backup job is stored in the configuration database and displayed in the Veeam Backup & Replication console.

The parent job runs in a regular manner — it starts by schedule or is started manually by the user. The archived log backup job is triggered by the parent backup job. This sequence ensures that the VM (and the database) restore point is present when you need to use archived logs to restore the database.

# Sessions of Archived Log Backup Jobs

The archived log backup job runs permanently in the background, shipping archived logs to the backup repository at a specific time interval (by default, every 15 minutes). A sequence of time intervals between sessions of the parent backup job makes up a session of the archived log backup job.

The archived log backup session starts and stops in the following way:

- The initial session starts when the parent backup job schedule is enabled. After that, the session starts with every new session of the parent backup job.
- The session ends before the next session of the parent backup job or when this parent backup job is disabled.
- When the session ends, Veeam Backup & Replication stops the non-persistent runtime components and uninstalls them from the VM guest OS. When a new session starts, the runtime components are deployed again.

# How Oracle Archived Log Backup Works

The archived log backup for Oracle VMs is performed in the following way:

1. Veeam Backup & Replication launches the parent backup job by schedule.
2. The parent backup job creates an image-level backup of the Oracle VM and stores this backup in the backup repository.
3. A new session of the archived log backup starts. Veeam Backup & Replication accesses the VM guest OS to perform guest processing, collect database information and handle archived log.

   If Oracle runs on a Microsoft Windows server, Veeam Backup & Replication accesses the VM guest OS over a guest interaction proxy. You can instruct Veeam Backup & Replication to select the guest interaction proxy automatically or assign it explicitly. For more information, see Guest Interaction Proxies.

   By default, Veeam Backup & Replication accesses the VM guest OS over the network:

   - For Linux VM guest OS — using SSH.
   - For Microsoft Windows VM guest OS — using RPC.

   If a network connection cannot be established, Veeam Backup & Replication accesses the VM guest OS over VIX API/vSphere Web Services.
4. Veeam Backup & Replication deploys the non-persistent runtime components or uses (if necessary, deploys) persistent agent components in the VM guest OS. The components scan the Oracle system and collect information about databases whose logs must be processed, including:

   - List of all databases
   - Database state — a database is on or off, in which logging mode it runs
   - Paths to all database files (configuration logs and so on) and other data required for backup

   Veeam Backup & Replication also detects whether it is possible to store logs in the backup repository through direct access or whether a log shipping server is required.

   The non-persistent runtime components or persistent agent components copy archived log files from the log archive destination (set by the Oracle administrator) to a temporary folder on the VM guest file system.

5. Veeam Backup & Replication maps information about the Oracle system collected at step 4 with information kept in the configuration database. This periodic mapping helps reveal databases for which Veeam Backup & Replication must ship archived logs to the backup repository during this time interval.

6. Veeam Backup & Replication transfers archived log backup files from the temporary location on the Oracle VM to the backup repository and saves them as VLB files. Veeam Backup & Replication transfers logs either directly or through the log shipping server. The source-side Veeam Data Mover compresses log data to be transferred according to its built-in settings. On the backup repository side, data is compressed according to the parent backup job settings.

Archived logs that, for some reason, were not processed during the log backup interval remain in the temporary folder and are processed during the next log backup interval. To detect these remaining logs, Veeam Backup & Replication enumerates log files in the temporary folder.

**NOTE**

If a new session of the archived log backup starts and the parent backup job has not created a new restore point yet, the archived log backup job will remain in the idle state, waiting for a new restore point to be created.

**IMPORTANT**

Before backup, Veeam Backup & Replication shuts down databases in the NOARCHIVELOG mode. For details, see the Backing Up a Database in NOARCHIVELOG Mode section in the Oracle Database Backup and Recovery User Guide.

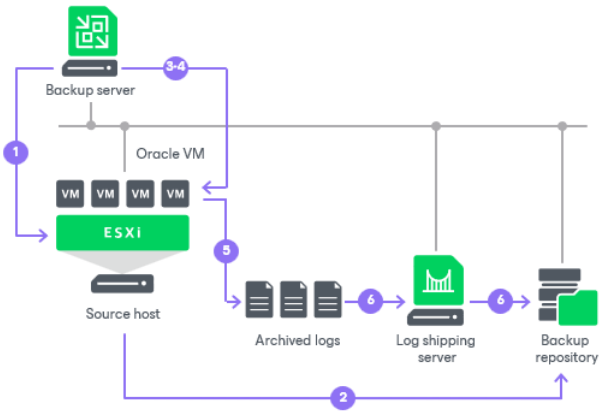


# Retention for Archived Log Backup

Archived log backups are stored in files of the proprietary Veeam format — VLB. Veeam Backup & Replication keeps archived log backups together with the VM image-level backup. The target location of VLB files depends on the type of the backup repository:

- If you store the VM image-level backup in a backup repository, Veeam Backup & Replication writes archived log backups to the same folder where files of the image-level backup reside.
- If you store the VM image-level backup in a scale-out backup repository, Veeam Backup & Replication writes archived log backups to the extent where the latest incremental backup file of the VM image-level backup is stored.

Veeam Backup & Replication removes archived log backups by retention. You can choose one of the following retention methods:

- Retain logs according to the image-level backup

- Retain logs for the specified number of days

# Retain Logs with Image-Level Backup

By default, Veeam Backup & Replication retains archived log backups together with the related image-level backup of the Oracle VM. Veeam Backup & Replication retains VM backup and log backups according to the short-term retention configured for VM backups. When Veeam Backup & Replication removes a restore point of the image-level backup from the backup chain, it also removes a chain of archived logs relating to this image-level backup. Note that even if long-term retention is configured for the VM backup, Veeam Backup & Replication retains archived log backups according to the short-term retention policy and deletes them after the short-term retention is exceeded.

This method allows you to have both the image-level backup and necessary archived log backups at hand. If you need to recover a database to some state, you can restore the Oracle VM from the necessary restore point and use archived logs to bring the database to the desired state.

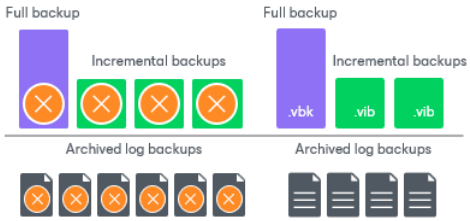


# Retain Logs for a Number of Days

You can instruct Veeam Backup & Replication to keep archived logs only for a specific period. This retention setting can be used, for example, if you want to save on storage space and plan to retain archived log backups for the last few days. In this case, you will be able to restore the database only to one of the most recent states.

If you select this retention method, you must make sure that retention policies for the image-level backup and archived log backup are consistent. The restore point of the image-level backup must always be preserved. If a backup of the database itself is missing, you will not be able to use archived logs.

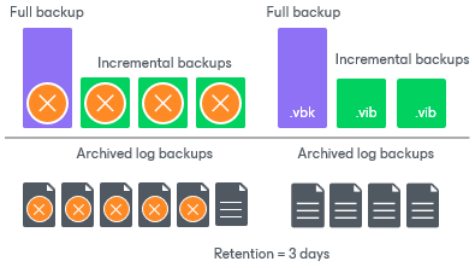


# Log Shipping Servers

For every Oracle VM whose archived logs you want to back up, Veeam Backup & Replication defines how to ship logs to the backup repository. Archived logs can be transported in the following ways:

- **Direct connection**. If it is possible to establish a direct connection between the VM guest OS and backup repository, log files will be shipped directly from the VM guest OS to the backup repository. This method is recommended — it does not involve additional resources and puts less load on the VM guest OS.

- **Over log shipping server**. If a direct connection is not possible, files will be shipped through log shipping servers. You can instruct Veeam Backup & Replication to choose a log shipping server automatically from the list of available ones or to use a specific server.

  Note that if a direct connection is possible, files will always be transferred from VM guest to repository directly (regardless of the configured log shipping server, as this server will not be involved). This approach helps to optimize performance at file transfer.

A log shipping server is a Microsoft Windows server added to the backup infrastructure. You can explicitly define what servers you want to use for log shipping or instruct Veeam Backup & Replication to choose an optimal log shipping server automatically. Veeam Backup & Replication chooses the log shipping server based on two criteria: possible data transfer methods and location of the Oracle VM and log shipping server. For more information, see Location of Log Shipping Server and VMs.

# Data Transfer Methods

Log shipping servers can transport data in two ways:

- **Over the network**. In this scenario, Veeam Backup & Replication obtains files from the VM guest OS and transfers them over the network.

  To offload the VM guest OS, logs are created one by one (not simultaneously). One log creation request is issued for every DB.

- **Over VIX API/vSphere Web Services**. In this scenario, Veeam Backup & Replication obtains transaction logs from the VM guest OS over the VIX API/vSphere Web Services, bypassing the network. For each Oracle Server instance one log creation request is created for all DBs (grouped by instance).

The default method is log shipping over the network.

# Location of Log Shipping Server and VMs

When choosing a log shipping server for the job, Veeam Backup & Replication considers the location of the Oracle VM and log shipping server. Veeam Backup & Replication uses the following priority rules to select the log shipping server:

1. Log shipping server is located on the same ESXi host as the Oracle VM.
2. Log shipping server and Oracle VM are located in the same network.
3. Log shipping server and Oracle VM are located in different networks (the production infrastructure is isolated from the backup infrastructure).

That is, when choosing a log shipping server, Veeam Backup & Replication will give the top priority to a VM that is located on the same ESXi host as the Oracle VM and that has a network connection to the Oracle VM

Log shipping servers are assigned per job session. When a new job session starts, Veeam Backup & Replication detects log shipping servers anew. Veeam Backup & Replication can also re-detect available servers during the job session. If a log shipping server becomes unavailable for some reason, Veeam Backup & Replication will fail over to another log shipping server.

> **IMPORTANT**
>
> If you do not want to use some servers for archived logs transport, you can manually define what server Veeam Backup & Replication must use as a log shipping server in the job settings. It is recommended that you assign the log shipping server role to a number of servers for availability purposes.

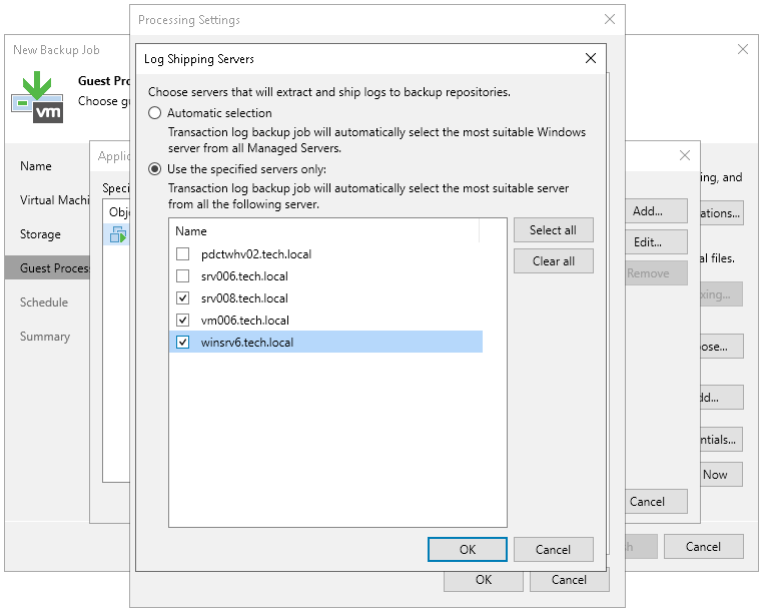


# Archived Log Backup Statistics

You can view the statistics of the archived log backup job in the **History** view or in the **Home** view in Veeam Backup & Replication.

In the statistics window, you can examine the overall statistics for the archived log backup job and view per-VM information.

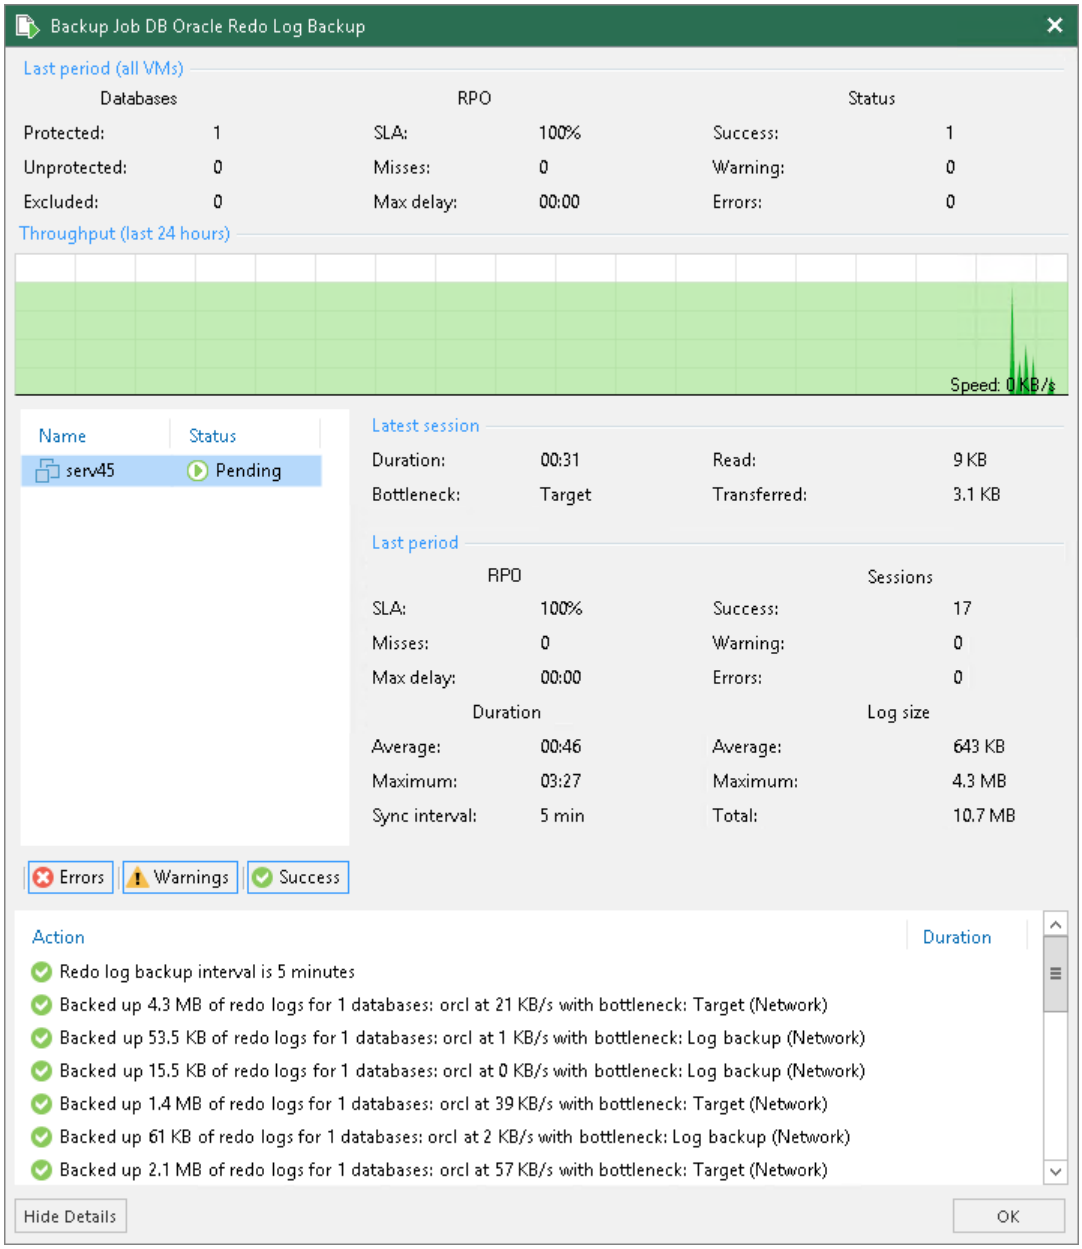


In the upper part of the statistics window, Veeam Backup & Replication displays information about the log backup job for all VMs included in the parent backup job.

The **Last period (all VMs)** section contains statistics data for the selected session of the backup job.

In the **Databases** column, you can view the following information:

- *Protected* — number of databases that were backed up at least once during the last session.
- *Unprotected* — number of databases that failed to be backed up during the last session.
- *Excluded* — databases excluded from processing. Databases may be excluded for the following reasons:
  - ARCHIVELOG mode is turned off for the database (the database is in NOARCHIVELOG mode).
  - The database was deleted after the latest full backup.
  - The database was added to the list of exclusions.

NOTE

Unprotected databases do not comprise **Excluded** databases, as they have different reasons for being non-processed.

In the **RPO** column, you can view the following information:

- *SLA* — how many log backup intervals were completed in time with successful log backup (calculated as a percentage of a total number of intervals).
- *Misses* — how many intervals failed to complete in time with successful log backup (number of intervals).
- *Max delay* — difference between the configured log backup interval and the time actually required for log backup. If exceeded, a warning is issued.

In the **Status** column, the following information is displayed (per job): number of VMs processed successfully, with warnings or with errors.

The **Latest session** section displays the following information for the latest log processing interval for the selected VM:

- *Duration* — duration of log shipment from the VM guest OS to the backup repository since the current log processing interval has started.
- *Bottleneck* — operation with the greatest duration in the last completed interval. The operation may have the following bottlenecks:

| Display Name | Slowing-down Operation |
|---|---|
| Log backup | Saving archived log files to a temporary location on VM guest OS (to work around, see the Veeam KB article: this Veeam KB article) |
| Network | Uploading log files to the log shipping server |
| Target | Saving files to the target repository |

- *Read* — amount of data read from the temporary folder on VM guest OS
- *Transferred* — amount of data transferred to the target repository

The **Last period** section displays the following statistics of log backups per VM for the latest session of the log backup job:

- The **RPO** column displays statistics on log processing interval.
- The **Sessions** column includes statistics of log backups per VM, calculated based on their status:
  - *Success* — number of intervals when all database logs were backed up successfully
  - *Warning* — number of sequential intervals with failed log processing (if not more than 4 intervals in a sequence)
  - *Errors* — number of sequential intervals with failed log processing (more than 4 intervals in a sequence)

- The **Duration** column includes the following information:
  - *Average* — average duration of log data transfer (through all intervals in the session)
  - *Maximum* — maximal duration of log data transfer (through all intervals in the session)
  - *Sync interval* — duration of periodic intervals specified for log backup in the parent job settings (default is 15 min)
- The Log size column displays the following information:
  - *Average* — average amount of data read from the VM guest OS through all intervals
  - *Maximum* — maximal amount of data read from the VM guest OS over all 15-min intervals
  - *Total* — total amount of data written to the backup repository

The pane at the bottom shows all actions performed during the job run. To filter out actions with a certain status, use the **Errors**, **Warnings** and **Success** buttons.

> **NOTE**
>
> Statistics on archived log processing are updated periodically, simultaneously for the VM backup job (parent) and archived log backup job (child job).

# Log Files

At each start of the Oracle backup job ('parent'), a new .VLB is created to store log backups in the repository:

- If the **Use per-machine backup files** option is selected for the repository, then Veeam will create a separate .VLB for each server processed by the job.
- If this option is cleared, then a single .VLB will be created for all servers processed by the job.

For example, if a job processes only one Oracle server, the repository will contain a number of .VLB files for it (a so-called chain).

As described in the section above, during database log backup ('child') job session, log archiving is performed by native means of the Oracle server. Archived logs are stored in a temporary folder on the Oracle VM guest file system. Then Veeam copies the archived log to the current .VLB in the repository. When the new 'parent' job session starts, another .VLB is created, and the archived log files that appear after that will be stored there during the 'child' job session. The resulting chain of .VLBs will look like the following one, depicted for a single Oracle VM1:

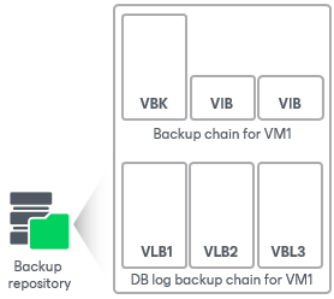


The total number of all archived log files stored at the moment in all VLBs is reported as **a number of restore points for the 'child' job** that backs up database logs. So, in the example above, the log backup job for Oracle VM1 has created 8 restore points by the moment.

# PostgreSQL WAL Files Backup

To protect VMs with PostgreSQL instances, you can instruct the backup job to create image-level VM backups and periodically back up write ahead log (WAL) files of PostgreSQL instances. Thus, you will create transactionally consistent backups of PostgreSQL instances that will contain backups of WAL files. If a PostgreSQL instance fails, you can use Application Item Restore to apply WAL files and recover PostgreSQL instances to the necessary state.

## Requirements and Limitations

Before you back up WAL files for PostgreSQL instances, consider the following requirements and limitations:

- Veeam Backup & Replication supports log backup and restore for PostgreSQL version 12, 13, 14, 15, 16, 17.
- Veeam Backup & Replication supports only backup of PostgreSQL instances configured on Linux-based VMs.
- Veeam Backup & Replication does not support backup of individual PostgreSQL databases.
- Veeam Backup & Replication does not support high availability cluster configurations and replication setups of PostgreSQL servers.
- Set the `archive_mode` parameter to the `on` mode in PostgreSQL instances.
- The `archive_command` parameter must not contain any values in PostgreSQL instances.
- Set the `wal_level` parameter to `replica` or `logical` in PostgreSQL instances.
- The temporary folder that keeps logs must have enough space and be accessible from Linux guest OS.
- To perform guest processing for PostgreSQL instances on Linux servers, make sure that the */tmp* directory is mounted with the *exec* option. Otherwise, you will get an error with the permission denial.

## Related Topics

- WAL Files PostgreSQL Backup Jobs
- How PostgreSQL WAL Files Backup Works
- Retention for PostgreSQL WAL Files
- Log Shipping Servers
- WAL Backup Statistics

## WAL Files PostgreSQL Backup Jobs

To maintain data consistency and integrity, Veeam Backup & Replication makes a backup of write ahead log (WAL) files for PostgreSQL instances. Veeam Backup & Replication uses these log files to recover PostgreSQL instances and bring them to the necessary consistent state.

To back up WAL files, you must create a backup job, add a PostgreSQL VM to it and specify advanced settings for WAL files in the job settings. The resulting job will consist of two jobs:

- Parent backup job — the job that creates an image-level backup of PostgreSQL VMs. You specify a parent backup job name when you define settings for a backup job, for example, *Daily DB Job*. For more information, see Specify Job Name and Description. For more information on how image-level backup of PostgreSQL VMs works, see Application-Aware Processing.
- Child job — the job that creates a backup of WAL files. To form a name of the child job, Veeam Backup & Replication adds a suffix to the name of the parent backup job: *<parent_job_name>* + *PostgreSQL Log Backup*, for example, *Daily Job PostgreSQL Log Backup*. When Veeam Backup & Replication detects a backup job that contains at least one PostgreSQL VM, it automatically creates the child job to back up WAL files. Veeam Backup & Replication keeps session data of the child job in the configuration database and displays it in the console. For more information on how the job creates a backup of WAL files, see How PostgreSQL WAL Files Backup Works.

The parent job runs in a regular manner — it starts by schedule or is started manually by the user. The child job is triggered by the parent backup job. This sequence ensures that the VM (and the instances) restore point is present when you need to use WAL files to restore the database.

# Sessions of Archived Log Backup Jobs

The child backup job runs permanently in the background, shipping WAL files to the backup repository at a specific time interval (by default, every 15 minutes). A sequence of time intervals between sessions of the parent backup job makes up a session of the child backup job.

The child backup session starts and stops in the following way:

- The initial session starts when the parent backup job schedule is enabled. After that, the session starts with every new session of the parent backup job.
- The session ends before the next session of the parent backup job or when this parent backup job is disabled.
- When the session ends, Veeam Backup & Replication stops the non-persistent runtime components and uninstalls them from the VM guest OS. When a new session starts, the runtime components are deployed again.

# How PostgreSQL WAL Files Backup Works

To maintain integrity and data consistency, Veeam Backup & Replication makes a backup of write ahead log (WAL) files that contain PostgreSQL instance logs. Veeam Backup & Replication uses these log files for instance recovery operations.

To back up WAL files, you must enable application-aware processing for PostgreSQL VMs. After this, during the job session Veeam Backup & Replication installs non-persistent runtime components or uses (if necessary, installs) persistent agent components on this VM to collect information about the instance and process WAL files according to job settings. For more information on how to configure application-specific settings for PostgreSQL WAL files backup, see PostgreSQL WAL Files Settings step of the backup job wizard — you can specify how logs should be backed up and deleted for PostgreSQL instances.

The WAL files backup for PostgreSQL VMs is performed in the following way:

1. Veeam Backup & Replication launches the parent backup job by schedule.
2. The parent backup job creates an image-level backup of the PostgreSQL VM and stores this backup in the backup repository.

3. A child job starts to back up WALs: Veeam Backup & Replication accesses the PostgreSQL VM guest OS over SSH. By default, Veeam Backup & Replication installs non-persistent components to the VM guest OS, and Veeam Backup & Replication will uninstall them after the job completes. You can also install a Linux management agent to the VM guest OS — in this case, the agent will remain installed on the VM and Veeam Backup & Replication will use it to access the VM guest OS. For more information, see Persistent Agent Components.
4. Veeam Backup & Replication scans the PostgreSQL system and gets the necessary information to back up backup and restore a PostgreSQL instance.
   - List of all databases added to the instance.
   - The instance state — if the instance is on or off, in which logging mode it runs.
   - Paths to all Instance files (configuration logs and so on) and other data required for backup.
5. The non-persistent runtime components or persistent agent components copy WALs from the log archive destination (set in the guest processing settings of a job) to a temporary directory on the VM guest file system.
6. Veeam Backup & Replication maps information about the PostgreSQL system collected at the step 4 with information kept in the configuration database. This periodic mapping helps to reveal instances for which Veeam Backup & Replication must ship WALs to the backup repository during this time interval.
7. Veeam Backup & Replication transfers WALs from the temporary location on the PostgreSQL VM to the backup repository and saves them as VLB files. Veeam Backup & Replication transfers WALs either directly or through the log shipping server. The source-side Veeam Data Mover compresses log data to be transferred according to its built-in settings. On the backup repository side, data is compressed according to the parent backup job settings.
8. Veeam Backup & Replication deletes WALs from the temporary directory on the VM guest file system.

> **IMPORTANT**
>
> Consider the following:
>
> - Veeam Backup & Replication removes only backed-up WALs from a temporary folder. The WALs that were not backed up during the log backup interval remain in the temporary folder. Veeam Backup & Replication will process these logs during the next log backup interval.
> - If a new session of the WALs backup starts and the parent backup job has not created a new restore point yet, the WALs backup job will remain in the idle state, waiting for a new restore point to be created.

# Retention for PostgreSQL WAL Files

Veeam Backup & Replication stores WAL files in the proprietary Veeam format — .VLB and keep them together with the VM image-level backup. For more information on the backup files, see Backup Chain.

The target location where Veeam Backup & Replication keeps .VLB files depends on the type of repository where you store VM image-level backup files:

- If it is the backup repository, Veeam Backup & Replication writes WAL files to the same folder where the VM image-level backup files reside.
- If it is the scale-out backup repository, Veeam Backup & Replication writes WAL files to the extent where the latest incremental backup file of the VM image-level backup resides.

  Note that Veeam Backup & Replication does not store .VLB files in capacity or archive extents added to the scale-out backup repository.

Veeam Backup & Replication removes WAL files by retention. You can choose one of the following retention methods:

- Retain WAL files according to the image-level backup
- Retain WAL files for the specified number of days

# Retain Logs with Image-Level Backup

By default, Veeam Backup & Replication retains WAL files together with the image-level backup of a PostgreSQL VM. Veeam Backup & Replication retains VM backup and log backups according to the short-term retention configured for VM backups. When Veeam Backup & Replication removes a restore point of the image-level backup from a backup chain, it also removes a chain of WAL files that relates to this image-level backup.

> **NOTE**
>
> If you configure long-term retention, Veeam Backup & Replication will retain WAL files according to the short-term retention policy and will delete them after the short-term retention is exceeded.

This method allows you to have both the image-level backup and necessary WAL files at hand. If you need to recover a PostgreSQL instance to a specific state, you can restore the PostgreSQL VM from the necessary restore point and use WAL files to bring the database to the necessary state.

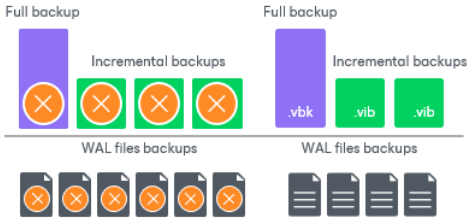


# Retain WAL Files for a Number of Days

You can instruct Veeam Backup & Replication to keep WAL files only for a specific period. You can use this option if you want to save storage space and plan to retain WAL files for the last few days. In this case, you will be able to restore a PostgreSQL instance only to one of the most recent states.

If you select this retention method, you must ensure that retention policies for the image-level backup and archived log backup are consistent. The restore point of the image-level backup must always be preserved. If a backup of the PostgreSQL instance itself is missing, you will not be able to use WAL files for recovery.

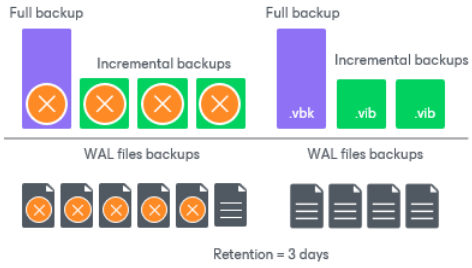

# WAL Files

At each start of the parent PostgreSQL backup job, Veeam Backup & Replication creates a new .VLB to store log backups in the repository:

- If you select the **Use per-machine backup files** option for the repository, Veeam Backup & Replication will create a separate .VLB for each server processed by the job.
- If you do not use the **Use per-machine backup files** option, Veeam Backup & Replication will create a single .VLB for all servers processed by the job.

For example, if a job processes only one PostgreSQL server, the repository will contain a number of .VLB files for it (a so-called chain).

Veeam Backup & Replication comprises native means of the PostgreSQL server to archive WAL files. During the WAL files backup job, the PostgreSQL server stores the files in a temporary folder on the PostgreSQL VM guest file system. After the parent backup job is completed, Veeam Backup & Replication copies WAL files to .VLB and stores them in a backup repository. When a new WAL files job session starts, the PostgreSQL server creates new WAL files and keeps them in the temporary folder. As a result, for a single PostgreSQL VM the job creates a chain of .VLBs that are located in a backup repository together with .VBK and .VIB files. A total number of all WAL files stored in .VLB sets a number of restore points. A single .VLB file can have several WAL files and, therefore, several restore points. For example. 1 .VLB can have 10 restore points.

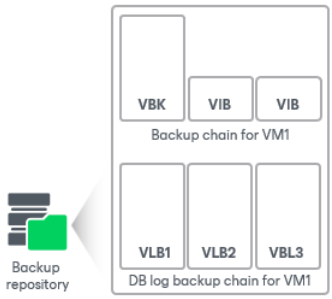


# Log Shipping Servers

For every PostgreSQL VM whose WAL files you want to back up, Veeam Backup & Replication defines how to ship logs to the backup repository. WAL files can be transported in the following ways:

- **Direct connection**. If it is possible to establish a direct connection between the VM guest OS and backup repository, WAL files will be shipped directly from the VM guest OS to the backup repository. This method is recommended — it does not involve additional resources and puts less load on the VM guest OS.
- **Over log shipping server**. If a direct connection is not possible, files will be shipped through log shipping servers. You can configure Veeam Backup & Replication to choose a log shipping server from the list of available ones or to use a specific server.

  Note that if a direct connection is possible, files will be always transferred from VM guest to repository directly (regardless of the configured log shipping server, as this server will not be involved). This approach helps to optimize performance at file transfer.

A log shipping server is a Microsoft Windows server added to the backup infrastructure. You can explicitly define what servers you want to use for log shipping or instruct Veeam Backup & Replication to automatically choose an optimal log shipping server. Veeam Backup & Replication chooses the log shipping server based on two criteria: possible data transfer methods and location of the PostgreSQL VM and log shipping server. For more information, see Location of Log Shipping Server and VMs.

# Data Transfer Methods

Log shipping servers can transport data in two ways:

- Over the network. In this scenario, Veeam Backup & Replication obtains files from the VM guest OS and transfers them over the network. This is a default method that uses log shipping servers.
- Over VIX API/vSphere Web Services. In this scenario, Veeam Backup & Replication obtains WAL files from the VM guest OS over the VIX API/vSphere Web Services, bypassing the network.

# Location of Log Shipping Server and VMs

When choosing a log shipping server for the job, Veeam Backup & Replication considers the location of the PostgreSQL VM and log shipping server. Veeam Backup & Replication uses the following priority rules to select the log shipping server:

1. The log shipping server is located on the same ESXi host as the PostgreSQL VM.
2. The log shipping server and PostgreSQL VM are located in the same network.
3. The log shipping server and PostgreSQL VM are located in different networks (the production infrastructure is isolated from the backup infrastructure).

That is, when choosing a log shipping server, Veeam Backup & Replication will give the top priority to a VM that is located on the same ESXi host as the PostgreSQL VM and that has a network connection to the PostgreSQL VM.

Veeam Backup & Replication assigns log shipping servers per job session. When a new job session starts, Veeam Backup & Replication detects a new log shipping server. Veeam Backup & Replication can also re-detect available servers during the job session. If the log shipping server becomes unavailable for some reason, Veeam Backup & Replication will fail over to another log shipping server.

> **IMPORTANT**
>
> If you do not want to use some servers for WAL files transport, you can manually define what server Veeam Backup & Replication must use as a log shipping server in the job settings. It is recommended that you assign the log shipping server role to a number of servers for availability purposes.

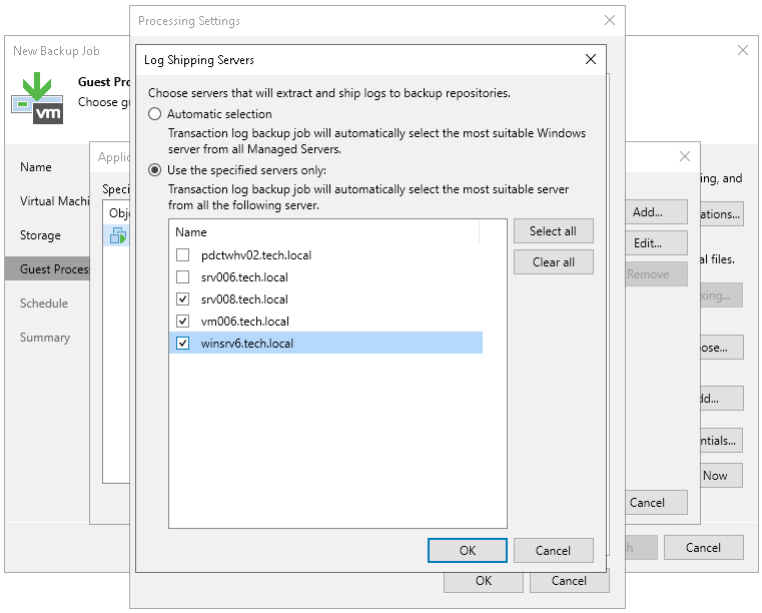

# WAL Backup Statistics

You can view the statistics of the WAL files backup job in the **History** view or in the **Home** view in Veeam Backup & Replication.

In the statistics window, you can examine the overall statistics for WAL files and view per-VM information.

In the upper part of the statistics window, Veeam Backup & Replication displays information about the WAL files backup job for all VMs included in the parent backup job.

The **Last period (all VMs)** section contains statistics data for the selected session of the backup job.

In the **Databases** column, you can view the following information:

- *Protected* — number of instances that were backed up at least once during the last session.
- *Unprotected* — number of instances that failed to be backed up during the last session.
- *Excluded* — instances excluded from processing. Instances may be excluded for the following reasons:
  - ARCHIVELOG mode is turned off for the database (the database is in NOARCHIVELOG mode).
  - The database was deleted after the latest full backup.
  - The database was added to the list of exclusions.

> **NOTE**
>
> Unprotected instances do not comprise **Excluded** instances, as they have different reasons for being non-processed.

In the **RPO** column, you can view the following information:

- *SLA* — how many log backup intervals were completed in time with successful log backup (calculated as a percentage of the total number of intervals).
- *Misses* — how many intervals failed to complete in time with successful log backup (number of intervals).
- *Max delay* — difference between the configured log backup interval and the time actually required for log backup. If exceeded, a warning is issued.

In the **Status** column, the following information is displayed (per job): number of VMs processed successfully, with warnings or with errors.

The **Latest session** section displays the following information for the latest log processing interval for the selected VM:

- *Duration* — duration of log shipment from the VM guest OS to the backup repository since the current log processing interval has started.
- *Bottleneck* — operation with the greatest duration in the last completed interval. The operation may have the following bottlenecks:

| Display Name | Slowing-down Operation |
|---|---|
| Log backup | Saving archived log files to a temporary location on VM guest OS (to work around, see the Veeam KB article: this Veeam KB article) |

| Display Name | Slowing-down Operation |
|---|---|
| Network | Uploading log files to the log shipping server |
| Target | Saving files to the target repository |

- *Read* — amount of data read from the temporary folder on VM guest OS.
- *Transferred* — amount of data transferred to the target repository.

The **Last period** section displays the following statistics of log backups for every VM for the latest session of the log backup job:

- The **RPO** column displays statistics on log processing interval.
- The **Sessions** column includes statistics of log backups per VM, calculated based on their status:
  - *Success* — number of intervals when all database logs were backed up successfully
  - *Warning* — number of sequential intervals with failed log processing (if not more than 4 intervals in a sequence)
  - *Errors* — number of sequential intervals with failed log processing (more than 4 intervals in a sequence)
- The **Duration** column includes the following information:
  - *Average* — average duration of log data transfer (through all intervals in the session)
  - *Maximum* — maximal duration of log data transfer (through all intervals in the session)
  - *Sync interval* — duration of periodic intervals specified for log backup in the parent job settings (default is 15 min)
- The **Log size** column displays the following information:
  - *Average* — average amount of data read from the VM guest OS through all intervals
  - *Maximum* — maximal amount of data read from the VM guest OS over all 15-min intervals
  - *Total* — total amount of data written to the backup repository

The pane at the bottom shows all actions performed during the job run. To filter out actions with a certain status, use the **Errors**, **Warnings** and **Success** buttons.

> **NOTE**
>
> Statistics on WAL files processing is updated periodically, simultaneously for the VM backup job (parent) and WAL files backup job (child job).

# Backup Job Scheduling

You can start backup jobs manually or schedule them to start automatically at a specific time. Veeam Backup & Replication lets you configure the following settings for the job:

- Automatic Startup Schedule
- Job retry
- Backup window

## Automatic Startup Schedule

To run a job periodically without user intervention, you can schedule the job to start automatically. The Veeam Backup Service running on the backup server continuously checks the configuration settings of all jobs configured on the backup server and starts them according to their schedule.

Veeam Backup & Replication lets you configure the following scheduling settings for jobs:

- You can schedule jobs to run at specific time every day or on selected days
- You can schedule jobs to run periodically at specific time intervals
- You can schedule jobs to run continuously
- You can chain jobs

## Jobs Started at Specific Time

You can schedule jobs to start at a specific time daily, on specific week days, or monthly on selected days.

This type of schedule requires you to define the exact time when the job must be started. For example, you can configure the job to start daily at 10:00 PM or every first Sunday of the month at 12:00 AM.

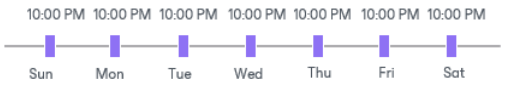


## Jobs Started at Specific Time Intervals

You can schedule jobs to start periodically throughout the day at a specific time interval. The time interval between job sessions can be defined in minutes or hours. For example, you can configure a job to start every 30 minutes or every 2 hours.

For periodically run jobs, the reference time is midnight (12:00 AM). Veeam Backup & Replication always starts counting defined intervals from 12:00 AM, and the first job session will start at 12:00 AM. For example, if you configure a job to run with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM, and so on.

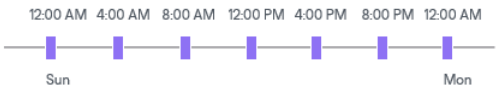

If necessary, you can specify an offset for periodically run jobs. The offset is an exact time within an hour when the job must start. For example, you can configure the job to start with a 4-hour interval and specify an offset equal to 15 minutes. In this case, the job will start at 12:15 AM, 4:15 AM, 8:15 AM, 12:15 PM, 4:15 PM, and so on.

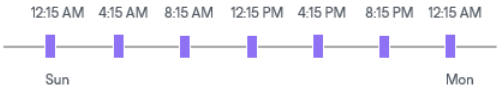


If a session of a periodically run job does not fit into the specified time interval and overlaps the next planned job session, Veeam Backup & Replication starts the next backup job session at the nearest scheduled interval. For example, you set up a job to run with a 4-hour interval. The first job session starts at 12:00 AM, takes 5 hours and is completed at 5:00 AM. In this case, Veeam Backup & Replication will start a new job session at 8:00 AM.

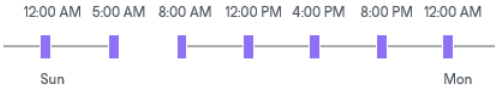


# Jobs Run Continuously

You can schedule the job to run continuously — that is, in a non-stop manner. A new session of a continuously running job starts as soon as the previous job session is completed. Continuously run jobs can help you implement near-continuous data protection (near-CDP) for the most critical applications installed on VMs.

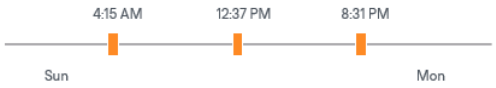


# Chained Jobs

In the common practice, data protection jobs configured in the virtual environment start one after another: when job *A* finishes, job *B* starts, and so on. You can create a chain of jobs using scheduling settings. To do this, you must define the start time for the first job in the chain. For other jobs in the chain, you must select the **After this job** option and choose the preceding job from the list.

Job chaining is not limited to jobs of a specific type only. You can create a chain of jobs of different types. For example, you can:

1. Set a backup job as the first job in the chain.
2. Configure a SureBackup job and chain with the backup job. In this case, Veeam Backup & Replication will automatically verify a backup file created by the backup job after the backup job is complete.

NOTE

Consider the following:

- When you start the initial job manually, Veeam Backup & Replication does not start chained jobs in the *Disabled* state. If at least one job chained directly to the initial job is in the *Enabled* state, Veeam Backup & Replication offers you to start the chained jobs as well. Click **Yes** to start the whole job chain or **No** to start only the first job in the chain.
- If you start the initial job manually and chain another job while the initial job is running, the chained job will start when the initial job completes.
- If you schedule active or synthetic full backups for the chained job, but the initial job does not run on these days, the active and synthetic full backups will not be created for the chained job.
- If a job in the chain fails or is canceled by a user, Veeam Backup & Replication still starts the next chained jobs. Note that if the failed or canceled job was started by a schedule, Veeam Backup & Replication will start the chained job only after all job retries.

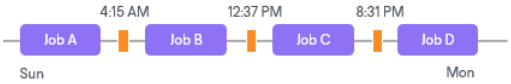


### Recommendations on Job Chaining

You should use job chaining wisely. Job chaining removes guesswork from job scheduling but has a number of drawbacks:

- You cannot predict precisely how much time the initial job will require and when jobs chained to it will start. Depending on the situation, the job schedule may shift, and some operations may not even be performed as planned.

  For example, you configure 2 jobs:

  - *Job 1* is scheduled to start at 10:00 PM daily and typically takes 1 hour.
  - *Job 2* is scheduled to start after *Job 1* daily. Synthetic full backup is scheduled on Saturday.

  Imagine that *Job 1* starts on Saturday and runs for 2.5 hours instead of 1 hour. *Job 2* will then start after midnight on Sunday, and the synthetic full backup planned on Saturday will not be created.

- Errors in job sessions may cause the job schedule to shift. For example, if the initial job in the chain fails, Veeam Backup & Replication will attempt to retry it, and the schedule for chained jobs will shift.

- Load on backup infrastructure resources may not be balanced. Some slots on backup proxies and backup repositories may be available but will not be used since jobs are queued to run one by one. And if you use a backup repository that supports multiple I/O streams, its resources will not be used efficiently.

Instead of job chaining, you can balance the load on backup infrastructure components. To do this, you must limit the number of concurrent tasks on backup proxies and backup repositories. For more information, see Limitation of Concurrent Tasks.

# Job Retry

You can instruct Veeam Backup & Replication to retry a job several times if the initial job run fails. By default, Veeam Backup & Replication automatically retries a failed job 3 times within one job session. You can specify a custom number of retries in the job settings.

Veeam Backup & Replication launches a retry for a job only if the previous job session failed and one or several tasks in the job were not processed. Veeam Backup & Replication does not perform the retry if a job session has finished with the *Success* or *Warning* status. During the job retry, Veeam Backup & Replication processes only failed tasks. If you change the job settings after it fails, consider starting the job manually.

The way Veeam Backup & Replication performs the retry differs depending on the backup chain format of the backup chain.

> **IMPORTANT**
>
> Veeam Backup & Replication does not perform automatic retry for jobs that were started or stopped manually.

# Job Retry for Per-Machine Backups

During one job session, Veeam Backup & Replication creates multiple backup files: one backup file for each VM in the job. If some VMs are not processed during a job run, Veeam Backup & Replication creates backup files only for those VMs that are successfully processed.

During a job retry, Veeam Backup & Replication attempts to process failed VMs. If the processing of failed VMs succeeds, Veeam Backup & Replication creates new backup files and writes data of the processed VMs to these files. If processing fails during all job retries, Veeam Backup & Replication processes the failed VMs during the next job sessions and writes VM data to the backup files created by the current job session.

For example, you have configured a job for 2 VMs: *VM 1* and *VM 2*. The job uses the forward incremental method.

During the first job session, Veeam Backup & Replication successfully processes *VM 1* and creates a full backup file for it. *VM 2* is not processed during all 3 job retries. On the next job session, Veeam Backup & Replication attempts to process the VMs. If processing succeeds, Veeam Backup & Replication creates an incremental backup for *VM 1* and a full backup for *VM 2*. As a result, you have 3 backups:

- Full and incremental backups containing *VM 1* data.
- Full backup containing *VM 2* data.

# Job Retry for Single-File Backups

During one job session, Veeam Backup & Replication creates one backup file for all VMs in the job. If some VMs are not processed during the first job run, Veeam Backup & Replication creates a backup file containing data for those VMs that are successfully processed.

During a job retry, Veeam Backup & Replication attempts to process failed VMs. If processing of failed VMs succeeds, Veeam Backup & Replication writes data of the processed VMs to the backup file that is created at the initial job run. If the processing fails during all job retries, Veeam Backup & Replication processes the failed VMs during the next job sessions and writes VM data to the backup file created by the current job session.

For example, you have configured a job for 2 VMs: *VM 1* and *VM 2*. The job uses the forward incremental method.

During the first job session, Veeam Backup & Replication successfully processes *VM 1* and creates a full backup file for it. *VM 2* is not processed during all 3 job retries. In this case, Veeam Backup & Replication attempts to process the failed *VM 2* within the next job session. Data for *VM 2* is written to the backup file created within this job session, which is an incremental backup. As a result, you have 2 files:

- Full backup file containing a full restore point for *VM 1.*
- Incremental backup file containing a full restore point for *VM 2* and an incremental restore point for *VM 1.*

# Backup Window

If necessary, you can specify a backup window for jobs. The backup window is a period of time on weekdays when jobs are permitted to run. If the job exceeds the allowed window, Veeam Backup & Replication will automatically stop this job. Also, if a job is in progress and enters the denied window, Veeam Backup & Replication will stop it immediately and reschedule according to the schedule settings.

The backup window can be helpful if you do not want data protection jobs to produce unwanted overhead for the production environment or do not want jobs to overlap production hours. In this case, you can define the time interval during which the job must not run.

> **IMPORTANT**
>
> The backup window affects only the data transport process and health check retry operations. Other transformation processes can be performed in the target repository outside the backup window. Linked jobs that process Microsoft SQL transaction logs and Oracle archived logs are not affected by the backup window settings.

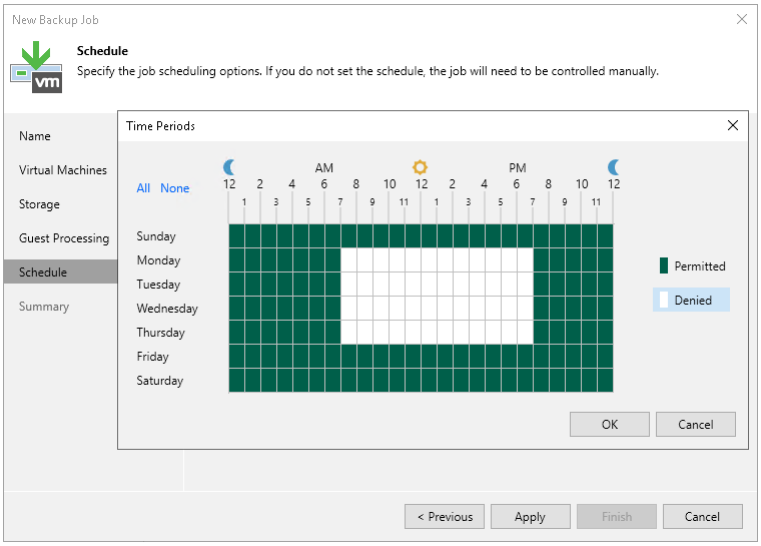


# Backup Window for Periodically Run Jobs

If you define the backup window for a job that runs periodically at specific time intervals, Veeam Backup & Replication will immediately start the job after the denied window is over. All subsequent backup job sessions will be performed according to specified scheduling settings.

For example, you have configured a job to run with a 4-hour interval with an offset of 15 minutes. The allowed backup window for the job is 7:00 PM to 8:00 AM. Veeam Backup & Replication will run this job in the following way:

1. The first job session will start at 12:15 AM (since midnight is a reference time for periodically run jobs).
2. The next job session will start at 4:15 AM.
3. The job session at 8:15 AM will not be performed as it falls into the denied period of the backup window.
4. The next job session will start immediately after the denied period is over: at 7:15 PM.
5. After that, Veeam Backup & Replication will run the job according to the defined schedule: at 8:15 PM, 12:15 AM and so on.

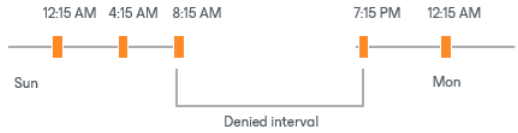


# Manual Start of Backup Jobs

You can start jobs manually if you need to capture VM data at a specific point in time and do not want to re-configure job scheduling settings. For example, you can start a job to create a VM backup before you install new software on a VM or enable a new feature.

When you start the job manually, Veeam Backup & Replication runs a regular job session that produces a new restore point in the backup chain in the backup repository.

To start and stop jobs configured on the backup server, you can use the **Start** and **Stop** buttons on the ribbon or the commands in the shortcut menu.

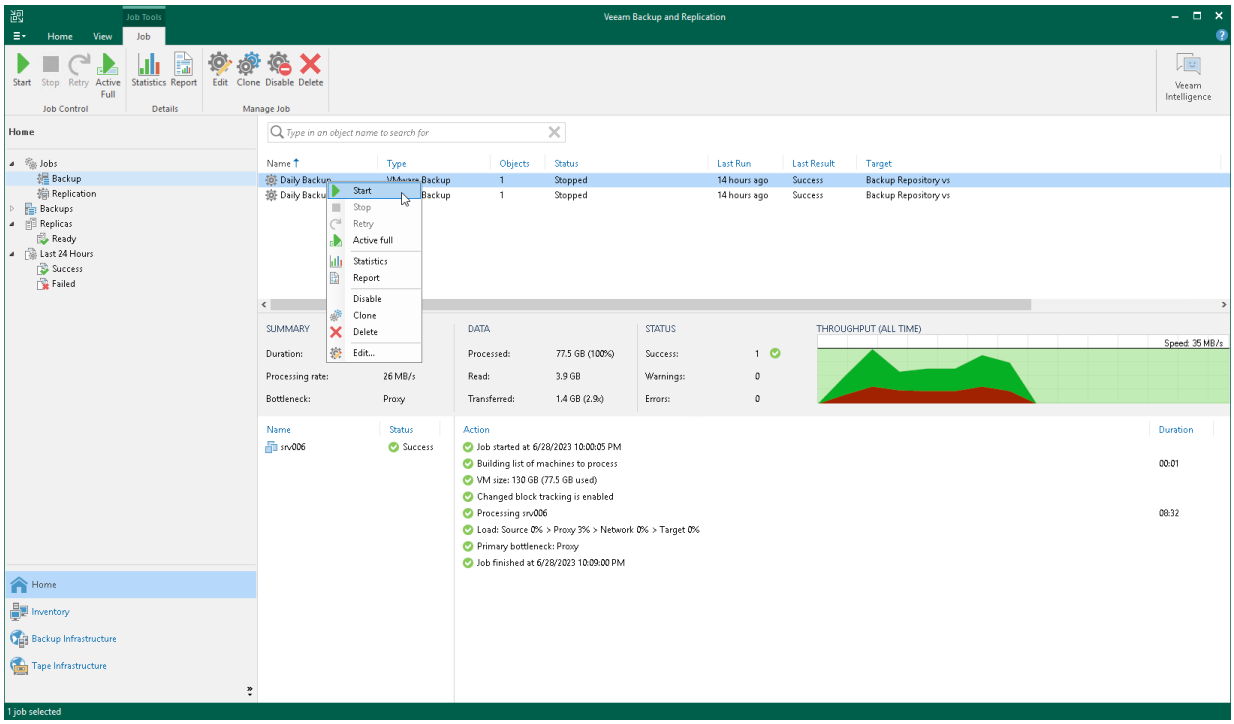

# Manual Stop of Backup Jobs

You can stop job execution at any moment. For example, you can stop a job if the job processes several VMs, but the workload appears to be greater than you expected. Or you can stop the job if there is not enough time to finish the job session.

You can stop a job in 2 ways:

- You can stop the job immediately. In this scenario, Veeam Backup & Replication terminates the job session and does not create a new restore point for VMs that are currently processed.
- You can stop the job gracefully. In this scenario, Veeam Backup & Replication creates a restore point for the VMs that are currently processed and then terminates the job session.

# Immediate Stop of Jobs

Immediate job stop terminates the job session instantly. The job finishes with the following results:

- VMs that Veeam Backup & Replication has succeeded in processing by the time you stop the job will have new restore points.
- VMs that Veeam Backup & Replication is currently processing and VMs that Veeam Backup & Replication has not started to process will not have new restore points.

When you stop a job session immediately, Veeam Backup & Replication performs the following operations:

1. If a snapshot for a VM has already been created, Veeam Backup & Replication instructs VMware vSphere to remove the snapshot.
2. Veeam Backup & Replication terminates all job processes and tasks. The job is finished with the *Failed* error.

All restore points created with the previous job sessions remain untouched. You can use them for restore operations.

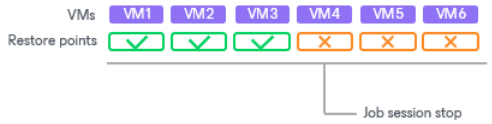


# Graceful Stop of Jobs

Graceful job stop instructs Veeam Backup & Replication that it must create restore points for VMs that are currently being processed, and then terminate the job. The job finishes with the following results:

- VMs that Veeam Backup & Replication has succeeded to process and VMs that are being processed will have new restore points.
- VMs that Veeam Backup & Replication has not started to process will not have new restore points.

You can use graceful job stop for the following types of jobs:

- Unstructured Data Backup jobs
- VM Backup jobs
- VM Copy jobs
- Replication jobs

You cannot use graceful job stop for the following types of jobs:

- File Copy jobs
- Backup Copy jobs
- Quick Migration job (during Quick Migration, Veeam Backup & Replication processes all VMs in one task)
- Restore operations

VMs added to the job are processed in the order defined in job settings. Information about VMs that have already been processed and VMs that are being processed is displayed in job details.

If you stop the job gracefully before Veeam Backup & Replication starts processing the first VM in the job, the job will be finished with the *Failed* error. You will see the message *Operation was canceled by user* in the job details.

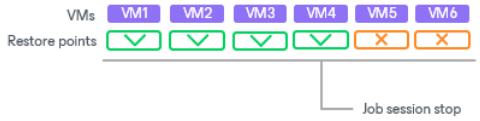

# Immutability for Backup Files

Veeam Backup & Replication allows you to prohibit the deletion of data from backup repositories by making that data temporarily immutable. It is done to increase security: immutability protects your data against loss due to attacks, malware activity or any other harmful actions.

> **NOTE**
>
> After you enable immutability, Veeam Backup & Replication will not support the deletion of immutable backups until the immutability expiration date comes.

You can enable the immutability feature for the following types of backup repositories:

- Dell Data Domain. For more information, see Immutability for Dell Data Domain.
- Hardened repository. For more information, see Hardened Repository.
- HPE StoreOnce. For more information, see HPE StoreOnce Supported Features.
- Object storage repository. For more information, see Immutability for Object Storage Repositories.
- Scale-out backup repository. For more information, see Immutability for Scale-out Backup Repository.

# Health Check for Backup Files

You can instruct Veeam Backup & Replication to periodically perform a health check for the latest restore point in the backup chain. During the health check, Veeam Backup & Replication performs a cyclic redundancy check (CRC) for metadata and a hash check for VM data blocks in the backup file to verify their integrity. The health check helps ensure that the restore point is consistent, and you can restore data from this restore point.

Health check can be performed for all types of backup chains:

- Forever forward incremental
- Forward incremental
- Reverse incremental backup chains

To run the health check periodically, you must enable the **Perform backup files health check** option in the backup job settings and define the health check schedule. By default, the health check is performed at 5:00 on the last Saturday of every month. You can change the schedule and run the health check weekly or monthly on specific days.

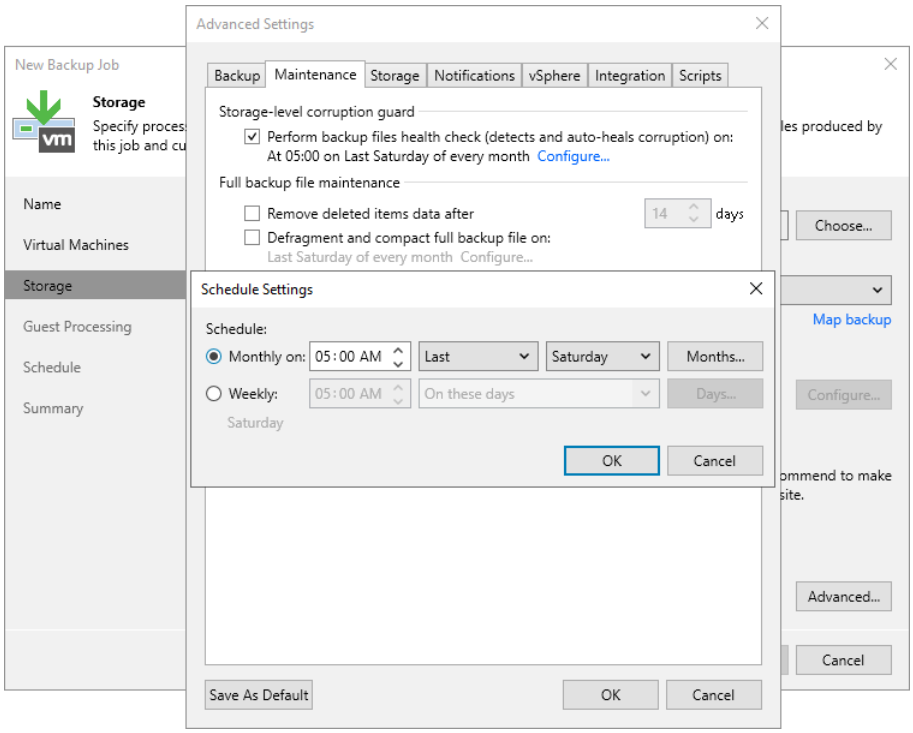


# Verification Content

The health check always verifies only the latest restore point in the backup chain. In case of forever forward incremental and forward incremental backup chains, if the latest restore point is incomplete, the health check verifies the restore point preceding the latest one.

Note that the health check procedure verifies not the latest backup file in the backup chain, but the latest restore point for a VM. The latest restore point corresponds to the state of the VM at the date and time when the latest backup file for this VM was created. Data blocks that are required to "compose" the VM latest state are typically spread out across several backup files in the backup chain. Therefore, to verify the latest state of the VM, Veeam Backup & Replication must open several backup files in the backup chain and read data blocks from these backup files. For this reason, the health check procedure may take long.

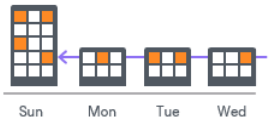


The health check verifies only those virtual disks of a VM that are available in the latest restore point. For example, you added a VM with 3 virtual disks to a backup job. The VM was backed up Sunday through Tuesday. On Wednesday, you removed 1 virtual disk, and Veeam Backup & Replication run the health check for the VM. During the health check, Veeam Backup & Replication will verify only the 2 remaining virtual disks.

The health check verifies only those VMs that are available in the latest restore point. For example, you added 2 VMs to a backup job and ran the job several times. The health check verified 2 VMs. If you remove 1 VM from the backup job, the next scheduled health check run will verify the latest unverified restore point for the removed VM, and the latest restore point for the remaining VM. In future, the health check will verify only the restore point for the remaining VM in the job.

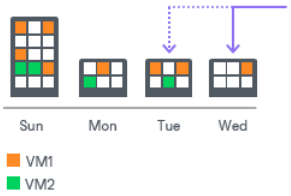


# Limitations for Health Check

Consider the following limitations for the health check:

- For data blocks located in the capacity tier, the health check works as described in the Health Check for Capacity Tier section.
- [For per-machine backup chains] If you add a new VM to an existing backup job that has been run for some time, Veeam Backup & Replication will perform the health check for it during the next incremental backup job session for the added VM.
- Linux immutable repositories do not support repair. If the health check detects corrupted data, Veeam Backup & Replication marks the restore point as corrupted in the configuration database and finishes the health check session. In that case, you need to perform the active full backup for this backup chain. Otherwise, Veeam Backup & Replication will finish every backup job session creating subsequent incremental restore points with the *Error* status.

# How Health Check Works

When Veeam Backup & Replication saves a new restore point to the backup repository, it calculates CRC values for backup metadata and hash values for data blocks of VM disk in the backup file, and saves these values in the metadata of the backup file, together with VM data. During the health check session, Veeam Backup & Replication uses these values to make sure that a verified restore point is consistent.

> **NOTE**
>
> If you perform health check for encrypted backup files, Veeam Backup & Replication will pass encryption keys to the regular backup repository or cloud repository. For more information on encryption, see Data Encryption.

Veeam Backup & Replication uses different mechanisms of health check for different types of backup chains:

- Forever forward incremental and forward incremental backup chains
- Reverse incremental backup chains

### Forever Forward Incremental and Forward Incremental Backup Chains

The health check for forward incremental backup chains is performed in the following way:

1. The health check starts according to the schedule. Depending on the storage type of the repository, the health check procedure can be performed on different infrastructure components:
   - For direct attached storage — on the repository itself.
   - For network attached storage — on the gateway server.
   - For deduplicating storage appliances — on the gateway server.

   > **NOTE**
   >
   > Consider the following:
   >
   > - The gateway server placement may influence the time needed for the health check procedure when using network attached storage and deduplicating storage appliances as repositories. As the health check is performed on the gateway server, we recommend to locate gateway servers as close to the backup repository as possible.
   > - The health check runs only after all jobs and operations that process a backup are completed (for example, a backup job, a restore job, synthetic operations with backups and so on). If some of these activities are started, the health check stops.
   > - If a health check session does not complete until the next scheduled run, this session will stop and a new session will start according to the schedule.
   > - If you add new VMs to a backup job which block data has been processed by the health check, the health check will verify only block data of new VMs and will skip backups of previously added VMs.

2. The health check calculates CRC values for backup metadata and hash values for VM disks data blocks in the backup file and compares them with the CRC and hash values that are already stored in the backup file.

   During the health check, Veeam Backup & Replication verifies the latest restore point in the backup chain (restore point created by the last backup job session). If the latest restore point in the backup chain is incomplete, Veeam Backup & Replication checks the restore point preceding the latest one.

3. If the health check does not detect data corruption, the backup job session completes in a regular way.

   If the health check detects corrupted data, Veeam Backup & Replication completes the backup job with the *Error* status and starts the health check retry process. The health check retry starts as a separate backup job session.

Depending on the revealed data corruption, Veeam Backup & Replication performs the following actions:

- If the health check has detected corrupted backup metadata in the full backup file, Veeam Backup & Replication marks the backup chain starting from this full restore point as corrupted in the configuration database. During the health check retry, Veeam Backup & Replication transports data blocks of the whole VM image from the source datastore, creates a new full backup file in the backup repository and saves transported data blocks to it.
- If the health check has detected corrupted backup metadata in the incremental backup file, Veeam Backup & Replication removes information about this incremental restore point and subsequent incremental restore points from the configuration database. During the health check retry, Veeam Backup & Replication transports new incremental data relative to the latest valid restore point in the backup chain from the source datastore, creates a new incremental backup file in the backup repository and saves transported data blocks to it.
- If the health check has detected corrupted VM disk blocks in the full or incremental backup file, Veeam Backup & Replication marks the restore point that includes the corrupted data blocks and subsequent incremental restore points as corrupted in the configuration database. During the health check retry, Veeam Backup & Replication transports data blocks from the source datastore. In addition, Veeam Backup & Replication transports data blocks that have changed since the backup job session that has triggered the health check. Veeam Backup & Replication stores these data blocks to the latest restore point that has been created by the current backup job session (session that has triggered the health check retry).

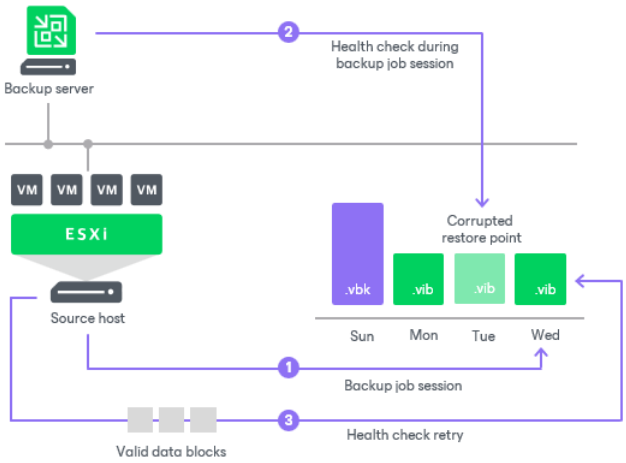


### Reverse Incremental Backup Chains

In the case of reverse incremental backup chains, the health check always verifies only the latest restore point in the backup chain, which is always a full backup file.

The health check is performed in the following way:

1. The health check starts according to the schedule. Depending on the storage type of the repository, the health check procedure can be performed on different infrastructure components:
   - For direct attached storage — on the repository itself.
   - For network attached storage — on the gateway server.
   - For deduplicating storage appliances — on the gateway server.
2. Veeam Backup & Replication verifies the full backup file. It calculates CRC values for backup metadata and hash values for VM disks data blocks in the full backup file, and compares them with the CRC and hash values that are already stored in the full backup file.

3. If the health check does not detect data corruption, the backup job session completes in a regular way.

   If the health check detects corrupted data, Veeam Backup & Replication completes the backup job with the *Error* status and starts the health check retry process. The health check retry starts as a separate backup job session.

   Depending on the revealed data corruption, Veeam Backup & Replication performs the following actions:

   - If the health check has detected corrupted backup metadata in the full backup file, Veeam Backup & Replication marks the whole backup chain (full backup file and preceding reverse incremental backup files) as corrupted in the configuration database. During the health check retry, Veeam Backup & Replication transports data blocks of the whole VM image from the source datastore, creates a new full backup file in the backup repository and saves transported data blocks to it.
   - If the health check has detected corrupted VM disk blocks in the full backup file, Veeam Backup & Replication marks the full backup file and preceding reverse incremental backup files as corrupted in the configuration database. During the health check retry, Veeam Backup & Replication transports data blocks from the source datastore. In addition, Veeam Backup & Replication transports data blocks that have changed since the backup job session that has triggered the health check. Veeam Backup & Replication stores these data blocks to the existing full backup file in the backup repository. Corrupted data blocks that have been replaced with data blocks from the source datastore are stored to an existing reverse incremental backup file preceding the full backup file.

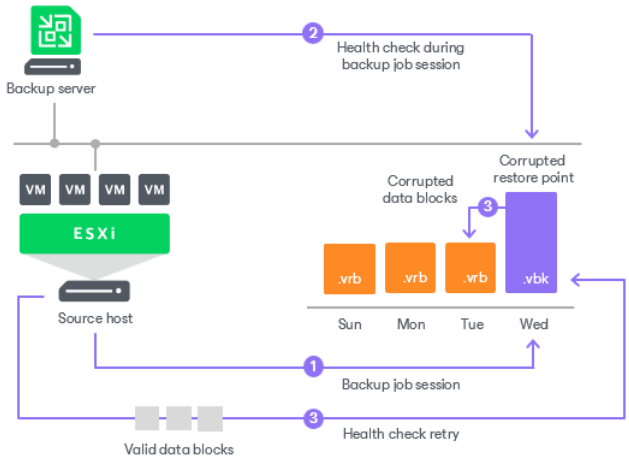


# Health Check Retry Mode

If the health check detects corrupted data, the backup job will switch to the *Retry* mode and will start the health check retry process. During the health check retry, Veeam Backup & Replication transports data blocks of the whole VM image from the source datastore, creates a new full backup file in the object storage repository and saves transported data blocks to it.

If Veeam Backup & Replication fails to fix the corrupted data during all health check retries, you must retry the job manually. In this case, Veeam Backup & Replication will transport the required data blocks from the source datastore, to fix the latest restore point. If the latest restore point in the backup chain is incomplete, Veeam Backup & Replication will attempt to fix the restore point preceding the latest one.

For scheduled jobs, the number of health check retries is equal to the number of job retries specified in the job settings. For jobs started manually, Veeam Backup & Replication performs 1 health check retry.

**IMPORTANT**

To allow Veeam Backup & Replication add the repaired data blocks to the latest restore point after completing the health check, the backup job must meet the following requirements:

- You must schedule the backups job to run automatically. For more information, see the Schedule step of the **New Backup Job** wizard.
- The target object storage is not set to the Maintenance mode.
- Backup window settings must allow a backup job to run after the health check completes.

If the backup job does not meet these requirements, Veeam Backup & Replication will add the repaired data blocks to a new restore point, created during a next run of the backup job.

# Compact of Full Backup File

The backup job constantly transforms the full backup file in the backup chain to meet retention policy settings. This transformation process, however, has a side effect. Over time, the full backup file becomes large and fragmented. Consequently, operations of reading and writing data from and to the backup file slow down.

To resolve this problem, Veeam Backup & Replication allows you to defragment and compact the full backup file periodically. During the file compact operation, Veeam Backup & Replication creates a new full backup file in the target repository: it copies existing data blocks from the old backup file, rearranges and stores them close to each other. As a result, the full backup file gets defragmented, its size reduced, and the speed of reading and writing from and to the file increases.

To compact the full backup file periodically, you must enable the **Defragment and compact full backup file** option in the backup job settings and define the compact operation schedule. By default, the compact operation is performed on the last Saturday of every month. You can change the compact operation schedule and instruct Veeam Backup & Replication to perform it weekly or monthly on specific days.

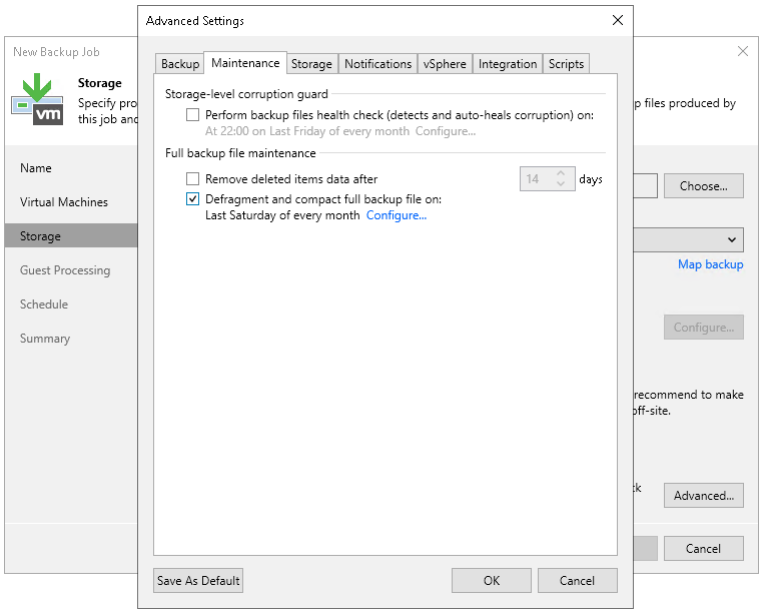


# Limitations for Full Backup File Compact

The full backup file compact has the following limitations:

- The **Defragment and compact full backup file** option works for forever forward incremental or reverse incremental backup chains. For this reason, you must not schedule active or synthetic full backups.

  Although you do not schedule active full backups for forever forward incremental or reverse incremental backup chains, you can create full backups. For example, you can create them manually or Veeam Backup & Replication can create them during the health check. On the day when active full backups are triggered, Veeam Backup & Replication does not create compact full backups. Veeam Backup & Replication will create them on another day during the backup job session.

- Veeam Backup & Replication does not compact full backup files that have been offloaded to cloud-based object storage. For more information, see Capacity Tier.

- The backup repository must have enough space to store a file of the full backup size. During the compact process, Veeam Backup & Replication creates auxiliary files that exist in the backup repository until the end of the compact operation.

- [For per-machine backup chains] If you add a new VM to an existing backup job that has been run for some time, Veeam Backup & Replication will perform the compact full operation for it during the next incremental backup job session for the added VM.
- If you change the block size in the backup job settings, Veeam Backup & Replication does not change the block size in the compacted backup file till the next full backup. However, if you change compression settings in the backup job settings, during the next compact file operation, Veeam Backup & Replication changes the compression level for the compacted backup file.

# Removal of Deleted VMs Data

During the compact operation, Veeam Backup & Replication does not copy all data blocks from the VBK file to the newly created file. It copies only data blocks of VMs whose information is stored in the configuration database. For example, if the VM is removed from the backup job, its data is not copied to the new full backup file. This approach helps reduce the size of the full backup file and remove unnecessary data from it.

# VM Data Take Out

If the full backup file contains data for a VM that has only one restore point and this restore point is older than 7 days, during the compact operation, Veeam Backup & Replication will extract data for this VM from the full backup file and write this data to a separate full backup file. Such backup will be displayed under the **Backups** > **Disk (Imported)** node in the **Home** view.

The mechanism works if the following conditions are met:

- The **Remove deleted VMs data** option is not enabled in the backup job settings.
- The **Use per-machine backup files** option is not enabled in backup repository settings.

# Backup Move

Veeam Backup & Replication allows you to move all backups of a backup job to another repository or to move specific workloads and their backups to another job.

Moving backups to another repository can be helpful if you are running out of free space on a repository and want to move all backups created by a backup job to another repository, target the backup job to this repository and continue the backup chain.

Moving workloads and their backups to another backup job can be helpful if you want to separate one backup job into multiple backup jobs or if you want to change backup settings for specific workloads. Also, you want to continue the existing backup chains for the workloads.

## How Moving to Another Repository Works

When moving backups to another repository and targeting the job to this repository, Veeam Backup & Replication performs the following:

1. Disables the backup job.
2. Copies backup files to a new location.
3. Targets the backup job to the new location.
4. Deletes source backup files.
5. After the successful move, Veeam Backup & Replication enables the backup job. The enabled backup job continues the backup chain for the existing and moved backups.

> **NOTE**
>
> The described algorithm is common. Nuances, requirements and limitations are described further in this section.

## How Moving to Another Job Works

When moving backups to another job, Veeam Backup & Replication performs the following:

1. Disables the source and target jobs.
2. If backups are moved within one repository without immutability, Veeam Backup & Replication invokes the native move method of the file system. This applies if the repository is Windows, Linux, CIFS, NFS, deduplicating storage appliance or a scale-out backup repository of one of the listed types. In other cases, Veeam Backup & Replication copies backup files of selected workloads to the repository where the target job saves backups.
3. Excludes the selected workloads from being processed by the source job.
4. Includes the selected workloads to be processed by the target job.
5. After the successful move, Veeam Backup & Replication enables the source and target jobs. The target backup job continues the backup chain for the existing and moved backups.

> **NOTE**
>
> The described algorithm is common. Nuances, requirements and limitations are described further in this section.

# Considerations and Limitations

This section lists considerations and known limitations for different types of jobs and different types of repositories where backups can be stored.

## Common Considerations

Consider the following:

- You can move individual workloads and their backups only if backups are per-machine backups with separate metadata files and only to repositories with the **Use per-machine backup files** check box enabled.
- You can move backups between repositories with the **Use per-machine backup files** check box selected and not selected. However, the move operation does not change the backup chain format (single-file backup, per-machine backup with single metadata file or per-machine backup with separate metadata files). The job will continue the backup chain of the original backup chain format. If you want to change the backup chain format of the chain, use the detach from job operation. For more information, see Upgrading Backup Chain Formats.

  [For moving to object storage repositories] You can move only per-machine backups with separate metadata files.
- [For VeeamZIP] You cannot move VeeamZIP backups to direct object storage repositories or when they are added as the performance tier of a scale-out backup repository. However, you can still move VeeamZIP backups to object storage repositories used as the capacity or archive tiers.
- After you move workloads and their backups between jobs, you must adjust the guest processing settings for the moved workloads. After the move, Veeam Backup & Replication uses the default guest processing setting. For more information on guest processing settings, see Guest Processing.
- You cannot move a workload and its backups to a job that already processes this workload.
- You cannot move backups from encrypted jobs to unencrypted jobs and vice versa.
- You cannot move backups imported without the metadata (.VBM) file.
- You cannot launch the restore operation for backups that are being moved at the moment.
- You cannot move backups from an HPE StoreOnce repository used as a target for a backup copy for HPE StoreOnce repositories.
- You cannot move backups created by backup copy jobs in the legacy periodic copy mode. To move backups, you first need to upgrade the backup format as described in section Upgrading Backup Chain Formats.
- After you move backups between backup copy jobs, Veeam Backup & Replication does not include and exclude workloads from the backup copy jobs. You should do that manually.
- You cannot move backups created by Veeam plug-ins for Enterprise applications, Veeam Cloud Plug-ins (Veeam Backup for AWS, Veeam Backup for Google Cloud, Veeam Backup for Microsoft Azure), Veeam Backup for Nutanix AHV, Veeam Backup for OLVM and RHV, Veeam Backup for Proxmox VE and Veeam Kasten.
- You cannot move Unstructured Data Backups.
- You cannot move backups created from the vSphere Self-Service Backup Portal.
- You cannot move backups created in the Veeam Cloud Connect repositories. For more information on Veeam Cloud Connect repositories, see the Cloud Repository section in the Veeam Cloud Connect Guide.

- The Move backup operations to, from and between cloud repositories are not supported. For more information on Veeam Cloud Connect's limitations, see the Veeam Cloud Connect Backup section in the Veeam Cloud Connect Guide.
- To use the move to another repository functionality and move backups to an object storage repository, check that defragment and compact full backup file, reverse incremental backup method and synthetic full backup are disabled for the job. Alternatively, you can use the move to another job functionality and move backups to a job already targeted to an object storage repository.
- Before moving workloads and their backups to a deduplicating storage appliance, Veeam Backup & Replication checks whether such a move can damage the backup chain. If so, Veeam Backup & Replication prohibits the move operation. After Veeam Backup & Replication completes the move operation, it checks that the target backup job is configured correctly. However, you must change the settings manually. You can find which backup settings need to be changed in the move session or in this Veeam KB article. For more information on how to view session statistics, see Viewing Job Session Results.
- You can move VMware Cloud Director backups of a whole job or individual vApps. You cannot move backups of VMs.
- If you move workloads and their backups between jobs linked to the same tape job, the tape job continues to create incremental backups within the existing backup chain. Otherwise, the tape job will create an active full backup for the moved workloads.
- Traffic throttling is not supported for backup move operations.

## Backup Jobs with Linked Backup Copy Jobs

If you move workloads and their backups from a backup job that is linked to a backup copy job, the following applies:

- If the source and target backup jobs are linked to backup copy jobs, Veeam Backup & Replication also moves backups and workloads between backup copy jobs. Note that Veeam Backup & Replication only performs the move if each backup job is linked to one backup copy job. For example, *Backup Job 1* processes *VM1* and *VM2* and is linked to *Backup Copy Job 1*. *Backup Job 2* processes *VM3* and *VM4* and is linked to *Backup Copy Job 2*. You move *VM1* and its backups to the *Backup Job 2*. Veeam Backup & Replication will also move the *VM1* and its backups from *Backup Copy Job 1* to *Backup Copy Job 2*. As a result, *Backup Job 2* and *Backup Copy Job 2* will process *VM3, VM4* and *VM1.*

  If a backup job is linked to multiple backup copy jobs, you need to move workloads and backups between the backup copy jobs manually. To do that, use the move to another job operation as described in section Moving Backups. If you do not move workloads between backup copy jobs, the backup copy job linked to the target backup job will create active full backup copies for the moved workloads.
- If the source backup job is still in progress, the move operation fails.
- Before moving backup copies, Veeam Backup & Replication waits till the source move backup process finishes successfully. Only then will Veeam Backup & Replication disable the linked backup copy job and move backup copies.

## Log Backups

If you move workloads and their backups from a backup job for which log backup is configured, the following applies:

- Veeam Backup & Replication moves log backups along with workloads.
- Veeam Backup & Replication disables the log backup jobs gracefully before moving log backups.
- [For SQL server Always On availability groups] You cannot move individual nodes from a job.

- [For Oracle Data Guard] You must move all servers from the job. If you move only one server, the application item restore stops working. To repair the application item restore functionality, you will need to move the server back to the original job.
- If guest processing is disabled in the target job, log backups are moved, however they will not be processed by the target job. To process log backups, make sure that guest processing is enabled for the target job as described in section Specify Guest Processing Settings.

## Backups with GFS Flags

If you move backups from a backup or backup copy job where long-term (GFS) retention is configured, the following applies:

- When you move all backups to another repository, Veeam Backup & Replication preserves all backups with GFS flags.
- When you move a workload and its backups to another job, Veeam Backup & Replication starts to retain backups according to the settings configured in the target job. Veeam Backup & Replication deletes all GFS backups that do not comply with the retention policy in the target job. For example, if yearly backups are not configured in the target job, Veeam Backup & Replication will remove them after the move operation finishes.

## Agent Backups

For more information on moving Veeam Agent backups, see the Moving Backups section in the Veeam Agent Management Guide.

## Immutable Repositories

If you move workloads and their backups from a repository for which immutability is enabled, the following applies:

- During the move operation, Veeam Backup & Replication copies the whole backup chain. After the move operation finishes, original backups stored in the immutable repository are moved to the node with the (**Orphaned**) postfix.
- Veeam Backup & Replication retains backups moved to the (**Orphaned**) node according to the retention period specified in the job from which the backups were moved. If the retention period is shorter than the immutability period specified in the repository settings where the backups were stored, Veeam Backup & Replication sets the retention period as equal to the immutability period.

  If the retention period is set in days, Veeam Backup & Replication deletes a backup after its retention period ends. If the retention period is set in restore points, Veeam Backup & Replication does not delete backups. You can delete these backups manually after the retention period ends.

If you move a backup chain to a repository with different immutability or retention settings, Veeam Backup & Replication will replace the source repository settings with the settings of the target repository for that specific backup chain.

## Scale-Out Backup Repositories

If you move workloads and their backups from a scale-out backup repository, the following applies:

- Veeam Backup & Replication moves backups only from the performance tier. If you want to move data from the capacity tier, you must first download it to the performance tier. For more information, see Downloading Data from Capacity Tier.

- Veeam Backup & Replication does not support moving backups between extents of a scale-out backup repository. To learn how to manage backups within the scale-out backup repository, see Scale-Out Backup Repositories.
- Backups stored in the capacity and archive tiers are moved to the node with the **(Orphaned)** postfix. The backups are retained according to the retention settings of the job from which the backups were moved. If the retention period is set in days, the Veeam Backup & Replication retains the backups according to the configured retention and deletes the backups from the repository after the retention period ends. If the retention period is set in restore points, Veeam Backup & Replication leaves backup files in a backup chain. The minimum number of backup files left equals the current retention period. You can delete these backup files manually as described in section Deleting Backups from Disk.
- Veeam Backup & Replication does not move backups from extents in the Maintenance mode.

# Resume on Disconnect

Veeam Backup & Replication can handle a situation of an unstable network during backup, backup copy and replication jobs. If a network connection drops for a short period during the data transport process, Veeam Backup & Replication automatically resumes the dropped network connection. The data transfer process starts from the point when the connection was lost. The resume on disconnect capability improves the reliability of remote data transfer, reduces the backup window and minimizes the network load.

Veeam Backup & Replication automatically re-establishes a connection between the following backup infrastructure components engaged in the data transfer process:

- Backup server
- Backup proxies
- Backup repository

Resume on disconnect works only for dropped network connections. Veeam Backup & Replication attempts to resume the connection at intervals of 15 seconds for 30 minutes. If the problem has any other nature, Veeam Backup & Replication retries the job in a regular manner.

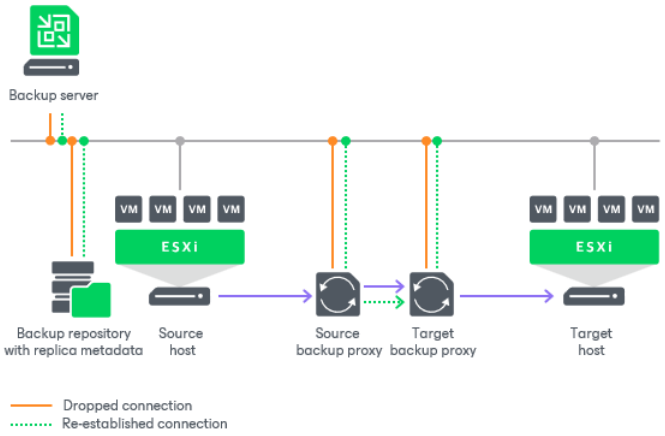


Veeam Backup & Replication does not create a new restore point on resume: VM data is written to the same restore point that was created for the current job session. When resuming the data transfer process, Veeam Backup & Replication regards VM disks, not the whole VM.

For example, a VM has 2 disks: *disk A* and *disk B*. Before the connection dropped, Veeam Backup & Replication managed to transfer 20 GB of *disk A* and did not start transferring *disk B*. After the connection is re-established, Veeam Backup & Replication will start transferring the data for *disk A* from the 20 GB point; data of the whole *disk B* will be transferred anew.

# Snapshot Hunter

The Snapshot Hunter is a Veeam technology used to detect and remove orphaned snapshots that may remain after backup or replication job sessions.

The Snapshot Hunter addresses the problem of "phantom" snapshots. Under some circumstances, VMware vSphere can report a successful removal of a snapshot, but the snapshot actually remains on the datastore.

Phantom snapshots can take substantial space on the datastore or impact VM performance. They can even cause the production VMs to stop if the datastore runs out of free space.

To solve the problem of phantom snapshots, Veeam Backup & Replication starts the Snapshot Hunter during each backup or replication job session. The Snapshot Hunter looks for snapshot files not registered in vSphere. If there are no orphaned files, the Snapshot Hunter stops. If orphaned snapshot files are detected, the Snapshot Hunter removes them in the background mode.

The Snapshot Hunter runs in jobs that use VMware VM snapshots:

- Backup jobs: regular backup and backup from storage snapshots
- Replication jobs (the source VM snapshot): regular replication, replication from storage snapshots
- VeeamZIP

> **NOTE**
>
> During Snapshot Hunter analysis, Veeam Backup & Replication skips VMware Cloud Director VMs.

## How Snapshot Hunter Works

A temporary snapshot of the VM is taken and then removed during every backup or replication job session. To remove the snapshot, Veeam Backup & Replication triggers the VMware snapshot consolidation mechanism that includes two steps:

1. VMware vSphere removes the snapshot from the VM snapshots list.
2. VMware vSphere consolidates the data written to the delta file with the VM disks.

The problem occurs when the snapshot was removed successfully, but the consolidation failed. This may happen, for example, if the files appear to be locked when VMware vSphere attempts to consolidate the snapshot files. In this case, the files remain on the datastore.

The Snapshot Hunter is started as a separate process scheduled within every job session. The discovery of phantom snapshots does not affect the job. If the phantom snapshots are discovered, the Veeam Backup Service schedules the snapshot consolidation, and the job runs normally.

Veeam Backup & Replication checks the datastore to discover orphaned snapshot files. To consolidate these files with the VM disks, Veeam Backup & Replication calls a consolidation algorithm. The algorithm consists of three steps, each representing a VMware method.

1. VMware Consolidate method

   As a first attempt, Veeam Backup & Replication calls the VMware Snapshot Consolidate method. This method is the same mechanism that VMware vSphere uses for VMs with the *Needs Consolidation* status.

2. Hard consolidation without quiesce

   If the first attempt fails, Veeam Backup & Replication creates a new snapshot and calls the VMware *Delete all snapshots* method. As a result, all VM snapshots and associated files are deleted. The snapshot is taken without quiescing the VM.

3. Hard consolidation with quiesce

   If the snapshot deletion still fails, Veeam Backup & Replication implies another VMware method that creates a quiesced snapshot and then removes all VM snapshots.

> **NOTE**
>
> Hard consolidation without quiesce and hard consolidation with quiesce are performed only if the VM does not have any user snapshots. In case there are one or more user snapshots, these steps will not be performed.

The 3-steps consolidation procedure is launched up to 4 times with a 4-hour interval.

In case all four attempts fail, Veeam Backup & Replication sends an e-mail notification informing that the user needs to manually troubleshoot the problem. Note that you need to have the global email notifications option enabled. For more information, see Specifying Email Notification Settings.

The Snapshot Hunter considers the backup window set for the job. If any of the attempts does not fit the backup window, Veeam Backup & Replication will not perform the consolidation and send the e-mail notification.

To view information on the Snapshot Hunter sessions, in the Veeam Backup & Replication console, open the **History** view and select **System** in the inventory pane.

In case no consolidation attempt could fit the backup window, the warning appears in the job statistics.

# Creating Backup Jobs

To back up VMs, you must configure a backup job. The backup job defines how, where and when to back up VM data. You can use one job to process one or more VMs. Jobs can be started manually or scheduled to run automatically at a specific time.

Before you create a backup job, check the prerequisites. Then, use the New Backup Job wizard to configure the backup job.

# Before You Begin

Before you create a backup job, check the following:

- Backup infrastructure components that will take part in the backup process must be added to the backup infrastructure and properly configured. These include ESXi hosts on which VMs are registered, backup proxy and backup repository.
- The backup repository must have enough free space to store created backup files. To receive alerts about low space in the backup repository, configure global notification settings. For more information, see Specifying Other Notification Settings.
- If you plan to map a backup job to a backup file that already exists in the backup repository, you must perform the rescan operations for this backup repository. Otherwise, Veeam Backup & Replication will not be able to recognize backup files in the backup repository. For more information, see Rescanning Backup Repositories.
- If you plan to use an HPE StoreOnce repository as a target for backup jobs, check the limitations and requirements for the repository.
- If you plan to configure a secondary destination for the backup job, you can create a backup copy job or backup to the tape job beforehand. While creating the backup copy or backup to tape job, you can leave the source empty and link the source to the job later. For more information, see the Creating Backup Copy Jobs and Creating Backup to Tape Jobs sections in the Veeam Backup & Replication User Guide.
- If you plan to enable guest processing for the job, check Guest Processing Requirements and Limitations.
- If you plan to use VM Guest OS File Indexing, you need enough disk space to store indexing data. For more information, see Requirements and Limitations for VM Guest OS File Indexing.
- If you plan to use pre-job and post-job scripts and pre-freeze and post-thaw scripts, you must create scripts before configuring the backup job.
- If you plan to perform maintenance operations with backup files periodically, consider the following limitations: Health Check for Backup Files, Retention Policy for Deleted Items and Compact of Full Backup File.
- If you assign the role of a backup proxy to a VM, you should not add this VM to the list of processed VMs in a job that uses this backup proxy. Such configuration may result in degraded job performance. Veeam Backup & Replication will assign this backup proxy to process other VMs in the job first, and processing of this VM itself will be put on hold. Veeam Backup & Replication will report the following message in the job statistics: *VM is a backup proxy, waiting for it to stop processing tasks*. The job will start processing this VM only after the backup proxy deployed on the VM finishes its tasks.
- If you use tags to categorize virtual infrastructure objects, check the limitations for VM tags. For more information, see VM Tags.
- If you plan to use the Backup from Storage Snapshots technology, check the requirements and limitations for backup proxies in the Storage System Snapshot Integration Guide. Also, check other requirements in the guide.
- If a job cannot be completed within 21 days period, it will be stopped with the *Failed* status.
- If you plan to protect VMware Cloud Director objects, consider using the dedicated backup job. For more information, see Backup for VMware Cloud Director.

# Step 1. Launch New Backup Job Wizard

To launch the **New Backup Job** wizard, do one of the following:

- On the **Home** tab, click **Backup Job** > **Virtual machine** > **VMware vSphere**.
- Open the **Home** view. In the inventory pane, right-click **Jobs** and select **Backup** > **Virtual machine** > **VMware vSphere**.
- Open the **Inventory** view. In the working area, select the VMs, click **Add to Backup** on the ribbon and select **New job** or right-click the VMs and select **Add to backup job** > **New job**.
  Veeam Backup & Replication will start the **New Backup Job** wizard and add the VMs to this job. You can add other VMs to the job later on, when you pass through the wizard steps.
- You can quickly add the VMs to an already existing job. To do this, open the **Inventory** view. In the working area, select the VMs and click **Add to Backup** > *name of the job* on the ribbon or right-click the VMs and select **Add to backup job** > *name of the job*.

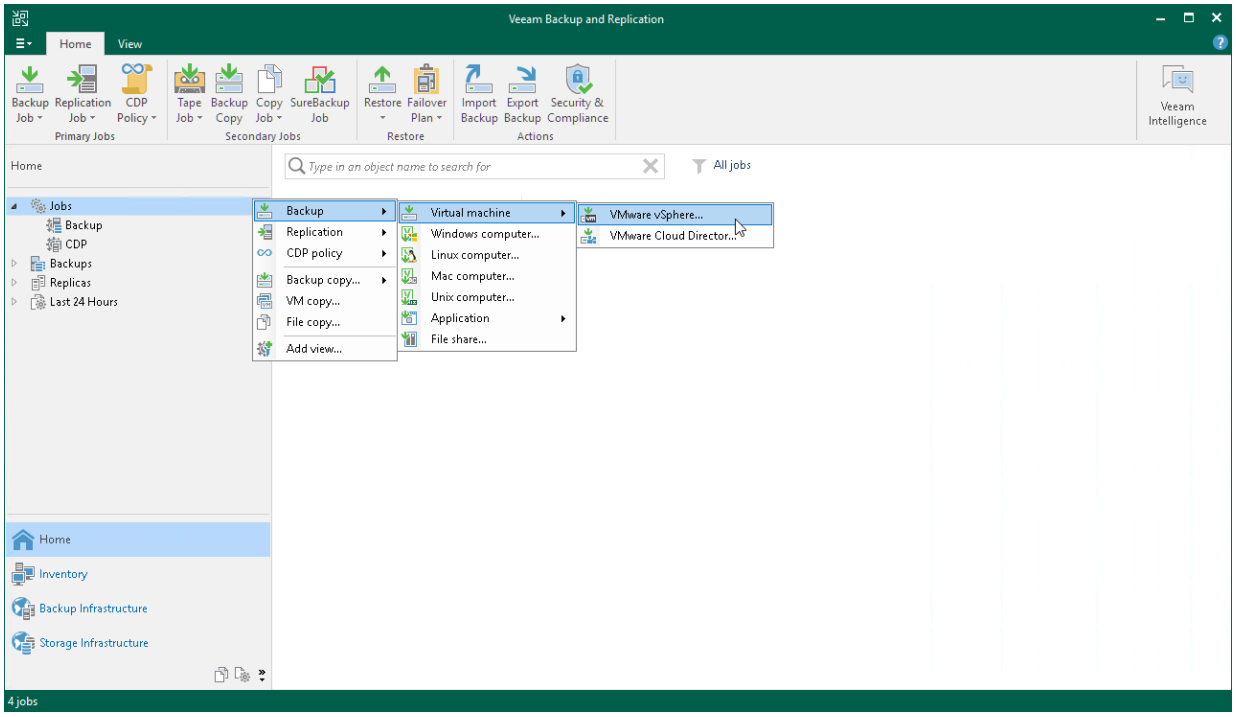

# Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a name and description for the backup job.

1. In the **Name** field, enter a name for the backup job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the job, date and time when the job was created.
3. Select the **High priority** check box if you want the resource scheduler of Veeam Backup & Replication to prioritize this job higher than other similar jobs and allocate resources to it in the first place. For more information on job priorities, see Job Priorities.

> **TIP**
>
> In the list of jobs in the Veeam Backup & Replication console, jobs with the **High priority** option enabled are marked with a red flag ( ).

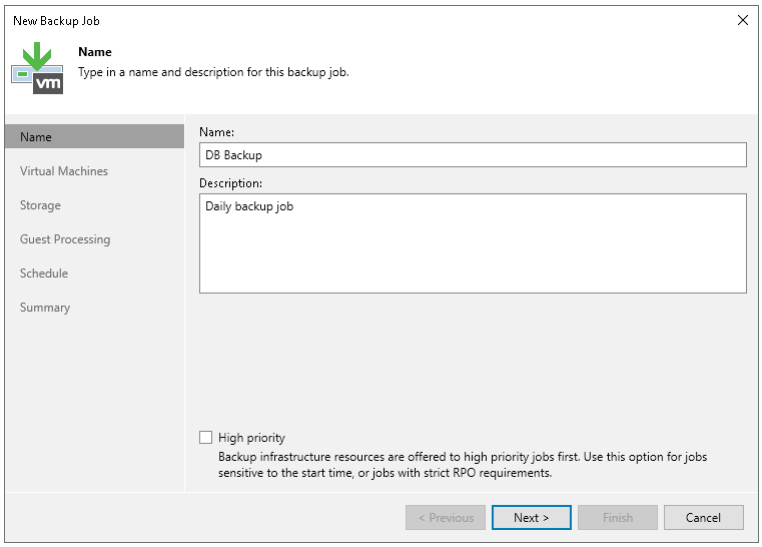

# Step 3. Select VMs to Back Up

At the **Virtual Machines** step of the wizard, select VMs and VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags) that you want to back up:

1. Click **Add**.
2. In the **Add Object** window, select the necessary VMs or VM containers and click **Add**. If you select VM containers and add new VMs to this container in the future, Veeam Backup & Replication will update backup job settings automatically to include these VMs.

   You can use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select the **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags, and only those VMs that have all the selected tags will be processed by the job.

   To quickly find the necessary VMs, you can use the search field at the bottom of the **Add Object** window. If you want to switch between the types of VMs you want to search through, use the button to the left of the search field.

> **NOTE**
>
> You can use a regular backup job to process VMs that are part of vApps created in the vCenter Server. To back up VMware Cloud Director vApps, you must use specifically developed VMware Cloud Director backup jobs. For more information, see Backup and Restore of vApps.

The total size of objects added to the job is displayed in the **Total size** field. Use the **Recalculate** button to refresh the total size value after you add a new object to the job.

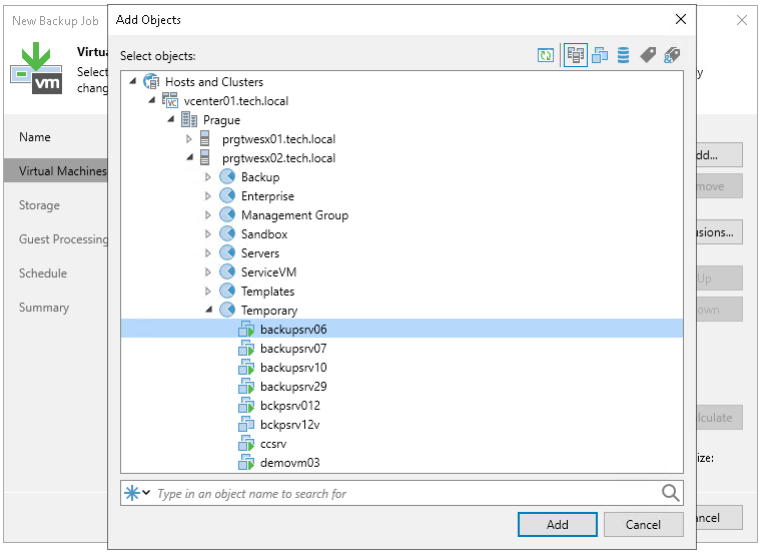

# Step 4. Exclude Objects from Backup Job

After you add VMs and VM containers to the job, you can specify which objects you want to exclude from the backup. You can exclude the following types of objects:

- VMs from VM containers
- Specific VM disks
- VM templates

> **NOTE**
>
> Veeam Backup & Replication automatically excludes VM log files from backup to make the backup process faster and reduce the size of the backup file.

To exclude VMs from a VM container:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. Click the **VMs** tab.
3. Click **Add**.
4. Use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only those VMs that have all the selected tags will be excluded from the job.
5. In the displayed tree, select the necessary object and click **Add**. Use the **Show full hierarchy** check box to display the hierarchy of all VMware Servers added to the backup infrastructure.
6. Click **OK**.

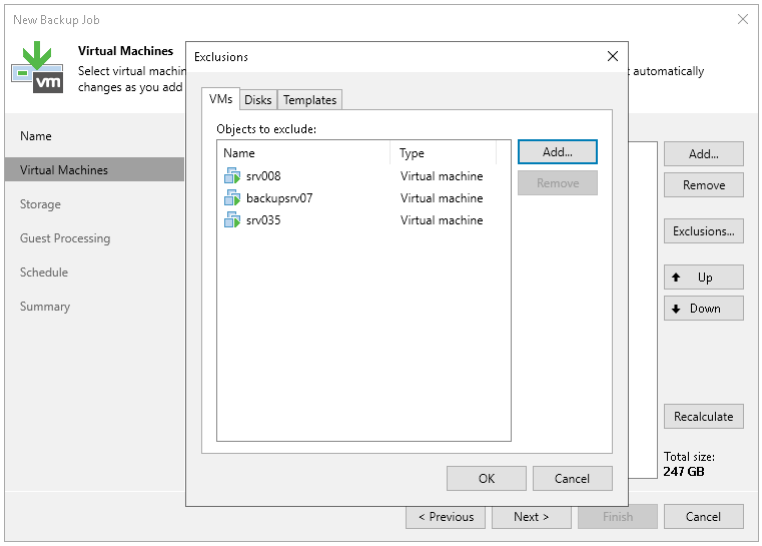


To exclude VM disks:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. Click the **Disks** tab.
3. Select the VM in the list and click **Edit**. If you want to exclude disks of a VM added as part of the container, click **Add** to include the VM in the list as a standalone object.

4. Choose disks that you want to back up. You can choose to process all disks, 0:0 disks (typically, system disks) or add to the list custom IDE, SCSI or SATA disks.
5. Select the **Remove excluded disks from VM configuration** check box. Veeam Backup & Replication will modify the VMX file of a backed-up VM to remove excluded disks from the VM configuration. If you restore this VM from the backup file to a location where excluded disks are not accessible with the original paths, you will not have to manually edit the VM configuration file to be able to power on the VM.

**NOTE**

Consider the following:

- If you exclude disks from a backup and enable application-aware processing, Veeam Backup & Replication will still perform application-aware processing for the excluded disks. This means that VSS will process disk data.
- When you exclude disks from a backup, they will only be removed from the backup operation. It is important to note that even if a disk is excluded from the backup, it might still be included in a snapshot created by the hypervisor.

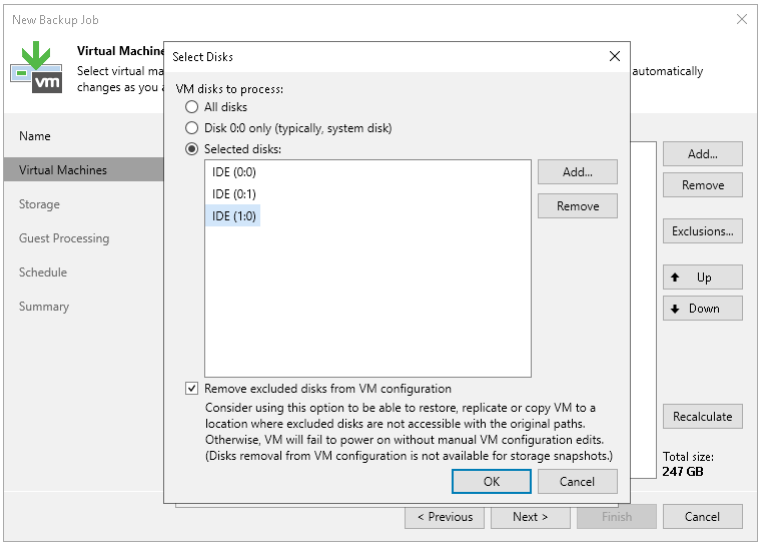


To exclude VM templates:

1. At the **Virtual Machines** step of the wizard, select a VM container and click **Exclusions**.
2. Click the **Templates** tab.
3. Clear the **Backup VM** templates check box.

4. If you want to include VM templates into the full backup only, leave the **Backup VM templates** check box selected and select the **Exclude templates from incremental backup** check box.

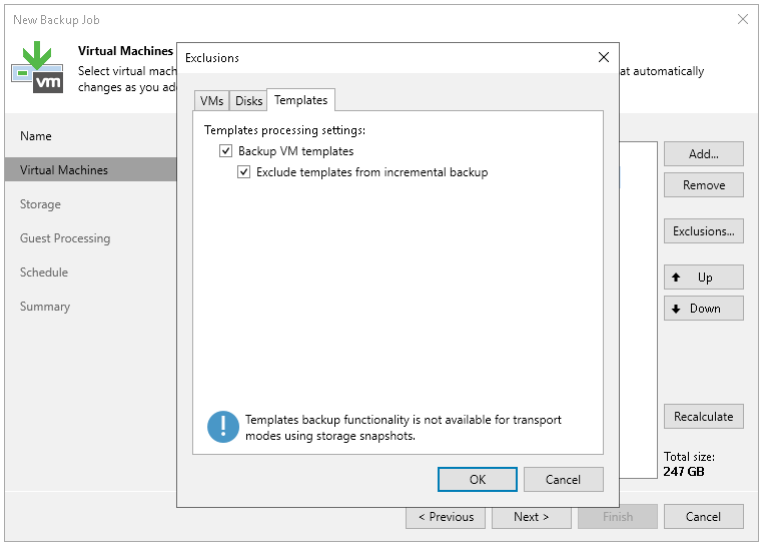

# Step 5. Define VM Backup Order

You can define the order in which the backup job must process VMs. Setting VM order can be helpful, for example, if you add some mission-critical VMs to the job and want the job to process them first. You can set these VMs first in the list to ensure their processing fits the backup window.

VMs inside a VM container are processed at random. To ensure that VMs are processed in the defined order, you must add them as standalone VMs, not as a part of the VM container.

To define the VM backup order:

1. At the **Virtual Machines** step of the wizard, select a VM or VM container.
2. Use the **Up** and **Down** buttons on the right to move the VM or VM container up or down in the list.

> **NOTE**
>
> VMs may be processed in a different order. For example, if backup infrastructure resources for a VM that is higher on the priority list are not available, and resources for a VM that is lower on the list are available, Veeam Backup & Replication will start processing the VM that is lower on the list first.

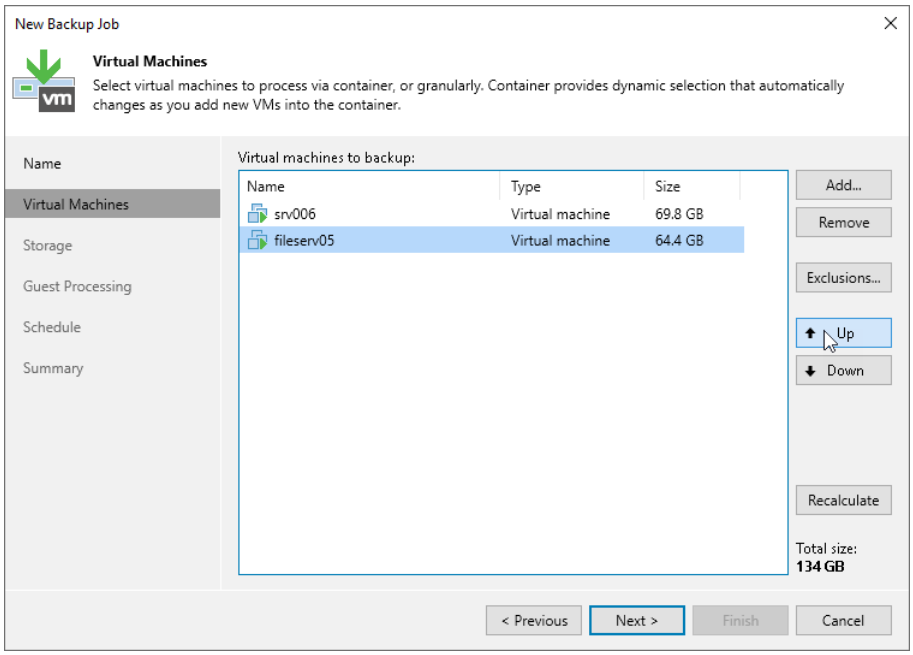

# Step 6. Specify Backup Storage Settings

At the **Storage** step of the wizard, select backup infrastructure components for the job — backup proxy and backup repository, and specify backup storage settings.

1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies with access to the source datastore and automatically assign an optimal backup proxy to process VMs in the job.

     Veeam Backup & Replication assigns backup proxies to VMs included in the backup job one by one. Before processing a new VM in the VM list, Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use to retrieve VM data and the current workload on the backup proxies to select the most appropriate one for VM processing.
   - If you choose **Use the selected backup proxy servers specified below**, you can explicitly select backup proxies that the job must use. It is recommended that you select at least two backup proxies to ensure that the backup job starts if one of the proxies fails or loses its connectivity to the source datastore.

   If you plan to use the Backup from Storage Snapshots technology, check the requirements and limitations for backup proxies in the Storage System Snapshot Integration Guide.
2. From the **Backup repository** list, select a backup repository where the created backup files must be stored. When you select a backup repository, Veeam Backup & Replication automatically checks how much free space is available in the backup repository.

   > **NOTE**
   >
   > Consider the following:
   >
   > - If you change the repository after the job has already run, Veeam Backup & Replication suggests that you move the existing backups to the new repository. If you want to move the backups, check the limitations and considerations in Backup Move.
   > - If you select an object storage repository or a scale-out backup repository which performance tier consists of object storage repositories, Veeam Backup & Replication will not provide the amount of free space in this repository since its capacity is constantly expanding.
3. You can map the job to a specific backup stored in the backup repository. Backup job mapping can be helpful if you have moved backup files to a new backup repository and want to point the job to existing backups in this new backup repository. You can also use backup job mapping if the configuration database gets corrupted and you need to reconfigure backup jobs.

   To map the job to a backup, click the **Map backup** link and select the backup in the backup repository. Backups can be easily identified by job names. To find the backup, you can also use the search field at the bottom of the window.
4. In the **Retention Policy** field, specify retention policy settings for restore points:
   - If you want to keep the last N restore points, select *restore points* from the drop-down list. Then, specify the number of restore points.
   - If you want to keep all restore points created during the last N days, select *days* from the drop-down list. Then, specify the number of days.

   When the specified number is exceeded, the earliest restore point is removed from the backup chain or merged with the next closest restore point. For more information, see Short-Term Retention Policy.

> **NOTE**
>
> If you enable the GFS retention, the short-term retention policy will not be able to delete and merge the GFS backup files. Thus, the backup chain will have more restore points than specified in the short-term retention policy.

5. If you want to archive backup files created by the backup job to a secondary destination, select the **Configure secondary destination for this job** check box. The **New Backup Job** wizard will include the **Secondary Target** step, where you will link the job to the required destination. You can use the following destinations:
   - Backup repository (by linking to a backup copy job)
   - Tape (by linking to a tape backup job)
   - Storage array (by linking to a storage array feature)

   You can enable this option only if a secondary destination job or feature is already configured. For more information, see Specify Secondary Target.

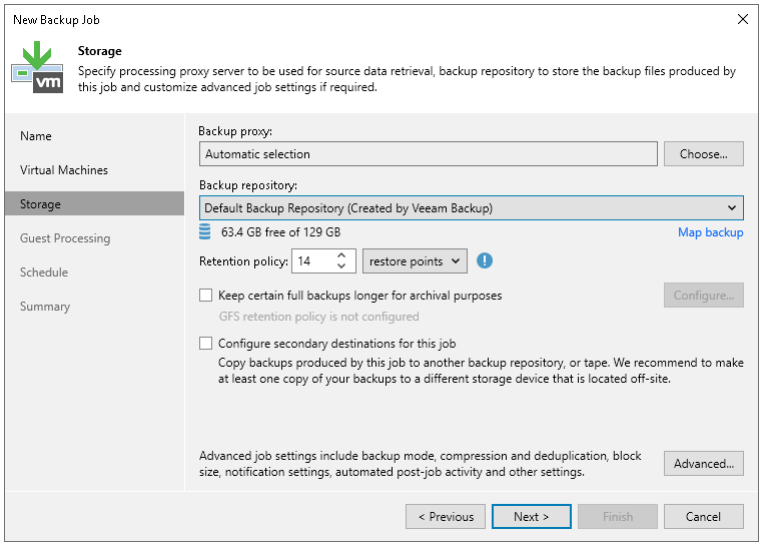

# Step 7. Configure Long-Term Retention

If you want to ignore the short-term retention policy for some full backups and store them for long-term archiving, you can configure a long-term retention policy (or GFS retention policy) for the backup job. For more information on GFS and its limitations, see Long-Term Retention Policy (GFS).

1. At the storage step of the wizard, select the **Keep certain full backups longer for archival purposes** check box.
2. Click **Configure**.
3. In the **Configure GFS** window, do the following:
   - If you want to create weekly restore points, select the **Keep weekly full backups for** check box. Then, specify the number of weeks during which you want to prevent restore points from being modified and deleted.

     From the **If multiple full backups exist, use the one from** drop-down list, select a week day when Veeam Backup & Replication must assign the weekly GFS flag to a full restore point.
   - If you want to create monthly restore points, select the **Keep monthly full backups for** check box. Then, specify the number of months during which you want to prevent restore points from being modified and deleted.

     From the **Use weekly full backup for the following week of a month** drop-down list, select a week when Veeam Backup & Replication must assign the monthly GFS flag to a full restore point. A week equals 7 calendar days; for example, the first week of May is days 1–7, and the last week of May is days 25–31.
   - If you want to create yearly restore points, select the **Keep yearly full backups for** check box. Then, specify the number of years during which you want to prevent restore points from being modified and deleted.

     From the **Use monthly full backup for the following month** drop-down list, select a month when Veeam Backup & Replication must assign the yearly GFS flag to a full restore point.

- If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

> **NOTE**
>
> If you select to assign multiple types of GFS flags, the flags begin to depend on each other. For more information on this dependency, see Algorithm for Multiple Flag Types.

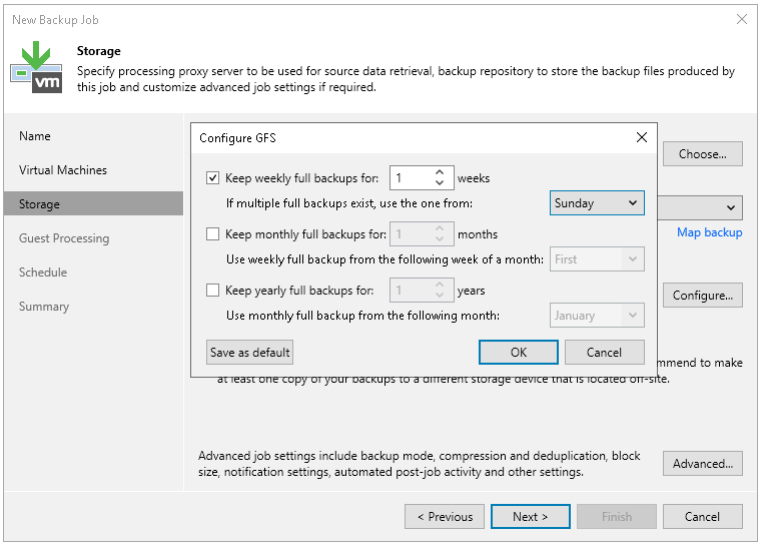
New Backup Job
Storage
Specify processing proxy server to be used for source data retrieval, backup repository to store the backup files produced by this job and customize advanced job settings if required.
Name
Virtual Machines
Storage
Guest Processing
Schedule
Summary
Configure GFS
Keep weekly full backups for: 1 weeks
If multiple full backups exist, use the one from: Sunday
Keep monthly full backups for: 1 months
Use weekly full backup from the following week of a month: First
Keep yearly full backups for: 1 years
Use monthly full backup from the following month: January
Save as default
OK
Cancel
Choose...
Map backup
Configure...
Advanced job settings include backup mode, compression and deduplication, block size, notification settings, automated post-job activity and other settings.
Advanced...
< Previous
Next >
Finish
Cancel

# Step 8. Specify Advanced Backup Settings

At the **Storage** step of the wizard, specify advanced settings for the backup job:

- Backup settings
- Maintenance settings
- Storage settings
- Notification settings
- vSphere settings
- Integration settings
- Script settings

> **TIP**
>
> After you specify necessary settings for the backup job, you can save them as default settings. To do this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new backup job, Veeam Backup & Replication will automatically apply the default settings to the new job.

## Backup Settings

To specify settings for a backup chain created by the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. On the **Backup** tab, select the backup method that you want to use to create the backup chain in the backup repository:
   - To create a reverse incremental backup chain, select **Reverse Incremental**.

     > **IMPORTANT**
     >
     > Consider the following:
     >
     > - Dell Data Domain and HPE StoreOnce do not support the reverse incremental backup method.
     > - Direct backup to object storage repositories does not support the reverse incremental backup method.

   - To create an incremental backup chain, select **Incremental** and enable synthetic full or active full backups, or both (see items 3-4).
   - To create a forever forward incremental backup chain, select **Incremental**. Also, check that synthetic full and active full backups are disabled (see items 3-4).

   For more information, see Backup Methods.
3. You can select to periodically create synthetic full backups if you choose the incremental backup method. Select the **Create synthetic full backups periodically on** check box and click **Configure** to schedule synthetic full backups on the necessary days of the week.

> **IMPORTANT**
>
> For backup jobs created in Veeam Backup & Replication version prior to 11, you could enable the Transform previous backup chains into rollbacks option. Starting from Veeam Backup & Replication version 11, this option was deprecated and the feature could no longer be selected in new jobs. However, jobs that were already using this feature will continue to do so after upgrading to Veeam Backup & Replication 11. Starting from Veeam Backup & Replication 12, this feature is fully deprecated, and the presence of any jobs still using this feature blocks the upgrade to Veeam Backup & Replication 12. For more details, see this Veeam KB article.

4. You can select to periodically create active full backups with any backup mode enabled. Select the **Create active full backups periodically on** check box and click **Configure** to define scheduling settings.

   Before you schedule periodic full backups, you must ensure you have enough free space in the backup repository. As an alternative, you can create active full backups manually when needed. For more information, see Active Full Backup.

> **IMPORTANT**
>
> Consider the following:
>
> - If you schedule the active full backup and synthetic full backup on the same day, Veeam Backup & Replication will perform only the active full backup. Synthetic full backup will be skipped.
> - If you schedule a job to start after another job (initial job), but the initial job does not run on days when the synthetic or active full backup is scheduled for the chained job, Veeam Backup & Replication will not create active or synthetic full backups. For more information on the job schedule options, see Define Job Schedule.
> - Synthetic full backups cannot be created independently for backup jobs targeted at object storage. To include synthetic full backups in the backup process, you must enable the GFS policy.

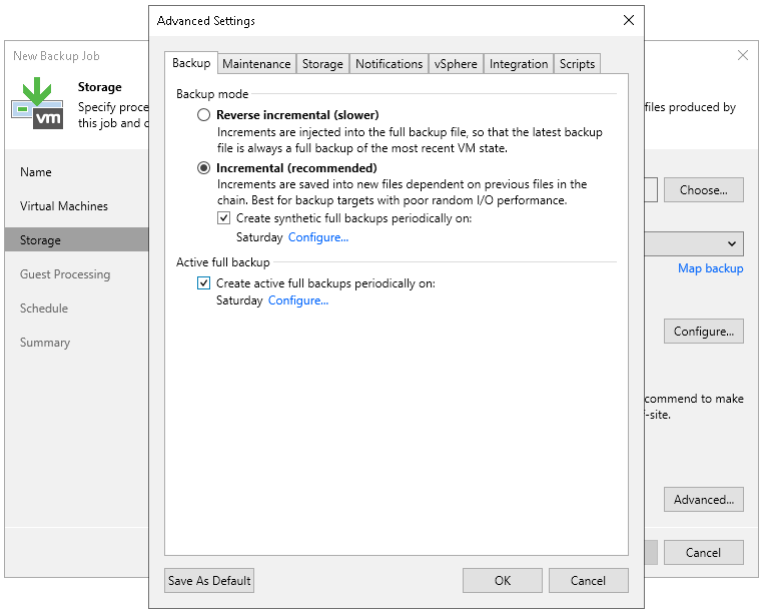

# Maintenance Settings

You can instruct Veeam Backup & Replication to periodically perform the health check for the latest restore point in the backup chain and run maintenance operations to make sure that the backup chain remains valid and consistent.

To specify the health check and maintenance settings, at the **Storage** step of the **New Backup Job** wizard, click **Advanced** and click the **Maintenance** tab.

# Specifying Health Check Settings

To specify health check settings for the backup job:

1. In the **Storage-level corruption guard** section, select the **Perform backup files health check** check box and specify the time schedule for the health check. By default, if the health check is enabled, it is performed monthly at 5:00 AM every last Saturday (at 10:00 PM every last Friday - for version 12).
2. To specify the health check schedule, click the **Configure** link.
3. In the **Schedule Settings** window, specify whether you want to perform the health check monthly or weekly and specify the schedule settings.

# Specifying Backup File Maintenance

To specify maintenance settings for the backup job:

1. To specify the retention period for deleted VMs, select the **Remove deleted items data after** check box and specify the number of days for which you want to keep backup data for deleted VMs.

   > **NOTE**
   >
   > Consider the following:
   >
   > - If a VM is no longer available (for example, it was deleted or excluded from the job), Veeam Backup & Replication will keep its data in the backup repository for the period that you have specified. When this period is over, data of the deleted VM will be removed from the backup repository.
   > - By default, the retention period for deleted VM data is 14 days. It is strongly recommended that you set the retention period to 3 days or more to prevent unwanted data loss. For more information, see Retention Policy for Deleted Items.

2. To periodically compact a full backup, select the **Defragment and compact full backup file** check box.

   > **IMPORTANT**
   >
   > Consider the following:
   >
   > - Direct backup to object storage repositories does not support the **Defragment and compact full backup file** option.
   > - The HPE StoreOnce backup repository does not support the **Defragment and compact full backup file** option.

3. To specify the schedule for the compact operation, click the **Configure** link.
4. In the **Schedule Settings** window, specify whether you want to compact a full backup monthly or weekly and specify the schedule settings.

5. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

> **IMPORTANT**
>
> If you schedule periodic full backups, the **Defragment and compact full backup file** check box does not apply.

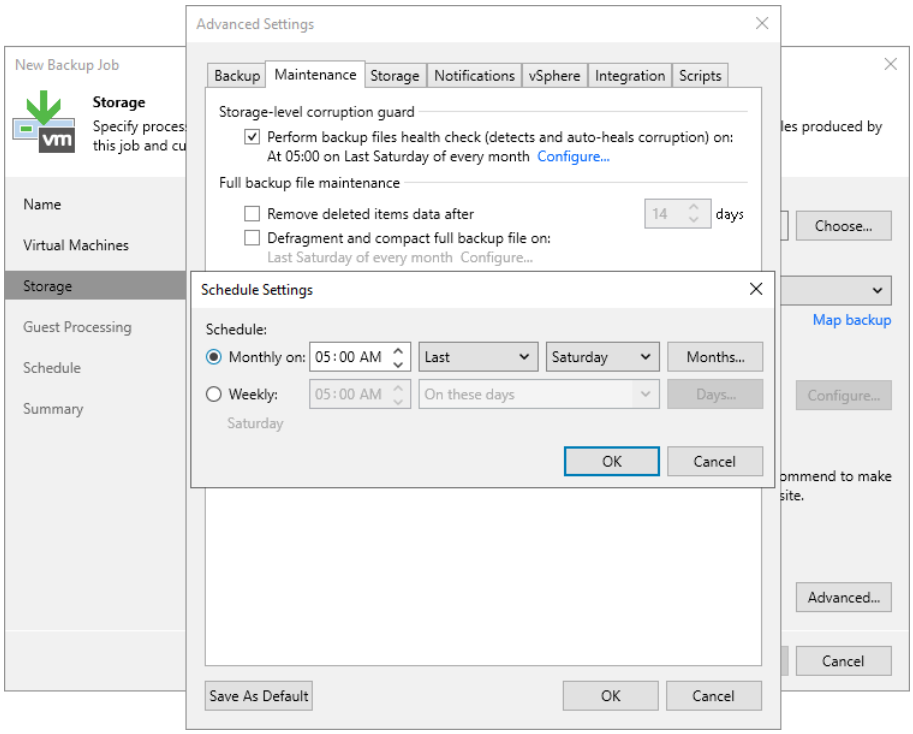


# Storage Settings

To specify storage settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Storage** tab.
3. By default, Veeam Backup & Replication deduplicates VM data before storing it in the backup repository. Data deduplication provides a smaller backup file size but may reduce the backup job performance. For more information on data deduplication, see Deduplication.

   To disable data deduplication, clear the **Enable inline data deduplication** check box. When you disable data deduplication, you also change the workflow for incremental backups. If Changed Block Tracking (CBT) is enabled for the job, Veeam Backup & Replication will save all data blocks marked by CBT as new to the destination storage without performing an additional check or using Veeam filtering mechanism. It will result in faster creation of incremental backups. However, backup files may increase in size.

   > **NOTE**
   >
   > The option that allows you to enable or disable data deduplication is not available for an object storage repository due to its underlying functionality. The backup to object storage applies various improvements to optimize storage consumption and make the backup process faster.
   >
   > If you have selected such a repository at the **Storage** step of the wizard, the **Enable inline data deduplication** check box will not be displayed.

4. By default, Veeam Backup & Replication checks the NTFS MFT file on VMs with Microsoft Windows OS to identify data blocks of the `hiberfil.sys` file (file used for the hibernate mode) and `pagefile.sys` file (swap file) and excludes these data blocks from processing. The swap file is dynamic and changes intensively between backup job sessions, even if the VM itself does not change much. Processing of service files reduces the job performance and increases the size of incremental backup files.

   If you want to include data blocks of the `hiberfil.sys` file and `pagefile.sys` file in the backup, clear the **Exclude swap file blocks** check box. For more information, see Swap Files.

5. By default, Veeam Backup & Replication does not copy deleted file blocks ("dirty" blocks on the VM guest OS) to the target location. This option lets you reduce the size of backup files and increase job performance. If you want to include dirty data blocks in the backup, clear the **Exclude deleted file blocks** check box. For more information, see Deleted File Blocks (BitLooker).

6. From the **Compression level list**, select a compression level for the backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*. For more information on data compression, see Data Compression and Deduplication.

7. In the **Storage optimization** section, select the block size that will be used to process VMs. For more information on the data block sizes and how they affect performance, see Storage Optimization.

8. To encrypt backup files, select the **Enable backup file encryption** check box. In the **Password** field, select the encryption method you want to use — a password that will generate secret keys or an external KMS server that will generate KMS keys. For more information about backup job encryption, see Encrypting Backup Jobs.

   If you have not created the password beforehand, click **Add** or use the *Manage passwords* link to specify a new password. For more information, see Password Manager.

   If you do not use password loss protection, Veeam Backup & Replication will display a warning about it. For more information, see Password Loss Protection.

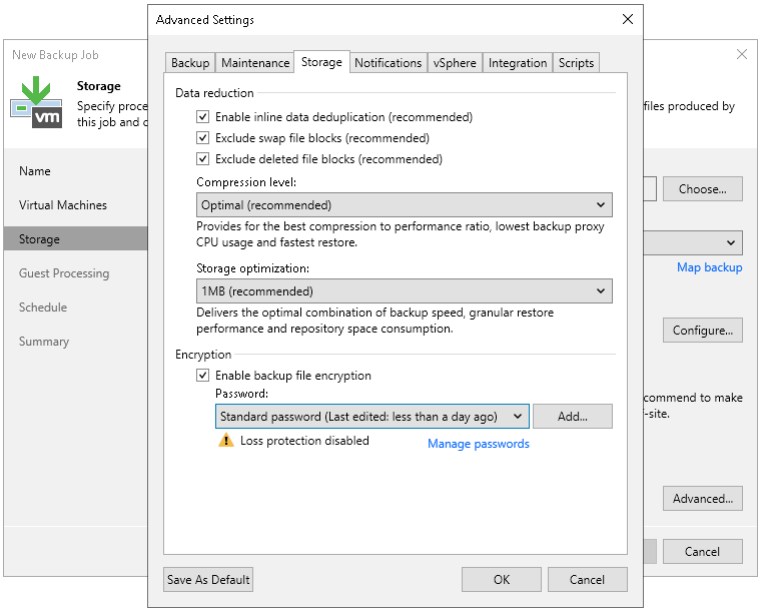


# Notification Settings

To specify notification settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.

3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully.

   SNMP traps will be sent if you specify global SNMP settings in Veeam Backup & Replication and configure software on the recipient's machine to receive SNMP traps. For more information, see Specifying SNMP Settings.

4. Select the **Send email notifications to the following recipients** check box if you want to receive notifications about the job completion status by email. In the field under the check box, specify the recipient's email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication. For more information, see Configuring Global Email Notification Settings.

5. You can choose to use global notification settings or specify custom notification settings.

   - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see Configuring Global Email Notification Settings.
   - To configure a custom notification for the job, select **Use custom notification settings specified below** check box. You can specify the following notification settings:

     i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%Time%* (completion time), *%JobName%*, *%JobResult%*, *%ObjectCount%* (number of VMs in the job) and *%Issues%* (number of VMs in the job that have finished with the *Warning* or *Failed* status).

     ii. Select the **Notify on success**, **Notify on warning** and **Notify on error** check boxes to receive email notification if the job completes successfully, fails or completes with a warning.

     iii. Select the **Suppress notifications until the last retry** check box to receive a notification about the final job status. If you do not enable this option, Veeam Backup & Replication will send one notification per every job retry.

> **NOTE**
>
> If you specify the same email recipient in both job notification and global notification settings, Veeam Backup & Replication will send two separate notifications only if the subject for the email message specified in the job notification is different from the subject specified in global notification settings. Otherwise, Veeam Backup & Replication will suppress global notification settings and will send job notifications only.

6. Select the **Set successful backup details to this VM attribute** check box to write information about successfully performed backup and backup results (backup date and time, backup server name and path to the backup file) to a VM attribute. In the field under the check box, enter the name of the attribute. If the specified attribute does not exist, Veeam Backup & Replication will create it.

7. Select the **Append to the existing attribute's value** check box to append information about successfully performed backup to an existing attribute's value. In this case, Veeam Backup & Replication will keep values added by the user in the attribute and will overwrite only the value added by the backup job. If you do not select this option, Veeam Backup & Replication will overwrite the existing attribute values (made both by the user and the backup job).

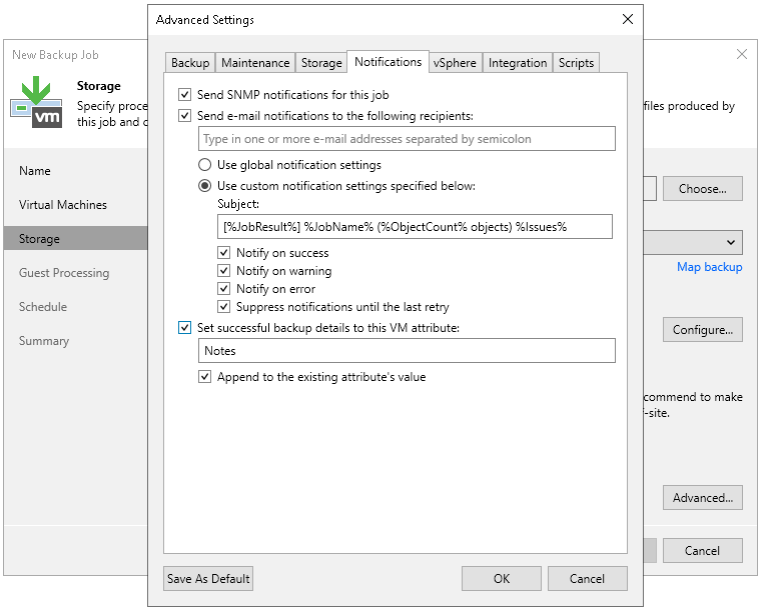


# vSphere Settings

To specify VMware vSphere settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **vSphere** tab.
3. Select the **Enable VMware tools quiescence** check box to freeze the file system of processed VMs during backup.

   Depending on the VM version, Veeam Backup & Replication will use the VMware FileSystem Sync Driver (vmsync) driver or VMware VSS component in VMware Tools for VM snapshot creation. These tools are responsible for quiescing the VM file system and bringing the VM to a consistent state suitable for backup. For more information, see VMware Tools Quiescence.
4. In the **Changed block tracking** section, configure VMware vSphere CBT:
   a. Ensure that the **Use changed block tracking data** check box is selected if you want to enable CBT.
   b. Ensure that the **Enable CBT for all processed VMs automatically** check box is selected if you want to force using CBT even if CBT is disabled in VM configuration.
   c. Ensure that the **Reset CBT on each Active Full backup automatically** check box is selected if you want to reset CBT before Veeam Backup & Replication creates active full backups.

      CBT reset helps avoid issues, for example, when CBT returns incorrect changed data.

   For more information on CBT, see Changed Block Tracking.

**IMPORTANT**

Consider the following:

- You can use CBT for VMs with virtual hardware version 7 or later. These VMs must not have existing snapshots.
- If you back up one VM with two different jobs, and the 1st job performs active full according to a schedule, the 2nd job, during an incremental run, will have to read the entire VMDK file of the processed VM. Therefore, the VM processing by a second job will take longer than during a normal incremental run. To avoid this behavior, disable the **Reset CBT on each Active Full backup automatically** option for both jobs.

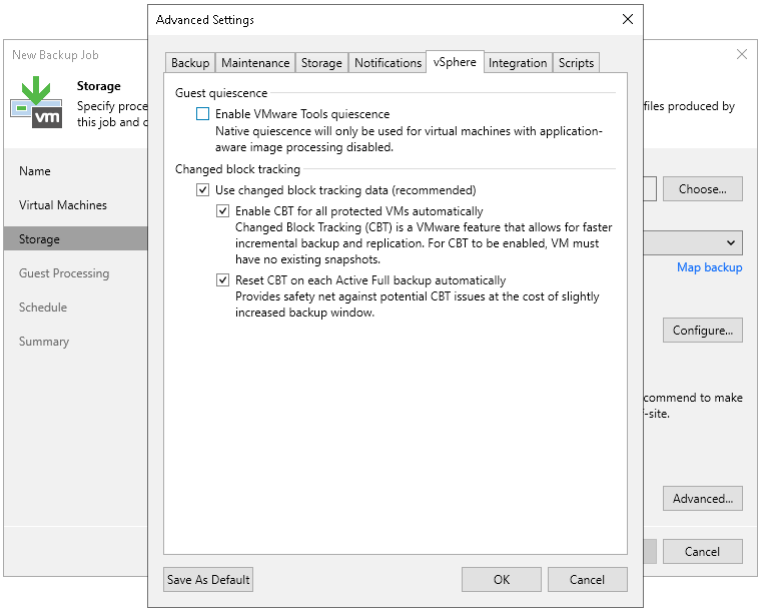


# Integration Settings

On the **Integration** tab, you can define whether you want to use the Backup from Storage Snapshots technology to create the backup. Backup from Storage Snapshots lets you leverage storage snapshots for VM data processing. The technology improves RPOs and reduces the impact of backup activities on the production environment.

Before you start using the Backup from Storage Snapshots technology, check the prerequisites. For more information, see the Requirements and Limitations section in the Storage System Snapshot Integration Guide.

To specify storage integration settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Integration** tab.
3. By default, the **Enable backup from storage snapshots** option is enabled. If you do not want to use Backup from Storage Snapshots, clear this check box. For more information, see the Configuring Backup from Storage Snapshots section in the Storage System Snapshot Integration Guide.

4. If you add to the job many VMs whose disks are located on the same volume or LUN, select the **Limit processed VM count per storage snapshot to <N>** check box and specify the number of VMs for which one storage snapshot must be created. Veeam Backup & Replication will divide VMs into several groups and trigger a separate storage snapshot for every VM group. As a result, the job performance will increase.

   For more information, see the Limitation on Number of VMs per Snapshot section in the Storage System Snapshot Integration Guide.

5. If Veeam Backup & Replication fails to create a storage snapshot, VMs whose disks are located in the storage system will not be processed by the job. To fail over to the regular VM processing mode and back up or replicate such VMs in the regular processing mode, select the **Failover to standard backup** check box.

6. [For secondary NetApp, HPE Nimble and HPE 3PAR storage systems] If Veeam Backup & Replication cannot create a storage snapshot on secondary storage arrays, the job will not back up VMs whose disks are located on the storage system. To fail over to Backup from Storage Snapshots on the production storage, select the **Failover to primary storage snapshot** check box. If Veeam Backup & Replication fails to create a storage snapshot on secondary storage arrays, it will trigger the storage snapshot on primary storage arrays and use it as a source for backup. Note, however, that Backup from Storage Snapshots on primary storage arrays will put an additional load on the production environment.

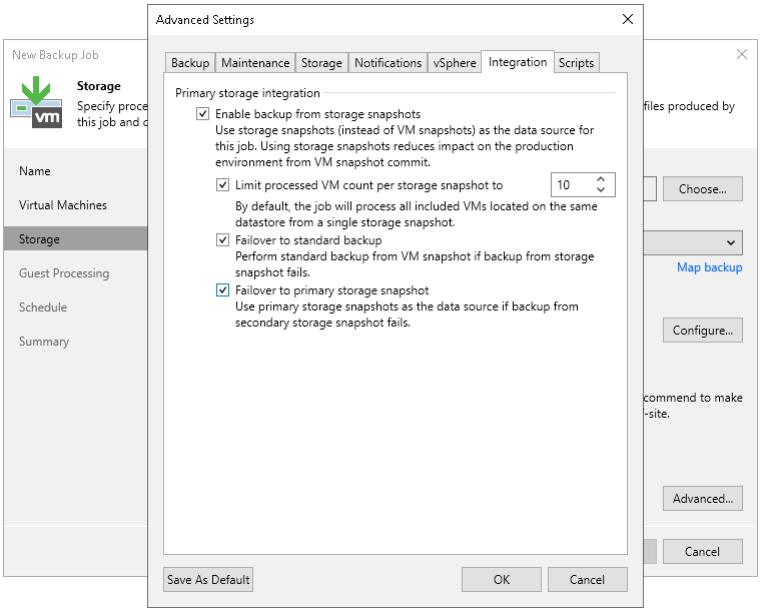


# Script Settings

To specify script settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Scripts** tab.
3. If you want to execute custom scripts before and after the backup job, select the **Run the following script before the job** and **Run the following script after the job** check boxes and click **Browse** to choose executable files from a local folder on the backup server. The scripts are executed on the backup server.

   You can select to execute pre- and post-backup actions after a number of backup sessions or on specific week days.

   - If you select the **Run scripts every <N> backup session** option, specify the number of the backup job sessions after which the scripts must be executed.

- If you select the **Run scripts on the selected days only** option, click **Days** and specify week days on which the scripts must be executed.

> **NOTE**
>
> Consider the following:
>
> - Custom scripts you define in the advanced job settings relate to the backup job itself, not the VM quiescence process. To add pre-freeze and post-thaw scripts for VM image quiescence, use the **Guest Processing** step of the wizard.
> - If you select the **Run scripts on the selected days only** option, Veeam Backup & Replication executes scripts only once on each selected day — when the job runs for the first time. During subsequent job runs, scripts are not executed.
> - To run the script, Veeam Backup & Replication uses the Service Account under which the Veeam Backup Service is running.

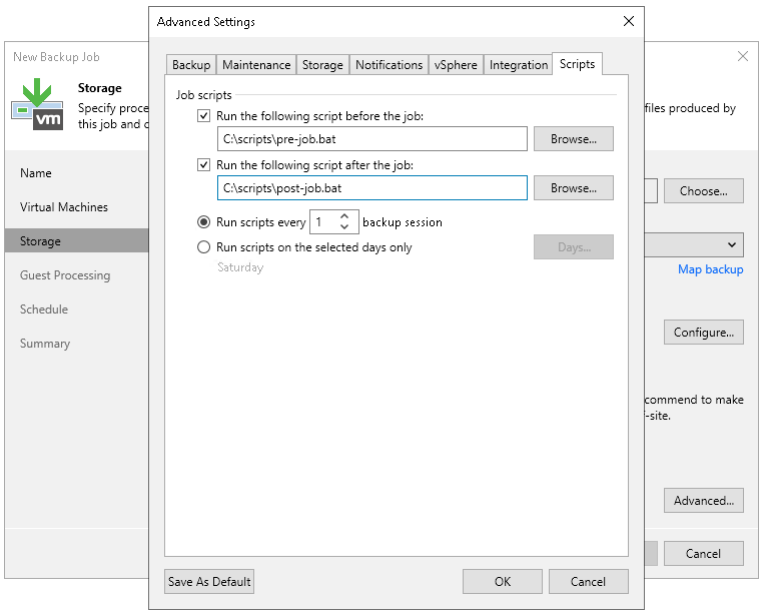

# Step 9. Specify Secondary Target

The **Secondary Target** step of the wizard is available if you have enabled the **Configure secondary destination for this job** option at the **Storage** step of the wizard.

If you plan to link the backup job to a storage array as the secondary destination, see the Backup from Storage Snapshots section in the Storage System Snapshot Integration Guide.

If you plan to link the backup job to a backup to tape or backup copy job, select the required job at the **Secondary Target** step of the wizard. As a result, the backup job will be added as a source to the backup to tape or backup copy job. Backup files created by the backup job will be archived to tape or copied to the secondary backup repository according to the secondary job schedule. For more information, see the Linking Backup Jobs to Backup Copy Jobs and Linking Backup Jobs to Backup to Tape Jobs sections in the Veeam Backup & Replication User Guide.

> **NOTE**
>
> The backup to tape job or backup copy job must be configured beforehand. You can create these jobs with an empty source. When you link the backup job to these jobs, Veeam Backup & Replication will automatically update the linked jobs to define the backup job as a source for these jobs.

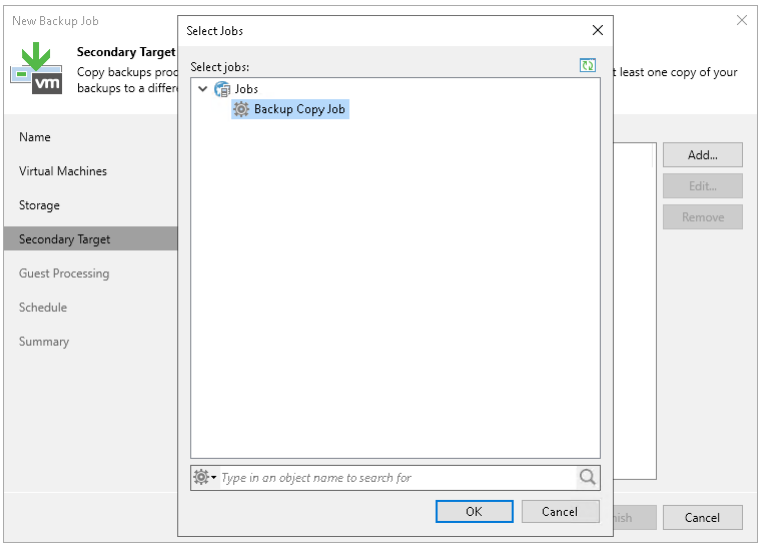

# Step 10. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, you can enable the following settings:

- **Enable application-aware processing** — to create a transactionally consistent backup of VMs running VSS-aware applications. The transactionally consistent backup guarantees the proper recovery of applications on VMs without data loss. For details on application-aware processing settings, see Application-Aware Processing.
- **Enable guest file system indexing and malware detection** — to create a catalog of files located on the guest OS. The catalog allows you to browse, search and perform 1-click restores of individual files. Guest indexing data in the catalog is scanned for suspicious file system activity and malware files. For more information, see Preparing for File Browsing and Searching section in the Enterprise Manager User Guide. For details on guest OS file indexing settings, see VM Guest OS File Indexing. For details on malware detection, see How Guest Indexing Data Scan Works.
- **Guest interaction proxy** — to specify interaction proxy settings that Veeam Backup & Replication will use to install non-persistent runtime components or use (if necessary, install) persistent agent components in each VM.
- **Guest OS credentials** — that allow Veeam Backup & Replication to connect to the VM guest OS.

To specify interaction proxy settings and credentials, do the following:

1. [Only for Microsoft Windows VMs] To connect a backup server to the processed VM guest OS, specify the guest interaction proxy. On the right of the **Guest interaction proxy** field, click **Choose** and select one of the following:
   - Leave **Automatic selection** to let Veeam Backup & Replication automatically select the guest interaction proxy.
   - Select **Use the selected guest interaction proxy servers only** to explicitly define which servers will perform the guest interaction proxy role. The list of servers contains Microsoft Windows servers added to the backup infrastructure.

   > **NOTE**
   >
   > The guest interaction proxy functionality is included in the Veeam Universal License. If you use the legacy socket-based license, you will require an Enterprise or higher edition.

2. To specify a user account that will be used to connect to the VM guest OS, from the **Guest OS credentials** drop-down list, select a user account that has enough permissions. For more information on the permissions and requirements for the user account, see Performing Guest Processing.

   If you have installed persistent agent components for VMs running Linux operating systems, you can select *Use management agent credentials* from the list. For more information, see Installing Persistent Agent Components on Linux VMs.

   > **NOTE**
   >
   > Management Agent credentials have root or elevated to root permissions. If you do not want to perform guest processing tasks under the account with such privileges, you can specify a non-root user account. This account or custom credentials added for specific VMs will also be used for SSH connection or networkless guest processing over VIX API/vSphere Web Services if the transport service connection fails.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add credentials. For more information on adding credentials, see the Credentials Manager section.

By default, Veeam Backup & Replication uses the **Log on as a batch job** policy to connect to guest OS. If the connection fails, Veeam Backup & Replication switches to **Interactive Logon**.

> **NOTE**
>
> If you use Kerberos authentication, consider the requirements and limitations described in section Kerberos Authentication.

3. By default, Veeam Backup & Replication uses the same credentials for all VMs in the job. If some VM requires a different user account, click **Credentials** and enter custom credentials for the VM.
4. To check if Veeam Backup & Replication can communicate with VMs added to the job, click **Test Now**. Veeam Backup & Replication will use the specified credentials to connect to all VMs in the list.

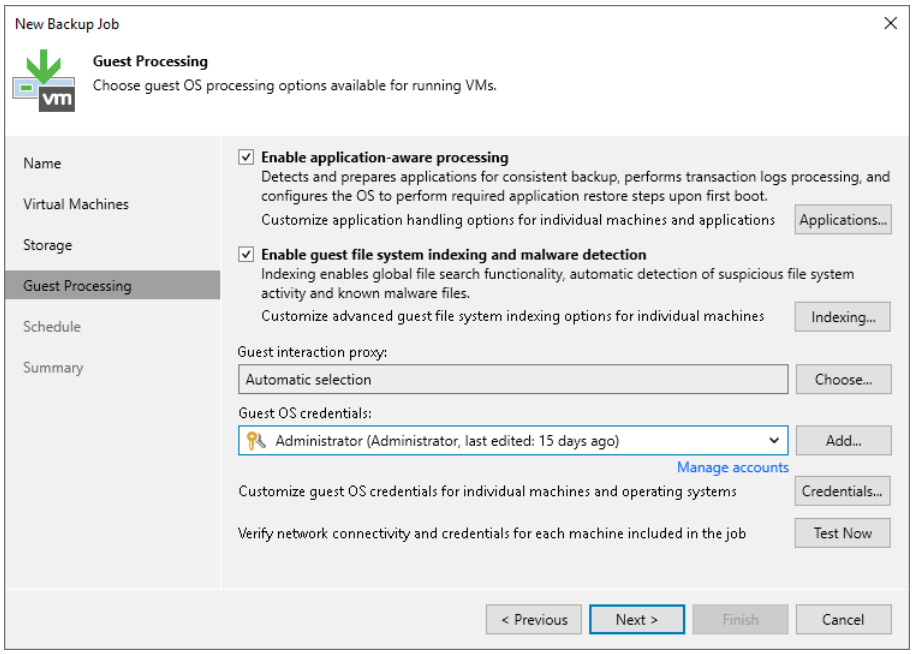


# Application-Aware Processing

To create transactionally consistent backups of VMs, you must enable application-aware processing. Application-aware processing allows you to define the method to process applications and application logs and select if you want to use the persistent agent.

To enable application-aware processing, at the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box and click **Applications**.

After that, define general application-aware processing settings:

1. In the **Application-Aware Processing Options** window, select the VM and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM to the list as a standalone object. To do this, click **Add** and choose the VM whose settings you want to customize. Then, select the VM in the list and define the necessary settings.
2. To specify the behavior scenario for application-aware processing, on the **General** tab in the **Applications** section, do one of the following:
   - If you want Veeam Backup & Replication to stop the backup process when an error occurs during application-aware processing, select **Require successful processing**.

- If you want to continue the backup process when an error occurs during application-aware processing, select **Try application processing, but ignore failures**. This option guarantees the completion of the backup job. In this case, Veeam Backup & Replication will create a crash-consistent backup instead of a transactionally consistent backup.

  > **IMPORTANT**
  >
  > If application-aware processing fails, the backup job will not process SQL transaction logs until Veeam Backup & Replication creates a new image-level backup of the Microsoft SQL Server VM.

- If you do not want to enable application-aware processing for the VM, select **Disable application processing**.

3. If you want Veeam Backup & Replication to process application logs or create copy-only backups, in the **VSS Settings** section, do one of the following:

   - [For Microsoft Exchange and Microsoft SQL VMs] If you want Veeam Backup & Replication to process application logs, select **Process transaction logs with this job** and specify settings on the SQL tab. For more information, see Microsoft SQL Server Transaction Log Settings.

     > **NOTE**
     >
     > [For Microsoft Exchange VMs] If you select this option, Veeam Backup & Replication will back up the Exchange database and its logs. The non-persistent runtime components or persistent components that run on the VM guest OS will wait for a backup job to complete successfully. After that, they will trigger truncation of transaction logs on a Microsoft Exchange server. If the backup job fails, the logs on this server will remain untouched.

   - [For Microsoft Exchange and Microsoft SQL VMs] If you use a third-party backup tool to perform VM guest level backup, and this tool maintains consistency of the database state, select **Perform copy only**. Veeam Backup & Replication will create a copy-only backup for the selected VM. The copy-only backup preserves the chain of full or differential backup files and transaction logs on the VM. For more information, see Microsoft Docs.

     Note that if you select this option, the **SQL** tab will not be available in the **Processing Settings** window.

   - [For Oracle VMs and PostgreSQL VMs] You must specify settings for application log handling on the **Oracle** and **PostgreSQL** tabs of the **VM Processing Settings** window. For more information, see Oracle Archived Redo Log Settings and PostgreSQL WAL Files Settings.

4. If you want Veeam Backup & Replication to use persistent guest agents on each protected VM for application-aware processing, select the **Persistent guest agent** check box.

   For more information, see Persistent Agent Components.

> **IMPORTANT**
>
> If both Microsoft SQL Server and Oracle Server are installed on one VM, and this VM is processed by a job with log backup enabled for both applications, Veeam Backup & Replication will back up only Oracle transaction logs. Microsoft SQL Server transaction logs will not be processed.

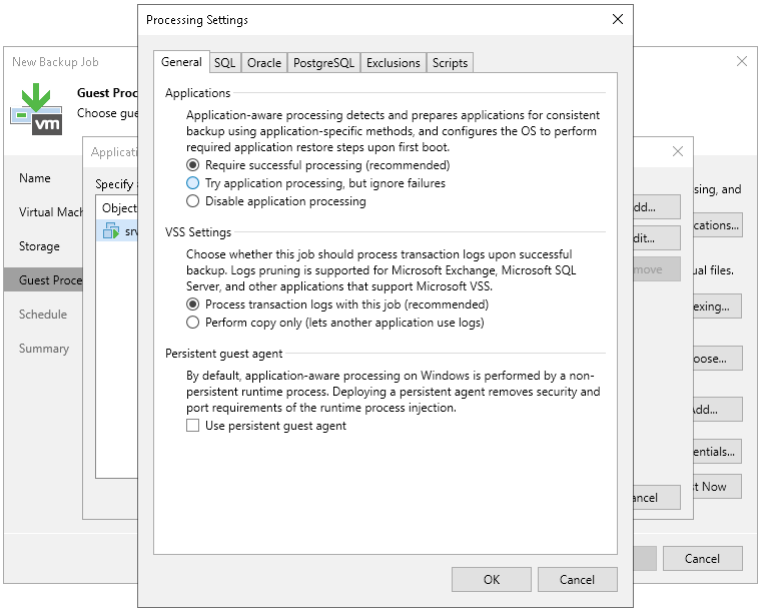


# Microsoft SQL Server Transaction Log Settings

To create a transactionally consistent backup of SQL VM, you must enable application-aware processing and define settings of transaction logs processing.

# Enabling Application-Aware Processing

Before configuring transaction logs processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the SQL VM and click **Edit**.

   By default, the **Processing Settings** window will display the **General** tab. In the **VSS Settings** section, select **Process transaction logs** with this job. For more information, see VSS settings.

# Specifying Transaction Logs Settings

To define how Veeam Backup & Replication will process transaction logs on this VM, specify the following settings:

1. In the **Processing Settings** window, click the **SQL** tab.
2. Specify how transaction logs must be processed. You can select one of the following options:
   - Select **Truncate logs** to truncate transaction logs after successful backup. The non-persistent runtime components or persistent components running on the VM guest OS will wait for the backup to complete successfully and then truncate transaction logs. If the job does not manage to back up the Microsoft SQL Server VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

> **NOTE**
>
> If the account specified at the Guest Processing step does not have enough rights, Veeam Backup & Replication tries to truncate logs using the local SYSTEM account for Microsoft SQL Server 2008 and 2008 R2. For other Microsoft SQL Server versions, Veeam Backup & Replication uses NT AUTHORITY\SYSTEM account.
>
> Make sure that these accounts have permissions listed in section Performing Guest Processing.

- Select **Do not truncate logs** to preserve transaction logs. When the backup job completes, Veeam Backup & Replication will not truncate transaction logs on the Microsoft SQL Server VM.

  It is recommended that you enable this option for databases that use the *Simple* recovery model. If you enable this option for databases that use the *Full* or *Bulk-logged* recovery model, transaction logs on the VM guest OS may grow large and consume all disk space. In this case, the database administrators must take care of transaction logs themselves.

- Select **Backup logs periodically** to back up transaction logs with Veeam Backup & Replication. Veeam Backup & Replication will periodically copy transaction logs to the backup repository and store them together with the image-level backup of the Microsoft SQL Server VM. During the backup job session, transaction logs on the VM guest OS will be truncated. For more information, see Microsoft SQL Server Log Backup.

To learn more about logging models and applicable options in Veeam Backup & Replication, see the Recovery Model section of the Veeam Explorers User Guide.

3. In the **Backup logs every <N> minutes** field, specify the frequency for transaction log backup. By default, transaction logs are backed up every 15 minutes. The maximum log backup interval is 480 minutes.
4. In the **Retain log backups** section, specify retention policy for transaction logs stored in the backup repository.

   - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for image-level backups and transaction log backups.

   - Select **Keep only last <N> days of log backups** to keep transaction logs for a specific number of days. By default, transaction logs are kept for 15 days. If you select this option, you must make sure that retention for transaction logs is not greater than retention for the image-level backups. For more information, see Retention for Transaction Log Backups.

5. In the **Log shipping servers** section, click **Choose** to select what log shipping server you want to use to transport transaction logs:

   - Select **Automatic selection** if you want Veeam Backup & Replication to choose an optimal log shipping server automatically. If the optimal shipping server is busy, Veeam Backup & Replication will direct the data flow to another shipping server so as not to lose data and comply with RPO. The process of transaction logs shipment does not require a dedicated server — Veeam Backup & Replication can use any Microsoft Windows server added to the backup infrastructure.

   - To define a log shipping server explicitly, select **Use the specified servers only** and select check boxes next to servers that you want to use as log shipping servers. The server list includes all Microsoft Windows servers added to the backup infrastructure.

     Ensure you select a server that is not engaged in other resource-consuming tasks. For example, you may want not to use a server that performs the WAN accelerator role as a log shipping server. For load balance and high availability purposes, it is recommended that you select at least 2 log shipping servers.

For more information on log shipping servers and how they are selected, see Log Shipping Servers.

**IMPORTANT**

Consider the following:

- Veeam Backup & Replication automatically excludes its configuration database from application-aware processing during backup if the database is hosted without using SQL Server Always On Availability Group. Transaction logs for the configuration database are not backed up.
- If the Veeam Backup & Replication configuration database is hosted using SQL Server Always On Availability Group, you should manually exclude this database from application-aware processing during backup, as described in this Veeam KB article. Otherwise, job processing will fail with the following error: Failed to freeze guest over network, wait timeout.

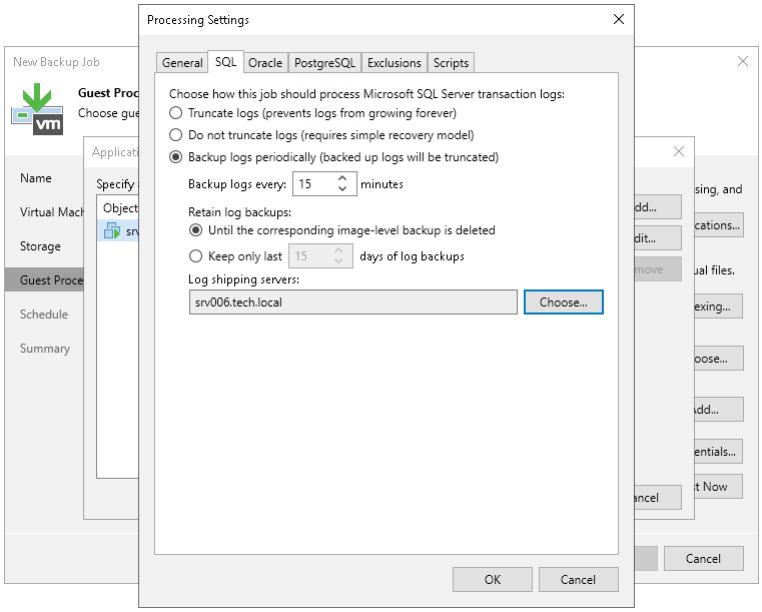


# Oracle Archived Redo Log Settings

To create transactionally consistent backups of a VM where Oracle Database is deployed, you must enable application-aware processing and define how Veeam Backup & Replication will process archived redo logs on this VM.

# Enabling Application-Aware Processing

Before configuring transaction logs processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Oracle VM and click **Edit**.

# Specifying Archive Redo Logs Settings

1. In the **VM Processing Settings** window, click the **Oracle** tab.

2. In the **Specify Oracle account with SYSDBA privileges** section, specify a user account that Veeam Backup & Replication will use to connect to the Oracle database. The account must have privileges for the Oracle database listed in the section Performing Guest Processing.

   You can select **Use guest credentials** in the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle database.

> **NOTE**
>
> [For Windows-based machines] Make sure you add the *%ORACLE_HOME%\bin* variable to the environmental path variable on the machine with your Oracle database.

3. In the **Archived logs** section, specify how Veeam Backup & Replication must process archived redo logs on the Oracle VM:

   - Select **Do not delete archived logs** if you want Veeam Backup & Replication to preserve archived logs on the VM guest OS. When the backup job completes, the non-persistent runtime components or persistent components will not delete archived logs.

     It is recommended that you select this option for databases in the NOARCHIVELOG mode. If the database is in the ARCHIVELOG mode, archived logs on the VM guest OS may grow large and consume all disk space. In this case, database administrators must take care of archived logs themselves.

   - Select **Delete logs older than <N> hours** or **Delete logs over <N> GB** if you want Veeam Backup & Replication to delete archived logs that are older than <N> hours or larger than <N> GB. The log size threshold refers not to the total size of all logs for all databases but to the log size of each database on the selected Oracle VM.

     When the parent backup job (job creating an image-level backup) runs, Veeam Backup & Replication will wait for the backup to complete successfully and then trigger archived log deletion on the Oracle VM over Oracle Call Interface (OCI). If the primary job does not manage to back up the Oracle VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

> **TIP**
>
> Veeam Backup & Replication removes redo logs only after the parent backup job session. To remove redo logs more often, you can schedule the parent job to run more often.

4. To back up Oracle archived logs with Veeam Backup & Replication, select the **Backup log every <N> minutes** check box and specify the frequency for archived log backup. By default, archived logs are backed up every 15 minutes. The maximum log backup interval is 480 minutes.

5. In the **Retain log backups** section, specify the retention policy for archived logs stored in the backup repository:

   - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for image-level backups and archived log backups.

   - Select **Keep only last <n> days** to keep archived logs for a specific number of days. By default, archived logs are kept for 15 days. If you select this option, you must ensure that retention for archived logs is not greater than retention for the image-level backups. For more information, see Retention for Archived Log Backup.

6. In the **Log shipping servers** section, click **Choose** to select what log shipping server you want to use to transport archived logs:

   - Select **Automatic selection** if you want Veeam Backup & Replication to select an optimal log shipping server automatically. The process of archived logs shipment does not require a dedicated server — Veeam Backup & Replication can use any Microsoft Windows server added to the backup infrastructure.

   - Select **Use the specified servers** only to define a log shipping server explicitly. In the **Log Shipping Servers** window, select check boxes next to the servers you want to use as log shipping servers. The server list includes all Microsoft Windows servers added to the backup infrastructure.

     Ensure you select a server that is not engaged in other resource-consuming tasks. For example, you may want not to use a server that performs the WAN accelerator role as a log shipping server. For load balance and high availability purposes, it is recommended that you select at least 2 log shipping servers.

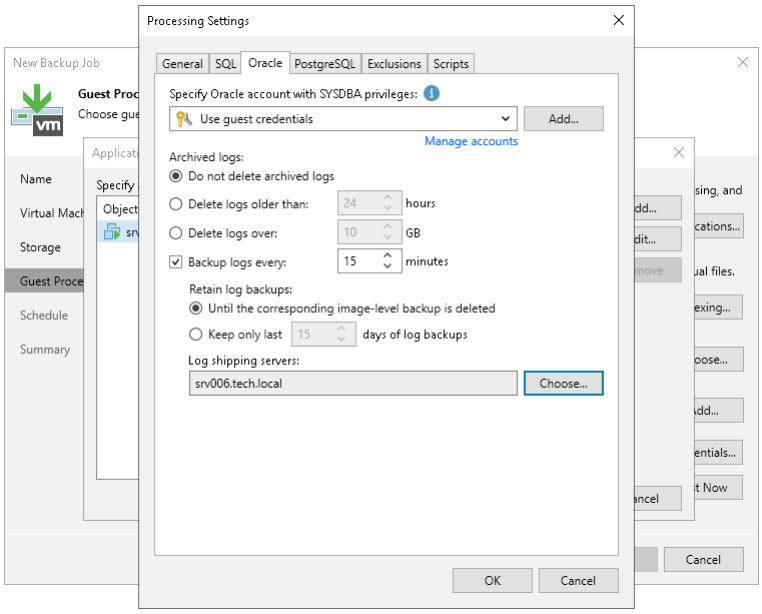


# PostgreSQL WAL Files Settings

To create a transactionally consistent backup of PostgreSQL VM, you must enable application-aware processing and define the settings of WAL files processing.

# Enabling Application-Aware Processing

Before configuring WAL file processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the PostgreSQL VM and click **Edit**.

NOTE

By default, Veeam Backup & Replication recursively scans the `/etc/postgresql`, `/var/lib/postgresql` and `/var/lib/pgsql` directories for the configuration files of PostgreSQL instances. If you keep configuration files in custom directories, the *pgsqlagent agent* will use its own `VeeamPostgreSQLAgent.xml` configuration file that is located in the `/etc/veeam/` directory. The pgsqlagent agent configuration file must be a single line XML.

To explicitly include or exclude specific configuration files from rescan, you can add the following commands to the `VeeamPostgreSQLAgent.xml` file:

- `ExcludeConfigDirs` — use this element to exclude configuration files.
- `AddConfigDirs` — use this element to include configuration files.

For example: `<config AddConfigDirs="/opt/psql/" ExcludeConfigDirs="/var/lib/postgresql/13/main45/,/var/lib/postgresql/13/maindd/" />`

# Specifying WAL Files Settings

To define how Veeam Backup & Replication will process WAL files on this VM, do the following:

1. In the **Processing Settings** window, click the **PostgreSQL** tab.
2. From the **Specify PostgreSQL account with superuser privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the PostgreSQL instance. The account must have privileges described in section Permissions. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add credentials.

   By default, the *Use guest credentials* option is selected in the list. With this option selected, Veeam Backup & Replication will connect to the PostgreSQL instance under the account. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the PostgreSQL instance.

   Note that if you plan to select the **System user without password file (peer)** authentication method at the step 3 of this procedure, you can add a user account in the Credentials Manager without specifying the password for the account.
3. In the **Specified user is** section, specify how the user will authenticate against the PostgreSQL instance:
   - Select **Database user with password** if the account you specified at the step 2 is a PostgreSQL account, and you entered the password for this account in the **Credentials Manager**.

     NOTE

     If you want Veeam Backup & Replication to use the user name map authentication, select **Database user with password** and leave the password field empty. Consider the following:

     - **Guest OS credentials** specified at the **Guest Processing** step of the wizard will be used as the System-Username.
     - PostgreSQL account specified at the step 2 will be used as the PG-Username.

     For more information about the user name maps, see PostgreSQL documentation.

- Select **Database user with password file (.pgpass)** if the password for the account you specified at the step 2 is defined in the `.pgpass` configuration file on the PostgreSQL VM. The password file must be located in the user's home directory. For more information about the password file, see PostgreSQL documentation.
- Select **System user without password file (peer)** if you want Veeam Backup & Replication to use the peer authentication method. In this case, Veeam Backup & Replication will use the account you specified at the step 2 as the OS account and as the PostgreSQL account to connect to PostgreSQL. For more information about the peer authentication method, see PostgreSQL documentation.

  > **IMPORTANT**
  >
  > If you have added a new PostgreSQL account and want to use it with the peer authentication method, make sure that you have added this account as a Linux user with sufficient permissions.

4. To back up PostgreSQL WAL files with Veeam Backup & Replication, select the **Backup log every <N> minutes** check box and specify the frequency for WAL files backup. By default, WAL files are backed up every 15 minutes. The maximum log backup interval is 480 minutes.
5. In the **Retain log backups** section, specify the retention policy for WAL files stored in the backup repository:
   - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for both image-level backups and WAL file backups.
   - Select **Keep only last <n> days** to keep WAL files for a specific number of days. By default, WAL files are kept for 15 days. If you select this option, you must make sure that retention for WAL files is not greater than retention for the image-level backups. For more information, see Retention for PostgreSQL WAL Files.
6. In the **Temporary location for archive logs** section, specify a path to the storage location where you want to keep WAL files.
7. In the **Log shipping servers** section, click **Choose** to select what log shipping server you want to use to transport WAL files:
   - Select **Automatic selection** if you want Veeam Backup & Replication to select an optimal log shipping server automatically. The process of WAL files shipment does not require a dedicated server — Veeam Backup & Replication can use any Microsoft Windows server added to the backup infrastructure.

- Select **Use the specified servers only** to define a log shipping server explicitly. In the **Log Shipping Servers** window, select check boxes next to the servers you want to use as log shipping servers. The server list includes all Microsoft Windows servers added to the backup infrastructure.

  Make sure that you select a server that is not used by other resource-consuming tasks. For example, you may want not to use a server that performs the WAN accelerator role as a log shipping server. For load balance and high availability purposes, it is recommended that you select at least 2 log shipping servers.

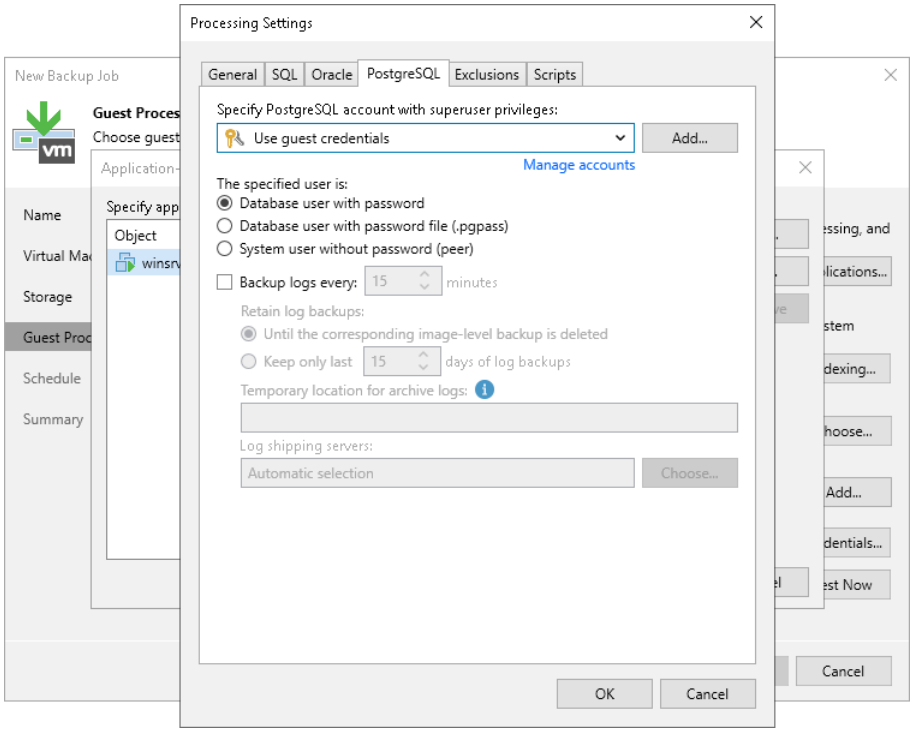


# VM Guest OS File Exclusion

If you do not want to back up specific files and folders on the VM guest OS, you can exclude them from the backup.

> **NOTE**
>
> VM guest OS file exclusion functionality is included in the Veeam Universal License. When using a legacy socket-based license, an Enterprise or higher edition is required.

To define what files and folders must be excluded:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the VM and click **Edit**.

   To define custom settings for a VM added as part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose a VM whose settings you want to customize. Then, select the VM in the list and define the necessary settings.
4. Click the **Exclusions** tab and specify what files must be excluded from the backup:
   - Select **Exclude the following files and folders** to remove the individual files and folders from the backup.

- Select **Include only the following files and folders** to leave only the specified files and folders in the backup.

5. Click **Add** and specify what files and folders you want to include or exclude. To form the list of exclusions or inclusions, you can use full paths to files and folders, environmental variables and file masks with the asterisk (*) and question mark (?) characters. For more information, see VM Guest OS Files.

> **NOTE**
>
> When you choose files to be included or excluded, consider the requirements and limitations listed in the section Requirements and Limitations for VM Guest OS File Exclusion.

6. Click **OK**.
7. Repeat steps 5-6 for every file or folder you want to exclude or include.

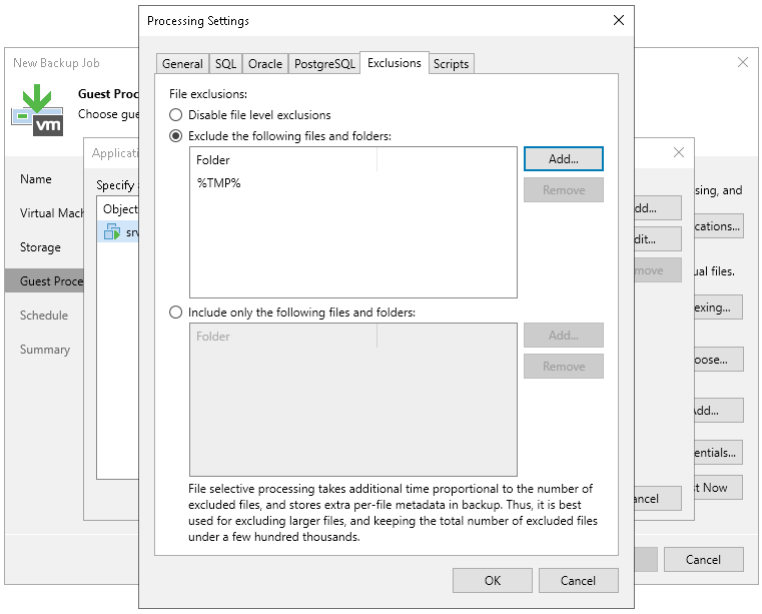


# Pre-Freeze and Post-Thaw Scripts

If you plan to back up VMs running applications that do not support VSS, you can specify what scripts Veeam Backup & Replication must use to quiesce the VM. The pre-freeze script quiesces the VM file system and application data to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script brings the VM and applications to their initial state.

To specify pre-freeze and post-thaw scripts for the job:

1. At the **Guest Processing** step, click **Applications**.
2. In the displayed list, select the VM and click **Edit**.
3. Click the **Scripts** tab.
4. In the **Script processing mode** section, specify the scenario for scripts execution:
   - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the backup process if the script fails.
   - Select **Ignore script execution failures** if you want to continue the backup process, even if script errors occur.

- Select **Disable script execution** if you do not want to run scripts for the VM.

5. In the **Windows scripts** section, specify paths to pre-freeze and post-thaw scripts for Microsoft Windows VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.
6. In the **Linux scripts** section, specify paths to pre-freeze and post-thaw scripts for Linux VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

   If you have added a VM container with Microsoft Windows and Linux VMs to the job, you can select to execute both Microsoft Windows and Linux scripts for the VM container. When the job starts, Veeam Backup & Replication will automatically determine what OS type is installed on the VM and use the required scripts to quiesce this VM.

> **TIP**
>
> Besides pre-freeze and post-thaw scripts for VM quiescence, you can instruct Veeam Backup & Replication to run custom scripts before the job starts and after the job completes. For more information, see Script Settings.

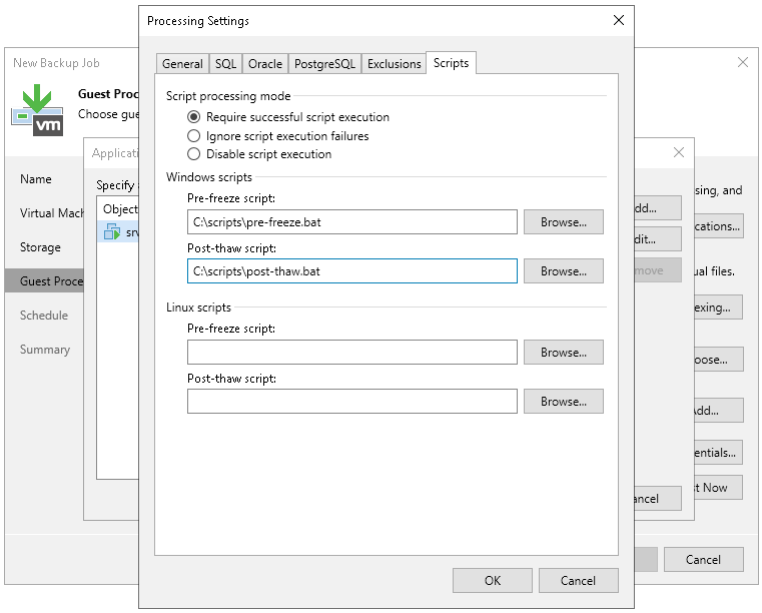


# VM Guest OS File Indexing

To specify VM guest OS indexing options for a VM:

1. At the **Guest Processing** step of the wizard, click **Indexing**.
2. Select a VM in the list and click **Edit** > **Windows indexing** or **Linux indexing**.
3. Specify the indexing scope:
   - Select **Disable indexing** if you do not want to index guest OS files of the VM.
   - Select **Index everything** if you want to index all VM guest OS files.
   - Select **Index everything except** if you want to index all VM guest OS files except those defined in the list. By default, system folders are excluded from indexing. You can add or delete folders using the **Add** and **Remove** buttons on the right. You can also use system environment variables to form the list, for example, *%windir%, %ProgramFiles%* and *%Temp%*.

     To reset the list of folders to its initial state, click **Default**.

- Select **Index only following folders** to define the folders you want to index. You can add or delete folders to index using the **Add** and **Remove** buttons on the right. You can also use system environment variables to form the list, for example, *%windir%, %ProgramFiles%* and *%Temp%*.

> **NOTE**
>
> To perform guest OS file indexing on Linux VMs, Veeam Backup & Replication requires several utilities to be installed on the Linux VM: openssh, gzip and tar. If these utilities are not found, Veeam Backup & Replication will prompt you to deploy them on the VM guest OS.

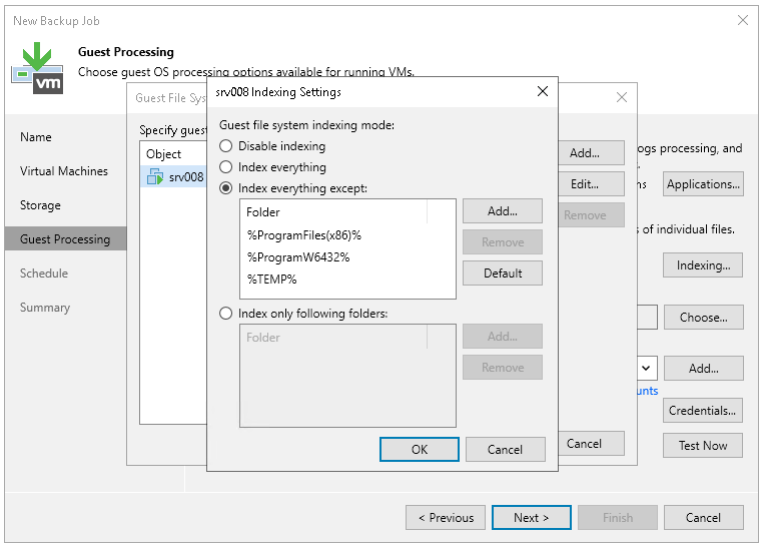

# Step 11. Define Job Schedule

At the **Schedule** step of the wizard, select to run the backup job manually or schedule the job to run regularly.

To specify the job schedule:

1. Select the **Run the job automatically** check box. If this check box is not selected, you must start the job manually to create the VM backup.
2. Define scheduling settings for the job:
   - To run the job at a specific time daily, on defined week days, or with specific periodicity, select **Daily at this time**. Use the fields on the right to configure the necessary schedule.
   - To run the job once a month on specific days, select **Monthly at this time**. Use the fields on the right to configure the necessary schedule.

     > **NOTE**
     >
     > When you configure the job schedule, consider possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a specific time interval, select **Periodically every**. In the field on the right, select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** and use the time table to define the permitted time window for the job. In the **Start time within an hour** field, specify the exact time when the job must start.

     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM, and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run it by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.

   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right. A new backup job session will start as soon as the previous backup job session finishes.
   - To chain jobs, use the **After this job** field. In common practice, jobs start one after another: when job *A* finishes, job *B* starts, and so on. If you want to create a chain of jobs, you must define the schedule for the first job in the chain. For the rest of the jobs in the chain, select the **After this job option** and choose the preceding job from the list.

     Note that chained jobs have some drawbacks and limitations. For more information, see Chained Jobs.
3. In the **Automatic retry** section, define whether Veeam Backup & Replication must attempt to run the backup job again if the job fails for some reason. During a job retry, Veeam Backup & Replication processes failed items only. Enter the number of attempts to run the job and define time intervals between them. If you select continuous backup, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job runs.

4. In the **Backup window** section, define the time interval within which the backup job must be completed. The backup window prevents the job from overlapping with production hours and ensures that the job does not provide unwanted overhead on the production environment. To set up a backup window for the job:

   a. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**.

   b. In the **Time Periods** window, define the allowed and prohibited hours for backup. If the job exceeds the allowed window, it will be automatically terminated.

> **NOTE**
>
> The **After this job** function will automatically start a job if the first job in the chain is started automatically by schedule. If you start the first job manually, Veeam Backup & Replication will display a notification. You will be able to choose whether Veeam Backup & Replication must start the chained job as well.

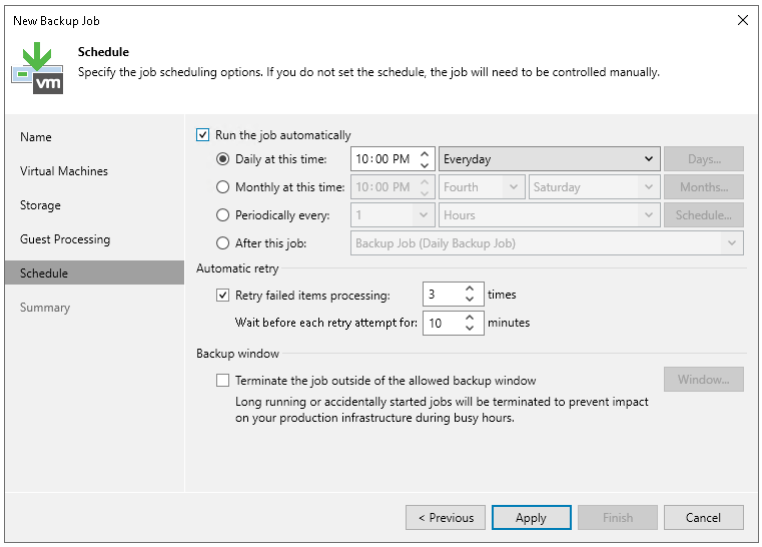

# Step 12. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of backup job configuration.

1. Review details of the backup job.
2. Select the **Run the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

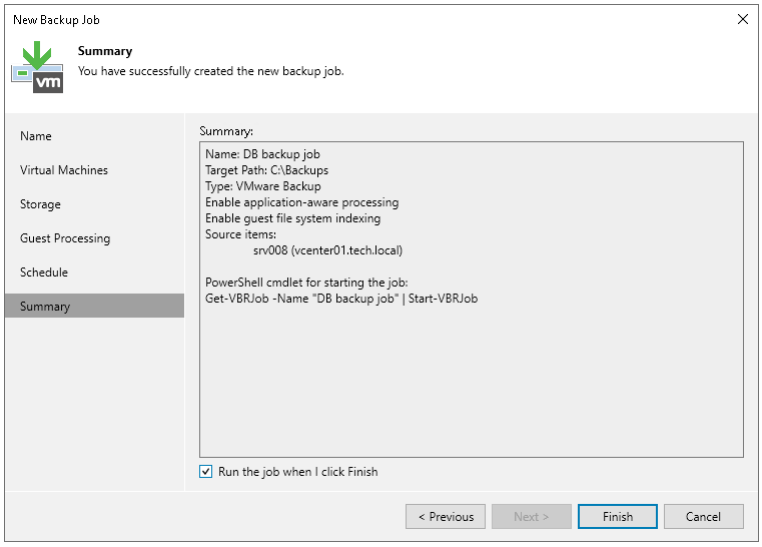

# Encrypting Backup Jobs

The encryption process for a backup job includes the following steps:

1. Job encryption is configured in the advanced job settings. When you enable the encryption option, you need to enter a password or select a KMS server.
2. If you enter a password, Veeam Backup & Replication generates a secret key. If you select a KMS server, Veeam Backup & Replication asks the KMS server to generate an asymmetric KMS key. A secret key or a KMS key will be used to encrypt data encryption keys during the job. For more information, see How Data Encryption Works.
3. When you start a job, Veeam Backup & Replication encrypts backup files with data encryption keys. By default, this operation is performed on the backup proxy.
4. Encrypted data is passed to the target backup repository and stored to a resulting backup file.

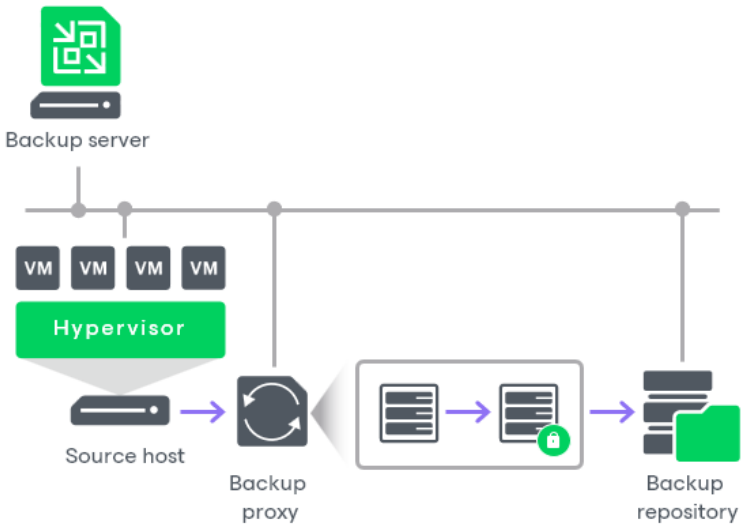


**NOTE**

If you plan to store backups in Adding Veeam Data Cloud Vault, you must enable job encryption.

If you update encryption settings for an existing backup job, consider the following:

- If you enable encryption, during the next job session Veeam Backup & Replication will automatically create a full backup file. The created full backup file and subsequent incremental backup files in the backup chain will be encrypted with the specified password or KMS key.

  Note that Veeam Backup & Replication does not encrypt the previous backup chain created by this job. If you want to start a new chain so that the unencrypted previous chain can be separated from the encrypted new chain, follow this Veeam KB article.
- If you change the password or start using KMS keys for the already encrypted job, during the next job session Veeam Backup & Replication will create a new incremental backup file. The created backup file and subsequent backup files in the backup chain will be encrypted with the new password or KMS key.
- If you disable encryption, during the next job session Veeam Backup & Replication will automatically create a full backup file.

# Performing Active Full Backup

You can create an ad-hoc full backup — active full backup, and add it to the backup chain in the backup repository. The active full backup resets the backup chain. All subsequent incremental backups use the active full backup as a starting point. The previously used full backup will remain in the backup repository until it is removed from the backup chain according to the retention policy.

## Performing Active Full Backup for All Workloads

To perform active full backup for all workloads in a backup job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the necessary job and click **Active Full** on the ribbon. Alternatively, you can right-click the job and select **Active Full**.

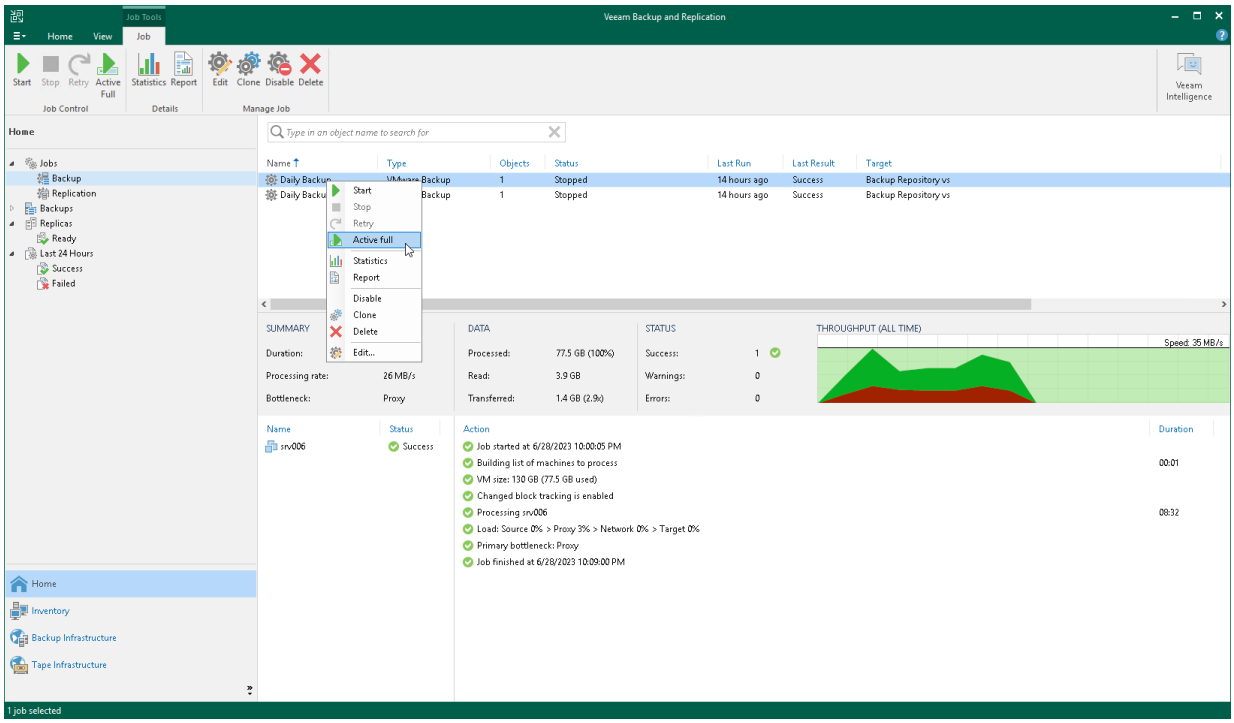


## Performing Active Full Backup for Individual Workloads

To perform active full backup for individual workloads:

1. Open real-time statistics or sessions results of the job.
2. Select workloads for which you want to perform active full backup.
3. Right-click one of the selected workloads and click **Active Full**. Note that you will be able to launch active full backup for other workloads in the job only after active full finishes for the selected workloads.

> **IMPORTANT**
>
> You can perform active full backup for individual workloads only if their backups are per-machine with separate metadata files.

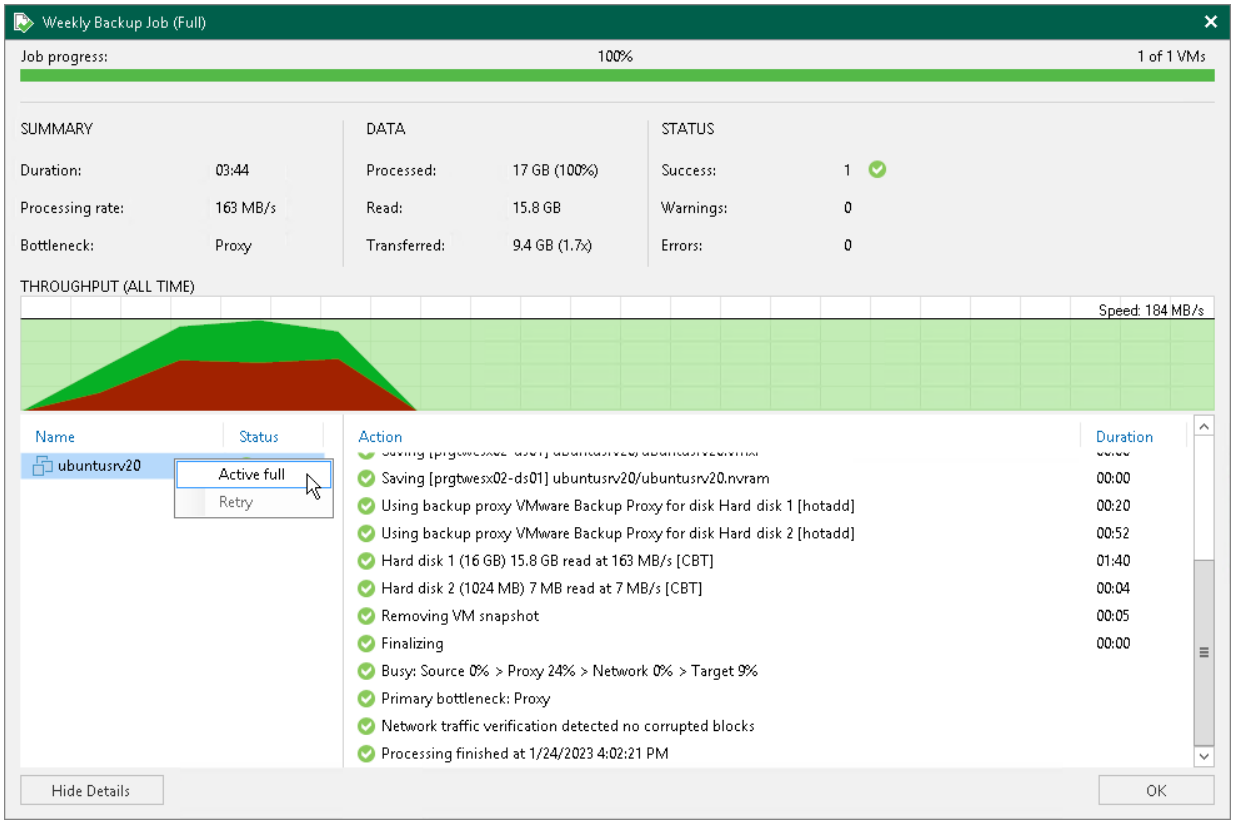
Weekly Backup Job (Full)
Job progress:
100%
1 of 1 VMs
SUMMARY
Duration:
03:44
Processing rate:
163 MB/s
Bottleneck:
Proxy
DATA
Processed:
17 GB (100%)
Read:
15.8 GB
Transferred:
9.4 GB (1.7x)
STATUS
Success:
1
Warnings:
0
Errors:
0
THROUGHPUT (ALL TIME)
Speed: 184 MB/s
Name
Status
Action
Duration
ubuntusrv20
Active full
Retry
Saving [prgtwesx02-ds01] ubuntusrv20/ubuntusrv20.nvram
00:00
Using backup proxy VMware Backup Proxy for disk Hard disk 1 [hotadd]
00:20
Using backup proxy VMware Backup Proxy for disk Hard disk 2 [hotadd]
00:52
Hard disk 1 (16 GB) 15.8 GB read at 163 MB/s [CBT]
01:40
Hard disk 2 (1024 MB) 7 MB read at 7 MB/s [CBT]
00:04
Removing VM snapshot
00:05
Finalizing
00:00
Busy: Source 0% > Proxy 24% > Network 0% > Target 9%
Primary bottleneck: Proxy
Network traffic verification detected no corrupted blocks
Processing finished at 1/24/2023 4:02:21 PM
Hide Details
OK

# Quick Backup

Quick backup lets you perform on-demand incremental backup for VMs. You can use quick backup if you want to produce an additional restore point for one or more VMs in a backup job and do not want to configure a new job or modify the existing one. Quick backup can be run for both incremental and reverse incremental backup chains.

Quick backup is an incremental backup task: Veeam Backup & Replication copies only changed data for selected VMs and saves this data to a new restore point in the backup chain. Similar to incremental backup, quick backup can only be run for VMs that have been successfully backed up at least once and have a full restore point. If there is no full restore point for a VM, quick backup cannot be performed.

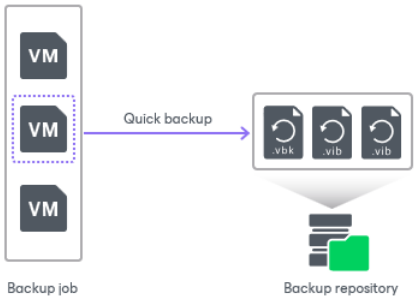


To perform a quick backup, Veeam Backup & Replication uses an existing backup job. When you start a quick backup task for a VM, Veeam Backup & Replication verifies that a backup job processing this VM exists on the backup server. If such a job is detected, Veeam Backup & Replication triggers a job and creates an incremental restore point for the VM. If a backup job for the VM does not exist, quick backup is terminated.

You can run a quick backup for one or multiple VMs simultaneously. If you start a quick backup for several VMs and these VMs are processed by different backup jobs, Veeam Backup & Replication triggers a set of backup jobs. Each triggered job creates a separate restore point and stores it in the correct backup chain.

In some cases, a VM may be processed by several backup jobs on the backup server. In this case, Veeam Backup & Replication starts the job that has created the most recent restore point for the VM.

For example, *VM01* is processed by 2 jobs:

- *Backup job 1* created the most recent restore point on Monday
- *Backup job 2* created the most recent restore point on Tuesday

When you start a quick backup for *VM01*, Veeam Backup & Replication will trigger *Backup job 2* to create a new incremental restore point.

> **NOTE**
>
> If the quick backup task overlaps the scheduled backup job, the backup job waits for the quick backup task to complete.

## Limitations for Quick Backup

You cannot perform a quick backup for VMware Cloud Director VMs processed with VMware Cloud Director jobs. However, if you process a VMware Cloud Director VM with a regular backup job, you can switch to the **Computer** view and start the quick backup operation for this VM.

# Retention Policy for Quick Backups

When you perform a quick backup, Veeam Backup & Replication creates a single VM incremental restore point. Unlike a regular incremental restore point that contains data for all VMs in a job, single VM incremental restore point contains data only for a specific VM.

A single VM restore point is not regarded as a full-fledged restore point in the backup chain. From the retention policy perspective, a single VM restore point is grouped with a regular restore point following it. When Veeam Backup & Replication needs to delete a single VM restore point by retention, it waits for the next regular restore point to expire — that is Veeam Backup & Replication increases the retention by one restore point for some time. After the next regular restore point expires, Veeam Backup & Replication deletes two restore points at once.

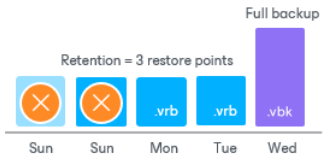


If the backup chain is stored in a repository with the Use per-machine backup files option enabled, the number of the retention points increase in a different way. As a rule, the number of retention points increase by the number of VMs from this chain for which quick backup was performed. It applies to the reverse incremental and forward incremental backup chains.

# Performing Quick Backup

You can create an ad-hoc incremental backup for one or more VMs — quick backup, and add it to the backup chain in the backup repository. A quick backup is useful when you want to produce an additional restore point for one or more VMs in the backup job and do not want to configure a new job or modify the existing one.

Quick backup can be performed for VMs that meet the following requirements:

- A backup job processing the VM exists on the backup server.
- A full backup file for the VM exists in the backup repository configured in the backup infrastructure.

To perform a quick backup:

1. Open the **Inventory** view.
2. In the infrastructure tree, select a host or VM container (host, cluster, folder, resource pool, VirtualApp, datastore or tag) where the VMs you want to back up reside.
3. In the working area, select the VMs and click **Quick Backup** on the ribbon. You can also right-click the VMs and select **Quick Backup**.

   Veeam Backup & Replication will trigger a backup job to create a new incremental restore point for each selected VM. Details of a running quick backup task are displayed in the job session window.

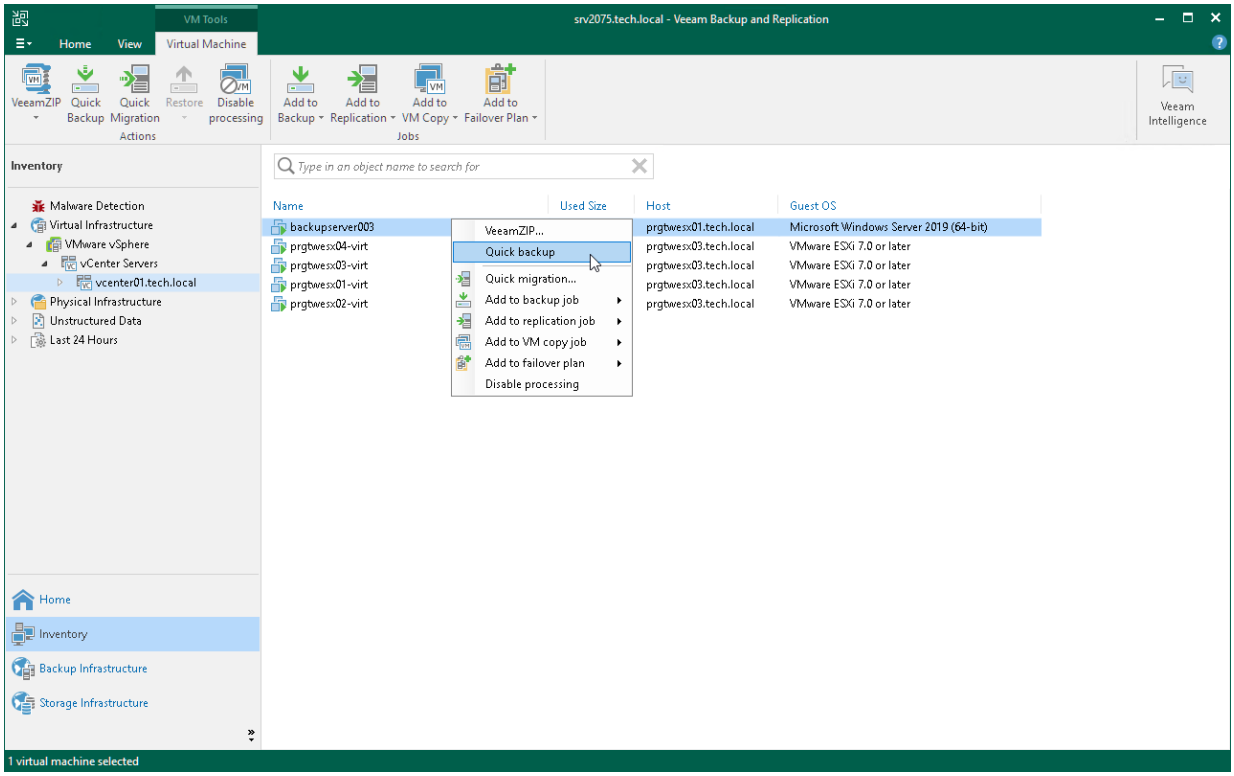

# Importing Backups Manually

You may need to import backups to Veeam Backup & Replication in the following situations:

- The backup server has failed and you have restored it in a new location. You want to restore VM data from backups created by the backup server that has failed.
- You want to restore VM data from backups created on another backup server.
- You want to restore VM data from backups in the backup repository that is not added to the backup infrastructure (for example, if you removed it earlier).
- You want to restore VM data from VeeamZIP files created on your backup server or another backup server.

The imported backup becomes available in the Veeam Backup & Replication console. You can use any restore operation to recover VM data from this backup.

## Considerations and Limitations

Before you import a backup, consider the following:

- In the Veeam Backup & Replication console, you can import backups from the following types of repositories: Microsoft Windows servers, Linux servers, HPE StoreOnce and Dell Data Domain deduplicating storage appliances. To import files from other types of repositories, use the `Import-VBRbackup` cmdlet as described in the Veeam PowerShell Reference.
- The server from which you plan to import backups must be added to the backup infrastructure. Otherwise, you will not be able to access backup files.
- To be able to restore VM data from previous backup restore points, ensure you have all required incremental backup files (forward or reverse) in the same folder where the full backup file resides.
- You can import backups created by a later Veeam Backup & Replication version. However, restore operations for this backup will not be available.
- Manually imported backups do not have an assigned repository. Veeam Backup & Replication server will perform the role of the mount server in case of the manual import.

## Importing Backups Manually

To import a backup to the Veeam Backup & Replication console:

1. On the **Home** tab, click **Import Backup**.
2. From the **Computer** list, select the server on which the backup you want to import is stored.
3. Click **Browse** and select the necessary VBM or VBK file. If you select the VBM file, the import process will be notably faster. It is recommended that you select the VBK file only if the VBM file is not available.
4. By default, index data of the VM guest OS file system is not imported with the backup to speed up the import process. If you want to import index data, select the **Import guest file system index** check box.
5. Click **OK** to import the backup. The imported backup will be displayed in the **Home** view, under the **Backups** > **Imported** node in the inventory pane. Backups are imported using the original name of the backup job with the `_imported` suffix appended.

### TIP

If you need to import all backups stored on a server, assign a backup repository role to it and enable import at the Review step of the wizard. If the repository is already added to the backup infrastructure, you can rescan it. Veeam Backup & Replication will automatically import backups. For more information, see Rescanning Backup Repositories.

Note that Veeam Backup & Replication will not be able to import backups if VBM files are unavailable. In this case, you will have to import backups manually using the VBK files.

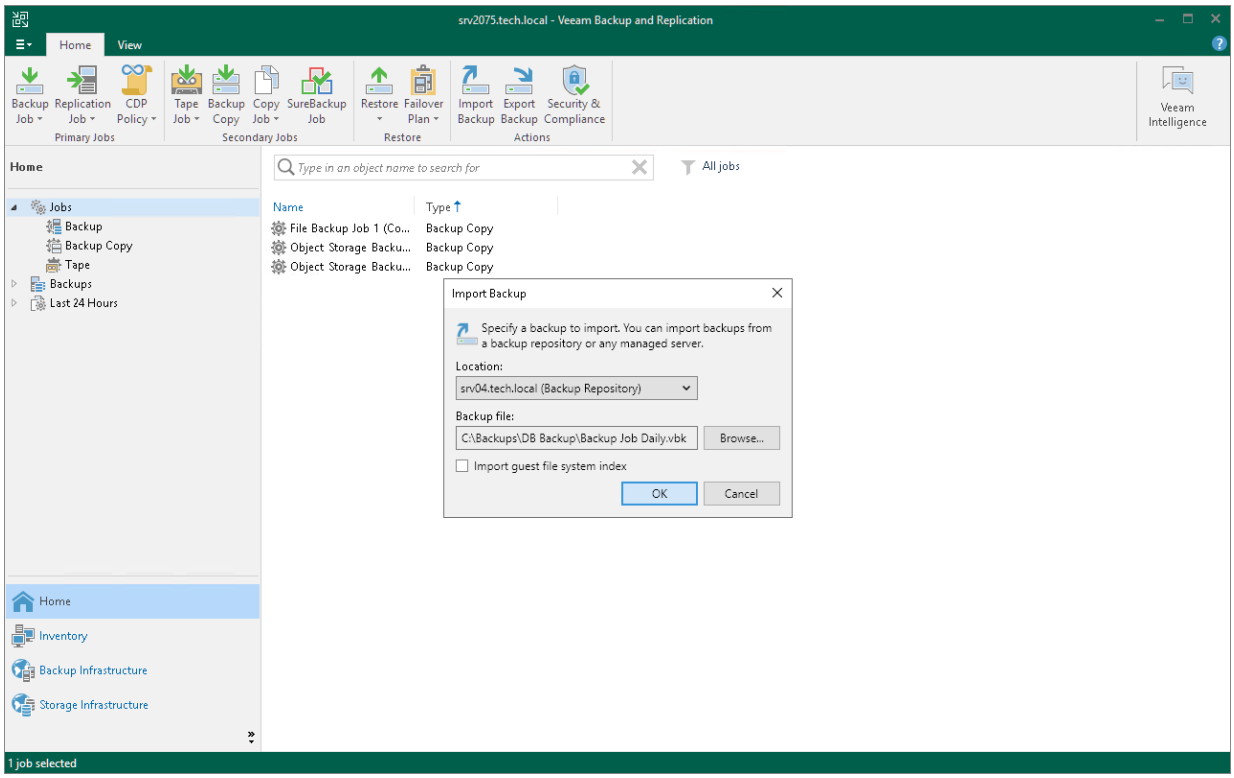

# Importing Encrypted Backups

You can import backups that were encrypted on this backup server or on another backup server.

To import an encrypted backup file:

1. On the **Home** tab, click **Import Backup**.
2. From the **Computer** list, select the host on which the backup you want to import is stored.
3. Click **Browse** and select the VBM or VBK file.
4. Click **OK**. The encrypted backup will appear under the **Backups** > **Disk (encrypted)** node in the inventory pane.
5. In the working area, select the imported backup and click **Specify Password** on the ribbon or right-click the backup and select **Specify password**.
6. In the **Password** field, enter the password for the backup file.

   If you changed the password one or several times while the backup chain was created, you must enter passwords in the following manner:

   - If you select a VBM file for import, you must specify the latest password that was used to encrypt files in the backup chain.
   - If you select a VBK file for import, you must specify the whole set of passwords that were used to encrypt files in the backup chain.

If you enter the correct passwords, Veeam Backup & Replication will decrypt the backup file. The backup will be moved under the **Backups** > **Disk (Imported)** node in the inventory pane.

> **NOTE**
>
> You can recover data from encrypted backups even if the password is lost. Restoring data without a password is included in the Veeam Universal License. When using a legacy socket-based license, an Enterprise or higher edition is required. Also, your backup server must be connected to Veeam Backup Enterprise Manager. For more information, see Decrypting Backups With Enterprise Manager Keys.

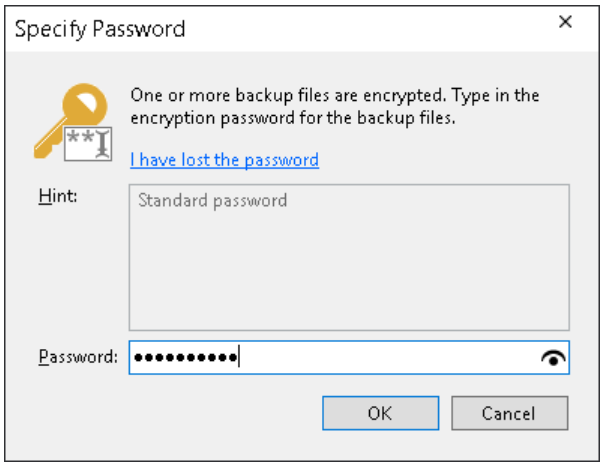

# Importing Transaction Logs

You cannot import transaction log backups without VM backups (as there will be no restore point to which the transaction logs can be applied).

To import a VM backup with transaction log backups, do either of the following:

- Import a backup metadata file (VBM). In this case, Veeam Backup & Replication will automatically import the database backup and log backups.
- Import a full backup file (VBK). In this case, Veeam Backup & Replication will browse the required log backups and import them.

# Importing Backup Files from Scale-Out Backup Repositories

You cannot import a backup directly from the scale-out backup repository. When you perform backup import, you cannot browse the whole scale-out backup repository. Veeam Backup & Replication lets you browse only through individual extents.

To import a backup from the scale-out backup repository, you must place backup files from all extents to one staging folder. The staging folder can reside on any server added to the backup infrastructure. After that, you can import the backup as usual.

> **TIP**
>
> If you need to import all backups stored in a scale-out backup repository, rescan the repository. In this case, you do not need to place files in one folder; Veeam Backup & Replication will import backups automatically. For more information, see Rescanning Backup Repositories.

# Managing Backups

You can perform the following operations with backups:

- Viewing Backup Properties
- Upgrading Backup Chain Formats
- Moving Backups
- Copying Backups
- Exporting Backups
- Detaching Backups from Jobs
- Removing Backups from Configuration
- Deleting Backups from Disk
- Deleting Backups from Object Storage Repositories
- Deleting Backups from Scale-Out Backup Repositories
- Removing Missing Restore Points
- Launching Background Retention
- Disabling Background Retention

# Viewing Backup Properties

You can view summary information about created backups. The summary information provides the following data:

- Available restore points.
- Date of restore points creation.
- Compression and deduplication ratios.
- Data size (amount of data before compression and deduplication), backup size (actual, physical amount of data stored in the repository after compression and deduplication), original size (size of the selected VM) and total size (sum of all the original sizes displayed in **Objects**).
- GFS retention policy applied to restore points (W — weekly; M — monthly; Y — yearly).
- Backup retention date. This column is available for backups created by VeeamZIP, export backup or copy backup and with the retention period specified.

In the **Backup Properties** window, you can see different icons whose meaning is described in the Infrastructure Icons section.

To view summary information for backups:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click the backup and select **Properties**.
4. To see the list of available restore points, select the required object from the **Objects** list.

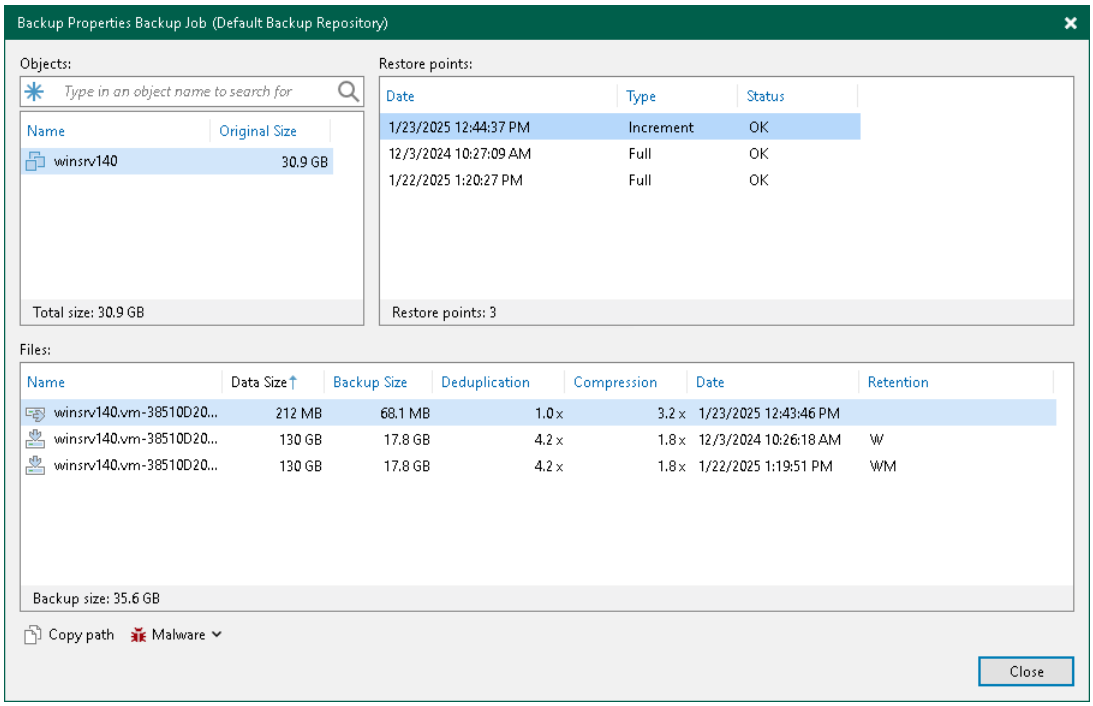

# Upgrading Backup Chain Formats

Veeam Backup & Replication supports the following ways to store backup files: per-machine backup with separate metadata files, per-machine backup with single metadata file and single-file backup. For more information, see Backup Chain Formats.

Changing the **Use per-machine backup files** option for existing backup repositories or moving backups to other repositories does not change the backup chain format. To change the format, follow the instructions in this section.

> **NOTE**
>
> Consider the following:
>
> - Change the backup chain format only for backups created successfully.
> - Before changing the backup chain format, Veeam Backup & Replication disables the job to which the backups belong. After the change, the job stays disabled. You need to enable it manually.
> - You must disable the log backup jobs manually before changing the format.
> - You can stop an upgrade session only for the backups for which the session has not started. If the session has started, Veeam Backup & Replication continues the upgrade to avoid backup corruption.
> - [For backups stored on repositories with immutability enabled] You can upgrade backups from per-machine backup with single metadata file to per-machine backup with separate metadata files. Other upgrades are not possible.
> - You cannot upgrade the backup chain format for backups stored on repositories with rotated drives.

## Per-Machine Backup with Single Metadata File to Per-Machine with Separate Metadata Files

To change the backup chain format from per-machine with single metadata file to per-machine with separate metadata files, use the upgrade functionality:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** node.
3. In the working area, select the necessary jobs. Note that the selected jobs must be targeted to the same repository.
4. Right-click one of the selected jobs and click **Upgrade backup chain format**. Alternatively, click **Upgrade Backup Chain Format** on the ribbon.

Veeam Backup & Replication will generate new metadata files for the existing backups. After the upgrade, the job will continue the backup chain and create per-machine backups with separate metadata files.

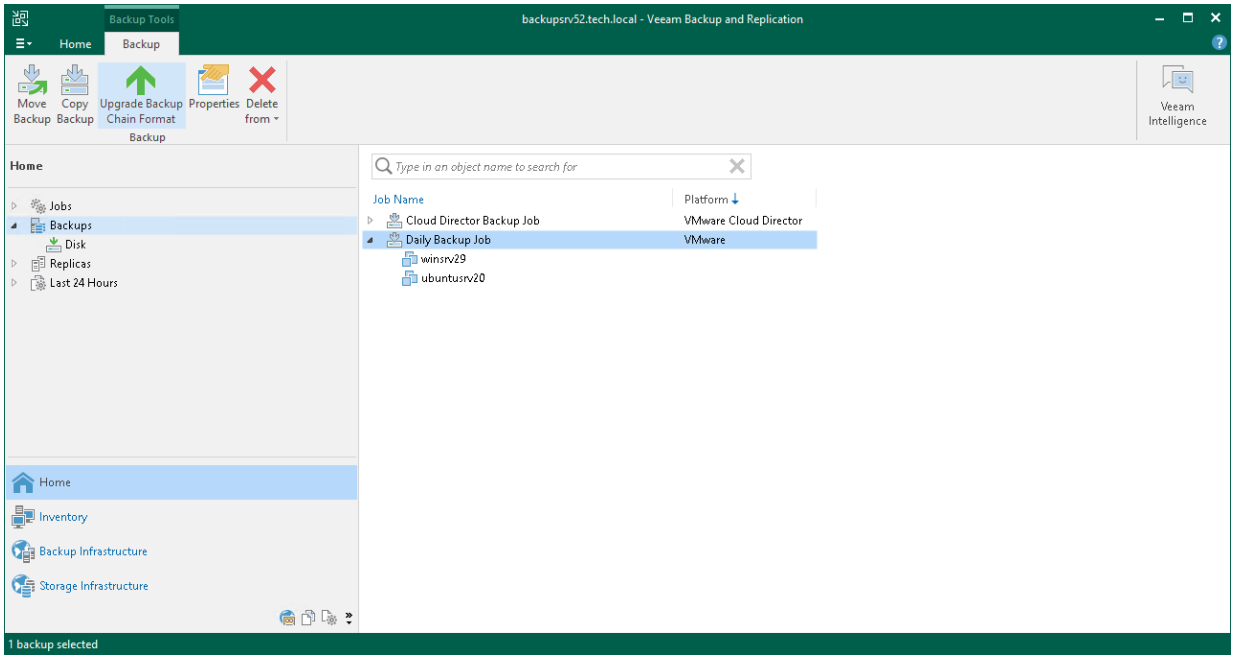


# Other Formats

To change the backup chain format, you must detach the existing backup chain from the job and then target the job to the required repository:

1. Detach backups from the job for whose backups you want to change the backup chain format. For more information, see Detaching Backups from Jobs.
2. Make sure that the job is targeted to a repository with the necessary state of the **Use per-machine backup files** check box.

   You can edit job settings and check the backup repository to which the job is targeted. Or, edit the repository settings to ensure the target repository has **Use per-machine backup files** check box selected. The selected check box is for per-machine backups with separate metadata files, the cleared check box is for single-file backups.

   > **NOTE**
   >
   > You cannot upgrade the backup chain to the per-machine backup with single metadata file as this format is obsolete since Veeam Backup & Replication version 12.

3. To create backup files of the required backup chain format, launch the job. The job will create active full backups.

   > **NOTE**
   >
   > If the job has a backup to tape job configured as a secondary target, the backup to tape job will archive a full backup during its next scheduled run.

# Related Topics

Upgrading Backup Chain Formats

# Moving Backups

Veeam Backup & Replication allows you to move all backups of a backup job to another repository or to move specific workloads and their backups to another job.

## Moving Backups to Another Repository

For information on the move operation, how it works and its limitations, see Backup Move.

To move backups to another repository and target a job to this repository, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** node.
3. In the working area, select the necessary job.
4. Right-click the job and select **Move backup**. Alternatively, click **Move Backup** on the ribbon.
5. In the **Move Backup to Another Location** window, select the repository to which you want to move backups.

   Veeam Backup & Replication will reconfigure and target the backup job to the selected repository.
6. Click **OK**.

Alternatively, you can change the repository in the job settings.

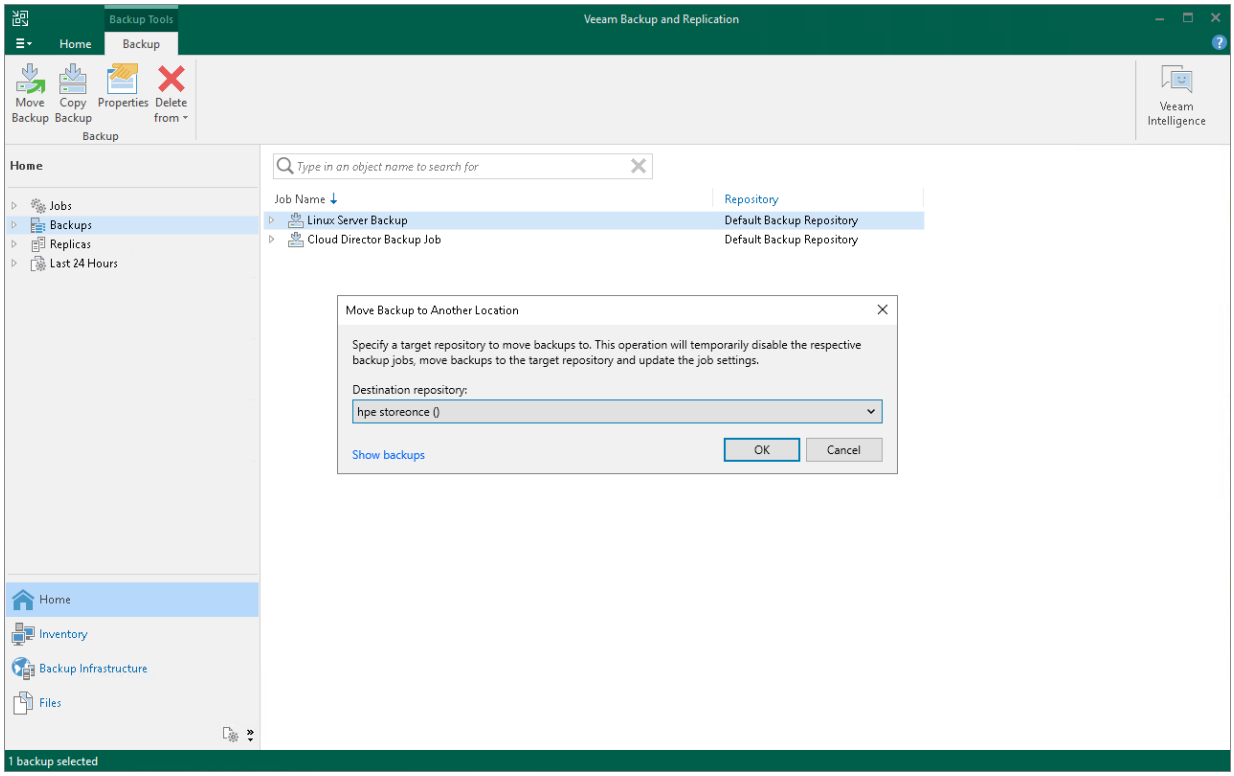


## Moving Backups to Another Job

For information on the move operation, how it works and its limitations, see Backup Move.

To move backups to another job:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** node.
3. In the working area, expand the necessary backup job and select workloads.

> **NOTE**
>
> You can move individual workloads and their backups only if backups are per-machine with separate metadata files for each workload.

4. Right-click one of the selected workloads and click **Move backup**. Alternatively, click **Move Backup** on the ribbon.
5. In the **Move Backup to Another Job** window, select the backup job to which you want to move backups.

   Veeam Backup & Replication will include workloads into the selected job and exclude workloads from the original job. Backups of the selected workloads will be moved to the repository to which the selected job is targeted.
6. Click **OK**.

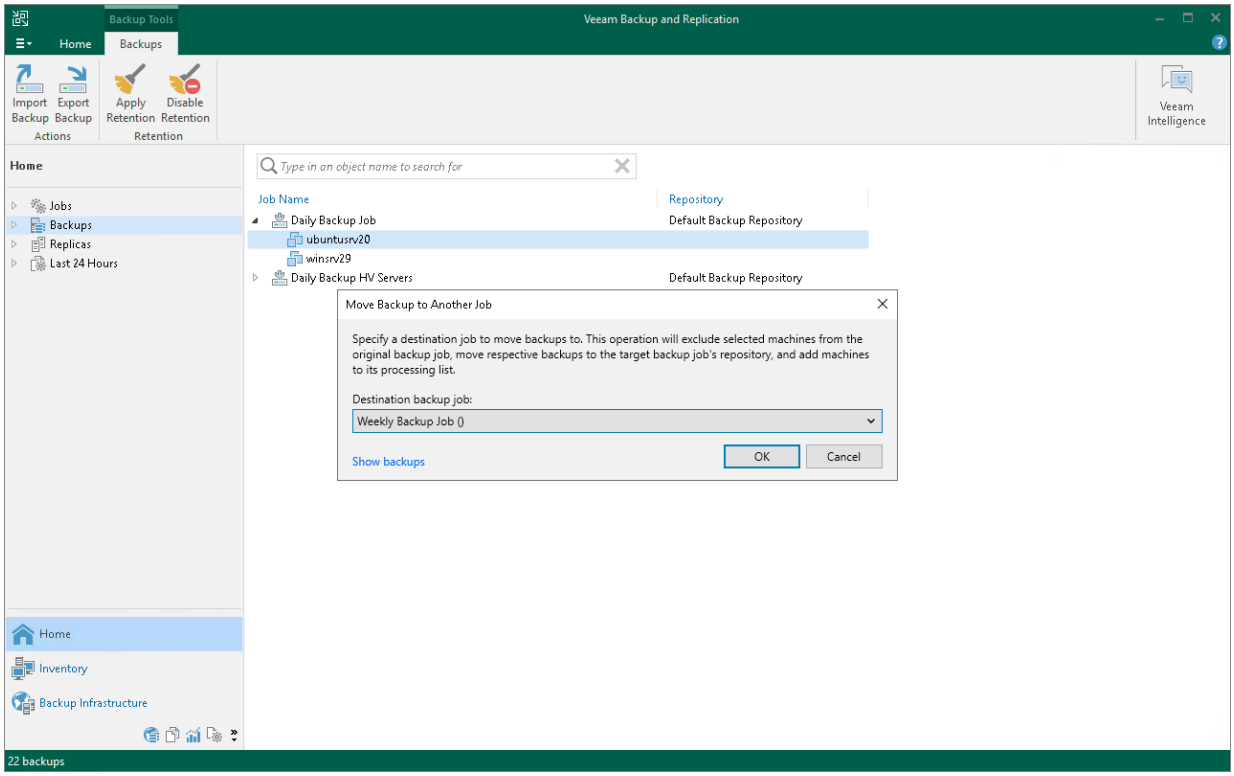


# Managing Failed Activities

If the move operation fails, Veeam Backup & Replication assigns the *User action required* status to it. In this case, you need to decide how to finish the operation: retry the move operation for failed workloads, cancel all changes, or move failed workloads but detach their backups. If you detach failed backups, the target job creates active full backups for failed workloads and continues backup chains for other workloads. The detached backups are shown in a node with the **(Orphaned)** postfix in the inventory pane.

> **NOTE**
>
> The original job will still be in the disabled state until you finalize the failed move operation.

To finalize the move operation:

1. Open the **Home** view.
2. In the inventory pane, select the **Last 24 Hours** node.
3. Right-click the failed move session and select the required action. Alternatively, select the required action on the ribbon.

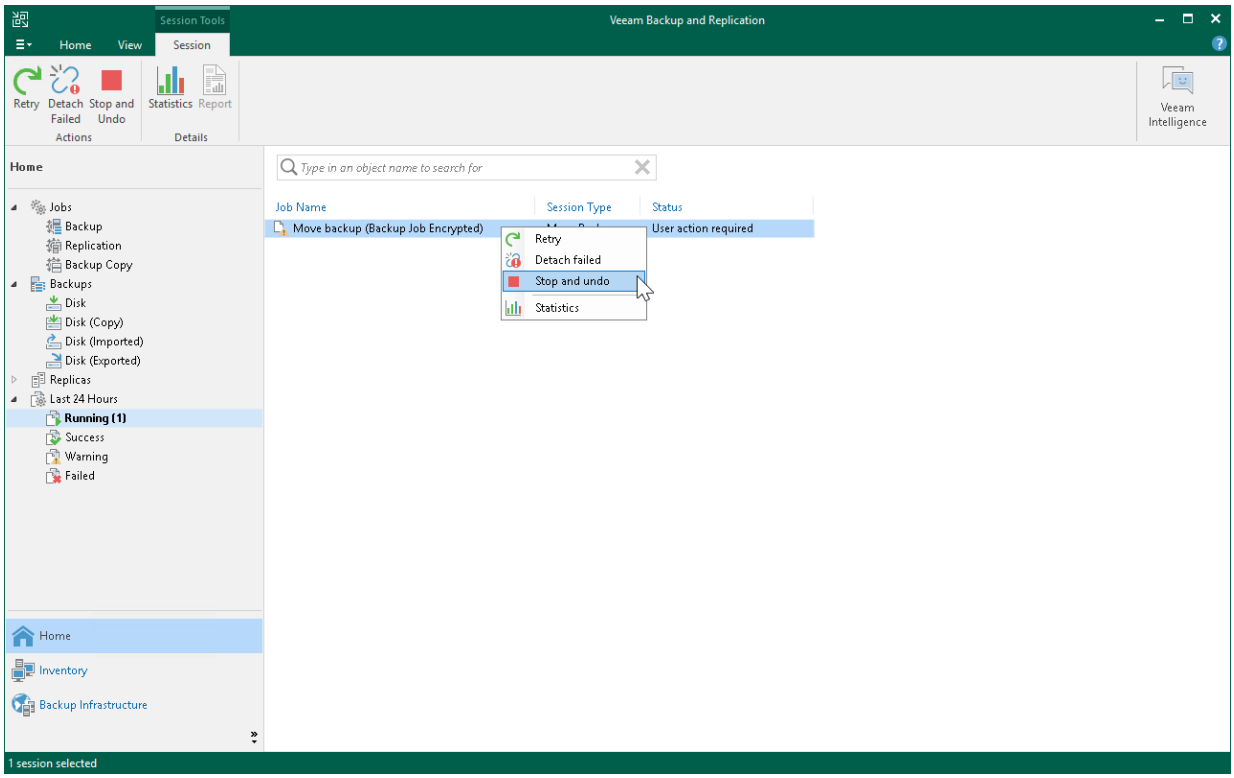

# Copying Backups

Copying backups can be helpful if you want to copy backups of a workload or backup job to another repository, local or shared folder. Veeam Backup & Replication copies the whole backup chain. If you want to convert a specific restore point into a single VBK file, use backup export. For more information, see Exporting Backups.

When Veeam Backup & Replication performs the copy operation, it disables the job, copies files to the target location and then enables the job. After the copy operation finishes, the copied backups are shown in a node with the **(Exported)** postfix in the inventory pane.

> **NOTE**
>
> This section is about one-time copy operation. If you want to copy backups on a schedule, create a backup copy job. For more information, see Backup Copy.

## Requirements and Limitations

Consider the following:

- The copy operation does not change the backup chain format (single-file backup, per-machine with single metadata file or per-machine with separate metadata files). If you copy backups between repositories with and without the **Use per-machine backup files** check box enabled, backups preserve their formats.
- If you copy backups from a scale-out backup repository and some backups are stored on extents in the Maintenance mode, such backups are not copied.
- Veeam Backup & Replication copies backups only from the performance tier of the scale-out backup repository. If you want to copy data from the capacity tier, you first must download it to the performance tier. For more information, see Downloading Data from Capacity Tier.
- You cannot copy backups between extents of a scale-out backup repository. To learn how to manage backups within the scale-out backup repository, see Scale-Out Backup Repositories.
- You cannot copy backups stored in Veeam Cloud Connect repositories. For more information on Veeam Cloud Connect repositories, see the Cloud Repository section in the Veeam Cloud Connect Guide.
- The Copy backup operations to, from and between cloud repositories are not supported. For more information on Veeam Cloud Connect's limitations, see the Veeam Cloud Connect Backup section in the Veeam Cloud Connect Guide.
- [For VMware Cloud Director] You can copy backups of a whole job or individual vApps. You cannot move backups of VMs.
- You cannot copy backups created by a backup job managed by Veeam Agent (backup policy).
- You cannot copy backups created by Veeam plug-ins for Enterprise applications, Veeam Cloud Plug-ins (Veeam Backup for AWS, Veeam Backup for Google Cloud, Veeam Backup for Microsoft Azure), and Veeam Kasten.
- Traffic throttling is not supported for backup copy operations.

## Copying Backups

To copy backups, do the following:

1. Open the **Home** view.

2. In the inventory pane, select the **Backups** node.
3. In the working area, select the necessary job or workload. Note that you can copy backups of an individual workload only if its backups are per-machine backups with separate metadata files.
4. Right-click the job and select **Copy backup**. Alternatively, click **Copy Backup** on the ribbon.
5. In the **Copy Backup to Another Location** window, choose where you want to copy backups — to a repository or to a local or shared folder.
6. If you want to delete the copied backups after a specific time period, select the **Delete this backup automatically in** check box and specify the time period.

   Backups that fall out of the specified retention policy will be removed automatically. If you do not specify the time period for deletion, copies will be stored until you remove them manually.

   > **TIP**
   >
   > You can customize retention period values in the drop-down list as described in this Veeam KB article.

7. Click **OK**.

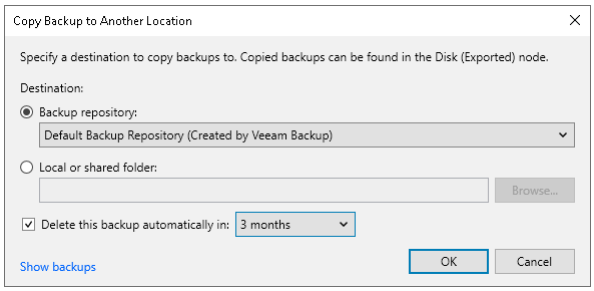


# Managing Failed Activities

If the copy operation fails, Veeam Backup & Replication assigns the *User action required* status to it. In this case, you need to decide how to finish the operation: retry the copy operation for failed backups, skip the failed backups or cancel all changes.

> **NOTE**
>
> The original job will still be in the disabled state until you finalize the failed copy operation.

To finalize the copy operation:

1. Open the **Home** view.
2. In the inventory pane, select the **Last 24 Hours** node.

3. Right-click the failed copy session and select the required action. Alternatively, select the required action on the ribbon.

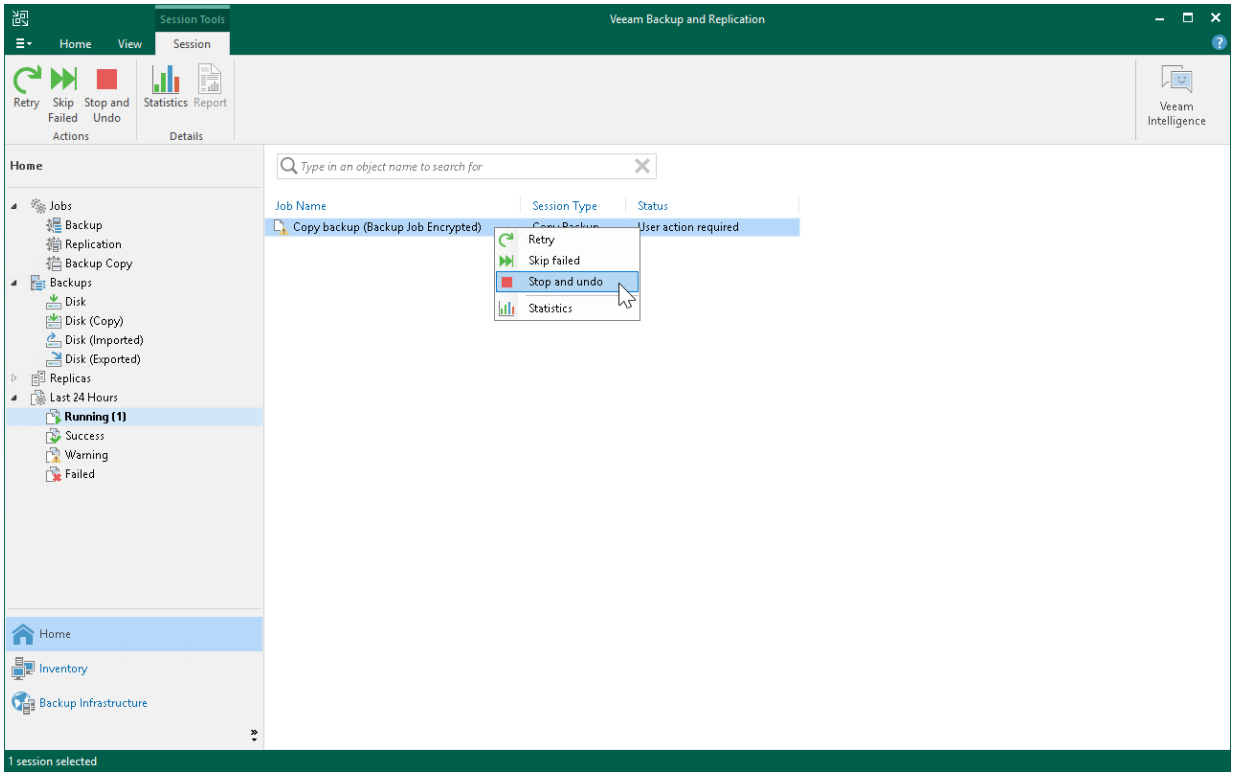

# Exporting Backups

Exporting backups allows you to synthesize a complete and independent full backup file out of selected restore points that are located in your backup repositories. You can transform any incremental or reverse-incremental backup chain (that is, all dependent .VBK, .VIB or .VRB files) into a full backup file and specify new retention policy settings for it. This option can be useful for legal hold and archiving purposes in case you want to prevent specific restore points from being deleted according to the retention policy of the primary backup job. It also allows you to provide a specific restore point on removable media.

> **TIP**
>
> Exporting backups allows you to export only a full backup file. If you want to export the whole backup chain, use the copy backup feature. For more information, see Copying Backups.

Export applies to *Full*, *Incremental* and *Reverse-incremental* restore points located in:

- Backup repositories.
- Object storage repositories.
- Scale-out backup repositories.
- Backup repositories of cloud service providers and their tenants.

Consider the following:

- Once export is complete, exported backup files will be displayed under the **Backups > Disk (Exported)** node.
- [For scale-out backup repositories] If you export backups to another scale-out backup repository, Veeam Backup & Replication will synthesize these backups in the performance tier. If you export data to the same scale-out backup repositories, Veeam Backup & Replication will synthesize these backups in the same tier from which it was exported:
  - If you export backups from the capacity tier, Veeam Backup & Replication synthesizes in the capacity tier. After a successful export, backups are displayed under the **Backups** > **Capacity Tier (Exported)** node.
  - If you export from the archive tier, Veeam Backup & Replication retrieves them from the archive tier and then synthesizes in the archive tier. After a successful export, backups are displayed under the **Backups** > **Archive Tier (Exported)** node.
- [For scale-out backup repositories] The target location to which you export backup defines where Veeam Backup & Replication synthesizes the backups:
  - If you export backups from the capacity tier to the same scale-out backup repositories, Veeam Backup & Replication synthesizes them in the capacity tier. After a successful export, backups are displayed under the **Backups** > **Capacity Tier (Exported)** node.
  - If you export from the archive tier to the same scale-out backup repositories, Veeam Backup & Replication retrieves backups from the archive tier and then synthesizes them in the archive tier. After a successful export, backups are displayed under the **Backups** > **Archive Tier (Exported)** node.
  - If you export backups to another scale-out backup repository, Veeam Backup & Replication will synthesize these backups in the performance tier.

> **NOTE**
>
> Note that the exported backup can be offloaded from the performance tier to the capacity or archive tier according to the settings of your scale-out backup repository.

- If a restore point that is being exported resides on the tenant side, a new full backup file will also be exported to the same repository (on the tenant side) from which the source restore point is being taken.
- If a tenant initiates the export of a restore point that resides in the *subtenant* directory, a new full backup file will be exported to the *tenant* directory.
- If you select an encrypted backup job, the exported backup file will be encrypted with the same password that you set in the advanced job settings.
- If you select a backup job consisting of multiple virtual machines, Veeam will synthesize a separate full backup file per each machine.
- When exporting VMs from VMware Cloud Director backups, all the VMs will be exported without vApps, that is, a new full backup file will be exported as a simple VMWare backup, not VMware Cloud Director backup. For more information about VMware Cloud Director backups, see Backup for VMware Cloud Director.
- Export session results are saved to the configuration database and available for viewing, as described in section Viewing Session Statistics.

# Performing Export

To export data, do the following:

1. Launch the Export Backup wizard.
2. Select restore points to export.
3. Specify a destination.
4. Specify an export reason.
5. Finish working with the wizard.

# Step 1. Launch Export Backup Wizard

To launch the **Export Backup** wizard, do either of the following:

- On the **Home** tab, click **Export Backup**.
- In the **Home** view under the **Backups > Disks** node, select a VM you want to transform into a full backup file and click **Export backup**.

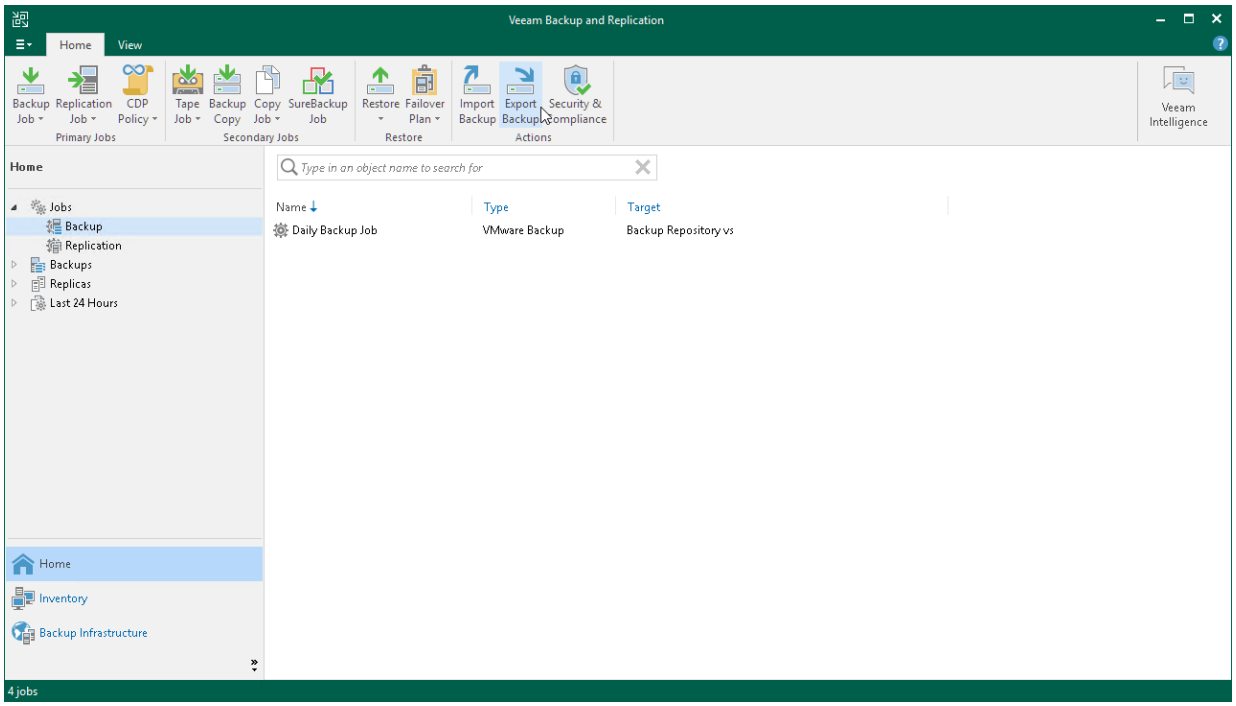

# Step 2. Select Restore Points to Export

At the **Restore Point** step of the wizard, do the following:

1. Click **Add** to select a VM, the restore points of which you want to transform into full backup files.
2. In the **Backups Browser** dialog box, select a backup job or virtual machine.

   When selecting a backup job consisting of multiple machines, each machine will be exported as an independent full backup file.
3. In the **Backups to export** list, select a VM and click **Point** to choose a restore point that you want to transform into a full backup file.

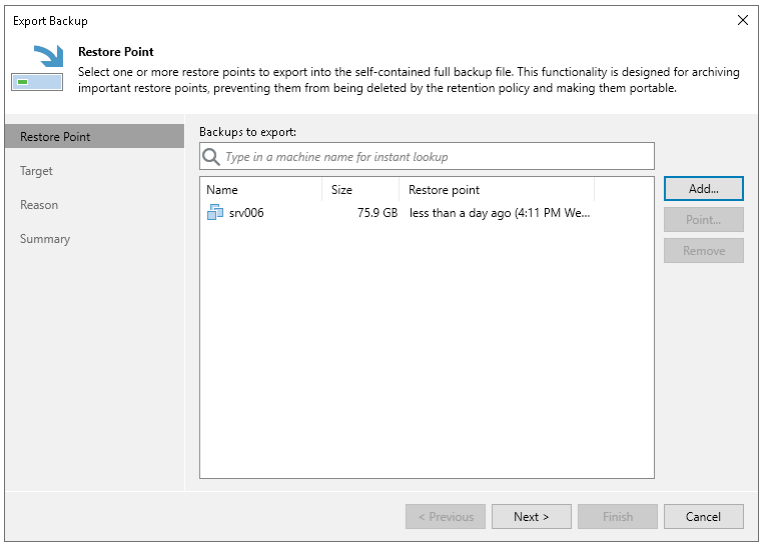

## Step 3. Specify Destination

At the **Target** step of the wizard, do the following:

1. In the **Export restore point to** area, choose whether you want to export restore points to a backup repository or to a local or shared folder.
2. If prompted, in the **Credentials** window, specify the credentials of the user account to access the target location.
3. To automatically delete the exported files after the specified period of time, select the **Delete exported backup file automatically** check box and specify the period. The exported files will be removed at 12:00:00 AM on the next day after the retention period ends.

> **TIP**
>
> You can customize retention period values in the drop-down list as described in this Veeam KB article.

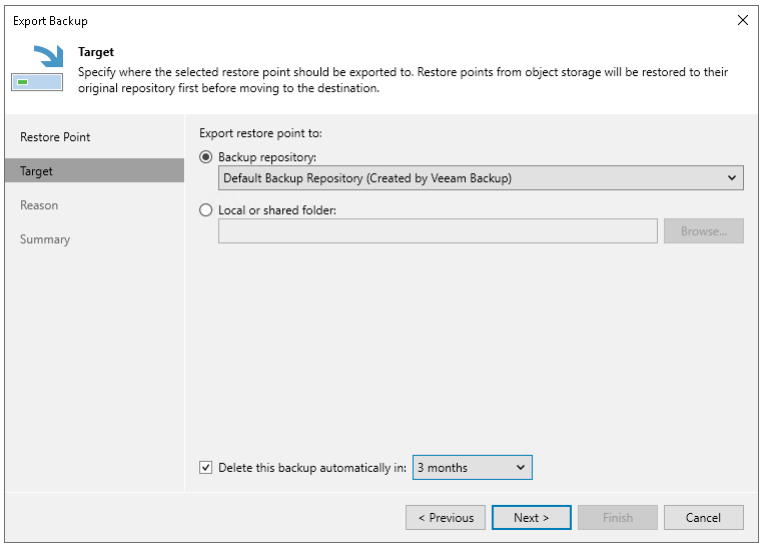

# Step 4. Specify Export Reason

At the **Reason** step of the wizard, provide the reason for restore.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

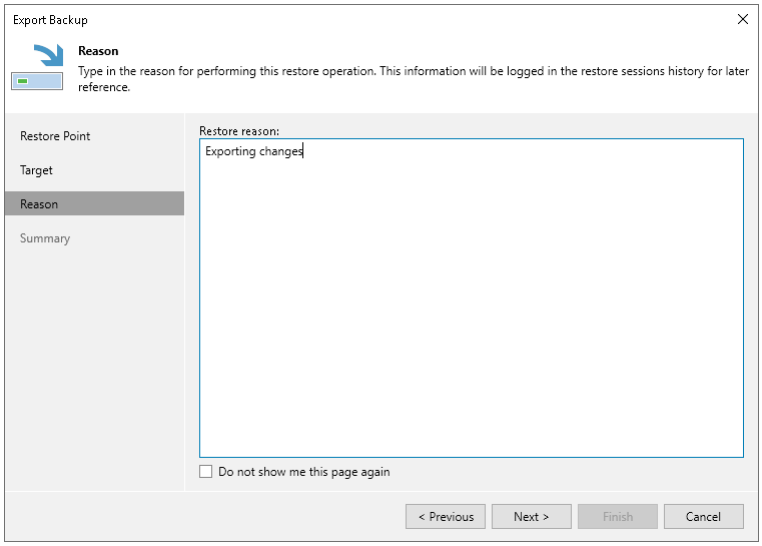

## Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information, click **Finish** and wait until the export session, which is described in section Viewing Session Statistics, is complete.

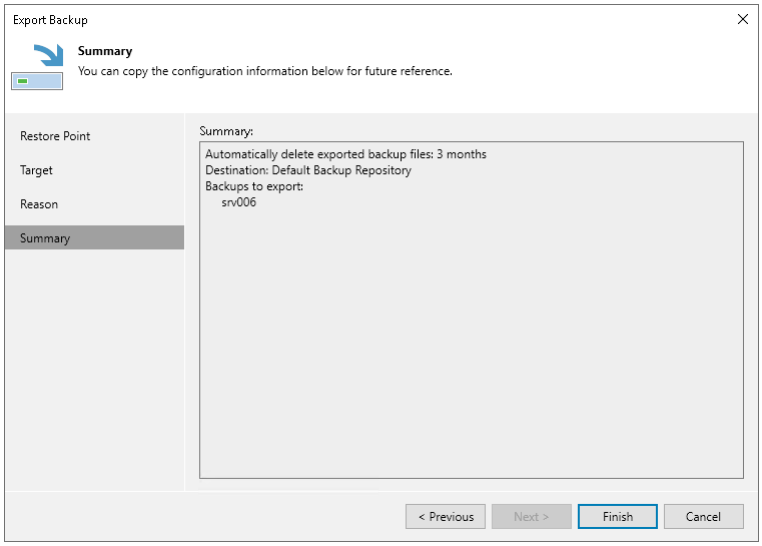


# Viewing Session Statistics

Once you invoke the export procedure, Veeam shows the **Export backup** progress dialog box that informs you of the current export status.

You can close the dialog box by clicking the **Close** button in the lower-right corner and let Veeam perform export in the background.

As each export session saves its results to the configuration database, you can review them at any time.

To review the export session results, do the following:

1. In the inventory pane, go to the **History** view and select the **Restore** > **Export** node.
2. In the working area, double-click a machine for which you want to review the session results or right-click a machine and select **Statistics**.

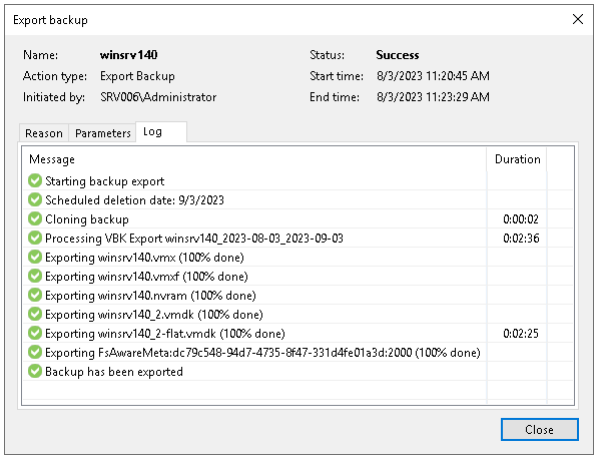


The **Export Backup** dialog box contains the following tabs:

- The **Reason** tab — shows you the reason for export you may have provided at the Specify Export Reason step of the wizard.

- The **Parameters** tab – shows you the date when the exported backup files will be removed due to the retention policy you may have configured at the Select Restore Points to Export step of the wizard. In this tab, you can also find a backup name and Date/time of a restore point that was synthesized into a full backup file.
- The **Log** tab – shows you the actual export progress.

# Canceling Session

To cancel a session, open the **Export Backup** dialog box, as described in steps 1-2 above, and click **Cancel restore task** in the upper-right corner of the dialog box.

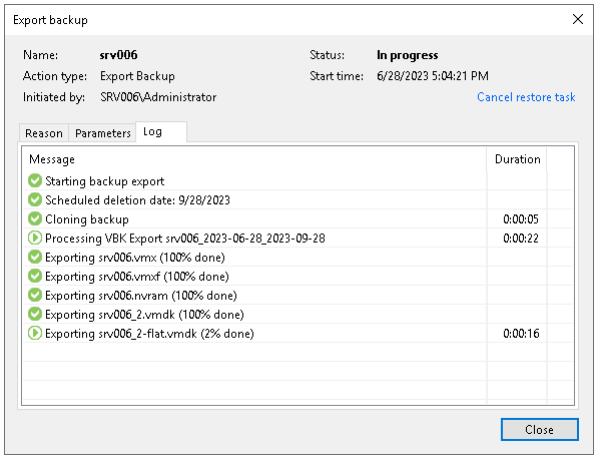

# Detaching Backups from Jobs

If you want to detach backups from a job, you can use the **Detach from job** operation.

When you detach backups from a job, the job stops processing these backup files. During the next run, the job will start a new backup chain; that is, will create active full backups.

The detached backup files remain in the backup repository and the Veeam Backup & Replication console. Veeam Backup & Replication shows the detached backups in the inventory pane in the node with the **(Orphaned)** postfix. These backups are retained according to the background retention process. For more information, see Background Retention.

To detach backups from a job:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** node.
3. In the working area, right-click the necessary backup and select **Detach from job**. Alternatively, click **Delete from** > **Job** on the ribbon.

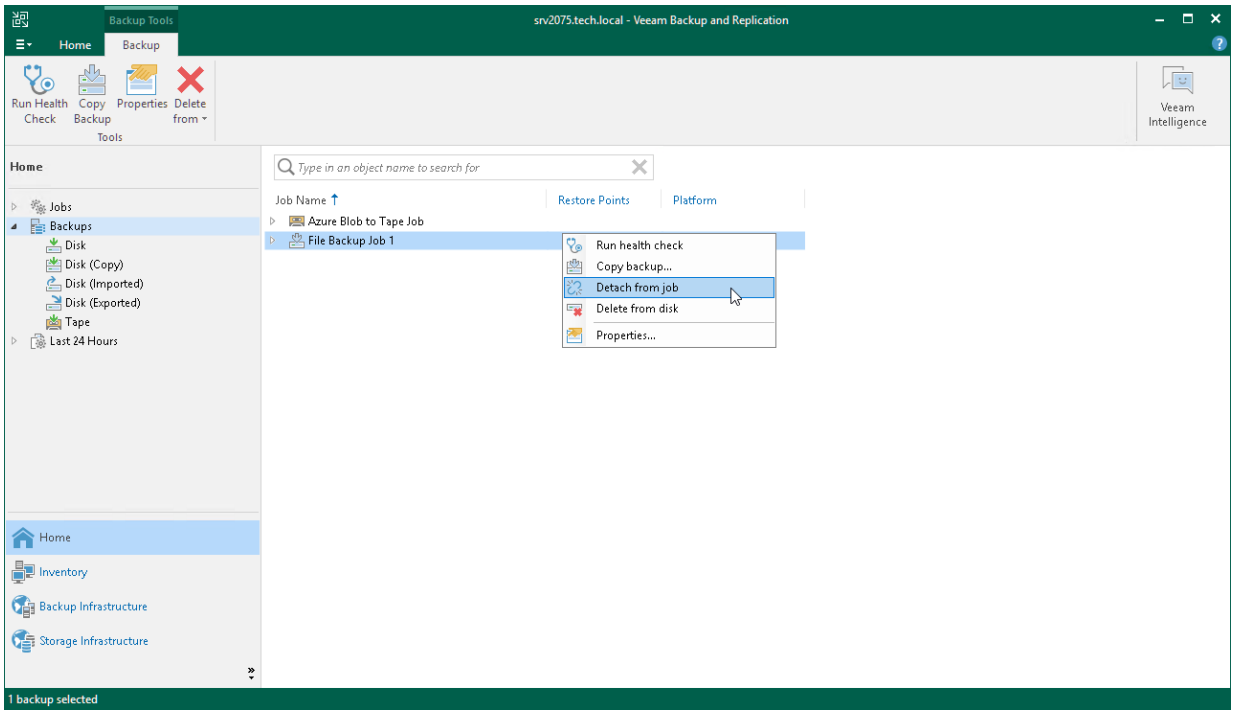

# Removing Backups from Configuration

If you want to remove records about backups from the Veeam Backup & Replication console and configuration database, you can use the **Remove from configuration** operation.

When you remove a backup from the configuration, backup files remain in the backup repository. You can import the backup to Veeam Backup & Replication at any time later and restore data from it.

When you remove an encrypted backup from the configuration, Veeam Backup & Replication removes encryption keys from the configuration database. If you import such a backup on the same backup server or another backup server, you will have to specify the password or unlock the backup with Veeam Backup Enterprise Manager. For more information, see Importing Encrypted Backups.

To remove a backup from the configuration:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** node.
3. In the working area, select the backup you want to remove, press and hold the [Ctrl] key, right-click the backup and select **Remove from configuration**.

> **IMPORTANT**
>
> Removing backups from configuration is designed for experienced users only. Consider using Deleting Backups from Disk or Detaching Backups from Jobs operations.
>
> Create an encrypted configuration backup before removing backups from the configuration.

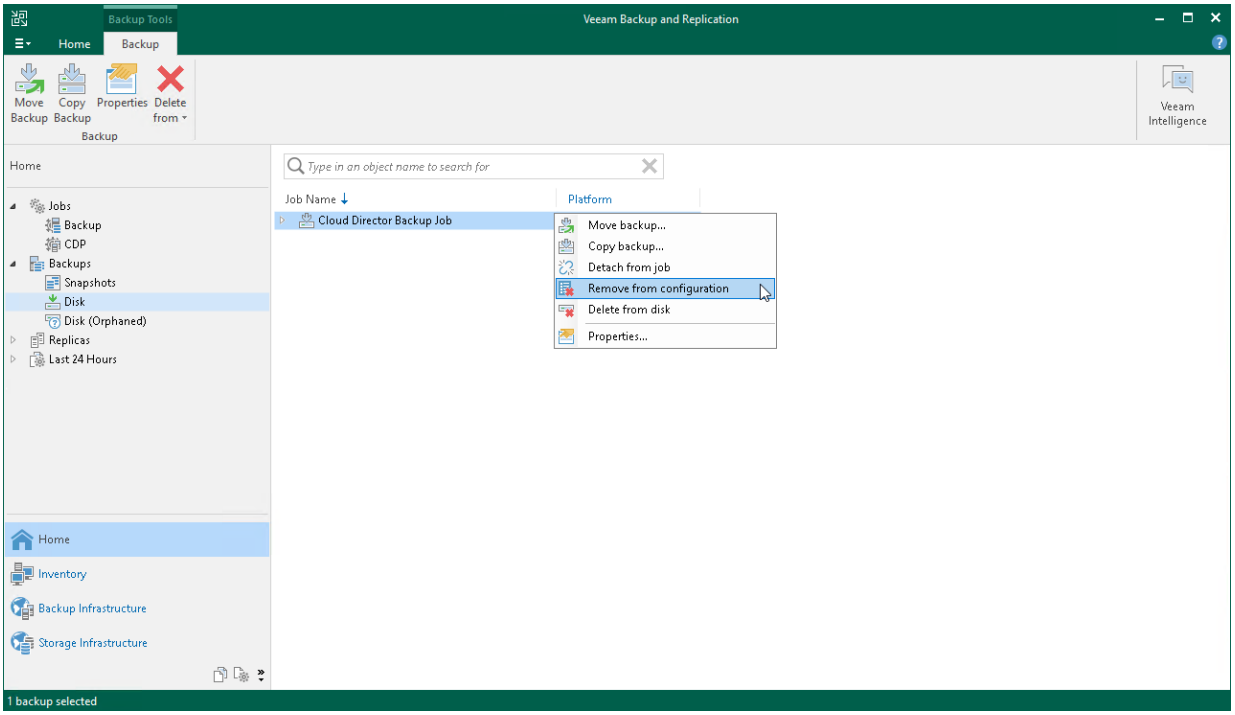

# Deleting Backups from Disk

If you want to permanently remove backup files from the backup repository and Veeam Backup & Replication console, you can use the **Delete from disk** operation. This option allows you to delete the following type of data:

- Backup files from the backup repository
- Separate VMs from backups

When you delete backup files from a disk, Veeam Backup & Replication deletes the whole chain from the backup repository. Thus, on the next backup job run, Veeam Backup & Replication will create full backups for VMs included in the job.

Consider the following:

- Do not delete backup files from the backup repository manually. Use the **Delete from disk** option instead. If you delete backup files manually, subsequent backup or replication job sessions will fail.
- If the per-machine functionality is enabled, you can perform the **Delete from disk** operation for separate VMs in the backup. If you delete backup files for one VM, on the next run of the backup job Veeam Backup & Replication will create a full backup for VMs whose backup files are deleted. For all other VMs, Veeam Backup & Replication will create increments.

  To learn more about per-machine backup files, see Backup Chain Formats.
- When you delete a separate VMs from a backup and if the per-machine functionality is enabled for this VM, Veeam Backup & Replication behaves differently:
  - [If per-machine is enabled] Veeam Backup & Replication deletes backup files of the selected VMs.

    On the next run of the job, Veeam Backup & Replication will create full backups for VMs whose backup files were deleted. For all other VMs, Veeam Backup & Replication will create increments.
  - [If per-machine is disabled] Veeam Backup & Replication only marks data blocks that belong to the deleted VMs as empty — the size of backup files does not change. However, Veeam Backup & Replication will use these data blocks during such operations as merging backup files. To reduce the size of full backup files in forever forward incremental and reverse incremental backup chains, you can compact full backup files. In forward incremental backup chains, files with blocks marked as empty will be deleted by retention.

    On the next run of the job, Veeam Backup & Replication will create full backups for VMs whose backup files were deleted. Note that full backups of these VMs will be stored in an incremental file. For all other VMs, Veeam Backup & Replication will create incremental backups.
- If you use the scale-out backup repository, keep in mind that the **Delete from disk** operation will remove the backups not only from the performance tier but also from the capacity and archive tier. If you want to remove backups from the performance tier only, you should move those backups to the capacity tier instead. For details, see Manually Moving Backups to Capacity Tier.

# Deleting Backups

To delete backup files from the backup repository, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups** or **Replicas**.
3. In the working area, select the backup and click **Delete from** > **Disk** on the ribbon. You can also right-click the backup and select **Delete from disk**.

4. To remove backups with GFS flags (weekly, monthly and yearly), select the **Remove GFS full backups** check box and click **Yes**.

# Deleting VMs from Backups

To delete a VM from a backup, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups** or **Replicas**.
3. In the working area, expand the necessary backup, select the VM you want to delete and click **Delete from** > **Disk** on the ribbon. You can also right-click the backup and select **Delete from disk**.
4. To remove backups with GFS flags (weekly, monthly and yearly), select the **Remove GFS full backups** check box and click **Yes**.

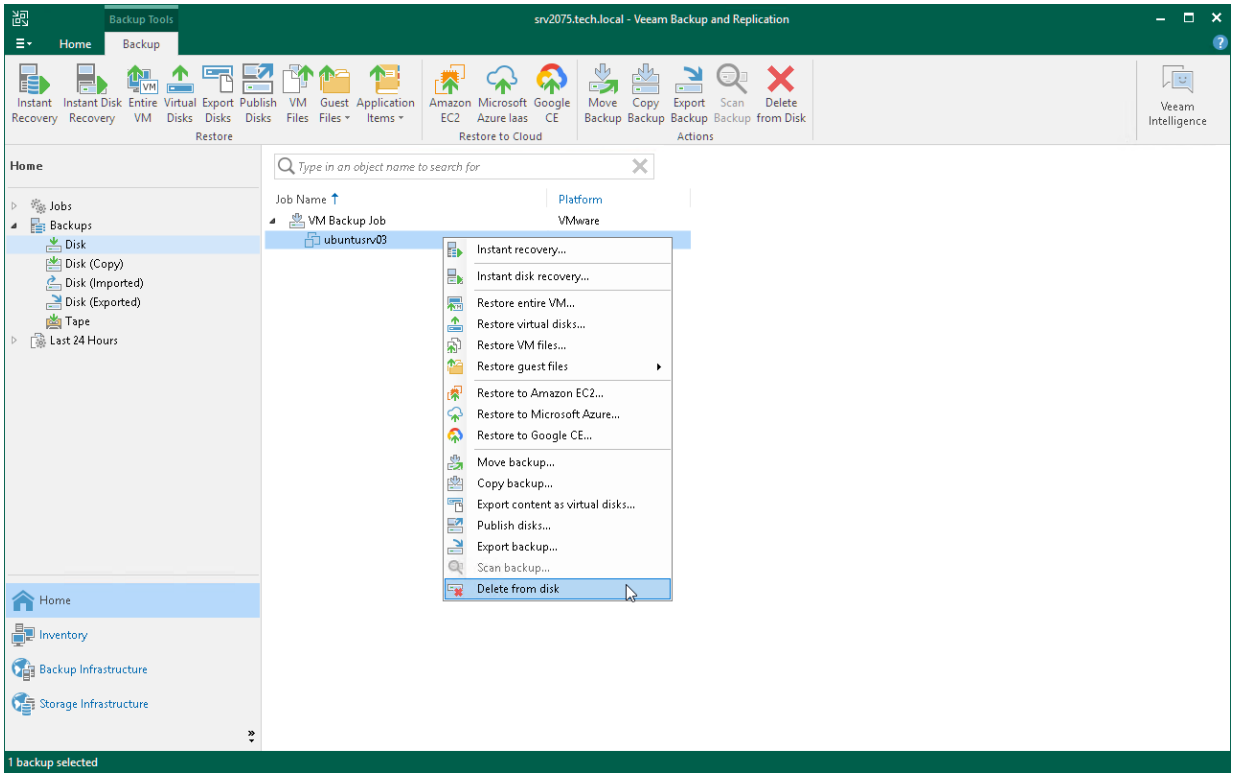

# Deleting Backups from Object Storage Repositories

This section describes how to delete backups from object storage repositories. To know how to delete backups from an object storage repository added as an extent of the scale-out backup repository, see Deleting Backups from Scale-Out Backup Repositories.

To delete a backup from object storage repositories, do the following:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Backups** > **Object Storage** node.
3. In the working area, select a backup or VM and click **Delete from Disk** on the ribbon. Alternatively, you can right-click a backup and select **Delete from disk**.

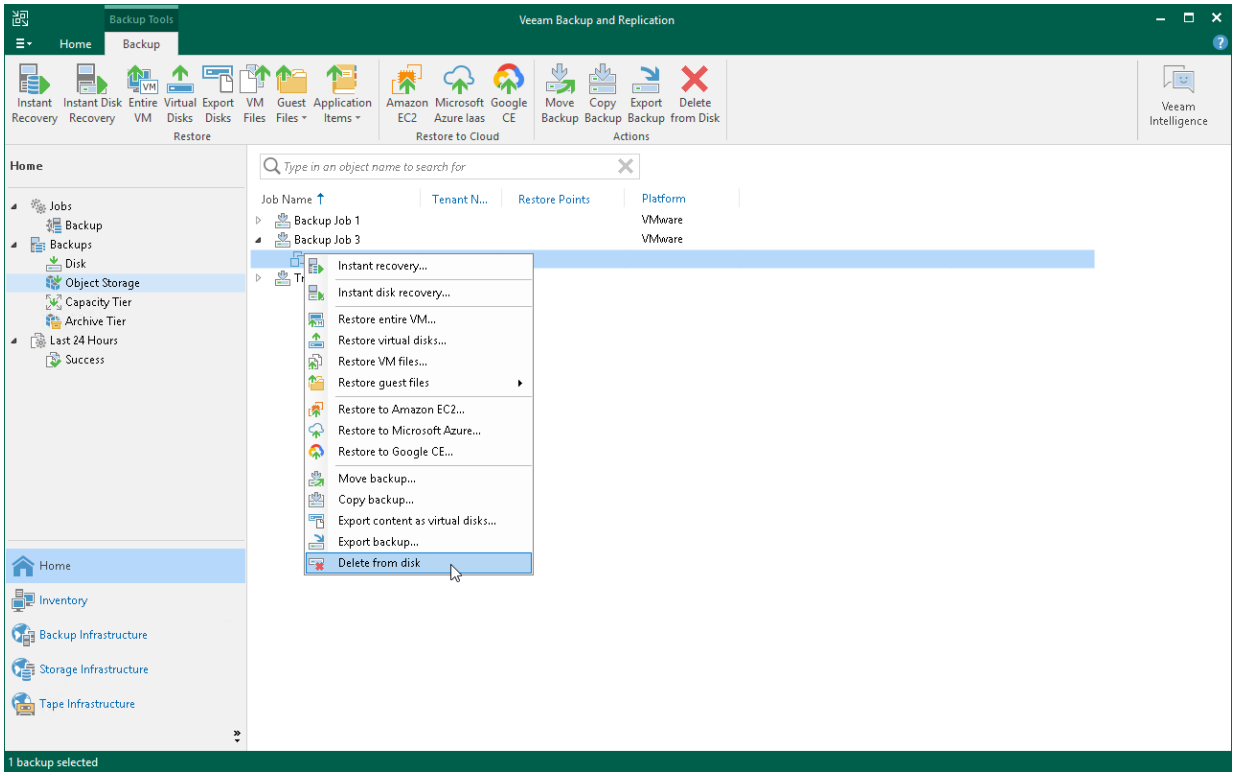

# Deleting Backups from Scale-Out Backup Repositories

To delete a backup from a scale-out backup repository, do the following:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Backups** and select one of the following nodes:
   - Select the **Object Storage** node if you want to delete a backup from the object storage repository added as a performance extent.

     > **IMPORTANT**
     >
     > Keep in mind that the **Delete from disk** operation will remove the backups not only from the performance tier but also from the capacity and archive tier. If you want to remove backups from the performance tier only, you should move those backups to the capacity tier instead. For details, see Manually Moving Backups to Capacity Tier.

   - Select the **Capacity Tier** node if you want to delete a backup from the capacity tier.
   - Select the **Archive Tier** node if you want to delete backups from the archive tier.
3. In the working area, select a backup or VM and click **Delete from Disk** on the ribbon. Alternatively, you can right-click a backup and select **Delete from disk**.

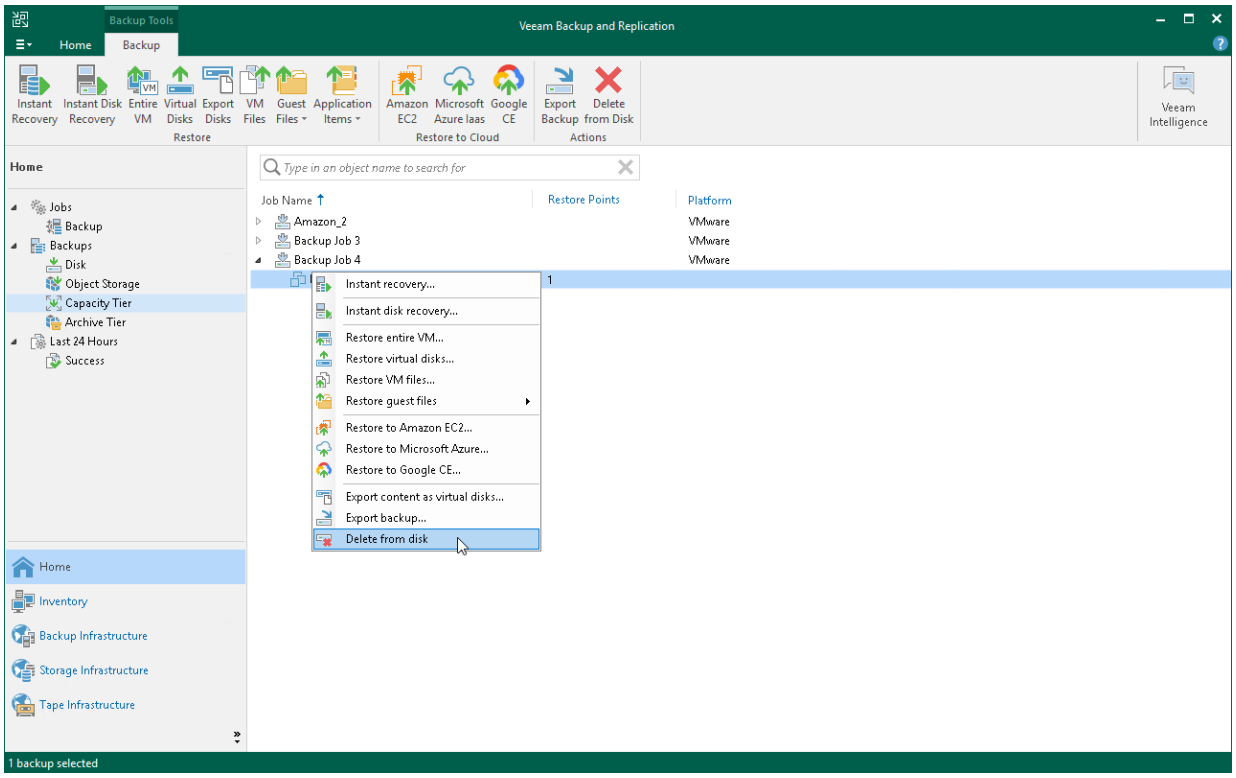

# Removing Missing Restore Points

In some cases, one or more restore points in the backup chain may be inaccessible. It can happen, for example, if the backup repository is put into Maintenance mode (for scale-out backup repositories), the backup repository is unavailable, or some backup file is missing in the backup chain. Backup chains that contain missing restore points get corrupted — you cannot perform backup or restore VM data from the missing restore point and restore points that depend on the missing restore point.

You can perform two operations with missing restore points:

- Forget — you can remove records about missing restore points from the configuration database. Veeam Backup & Replication will "forget" about missing restore points and will not display them in the console. The actual backup files will remain on disk (if backup files are still available).
- Remove — you can remove records about missing restore points from the configuration database and delete backup files from disk (if backup files are still available).

> **NOTE**
>
> Consider the following:
>
> - The **Forget** and **Remove from disk** options are available only for the restore points missing from the backup chain or points that depend on missing ones. If the restore point is available in the backup chain and does not depend on a missing restore point, you will not be able to use the **Forget** and **Remove from disk** options for it.
> - You can manually update information about missing restore points. For this, disable a backup job and rescan the backup repository that is the target for the job. For more information, see Disabling and Deleting Jobs and Rescanning Backup Repositories.
>
>   Manual update can be required because Veeam Backup & Replication requires some time to update information in the configuration database for restore points that were removed from a backup chain or became inaccessible. That is why such restore points may not be displayed in the console as missing restore points.
> - Veeam Backup & Replication does not track missing restore points in backups that reside in the cloud repository.
> - If you apply the **Forget** or **Remove from disk** options to a missing restore point in a scale-out backup repository, the backup file associated with the missing restore point will be deleted from the capacity tier and archive tier on the next offload and archiving job run.

To remove records about missing restore points from the configuration database:

1. Open the **Home** view.
2. In the inventory pane, select **Disk** under **Backups**.
3. In the working area, select the backup and click **Properties** on the ribbon or right-click the backup and select **Properties**.
4. In the **Backup Properties** window, right-click the missing restore point and select **Forget**.
   - To remove only the selected restore point and restore points that depend on it (that is, a part of the backup chain starting from this restore point), select **This and dependent backups**.

- To remove all missing restore points, select **All unavailable backups**.

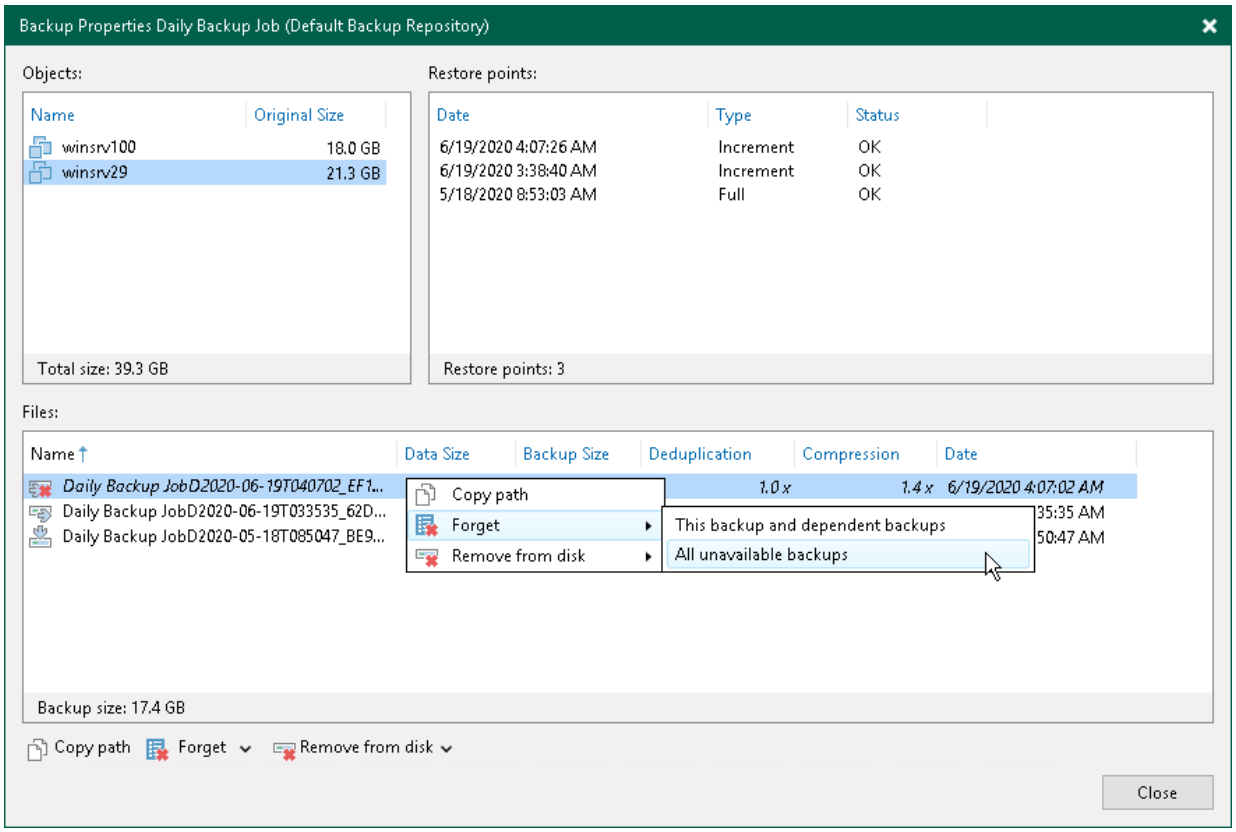


To remove missing restore points from the configuration database and disk:

1. Open the **Home** view.
2. In the inventory pane, click **Disk** under **Backups**.
3. In the working area, select the backup and click **Properties** on the ribbon or right-click the backup and select **Properties**.
4. In the **Backup Properties** window, right-click the missing restore point and select **Remove from disk**.
   - To remove only the selected restore point and restore points that depend on it (that is, a part of the backup chain starting from this restore point), select **This and dependent backups**.

- To remove all missing restore points, select **All unavailable backups**.

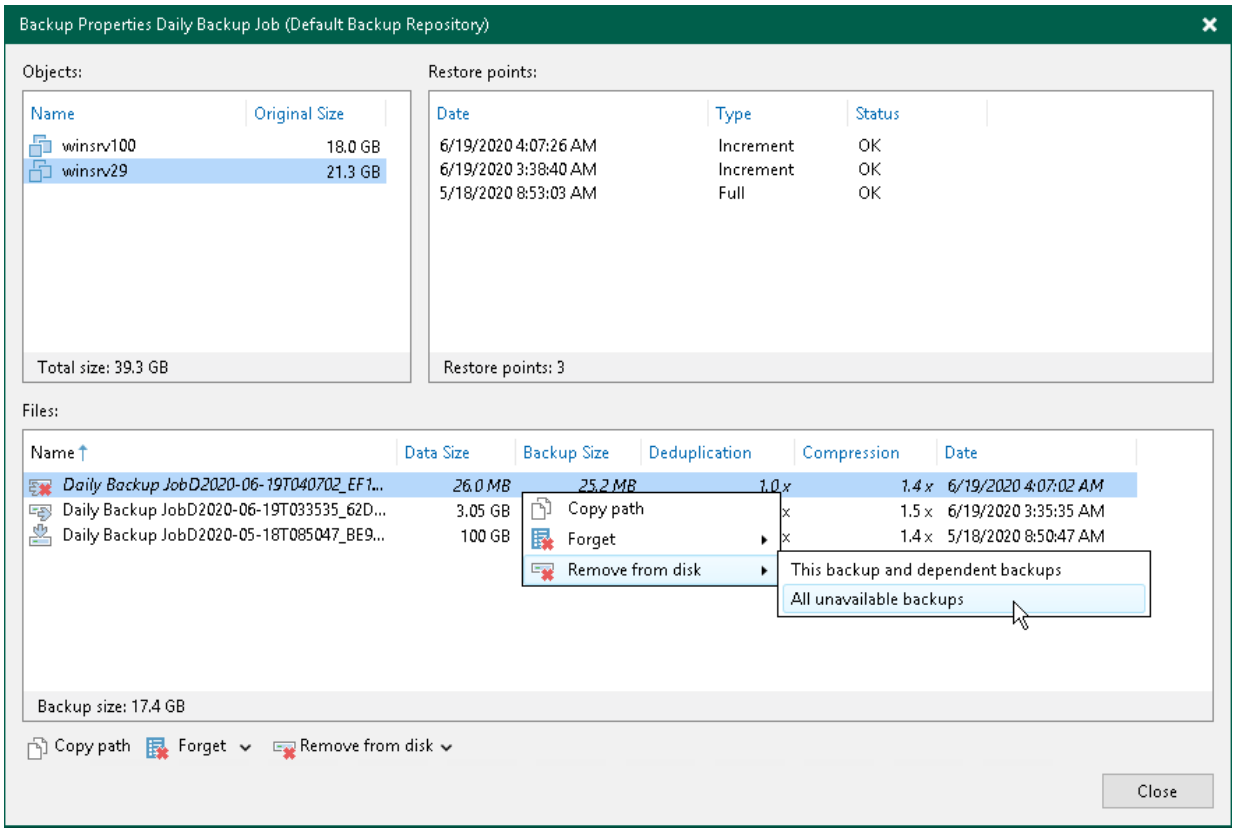

# Launching Background Retention

Veeam Backup & Replication launches background retention every 24 hours at 00:30. If you need to free up disk space from outdated backups earlier, you can launch the background retention manually.

> **NOTE**
>
> Be careful when launching background retention to avoid losing necessary data. For more information, see Background Retention.

To launch the background retention outside its automatic schedule, do the following:

1. Open the **Home** view.
2. In the inventory pane, right-click the **Backups** node.
3. Select **Apply retention now** or click **Apply Retention** on the ribbon.

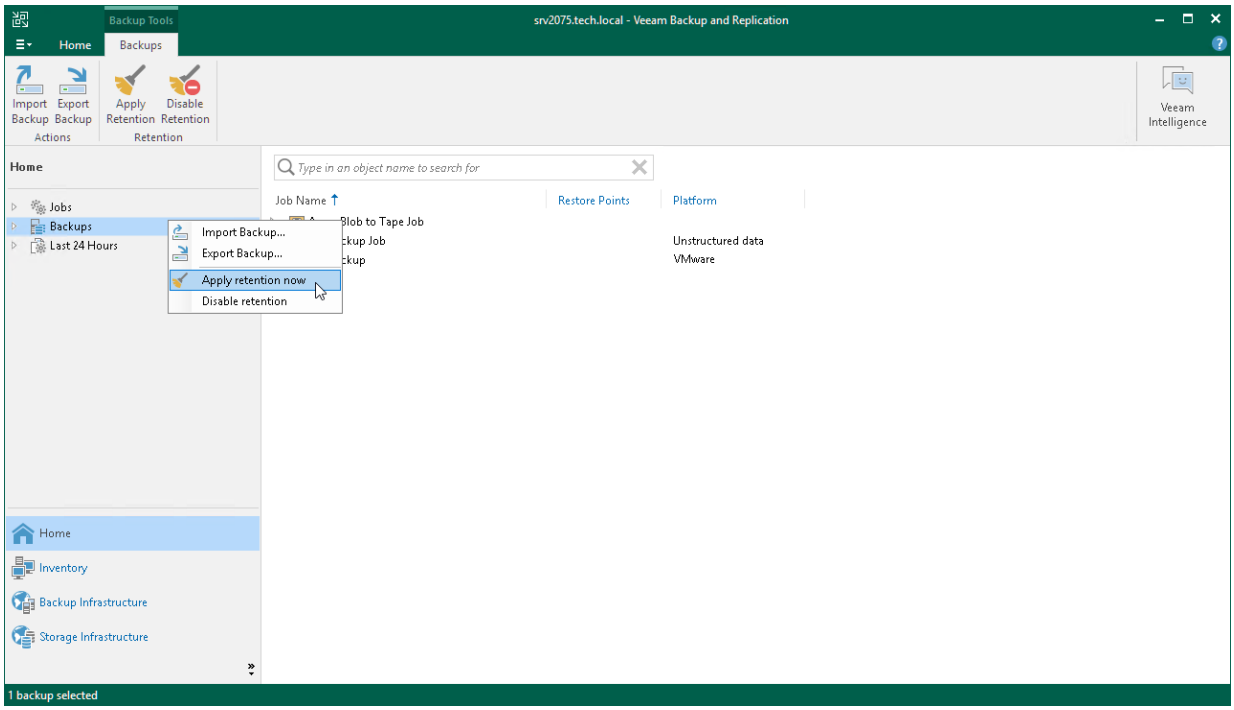

# Disabling Background Retention

You can manually disable background retention, which is launched every 24 hours at 00:30. You can enable disabled background retention at any time.

If you turn off background retention, you will receive a daily email report reminding you that it is disabled. Veeam Backup & Replication will send these reports only after you have enabled and configured email notifications, as described in the section Configuring Global Email Notification Settings.

If background retention is turned off while backup jobs are still enabled, auxiliary data cannot be cleaned from object storage repositories. This can decrease backup job performance over time.

To disable background retention:

1. Open the **Home** view.
2. In the inventory pane, right-click the **Backups** node.
3. Select **Disable retention** or click **Disable Retention** on the ribbon.

To enable disabled background retention, select the **Backups** node and click **Disable Retention** on the ribbon again. Alternatively, you can right-click the **Backups** node and select **Disable retention**.

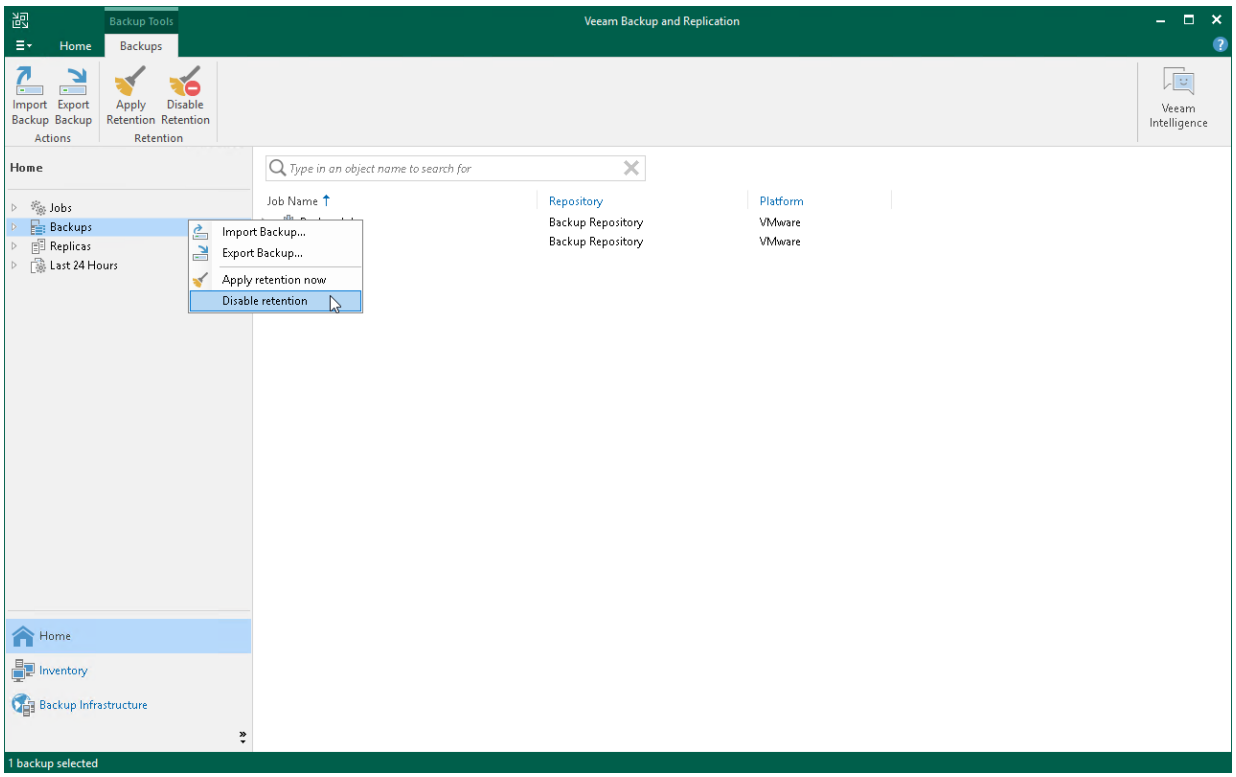

# Managing Backup Jobs

To view all jobs configured on the backup server, open the **Home** view and select the **Jobs** node in the inventory pane. The list of available jobs is displayed in the working area. You can edit job properties, start and stop jobs, restart failed jobs, clone jobs, view job statistics and delete unnecessary jobs.

# Editing Job Settings

After you add a backup job, you can edit its settings at any time. For example, you may want to change the scheduling settings or add VMs to the job.

## Editing General Settings

To edit job settings:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs** > **Backup**.
3. In the working area, select the job and click **Edit** on the ribbon or right-click the job and select **Edit**.

You will follow the same steps you followed when creating the job and can change job settings as required.

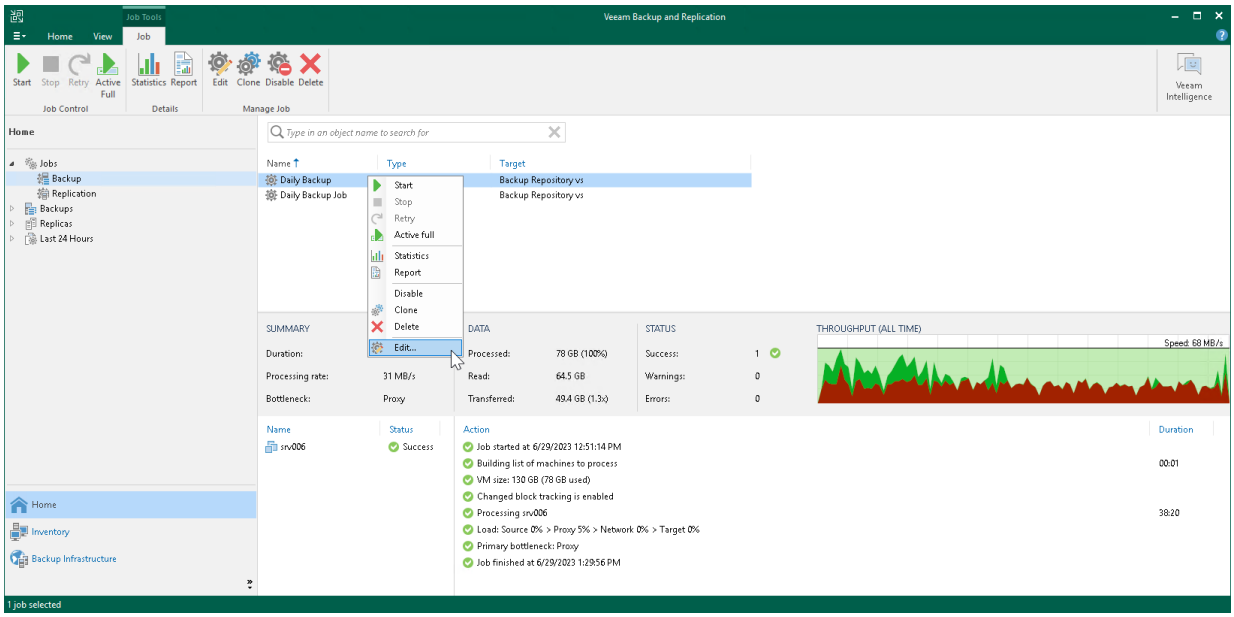


## Disabling GFS Scheme

If you disable the **Keep certain full backups longer for archival purposes** option at the **Storage** step of the **Edit Backup Job** wizard, and there are archive full backups in the target backup repository, Veeam Backup & Replication offers to remove them.

- Click **Yes** to remove archive full backups from the target backup repository. These backups will be removed during the next retention cycle (next backup session). The backup job will not create archive full backups.
- Click **No** to cancel the disable operation.

> **NOTE**
>
> If you disable the **Keep certain full backups longer for archival purposes** option and enable it again later, archive full backups that remained on disk will not be linked to the backup job. They will still be displayed under the **Backups** > **Disk (Imported)** node in the Veeam Backup & Replication console.

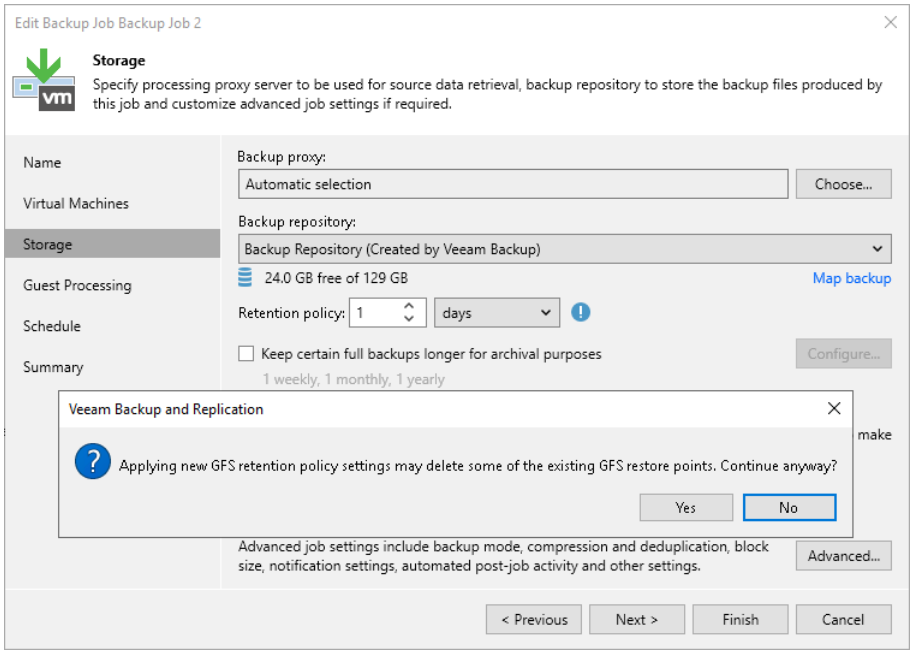
Edit Backup Job Backup Job 2
Storage
Specify processing proxy server to be used for source data retrieval, backup repository to store the backup files produced by this job and customize advanced job settings if required.
Name
Virtual Machines
Storage
Guest Processing
Schedule
Summary
Backup proxy:
Automatic selection
Choose...
Backup repository:
Backup Repository (Created by Veeam Backup)
24.0 GB free of 129 GB
Map backup
Retention policy:
1
days
Keep certain full backups longer for archival purposes
Configure...
1 weekly, 1 monthly, 1 yearly
Veeam Backup and Replication
Applying new GFS retention policy settings may delete some of the existing GFS restore points. Continue anyway?
Yes
No
make
Advanced job settings include backup mode, compression and deduplication, block size, notification settings, automated post-job activity and other settings.
Advanced...
< Previous
Next >
Finish
Cancel

# Cloning Jobs

You can create new jobs using job cloning. Job cloning allows you to create an exact copy of any job with the same job settings. Configuration information of the created job copy is written to the configuration database that stores information about the original job.

To create multiple jobs with similar settings, you can configure a set of jobs that will be used as 'job templates'. You can then clone these 'job templates' and edit the settings of cloned jobs as required.

The name of the cloned job is formed by the following rule: *<job_name_clone1>*, where *job_name* is the name of the original job and *clone1* is a suffix added to the original job name. If you clone the same job again, the number in the name will be incremented, for example, *job_name_clone2*, *job_name_clone3*, and so on.

When cloning a job, Veeam Backup & Replication can change some job settings so that cloned jobs do not hinder original jobs.

- If the original job is scheduled to run automatically, Veeam Backup & Replication disables the cloned job. To enable the cloned job, select it in the job list and click **Disable** on the ribbon or right-click the job and select **Disable**.
- If the original job is configured to use a secondary target, the cloned job is created without the secondary target settings.

To clone a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Clone** on the ribbon or right-click the job and select **Clone**.
4. After a job is cloned, you can edit all its settings, including the job name.

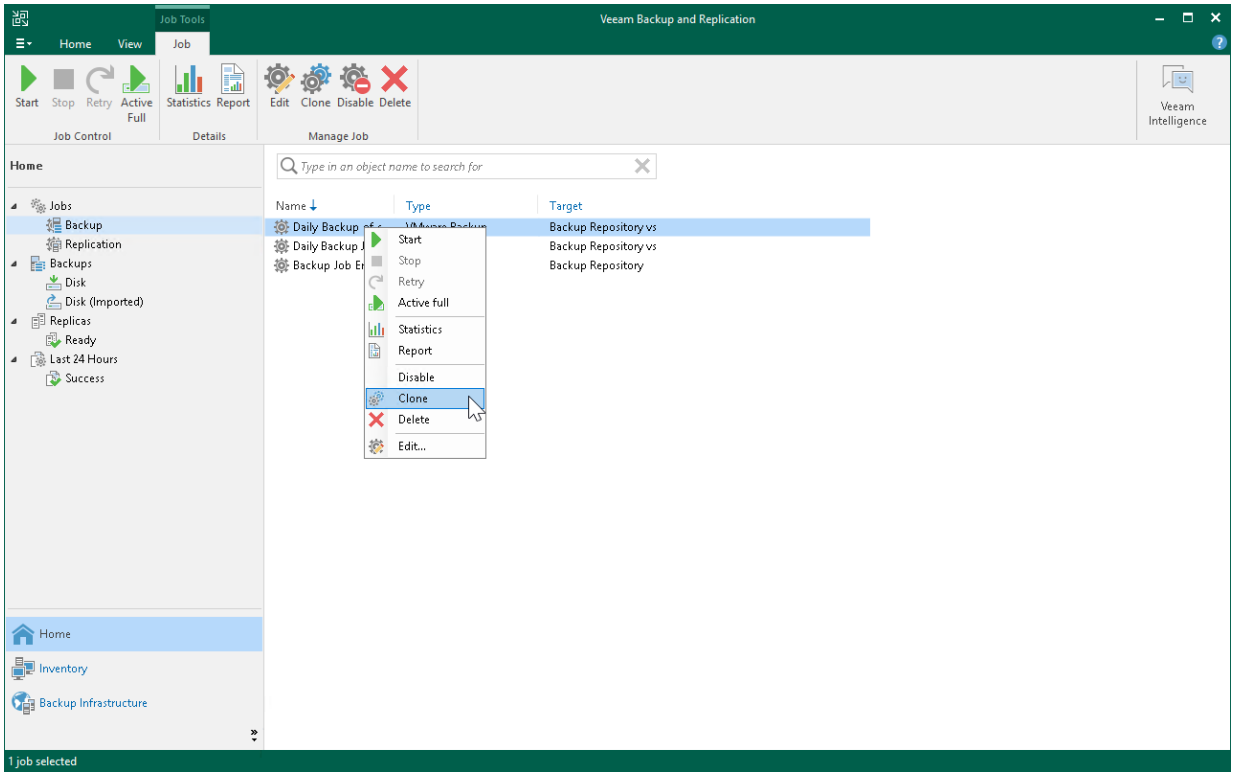

# Disabling and Deleting Jobs

You can temporarily disable scheduled jobs. The disabled job is paused for some period of time and is not run by the specified schedule. You can enable a disabled job at any time. You can also permanently delete a job from Veeam Backup & Replication and the configuration database.

## Disabling Job

To disable a job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Backup** node.
3. In the working area, select the job and select **Disable** on the ribbon or right-click the job and select **Disable**.

To enable a disabled job, select it in the list and click **Disable** on the ribbon once again.

> **NOTE**
>
> If the health check is enabled in the job settings, it will continue to run according to its own schedule, regardless of the job status.

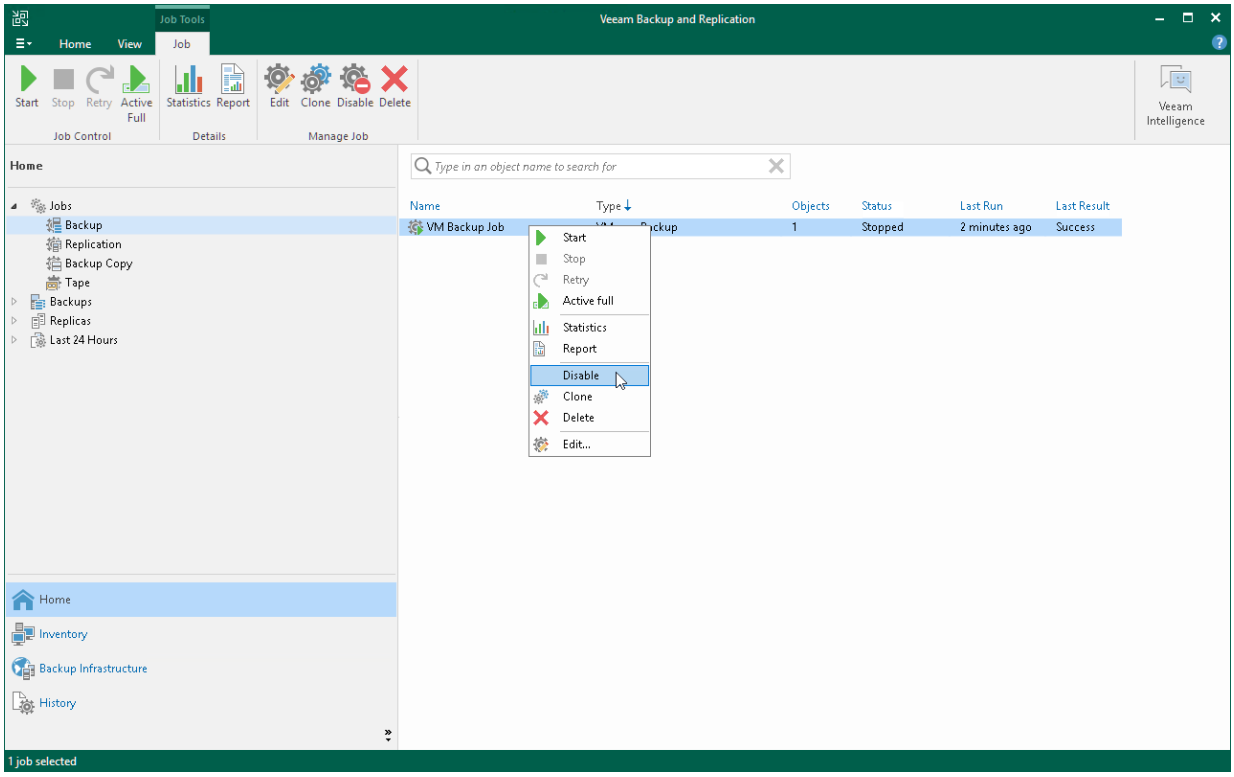


## Deleting Job

To delete a job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Backup** node.

3. In the working area, select the job and click **Delete** on the ribbon or right-click the job and select **Delete**.

After the job is deleted, the backups created by this job are displayed under the **Backups** > **Disk (Orphaned)** node. If the backup files created by this job were also stored in capacity tier or archive tier, they will also be displayed under the **Backups** > **Object Storage (Orphaned)** or **Backups** > **Archive (Orphaned)** nodes.

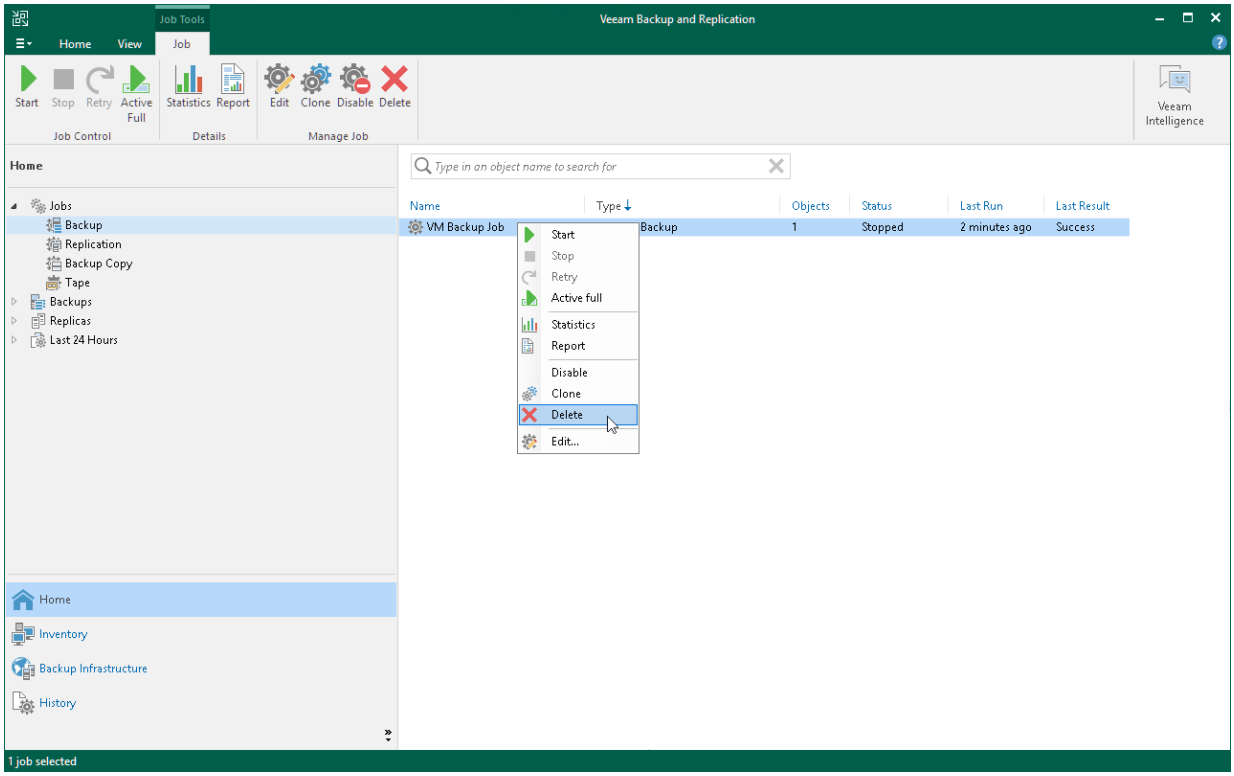

# Starting and Stopping Jobs

You can start a job manually if, for example, you want to create an additional restore point for a VM backup or replica without changing the job schedule. You can also stop a job if, for example, VM processing is about to take a long time, and you do not want the job to produce workload on the production environment during business hours.

## Starting Jobs

To start a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Start** on the ribbon or right-click the job and select **Start**.

## Stopping Jobs

You can stop a job in one of the following ways:

- Stop the job immediately. In this case, Veeam Backup & Replication will produce a new restore point only for those VMs that have already been processed when you stop the job.
- Stop the job after the current VM. In this case, Veeam Backup & Replication will produce a new restore point only for those VMs that have already been processed and for VMs that are being processed at the moment.

To stop a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Stop** on the ribbon or right-click the job and select **Stop**. In the displayed window, click **Immediately**.

To stop the job after the current VM:

1. Open the **Home** view.
2. In the inventory pane, click **Jobs**.

3. In the working area, right-click the job and select **Stop**. In the displayed window, click **Gracefully**.

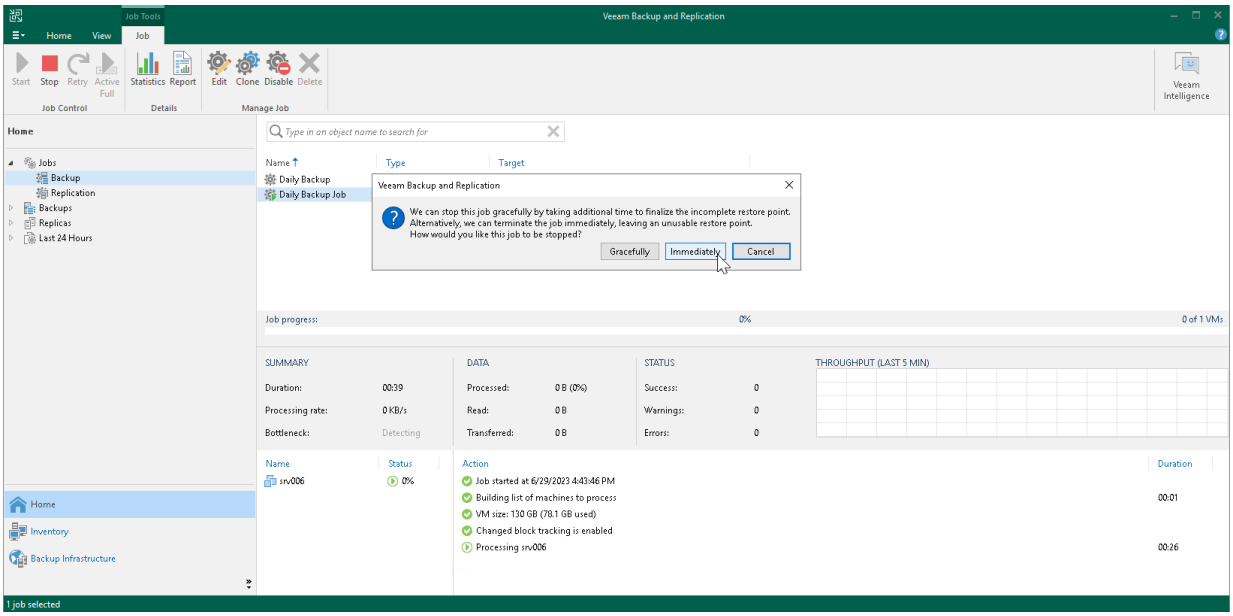

# Starting and Stopping Transaction Log Backup Jobs

If you create a backup job and instruct it to ship transaction logs, the backup job comprises 2 jobs:

1. A parent backup job creates an image-level backup of the VM on which the database runs. This job is named like a regular backup job, for example, *Daily Job*.
2. A transaction log backup job is responsible for shipping transaction logs to the backup repository. This job is named according to the following pattern:
   - For MS SQL: *<job_name> SQL Server Transaction Log Backup*. For example, *Daily Job SQL Server Transaction Log Backup*.
   - For Oracle: *<job_name> Oracle Redo Log Backup*. For example, *Daily Job Oracle Redo Log Backup*.

The transaction log backup job is created automatically by Veeam Backup & Replication if it detects that you have added to the backup job at least one Microsoft SQL Server or Oracle VM, enabled application-aware processing and instructed Veeam Backup & Replication to back up transaction logs periodically.

## Starting Transaction Log Backup Jobs

A parent backup job is manually started when you click **Start** on the toolbar or automatically by schedule. The transaction log backup job is initially started when you enable the schedule for the parent backup job. The transaction log backup works continuously in the background. A new session of the transaction log backup job starts every time the parent backup job is launched.

## Stopping Transaction Log Backup Jobs

You can stop transaction log processing in one of the following ways:

- Disable transaction log shipping
- Disable the parent backup job

If you want the backup job to create image-level backups of the VM but do not want it to ship transaction logs anymore, you can disable transaction log backup in the backup job settings.

To disable transaction log shipping:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Edit** on the ribbon or right-click the backup job and select **Edit**.
4. Pass to the **Guest Processing** step of the wizard and click **Applications**.
5. In the **Application-Aware Processing Options** window, select the VM and click **Edit**.
6. On the **SQL** or **Oracle** tab of the **VM Processing Settings** window, disable transaction log backup.
7. Click **Finish** to save the job settings.

If you do not want to create image-level backups of the VM and back up database transaction logs, you can disable scheduling for the parent backup job. Veeam Backup & Replication will instruct the transaction log backup job to complete log processing for all VMs added to the parent backup job and will switch the parent backup job to the non-scheduled mode. The parent backup job will no longer be started automatically by schedule — you will have to run it manually.

To disable scheduling for the parent backup job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Edit** on the ribbon. Alternatively, you can right-click the backup job and select **Edit**.
4. Pass to the **Schedule** step of the wizard and clear the **Run the job automatically** check box.
5. Click **Finish** to save the job settings.

Alternatively, you can disable the parent backup job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Disable** on the ribbon or right-click the job and select **Disable**.

To re-activate transaction log processing for all VMs in the parent backup job, select the job in the list and click **Disable** on the ribbon again.

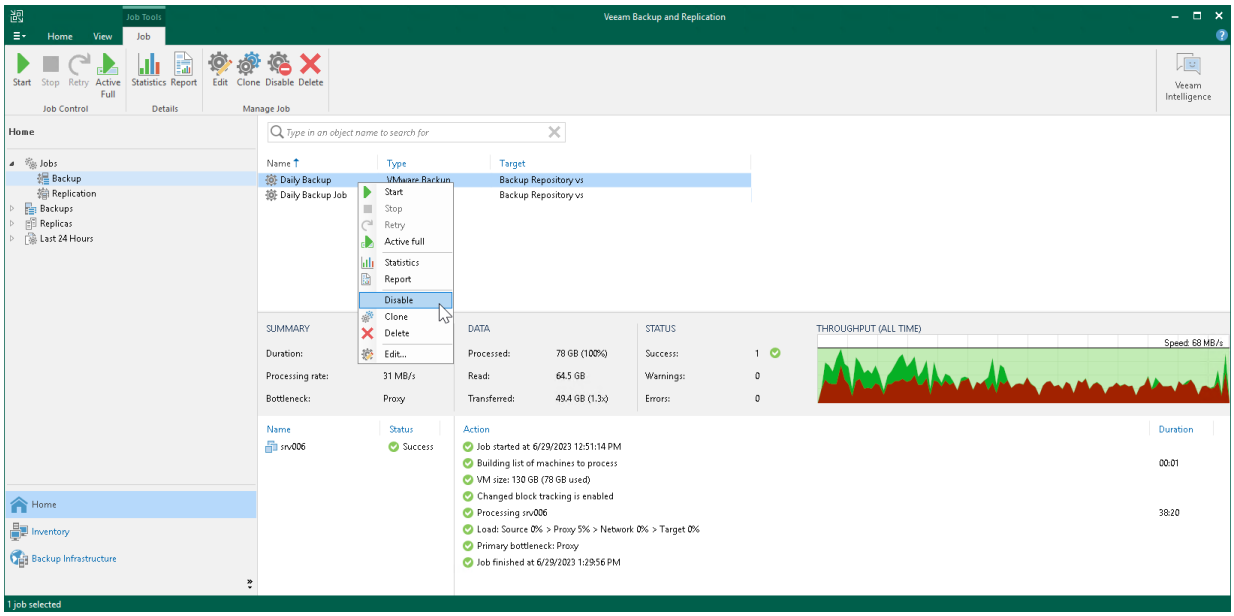

# Retrying Jobs

The retry option is necessary if a job fails and you want to retry this operation again. When you perform a retry, Veeam Backup & Replication restarts the operation only for the failed workloads added to the job and does not process VMs that have been processed successfully. As a result, the retry operation takes less time than running the job for all workloads.

## Retrying Job for All Failed Workloads

To perform retry for all workloads in a backup job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the necessary job and click **Retry** on the ribbon. Alternatively, you can right-click the job and select **Retry**.

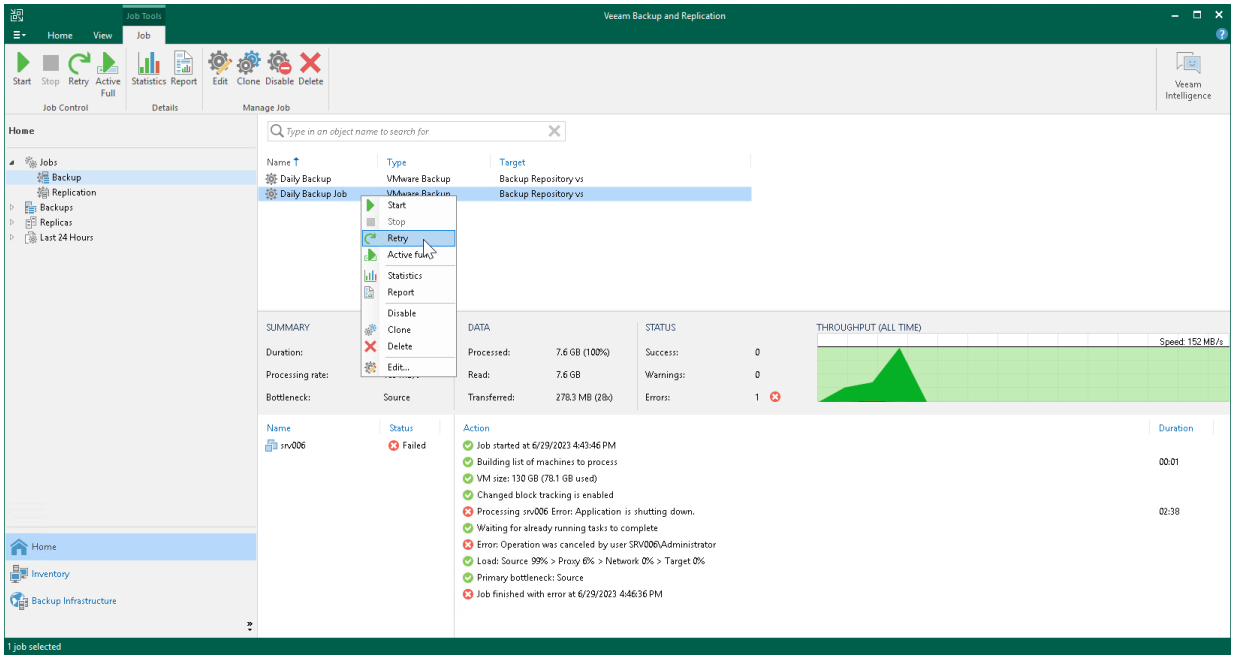


## Retrying Job for Individual Workloads

To perform retry for individual workloads:

1. Open real-time statistics or sessions results of the job.
2. Select failed workloads for which you want to perform retry.
3. Right-click one of the selected workloads and click **Retry**. Note that you will be able to launch retry for other workloads in the job only after retry finishes for the selected workloads.

> **IMPORTANT**
>
> You can perform retry for individual workloads only if their backups are per-machine with separate metadata files.

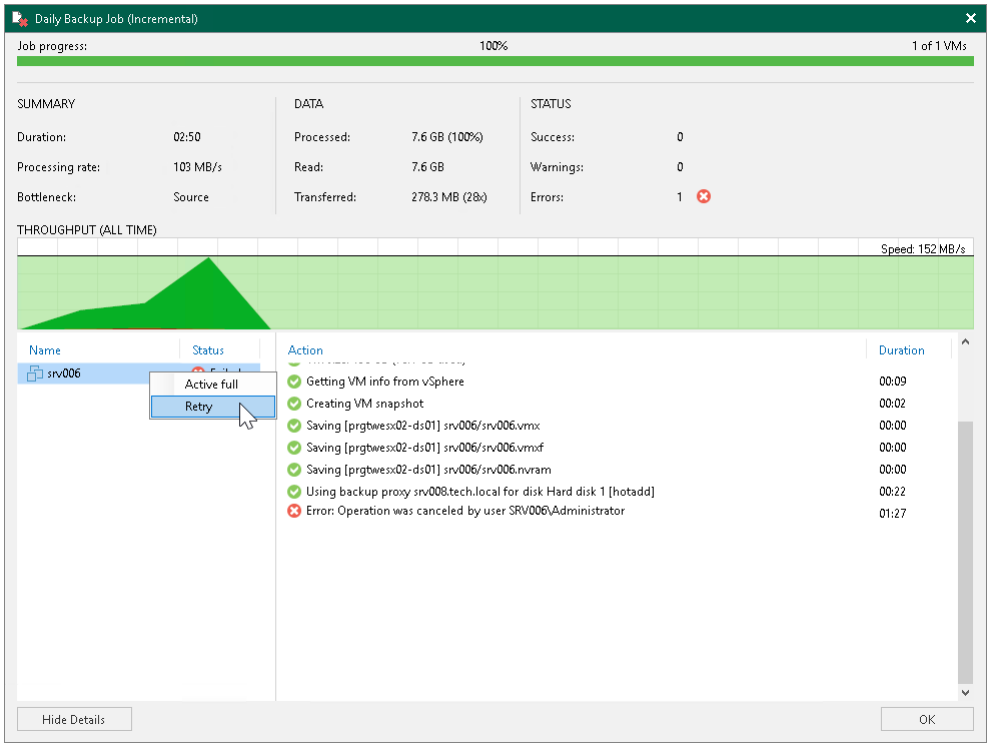
Daily Backup Job (Incremental)
Job progress:
100%
1 of 1 VMs
SUMMARY
Duration:
02:50
Processing rate:
103 MB/s
Bottleneck:
Source
DATA
Processed:
7.6 GB (100%)
Read:
7.6 GB
Transferred:
278.3 MB (28x)
STATUS
Success:
0
Warnings:
0
Errors:
1
THROUGHPUT (ALL TIME)
Speed: 152 MB/s
Name
Status
Action
Duration
srv006
Active full
Retry
Getting VM info from vSphere
00:09
Creating VM snapshot
00:02
Saving [prgtwesx02-ds01] srv006/srv006.vmx
00:00
Saving [prgtwesx02-ds01] srv006/srv006.vmxf
00:00
Saving [prgtwesx02-ds01] srv006/srv006.nvram
00:00
Using backup proxy srv008.tech.local for disk Hard disk 1 [hotadd]
00:22
Error: Operation was canceled by user SRV006\Administrator
01:27
Hide Details
OK

# Reconfiguring Jobs with Microsoft SQL Server VMs

In some situations, you may need to reconfigure a backup job that processes a Microsoft SQL Server VMs and ships transaction logs. For example, you may want to create a separate backup job to process the virtualized database and delete the VM running the database from the previously created job.

When you configure a new job, mind the restriction on transaction log shipping. By default, the new backup job that processes the VM will not ship transaction logs if transaction logs for this VM have been shipped for the last 7 days by another backup job on the same backup server.

You can overcome this restriction with registry values. For more information, contact Veeam Customer Support.

# Targeting Jobs to Another Repository

Veeam Backup & Replication offers you 2 ways to target a job to another repository:

- You can use the move to another repository operation described in the Moving Backups to Another Repository section. In this case, Veeam Backup & Replication moves all backups to a new repository and automatically reconfigures the job to target to the new location.
- You can edit job settings and select a new repository at the **Storage** step of the wizard. In this case, Veeam Backup & Replication will prompt you to choose whether to move all the existing backups to the new repository or start a new backup chain on the new repository. If you start a new backup chain, the previously created backups are detached from the job and put into the node with the **(Orphaned)** postfix.

For more information on how the move to another repository operation works and its considerations, see Backup Move.

# Reporting

When you run a job, Veeam Backup & Replication saves the job statistics and operation data to the configuration database. You can view real-time statistics for any performed job and generate reports with statistics data for any job or separate job session.

# Viewing Real-Time Statistics

To view real-time statistics for a job, do one of the following:

- Open the **Home** view. In the inventory pane, select one of the following nodes: **Jobs**, **Last 24 hours** or **Running**. In the working area, right-click the job and select **Statistics**.
- Open the **Home** view. In the inventory pane, select one of the following nodes: **Jobs**, **Last 24 hours** or **Running**. In the working area, double-click the running job.

The real-time statistics provide detailed data on job sessions: job progress, duration, processing rate, performance bottlenecks, amount of processed, read and transferred data, and details of the session performance, for example, warnings and errors that have occurred in the process of operation.

In addition to overall job statistics, real-time statistics provide information on each object processed with the job. To view the processing progress for a specific object, select it in the list on the left.

> **TIP**
>
> Consider the following:
>
> - To collapse and expand the real-time statistics window, use **Hide Details** and **Show Details** buttons at the bottom left corner of the window.
> - To switch between the job sessions backward and forward, use left and right arrow keys on the keyboard.

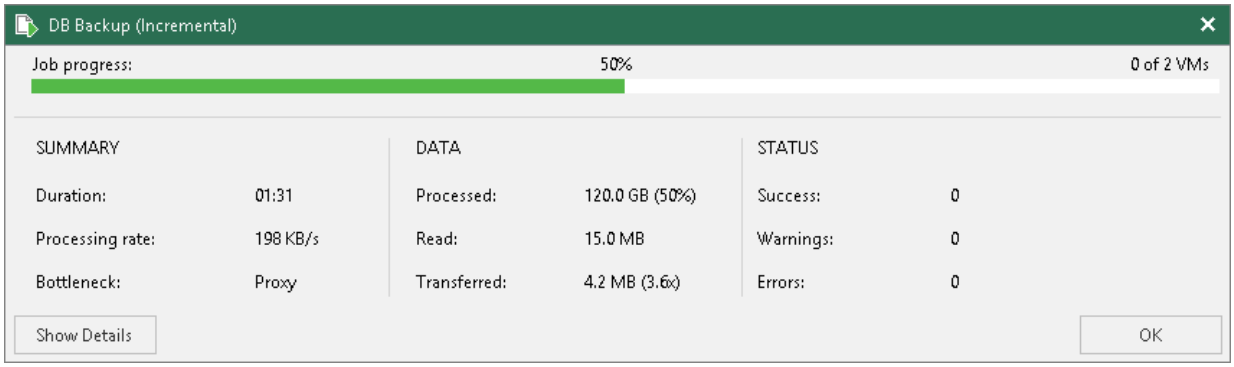


# Statistics Counters

Veeam Backup & Replication displays job statistics for the following counters:

- The **Job progress** bar shows the percentage of job completion.
- The **Summary** box shows general information about the job:
  - **Duration** – time from the job start till the current moment or job end.
  - **Processing rate** – average speed of VM data processing. This counter is a ratio between the amount of data that has actually been read and the time it took to process the data. Note that only the data transfer time is used for the calculation, and the job's runtime is irrelevant.
  - **Bottleneck** – a bottleneck in the data transmission process. To learn about job bottlenecks, see Performance Bottlenecks.
- The **Data** box shows information about processed VM data:
  - **Processed** – total size of all VM disks processed by the job.

- **Read** — the amount of data read from the datastore by the source-side Data Mover prior to applying compression and deduplication. For incremental job runs, the value of this counter is typically lower than the value of the **Processed** counter. Veeam Backup & Replication reads only data blocks that have changed since the last job session, processes and copies these data blocks to the target.
- **Transferred** — the amount of data transferred from the source-side Veeam Data Mover to the target-side Veeam Data Mover after applying compression and deduplication. This counter does not directly indicate the size of the resulting files. Depending on the backup infrastructure and job settings, Veeam Backup & Replication can perform additional activities with data: deduplicate data, decompress data prior to writing the file to disk, and so on. The activities can impact the size of the resulting file.

- The **Status** box shows information about the job results. This box informs how many tasks have been completed with the *Success*, *Warning* and *Error* statuses (1 task per 1 VM):
  - *Success* — the task is completed successfully.
  - *Warning* — the task is completed with minor errors. Depending on the nature of the errors, the backup data may not be consistent.
  - *Error* — the task is not completed due to a blocking error.
- The pane in the lower left corner shows a list of objects processed by the job.
- The pane in the lower right corner shows a list of operations performed during the job. To see a list of operations for a specific object included in the job, click the object in the pane on the left. To see a list of operations for the whole job, click anywhere on the blank area in the left pane.

# Colored Graph

To visualize the data transfer process, Veeam Backup & Replication displays a colored graph in the real-time statistics window:

- The green area defines the amount of data read from the source.
- The brown area defines the amount of data transported to the target.
- The horizontal line defines the current data processing speed.

If the job session is still being performed, you can click the graph to view the data rate for the last 5 minutes or the whole processing period. If the job session has already ended, the graph will display information for the whole processing period only.

The colored graph is displayed only for the currently running job session or the latest job session. If you open real-time statistics for past sessions other than the latest one, the colored graph will not be displayed.

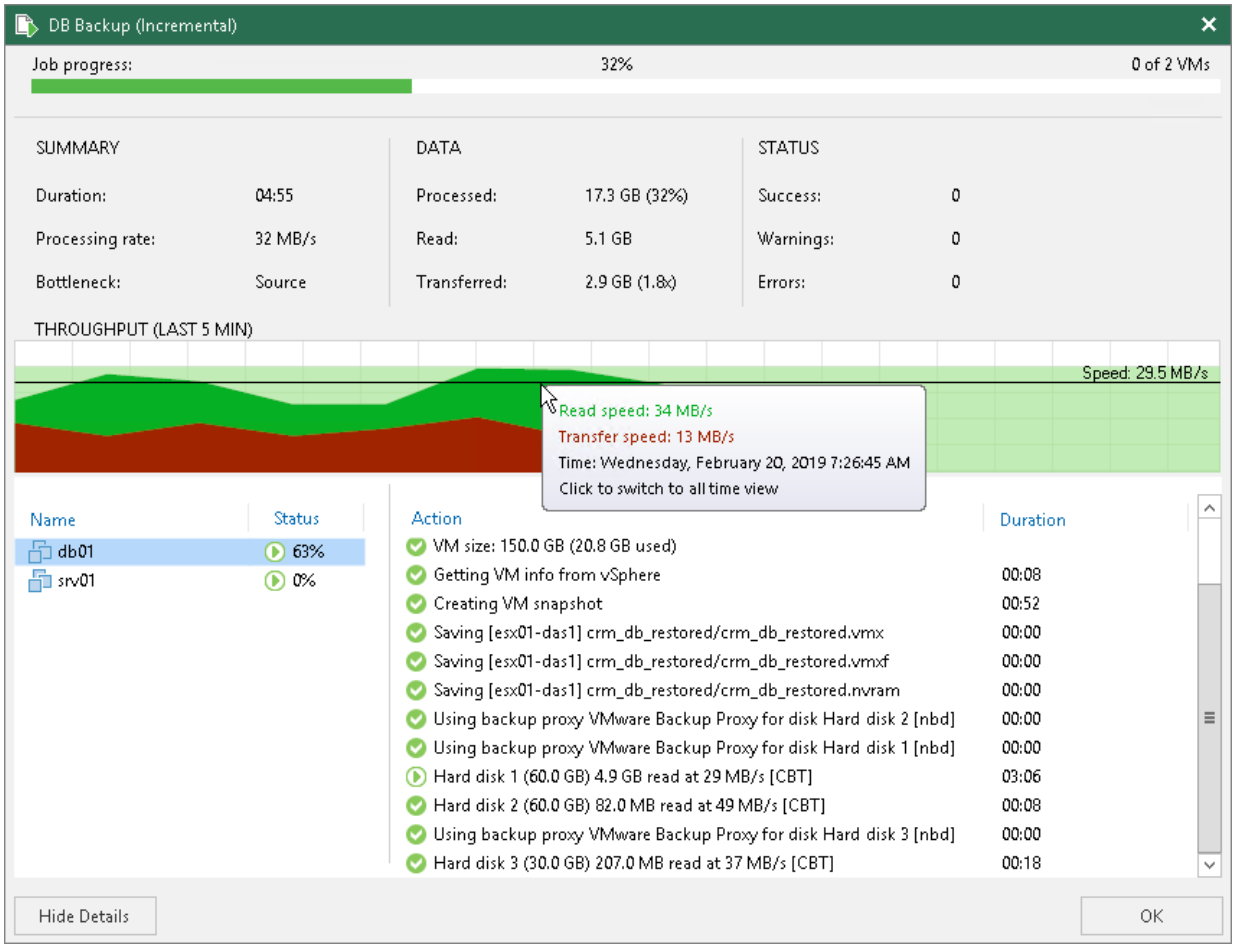

# Viewing History Statistics

The **History** view displays statistics for operations performed with Veeam Backup & Replication: backup and restore jobs, system operations, retention jobs and malware detection sessions. The **History** view shows data for all sessions stored in the configuration database.

To view the history of jobs and operations performed by Veeam Backup & Replication:

1. Click the arrow icon (») at the bottom of the navigation pane.
2. Click **History** in the list.
3. Select one of the following nodes: **Jobs**, **Restore**, **System** or **Malware Detection**.

The **History** view provides overall session statistics: name, status, start and end time, who initiated the session, and session type. To view detailed data on each session, double-click the session in the working area or right-click it and select **Statistics**. For backup and backup copy jobs, you can switch between **Compact View** and **Full View** modes within the **View** tab on the ribbon to hide or show job details.

> **TIP**
>
> Consider the following:
>
> - The **History** view can be shown as an icon if it does not fit into the pane. To show the **History** view in the full size, drag and drop the upper border of the pane.
> - You can configure the period for which Veeam Backup & Replication shows sessions as described in Specifying Session History Settings.

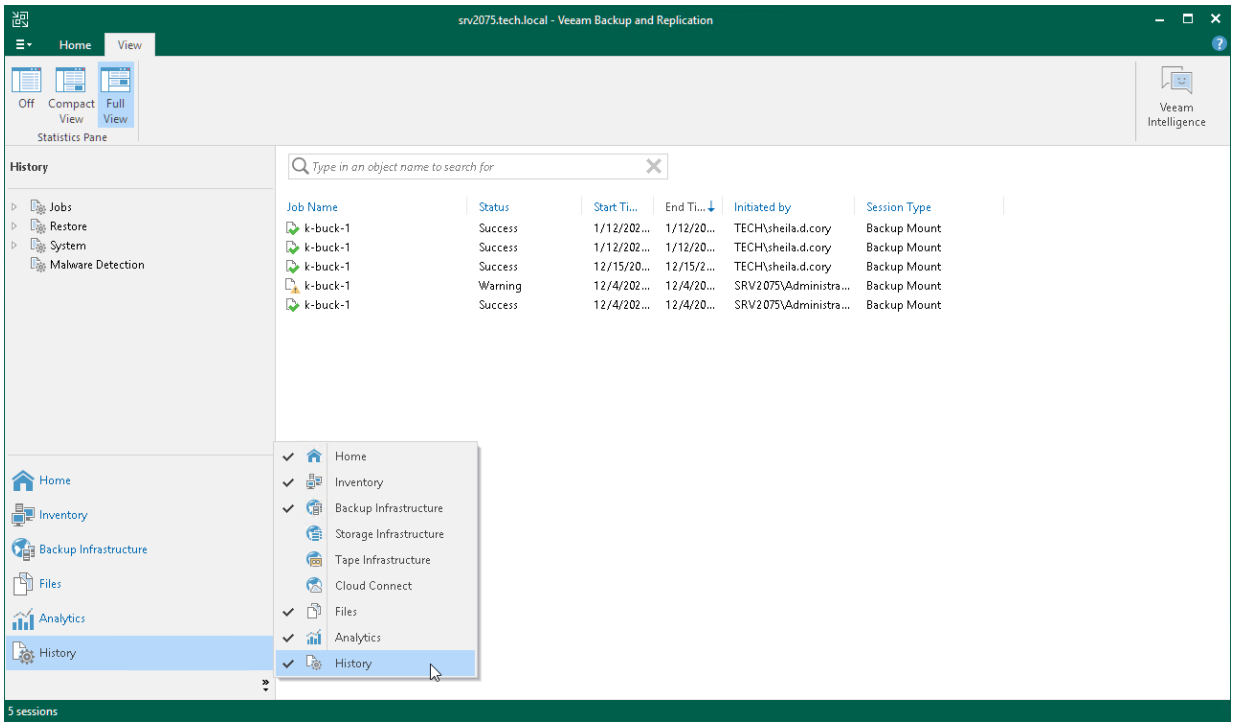

# Viewing Job Session Results

You can view detailed statistics on every job session.

To view statistics for a selected job session, do either of the following:

- Open the Viewing History Statistics view. In the inventory pane, select **Jobs**. In the working area, double-click the necessary job session.
- Open the Viewing History Statistics view. In the inventory pane, select **Jobs**. In the working area, right-click the necessary job session and select **Statistics**.

> **TIP**
>
> To switch between past job sessions, use left and right arrow keys on the keyboard.

## Statistics Counters

Veeam Backup & Replication displays backup job statistics for the following counters:

- The **Job progress** bar shows the percentage of job completion.
- The **Summary** box shows general information about the job:
  - **Duration** — duration of the job session.
  - **Processing rate** — average speed of VM data processing. This counter is a ratio between the amount of data that has actually been read and the time it took to process the data. Note that only the data transfer time is used for the calculation and the job's runtime is irrelevant.
  - **Bottleneck** — a bottleneck in the data transmission process. To learn more about bottlenecks, see Performance Bottlenecks.
- The **Data** box shows information about processed VM data:
  - **Processed** — total size of all VM disks processed by the job.
  - **Read** — the amount of data read from the datastore before compression and deduplication. For incremental job runs, the value of this counter is typically lower than the value of the **Processed** counter. Veeam Backup & Replication reads only data blocks that have changed since the last job session, processes and copies these data blocks to the target.
  - **Transferred** — the amount of data transferred from the source-side Veeam Data Mover to the target-side Veeam Data Mover after applying compression and deduplication. This counter does not directly indicate the size of the resulting files. Depending on the backup infrastructure and job settings, Veeam Backup & Replication can perform additional activities with data: deduplicate data, decompress data prior to writing the file to disk, and so on. The activities can impact the size of the resulting file.
- The **Status** box shows information about the job results. This box informs how many tasks have been completed with the *Success*, *Warning* and *Error* statuses (1 task per 1 VM):
  - *Success* — the task is completed successfully.
  - *Warning* — the task is completed with some minor errors. Depending on the nature of the errors, the backup data may not be consistent.
  - *Error* — the task is not completed due to a blocking error.

- The pane in the lower-left corner shows a list of objects processed by the job.
- The pane in the lower-right corner shows a list of operations performed during the session. To see a list of operations for a specific object, click the object in the pane on the left. To see a list of operations for the whole copy session, click anywhere on the blank area in the left pane.

# Colored Graph

To visualize the data transfer process, Veeam Backup & Replication displays a colored graph in the real-time statistics window:

- The green area defines the amount of data read from the source.
- The brown area defines the amount of data transported to the target.
- The horizontal line defines the current data processing speed.

If the job session is still being performed, you can click the graph to view the data rate for the last 5 minutes or the whole processing period. If the job session has already ended, the graph will display information for the whole processing period only.

The colored graph is displayed only for the currently running job session or the latest job session. If you open real-time statistics for past sessions other than the latest one, the colored graph will not be displayed.

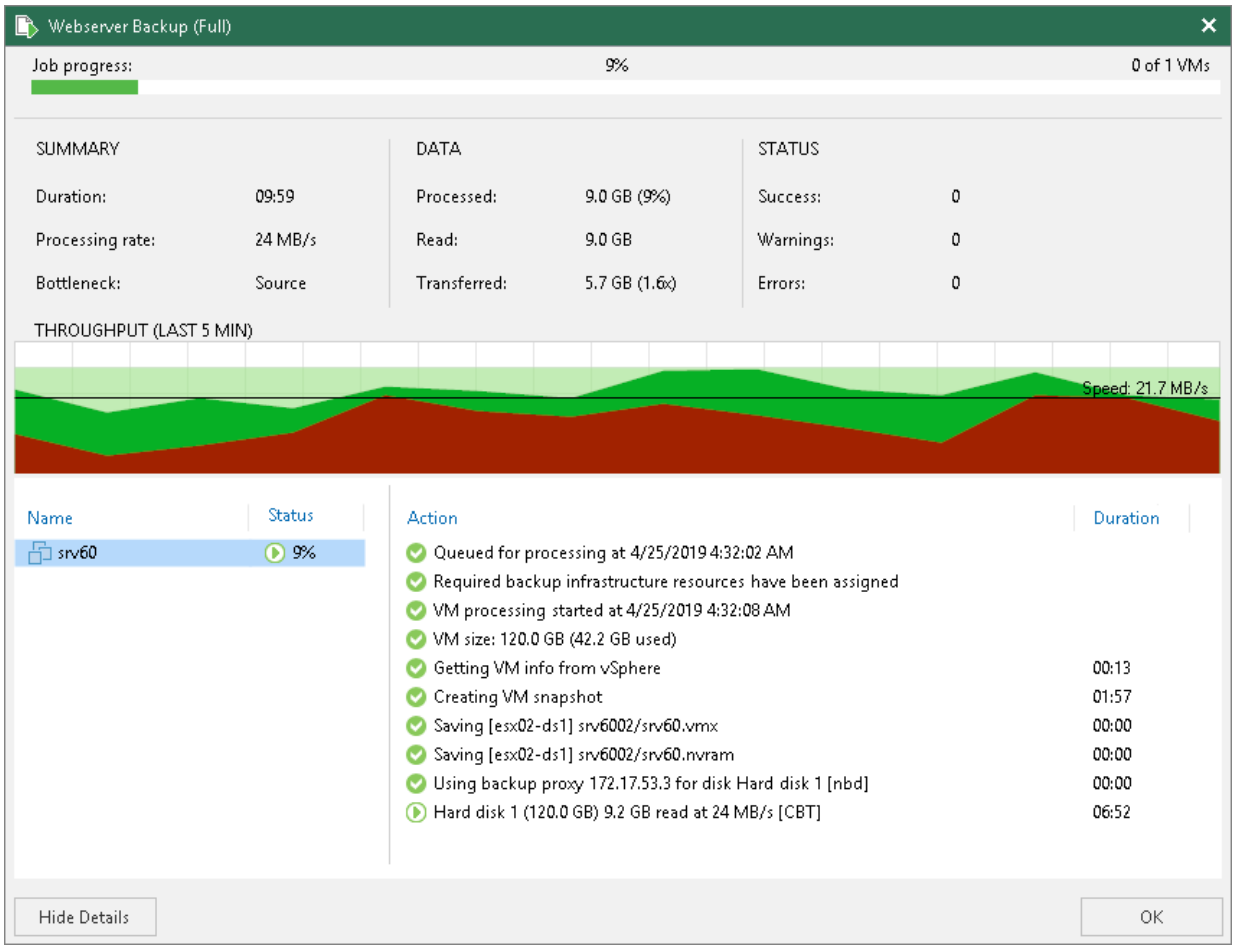

# Viewing Job and Job Session Reports

You can generate reports with details about all sessions of a job or a single session only.

## Job Report

The job report contains data on all sessions initiated for a specific job, that is, job history. The report shows data for all sessions stored in the configuration database.

To generate a job report:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the necessary job and click **Report** on the ribbon. You can also right-click the job and select **Report**.

For more information on counters in the report, see Report Counters.

> **TIP**
>
> Generated reports are stored in the `C:\Users\<username>\AppData\Local\Temp` folder.

## Session Report

To generate a report for a single session:

1. Open the Viewing History Statistics view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the necessary session and click **Report** on the ribbon. You can also right-click the necessary session and select **Report**.

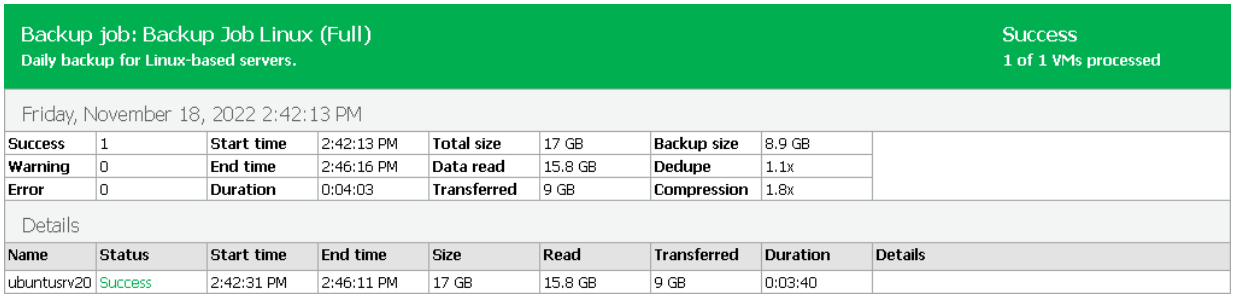

Backup job: Backup Job Linux (Full)
Daily backup for Linux-based servers.

Success
1 of 1 VMs processed

Friday, November 18, 2022 2:42:13 PM

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| Success | 1 | Start time | 2:42:13 PM | Total size | 17 GB | Backup size | 8.9 GB |
| Warning | 0 | End time | 2:46:16 PM | Data read | 15.8 GB | Dedupe | 1.1x |
| Error | 0 | Duration | 0:04:03 | Transferred | 9 GB | Compression | 1.8x |

Details

| Name | Status | Start time | End time | Size | Read | Transferred | Duration | Details |
|---|---|---|---|---|---|---|---|---|
| ubuntusrv20 | Success | 2:42:31 PM | 2:46:11 PM | 17 GB | 15.8 GB | 9 GB | 0:03:40 | |

## Report Counters

Veeam Backup & Replication displays the following counters in reports:

- The **Success**, **Warning** and **Error** counters show how many workloads were processed with the *Success*, *Warning* and *Error* statuses.
- The **Start time** and **End time** counters show when the job started and completed.
- The **Duration** counter shows the time from the job start till the current moment or job end.

- The **Total size** counter shows the provisioned size (the maximum configured size) of all workload disks in the job.
- The **Data read** counter shows the amount of data read from the datastore before compression and deduplication. The value of this counter is typically lower than the value of the **Total size** counter. Veeam Backup & Replication reads only data blocks that have changed since the last job session, processes and copies these data blocks to the target.
- The **Transferred** counter shows the amount of data transferred from the source side to the target side after applying compression and source-side deduplication. This counter does not directly indicate the size of the resulting files. Depending on the backup infrastructure and job settings, Veeam Backup & Replication can perform additional activities with data: perform target-side deduplication, decompress data prior to writing the file to disk, and so on. The activities can impact the size of the resulting file.
- The **Backup size** counter shows the resulting backup file size.
- The **Dedupe** counter shows the deduplication level.
- The **Compression** counter shows the compression level.

In the **Details** section, you can see similar counters for each workload the job processed.

> **NOTE**
>
> A synthetic full backup that is part of an incremental backup session is synthesized directly on the backup repository. Therefore, the **Transferred** and **Backup size** counters do not include data processed during the synthetic backup. They cover only data processed during the incremental run. Besides, the report header for such a backup does not have the **(Full)** mark as Veeam Backup & Replication considers it incremental. For more information, see the Synthetic Full Backup section.

# Replication

Replication is a technology that helps you protect mission-critical VMware virtual machines. When you replicate a VM, Veeam Backup & Replication creates an exact copy of the VM in the native VMware vSphere format on the target host. Veeam Backup & Replication maintains this copy in sync with the source VM. Replication provides minimum recovery time objective (RTO) in case a disaster strikes because VM replicas are in a ready-to-start state.

We recommend you to replicate VMs for which recovery point objective (RPO) of hours is required. If you need RPO of seconds, consider continuous data protection (CDP).

## Data Replication

To replicate VMs, Veeam Backup & Replication leverages VMware vSphere snapshot capabilities. During replication, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot can be thought of as a point-in-time copy of a VM that includes virtual disks, system state, configuration and so on. Veeam Backup & Replication uses the snapshot as a source of data for replication.

During the first replication cycle, Veeam Backup & Replication copies data of the source VM running on the source host, and creates its full replica on the target host. Unlike backup files, replica virtual disks are stored decompressed in their native format. All subsequent replication cycles are incremental.
Veeam Backup & Replication copies only those data blocks that have changed since the last replication job session. To keep track of changed data blocks, Veeam Backup & Replication uses different approaches. For more information, see Changed Block Tracking.

Veeam Backup & Replication lets you perform on-site replication for high availability scenarios and remote (off-site) replication for disaster recovery scenarios. To facilitate replication over the WAN or slow connections, Veeam Backup & Replication optimizes traffic transmission. It filters out unnecessary data blocks such as duplicate data blocks, zero data blocks, blocks of swap files and blocks of excluded VM guest OS files, and compresses replica traffic. Veeam Backup & Replication also allows you to use WAN accelerators and apply network throttling rules to prevent replication jobs from consuming the entire network bandwidth.

To replicate a VM, you need to configure required backup infrastructure components and create a replication job.

## Recovery

If a disaster strikes and the production VM stops working properly, you can fail over to its replica.

When you fail over to a replica, the replica takes over the role of the source VM. After your source VM is repaired, you can fail back to it and transfer all changes that occurred to replica to the source VM. If your source VM cannot be repaired, you can perform permanent failover, that is, permanently switch from the source VM to the VM replica and use this replica as the source VM. For more information, see Failover and Failback for Replication.

# Considerations and Limitations

Replication has the following requirements and limitations:

- Due to VMware vSphere limitations, if you change the size of VM disks on the source VM, Veeam Backup & Replication deletes all available restore points (represented as VM snapshots) on the VM replica during the next replication job session. For more information, see this VMware KB article.
- Due to a change in ESXi 6.0 Update 1, replication and Quick Migration to vVol datastores are not possible with either Veeam or native VMware vSphere replication.
- The target host must support the hardware version of the VM that you plan to replicate. For the compatibility table, see this VMware KB article.
- If you assign the role of a backup proxy to a VM, you should not add this VM to the list of processed VMs in a job that uses this backup proxy. Such configuration may result in degraded job performance. Veeam Backup & Replication will assign this backup proxy to process other VMs in the job first, and processing of this VM itself will be put on hold. Veeam Backup & Replication will report the following message in the job statistics: *VM is a backup proxy, waiting for it to stop processing tasks*. The job will start processing this VM only after the backup proxy deployed on the VM finishes its tasks.
- Replication of VM templates is not supported.
- If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.
- How Veeam Backup & Replication requests to assign MAC addresses to replicas depends on the MAC allocation mode configured on the vCenter Server:
  - Automatic allocation. If replica is located on the same ESXi host or vCenter Server, new MAC address is assigned. If replica is located on another ESXi host or vCenter Server, MAC address is preserved.
  - Manual allocation. MAC address is preserved.
- Due to Microsoft limitations, you cannot use Microsoft Entra ID (formerly Azure Active Directory) credentials to perform application-aware processing on VMs running Microsoft Windows 10 (or later).
- If a job is unable to complete within 21 days period, it will be stopped with the *Failed* status.

# Backup Infrastructure for Replication

Veeam Backup & Replication uses the following components for the replication process:

- Backup server
- Source and target hosts
- Backup proxies
- Backup repository
- [Optional] WAN accelerators

The amount and placement of these components depend on a replication scenario you use. For more information, see Replication Scenarios.

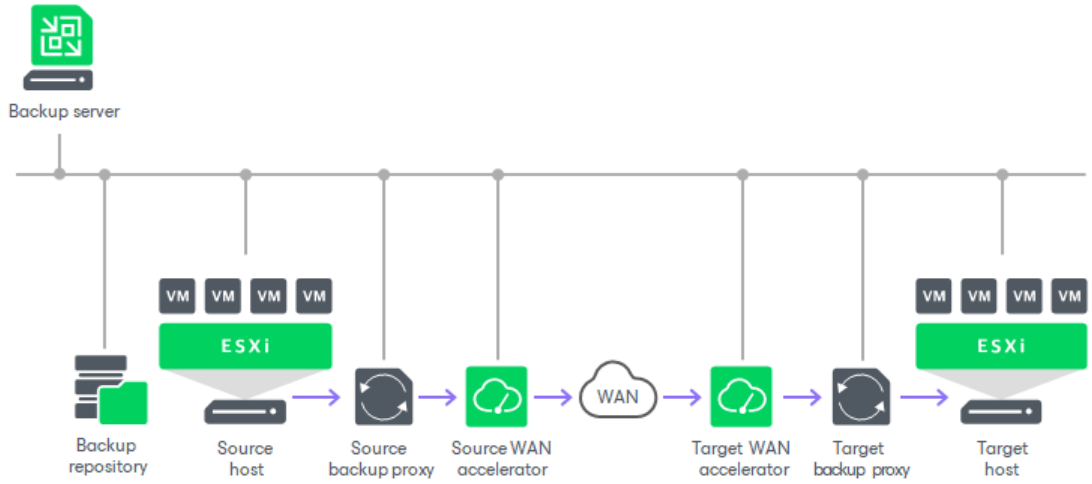


## Backup Server

During the replication process, the backup server coordinates replication tasks, controls resource allocation and replica job scheduling. The backup server runs the Veeam Backup Service and Veeam Broker Service that coordinate and interact with the virtual infrastructure.

For more information on the backup server, see Backup Server.

## Source and Target Hosts

The source host is the host where VMs that you plan to replicate are located. The target host is the host where VM replicas will be created and maintained in the ready-to-start state.

The role of a target host can be assigned to a standalone ESXi host or ESXi cluster. If you assign a cluster or vCenter Server as a target, the replication process becomes more sustainable — the replication process will not fail if there is at least one available host in the cluster.

To replicate data from and to hosts, they must be first added to the Veeam Backup & Replication infrastructure. For more information on how to add hosts, see the Adding VMware vSphere Servers section.

# Backup Proxies

A backup proxy collects, transforms and transports VM data during the replication process. For more information on backup proxies, requirements and limitations for them, see Backup Proxy.

For replication, you can deploy backup proxies on the following machines:

- **Physical machines**. In this case, Veeam Backup & Replication uses the Network transport mode to populate replica disk files. For more information, see Network Mode.
- **VMs**. The virtual backup proxy must be registered on an ESXi host that has a direct connection to the target datastore. In this case, the backup proxy will be able to use the Virtual appliance transport mode to populate replica disk files. This results in increased writing speed and fail-safe replication to ESXi targets. For more information, see Virtual Appliance (HotAdd). Note that if the Virtual appliance transport mode cannot be used, the backup proxy can fail over to the network mode if you configure it while adding a backup proxy.

We recommend you to use at least two backup proxies to ensure that the job will be performed if one of backup proxies fails or loses its connectivity to the source datastore. For more information on how assign the role of a backup proxy, see Adding VMware Backup Proxies. For more information on how to assign proxies to a replication job, see Specify Data Transfer Settings.

# Backup Repository

The backup repository stores replica metadata that contains information on the read data blocks (such as checksums and digests). Metadata is required when Veeam Backup & Replication performs incremental replication or if you fail back from a VM replica to the source VM in the original location using quick rollback. Veeam Backup & Replication uses metadata to quickly detect changed data blocks between two replica states.

The backup repository must have access to the source backup proxy. We recommend to deploy the backup repository as close to the source backup proxy as possible.

For more information, see Backup Repository.

# WAN Accelerators

WAN accelerators are optional components in the backup infrastructure. You can use WAN accelerators if you replicate VMs over a slow connection or over WAN.

In the replication process, WAN accelerators are responsible for global data caching. To use WAN acceleration, you must deploy two WAN accelerators in the following way:

- The source WAN accelerator must be deployed in the source side, close to the backup proxy.
- The target WAN accelerator must be deployed in the target side, close to the backup proxy.

For more information, see WAN Acceleration.

# Replication Scenarios

Veeam Backup & Replication supports a number of replication scenarios that depend on the location of the target host:

- **On-site replication**. If the target host is located in the same site as the source host, use the on-site replication scenario.
- **Off-site replication**. If the target and source hosts are located in different sites, use the off-site replication scenario.

Depending on the scenario you are planning to use, different backup infrastructure components will be involved in the replication process. For more information on the components, see Backup Infrastructure for Replication.

## On-Site Replication

On-site replication requires the following backup infrastructure components:

- Source and target hosts.
- Backup proxy. In the on-site replication scenario, the source Veeam Data Mover and target Veeam Data Mover are started on the same backup proxy. The backup proxy must have access to the backup server, source host, target host and backup repository that stores replica metadata.
- Backup repository for storing replica metadata.

In the on-site replication scenario, Veeam Backup & Replication does not perform data compression. Replication traffic is transferred decompressed between the two Veeam Data Mover started on the same backup proxy.

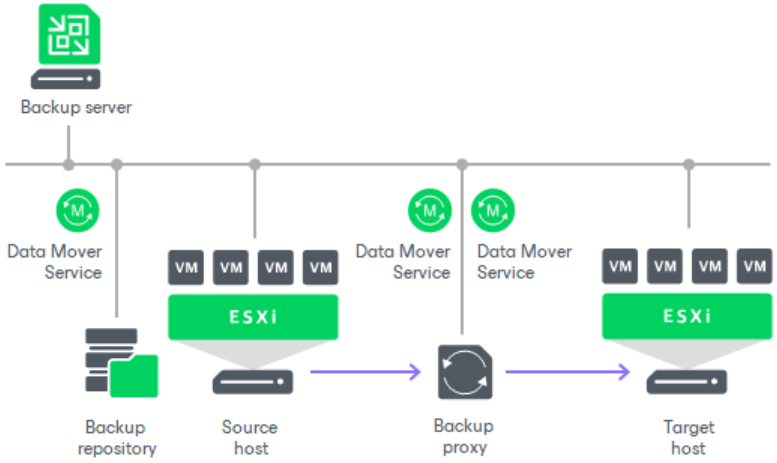


## Off-Site Replication

For the off-site replication scenario, you must use at least two backup proxies:

- One backup proxy in the production site, closer to the source host.
- One backup proxy in the remote DR site, closer to the target host.

Depending on the network speed between the production and DR sites, off-site replication can run in two ways:

- **Directly over backup proxies**. Transfer data directly over backup proxies if the connection between two sites is fast and stable.

- **Over a pair of WAN accelerators**. Use WAN accelerators if the connection between the sites is slow.

> **NOTE**
>
> With off-site replication, you can also use technologies that help reduce the amount of replication traffic and streamline replica configuration: replica seeding and mapping, network mapping and re-IP rules.

## Replication Directly over Backup Proxies

Off-site replication directly over backup proxies requires the following backup infrastructure components:

- Source and target hosts.
- At least one backup proxy in the production site. The backup proxy must have access to the backup server, source host, backup proxy in the target site and backup repository that stores replica metadata.
- At least one backup proxy in the DR site. The backup proxy must have access to the backup server, target host and backup proxy in the production site.
- Backup repository for storing replica metadata.

In the off-site replication scenario, Veeam Backup & Replication uses data compression. Veeam Data Mover on the source backup proxy compresses VM data blocks and sends them to the target backup proxy in the compressed format. Veeam Data Mover on the target backup proxy decompresses VM data and stores it to a datastore in a native VMware vSphere format.

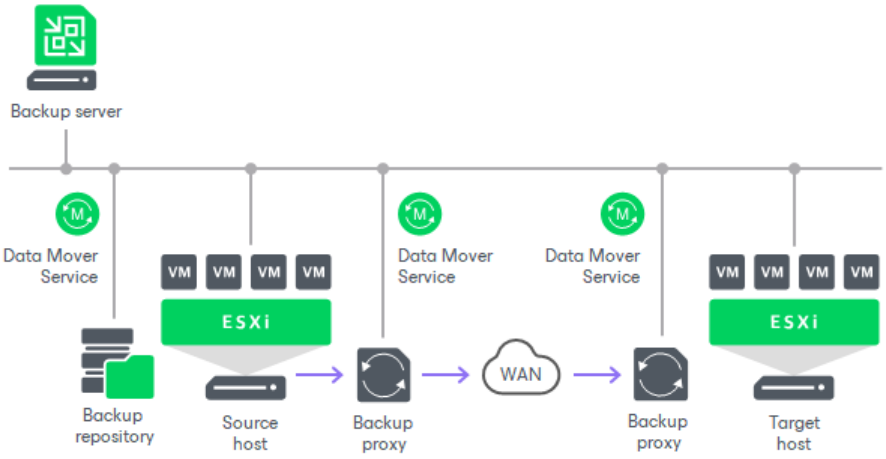


## Replication over WAN Accelerators

Off-site replication over WAN accelerators requires the following backup infrastructure components:

- Source and target hosts.
- A pair of WAN accelerators at each end of the WAN link:
  - Source WAN accelerator in the production site. The source WAN accelerator must have access to the backup server, source backup proxy and target WAN accelerator.
  - Target WAN accelerator in the DR site. The target WAN accelerator must have access to the backup server, source WAN accelerator and target backup proxy.
- At least one backup proxy in the production site. The backup proxy must have access to the backup server, source host, source WAN accelerator and backup repository that stores replica metadata.

- At least one backup proxy in the DR site. The backup proxy must have access to the backup server, target host and target WAN accelerator.
- Backup repository for storing replica metadata. The backup repository must be located in the production site, closer to the source backup proxy, and must have access to it.

In the off-site replication scenario using WAN accelerators, Veeam Backup & Replication compresses VM data. VM data blocks are compressed on the source WAN accelerator, transported to the DR site in the compressed format and decompressed on the target WAN accelerator.

# How Replication Works

Veeam Backup & Replication performs VM replication in the following way:

1. When a new replication job session starts, Veeam Backup & Replication reads job settings from the configuration database, checks what backup infrastructure resources are available and prepares the backup infrastructure for replication.
2. Veeam Backup & Replication queries information about VMs and virtualization hosts from the vCenter Server.
3. If application-aware image processing is enabled for the job, Veeam Backup & Replication connects to VM guest OSes, deploys non-persistent runtime components or uses (if necessary, deploys) persistent agent components on VM guest OSes and performs in-guest processing tasks.
4. Veeam Backup & Replication requests vCenter Server or ESXi host to create a VM snapshot. VM disks are put to the read-only state, and every virtual disk receives a delta file. All changes that the user makes to the VM during replication are written to delta files.
5. The source backup proxy reads the VM data from the read-only VM disk and copies it.

   During incremental job sessions, the source proxy uses changed block tracking (CBT) to retrieve only those data blocks that have changed since the previous job session. If CBT is not available, the source proxy interacts with the backup repository to obtain replica metadata, and uses this metadata to detect blocks that have changed since the previous job session.

   While copying VM data, the source proxy performs additional processing. It filters out zero data blocks, blocks of swap files and blocks of excluded VM guest OS files.
6. The source proxy compresses VM data and transports it to the target proxy.
7. The target proxy decompresses VM data and writes the result to the destination datastore.
8. After the backup proxy finishes reading VM data, Veeam Backup & Replication requests the vCenter Server or ESXi host to commit the VM snapshot and delta disks into the VM.

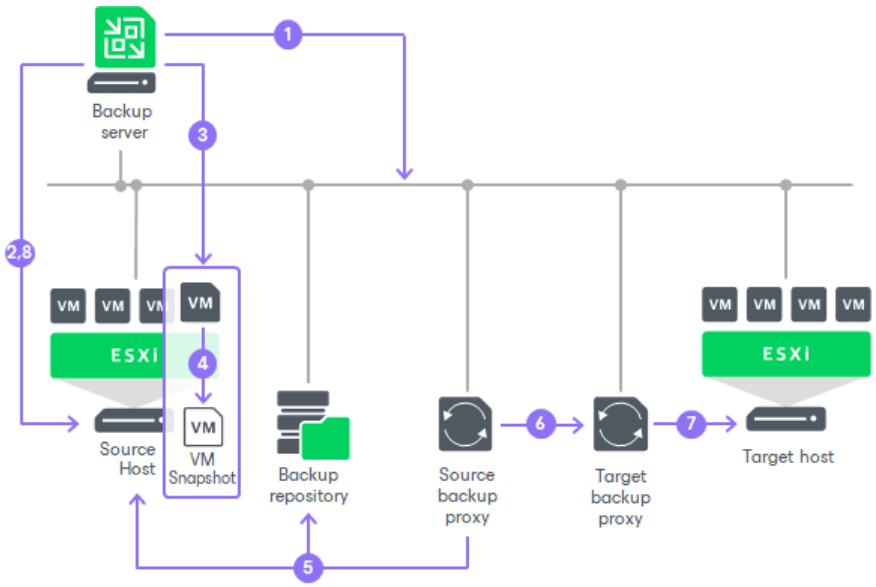

# Replication Chain

For every VM replica, Veeam Backup & Replication creates a replication chain that consists of restore points. Veeam Backup & Replication utilizes VMware ESXi snapshot capabilities to create and manage replica restore points.

Veeam Backup & Replication creates a restore point during every replication job session. During the first replication job session, Veeam Backup & Replication creates a copy of the source VM on the target host. During every subsequent replication job session, it adds a new snapshot to the replication chain for the VM replica. Blocks of data that have changed since the last job run are written to the snapshot delta file, and the snapshot delta file acts as a restore point. You can view a replication chain created for a VM using VMware vSphere client.

You can specify how many restore points you want to store in the replication chain. For this, configure retention policy settings for the replication job. For more information, see Specify Replication Job Settings.

VM replica restore points are stored in a native VMware vSphere format next to replica virtual disk files, which allows Veeam Backup & Replication to accelerate failover operations. To fail over to the necessary point of the VM replica, Veeam Backup & Replication does not need to apply rollback files. Instead, it uses a native VMware vSphere mechanism of reverting to a snapshot.

> **IMPORTANT**
>
> We recommend you against switching restore points for replicas and powering on replicas using VMware vSphere client. This may disrupt further replication operations in Veeam Backup & Replication or cause loss of important data. Instead, use Veeam Backup & Replication to perform failover operations. For more information on how to fail over to a VM replica, see Failover.

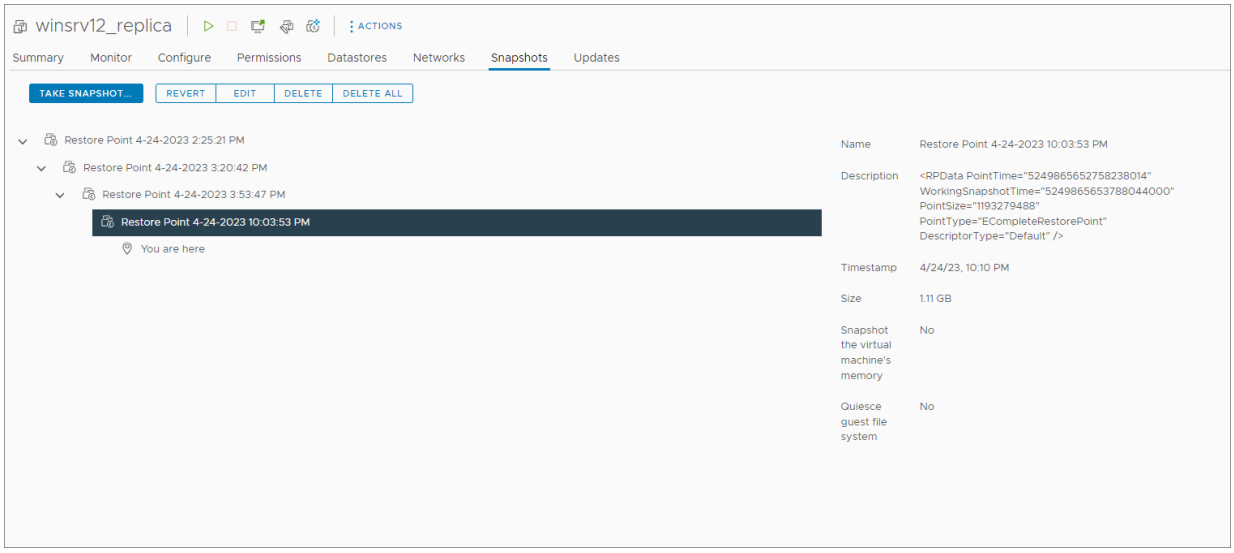

# Replica Seeding and Mapping

Replica seeding and mapping are technologies that help reduce the amount of traffic sent over a network. With these technologies, Veeam Backup & Replication do not have to transfer all of VM data from the source host to the target host across the sites during the first session of a replication job (during the initial replication).

You can use seeding and mapping in the following scenarios:

- **Seeding**

  Configure replica seeding if, in a backup repository located in the disaster recovery (DR) site, you have backups of VMs that you plan to replicate. During replication, Veeam Backup & Replication will restore VMs from these backups and will synchronize the state of the restored VMs with the latest state of the source VMs. Then Veeam Backup & Replication will use these restored VMs as replicas.

  For more information on how to create backups that can be used as "seeds" for replica, see Creating Replica Seeds.

- **Mapping**

  Configure replica mapping if, on the host in the DR site, you have ready-to-use copies of the source VMs. These can be restored VMs or replicas created by other replication jobs. Veeam Backup & Replication will synchronize the state of these ready-to-use VMs with the latest state of the source VMs and will use these VMs as replicas. You can also use replica mapping if you need to reconfigure or recreate replication jobs, for example, split one replication job into several jobs.

You can also configure both replica seeding and replica mapping in the same replication job. For example, if a replication job includes 2 VMs, you can use seeding for one VM and map the other VM to an existing VM.

> **IMPORTANT**
>
> If seeding or mapping is enabled in a replication job, all VMs in the job must be covered with seeding or mapping. If a VM neither has a seed, nor is mapped to an existing VM, it will be skipped from processing.

# Algorithm for Seeding

Replica seeding includes the following steps:

1. As a preparatory step for replica seeding, you need to create a backup of a VM that you plan to replicate. For more information on how to create a backup that will be used as a "seed" for replica, see Creating Replica Seeds.
2. When you create a replication job, you should point it to a backup repository in the DR site. During the initial synchronization, Veeam Backup & Replication accesses the backup repository where the replica seed is located, and restores the VM from the backup. The restored VM is registered on the target host in the DR site. Files of the restored VM are placed to the location you specify as the replica destination datastore.

   Virtual disks of a replica restored from the backup preserve their format (that is, if the source VM used thin provisioned disks, virtual disks of the VM replica are restored as thin provisioned).
3. Veeam Backup & Replication synchronizes the restored VM with the latest state of the source VM.

   After successful synchronization, in the **Home** view in the Veeam Backup & Replication console, under **Replicas** node you will see a VM replica with two restore points. One point will contain the state of the VM from the backup file; the other point will contain the latest state of the source VM you want to replicate.

4. During incremental synchronization, Veeam Backup & Replication transfers only incremental changes in a regular manner.

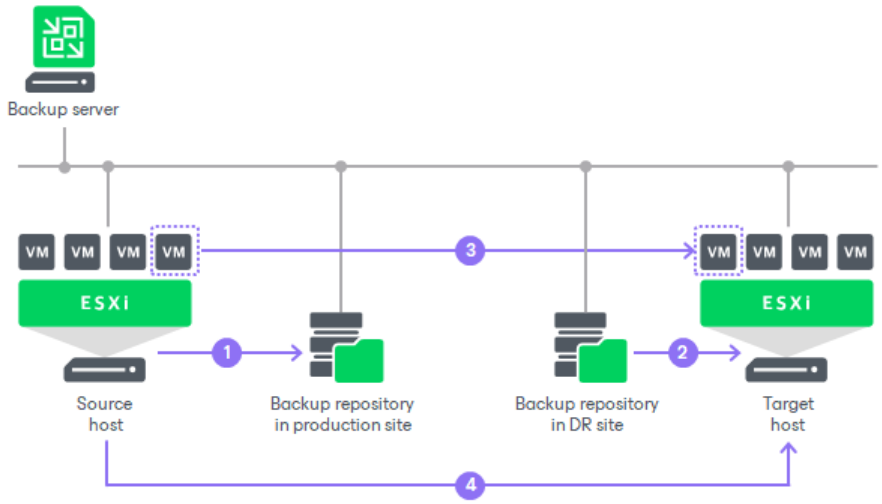


Replica seeding dramatically reduces traffic sent over WAN or slow connections because Veeam Backup & Replication does not send the full contents of the VM image. Instead, it transmits only differential data blocks.

> **TIP**
>
> If you add new VMs to an already existing replication job, you can enable replica seeding settings for these VMs. In this case, the newly added VMs will be seeded from the selected backups at the next pass of the replication job. VMs that have already been processed by the job by the time you add new VMs will be processed in a regular manner.

# Algorithm for Mapping

Replication to a mapped VM is performed in the following way:

1. During the first run, the replication job calculates the differences between the source and mapped VM. Instead of copying and transferring all data of the source VM, the replication job transfers only incremental changes to synchronize the state of the mapped VM with the state of the source VM.

   After successful synchronization, in the **Home** view of Veeam Backup & Replication, under **Replicas** node you will see a VM replica with 2 restore points:

   - One restore point will contain the latest state of the mapped VM.
   - The other restore point will contain the latest state of the source VM on the source host.

2. All subsequent runs of the replication job will be performed in a regular manner: Veeam Backup & Replication will transfer only incremental changes to the target host.

> **NOTE**
>
> If a VM replica to which the source VM is mapped has any snapshots, these snapshots will be removed during the run of the replication job.

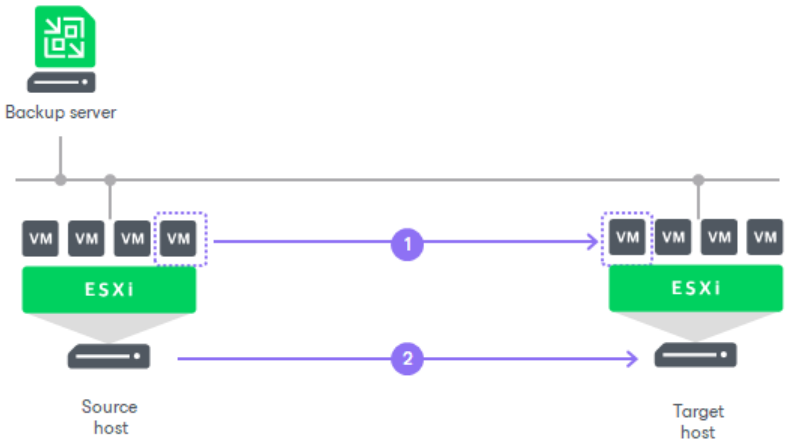
Backup server
VM
VM
VM
VM
ESXi
1
VM
VM
VM
VM
ESXi
2
Source host
Target host

# Replica from Backup

Disaster recovery plans often require that you back up and replicate the same VM for disaster recovery (DR) and high availability (HA) purposes. As a rule, this doubles the workload on the virtual infrastructure: two VM snapshots need to be created independently from one another, and VM data need to be transferred from the production site twice.

To minimize the use of compute, storage and network resources, you can use the replica from backup option. You can use this option in both on-site and off-site replication scenarios.

When you perform replication from backup, Veeam Backup & Replication does not address hosts and storage in the production environment to read VM data. As a source of data, Veeam Backup & Replication uses a backup chain that already exists in a backup repository. As a result, Veeam Backup & Replication creates only one snapshot and transfers VM data only once. Veeam Backup & Replication retrieves VM data only while a backup or backup copy job is running. The replication job re-uses retrieved data to build VM replica restore points.

## Differences between Seeding and Replica from Backup

Although replica from backup may resemble replica seeding, there is difference between these options:

- Replica seeding uses a backup file only during the first run of a replication job. To further build VM replica restore points, the replication job addresses the production environment and reads VM data from the source storage.
- Replica from backup uses a backup chain in a backup repository as the only source of data. When building a new VM replica restore point, Veeam Backup & Replication always reads data from the latest restore point in the backup chain, either full or incremental. The backup chain in the backup repository may be created by a backup job or a backup copy job.

## Requirements and Limitations for Replica from Backup

Consider the following requirements and limitations:

- You can perform replication only from backups of VMware vSphere virtual machines created by Veeam Backup & Replication. Replication from backups of VMware Cloud Director virtual machines is not supported.
- Replica from backup processes disks sequentially, not in parallel.
- Replica from backup can use the following backups only:
  - Backups that backup jobs or backup copy jobs create. These jobs must be configured on the same backup server where you configure the replication job.
  - Backups to which backup jobs or backup copy jobs are mapped. These jobs must be configured on the same backup server where you configure the replication job.
- The jobs mentioned in the previous list items must run periodically and produce new restore points. Otherwise, the replication job will have no data to retrieve and replicas will be in an outdated state.
- Replica from backup cannot use imported backups. However, there may be a situation when you need to use backups created on another backup server. In this case, use the instructions provided in Using Backups Created on Crashed Backup Server. Note that the backup server from which you import backups must not operate anymore, otherwise replica from backup may behave in an unexpected way.
- [For backups stored on scale-out backup repositories] Replica from backup ignores backups stored in a tier other than the performance tier.

# How Replica from Backup Works

Replica from backup is performed along with a regular replication job. When you set up a replication job, you define a backup repository with VM backups as a source of data. If the backups for this VM are available in different backup repositories, you can select several backup repositories as a source. In this case, Veeam Backup & Replication will look for the latest VM restore point across these backup repositories.

For example, you configure two backup jobs that process the same VM, and target these jobs at two different backup repositories. The backup jobs create the following backup files:

- *Backup job 1* creates 2 restore points in *Backup repository 1*: full backup file on Sunday and incremental backup file on Tuesday.
- *Backup Job 2* creates 1 restore point in *Backup repository 2*: full backup file on Monday.

The replication job is configured to retrieve VM data from backups and is scheduled to run daily. In this case, the replication job retrieves VM data from backups in the following way:

1. On Sunday, the replication job retrieves VM data from the full backup file in *Backup repository 1*.
2. On Monday, the replication job retrieves VM data from the full backup file in *Backup repository 2*.
3. On Tuesday, the replication job retrieves VM data from the incremental backup file in *Backup repository 1*.
4. Till next Sunday, the replication job does not retrieve any VM data because backup files are not created.

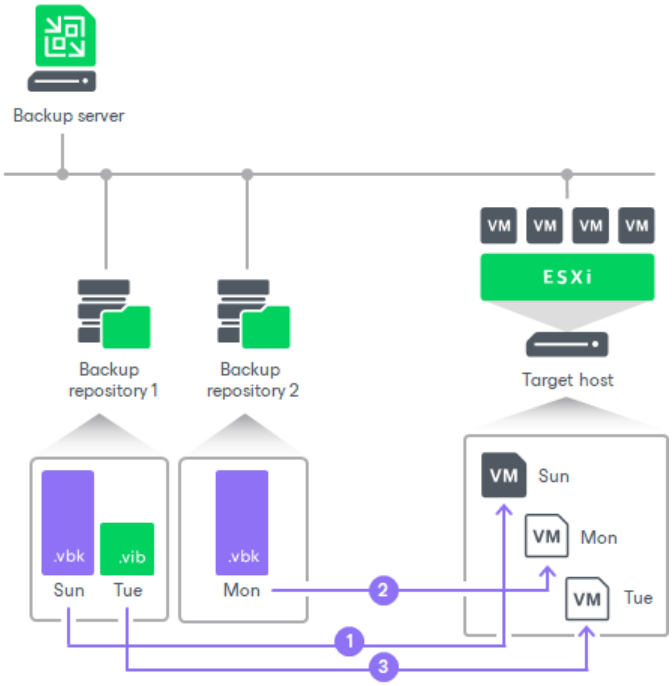


In some cases, a new restore point in the backup repository may not be created by the time a replication job starts. In this case, Veeam Backup & Replication displays a warning notifying that the latest restore point has already been replicated. The replication job session finishes with the *Warning* status.

In some cases, Veeam Backup & Replication can resume the replication process if data transfer was not finished, for example, because of the network disconnection. On the next run, Veeam Backup & Replication will continue data transfer for those disks for which Veeam Backup & Replication has started data processing and created snapshots during the current run.

NOTE

When you replicate a VM over a production network, Veeam Backup & Replication retrieves VM data as of the latest VM state. When you replicate a VM from backup, Veeam Backup & Replication retrieves VM data as of the point in time when the backup was created. The VM replica restore point has the same timestamp as a VM backup restore point, not the time when the replica job session is run.

# Using Backups Created on Crashed Backup Server

There may be a situation when you created backups on one backup server, the server crashed and you want to use these backups for replica from backup on another backup server.

Veeam Backup & Replication considers backups created on other backup servers as imported backups. Replica from backups cannot use imported backups, that is why you need to perform the following steps to use backups created on the crashed server:

1. Import the backups to the backup server where you create the replication job. You have several options:
   - You can connect the repository where the backups are stored to the backup server and then rescan the repository.
   - You can copy backups to a backup repository already added to the backup server and then rescan the repository.
   - You can copy backups to a backup repository already added to the backup server. Then edit repository settings and select the **Search the repository for existing backups and import them automatically check box** at the Review step of the wizard.
2. Create a new backup job or a backup copy job and map the imported backups to it. For more information on how to map backups, see Specify Backup Storage Settings.

   After you map a backup to a job, Veeam Backup & Replication stops considering the backup as imported.
3. Create a replication job. In the job settings, specify that you want to use backups as a data source and select the backup repository where the imported backups reside. For more information, see Specify Data Source.

IMPORTANT

Consider the following:

- The backup job or backup copy job to which you map the imported backup file must run periodically and produce new restore points. Otherwise, the replication job will have no data to retrieve and replicas will be in an outdated state.
- No other running backup servers must use the imported backups.

# Creating Replication Jobs

To create VM replicas, you must configure a replication job. The replication job defines how, where and when to replicate VM data. One job can be used to process one VM or multiple VMs.

To create a replication job, use the **New Replication Job** wizard.

# Before You Begin

Before you create a replication job, check the following prerequisites and limitations:

- Check limitations for replication described in section Considerations and Limitations.
- Backup infrastructure components that will take part in the replication process must be added to the backup infrastructure and properly configured. These include source and target ESXi hosts, one backup proxy for on-site replication scenario or two backup proxies for off-site replication scenario and backup repository for storing replica metadata. For more information, see Backup Infrastructure for Replication.
- The backup server must be able to resolve short names and connect to source and target virtualization hosts.
- The target datastore must have enough free space to store disks of replicated VMs. To receive alerts about low space on the target datastore, configure global notification settings. For more information, see Specifying Other Notification Settings.
- If you plan to replicate VMs using WAN accelerators, source and target WAN accelerators must be added to the backup infrastructure and properly configured. For more information, see Adding WAN Accelerators.
- If you plan to replicate VMs using WAN accelerators, it is recommended that you pre-populate global cache on the target WAN accelerator before you start the replication job. Global cache population helps reduce the amount of traffic transferred over WAN. For more information, see Manually Populating Global Cache.
- If you plan to replicate VMs from a backup, the backup job that you plan to use as the source must be configured beforehand. For more information, see Replica from Backup.
- If you plan to use pre-job and post-job scripts and pre-freeze and post-thaw scripts, you must create scripts before you configure the replication job. For the supported script format, see Pre-Freeze and Post-Thaw Scripts.
- If you plan to protect VMware Cloud Director objects, consider using the dedicated replication job. For more information, see Replication for VMware Cloud Director.

# Step 1. Launch New Replication Job Wizard

To launch the **New Replication Job** wizard, do one of the following:

- On the **Home** view. On the ribbon, click **Replication Job** > **Virtual machine** > **VMware vSphere**.
- Open the **Home** view. In the inventory pane right-click the **Jobs** node and select **Replication** > **Virtual machine** > **VMware vSphere**.
- Open the **Inventory** view. In the working area, select VMs that you want to replicate and right-click one of them. Select **Add to replication job** > **New job** if you want to create a new replication job, or **Add to replication job** > *<Job Name>* if you want to add VMs to an existing replication job.

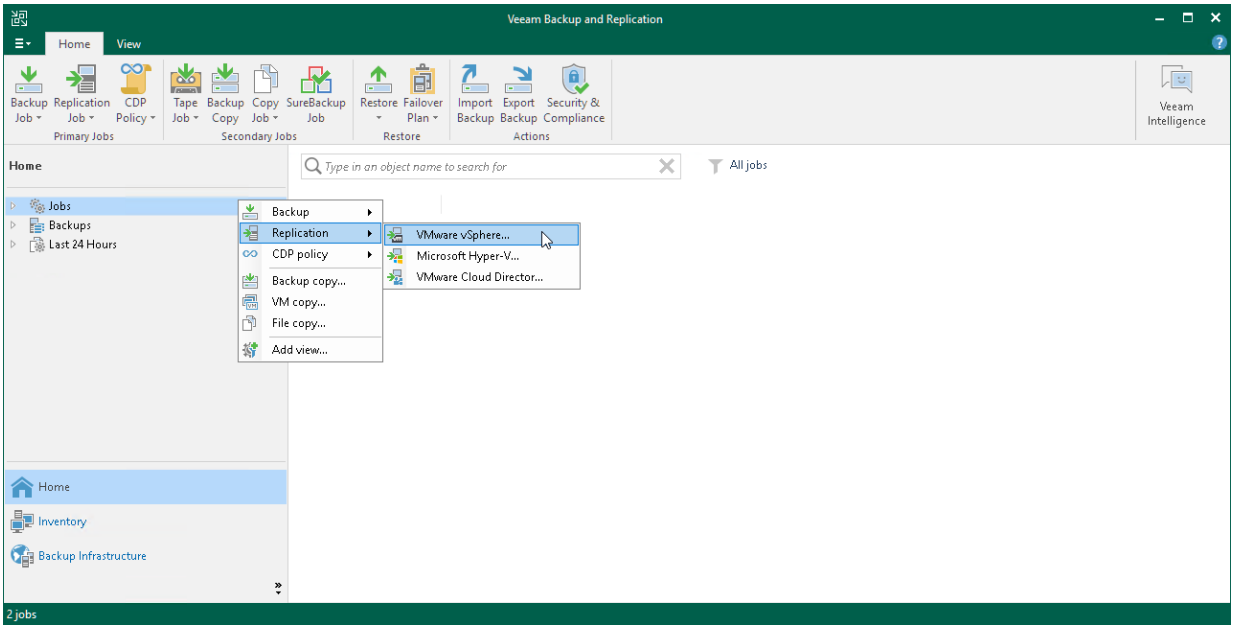

# Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a job name and description, and configure advanced settings for the replication job:

1. In the **Name** field, enter a name for the replication job.
2. In the **Description** field, provide a description for future reference.
3. If a network between your production and disaster recovery (DR) sites has low bandwidth, and you want to reduce the amount of traffic sent during the first run of the replication job, select the **Replica seeding (for low bandwidth DR sites)** check box.

   When selected, this check box enables the **Seeding** step where you will have to configure replica seeding and mapping. For more information on seeding and mapping, see Replica Seeding and Mapping.
4. If your DR site networks do not match your production site networks, select the **Network remapping (for DR sites with different virtual networks)** check box.

   When selected, this check box enables the **Network** step where you will have to configure a network mapping table.
5. If the IP addressing scheme in your production site differs from the scheme in the DR site, select the **Replica re-IP (for DR sites with different IP addressing scheme)** check box.

   When selected, this check box enables the **Re-IP** step where you will have to configure replica re-IP rules.
6. If you want the resource scheduler of Veeam Backup & Replication to prioritize this job higher than other similar jobs and to allocate resources to it in the first place, select the **High priority** check box. For more information on job priorities, see Job Priorities.

> **TIP**
>
> In the list of jobs in the Veeam Backup & Replication console, jobs with the **High priority** option enabled are marked with a red flag ( ).

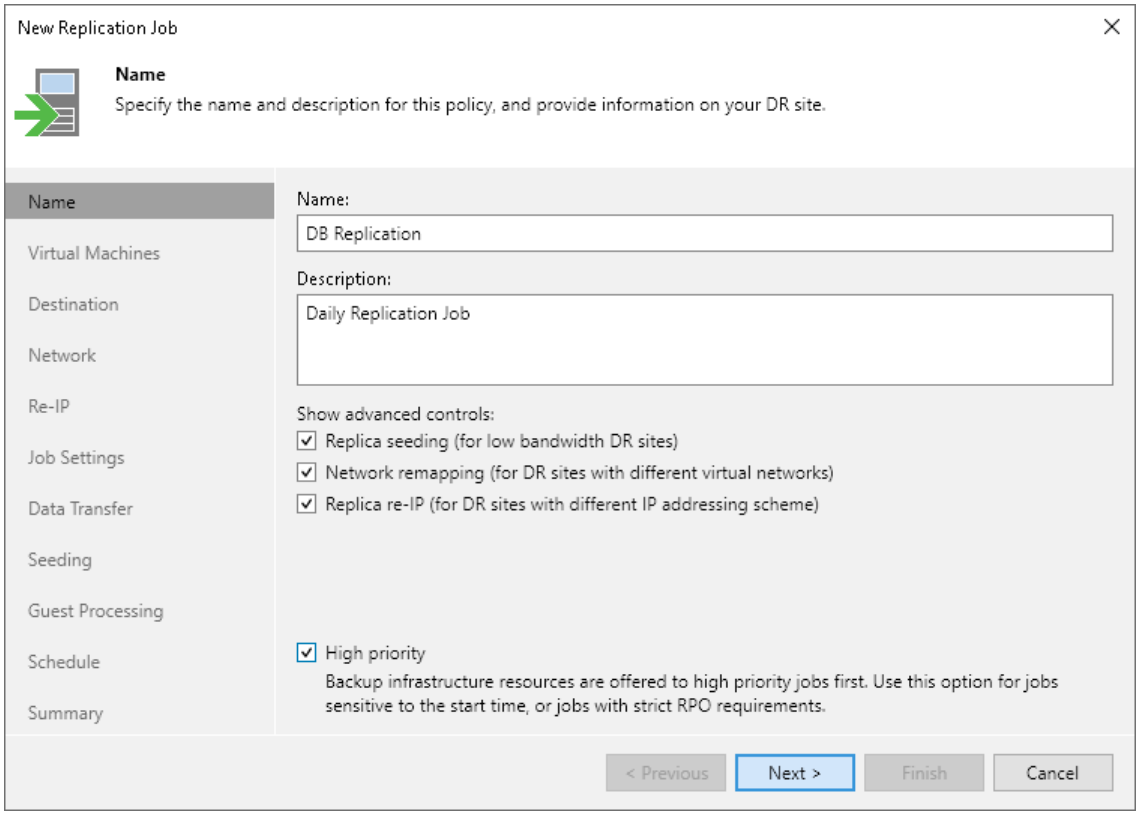
New Replication Job
Name
Specify the name and description for this policy, and provide information on your DR site.
Name
Virtual Machines
Destination
Network
Re-IP
Job Settings
Data Transfer
Seeding
Guest Processing
Schedule
Summary
Name:
DB Replication
Description:
Daily Replication Job
Show advanced controls:
Replica seeding (for low bandwidth DR sites)
Network remapping (for DR sites with different virtual networks)
Replica re-IP (for DR sites with different IP addressing scheme)
High priority
Backup infrastructure resources are offered to high priority jobs first. Use this option for jobs sensitive to the start time, or jobs with strict RPO requirements.
< Previous
Next >
Finish
Cancel

# Step 3. Select VMs to Replicate

At the **Virtual Machines** step of the wizard, select VMs and VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags) that you want to replicate:

1. Click **Add**.
2. In the **Add Object** window, select the necessary VMs or VM containers and click **Add**. If you select VM containers and add new VMs to this container in future, Veeam Backup & Replication will update replication job settings automatically to include these VMs.

   You can use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only VMs that have all the selected tags will be processed by the job.

   To quickly find the necessary VMs, you can use the search field at the bottom of the **Add Object** window. If you want to switch between types of VMs you want to search through, use the button to the left of the search field.

The total size of objects added to the job is displayed in the **Total size** field. Use the **Recalculate** button to refresh the total size value after you add a new object to the job.

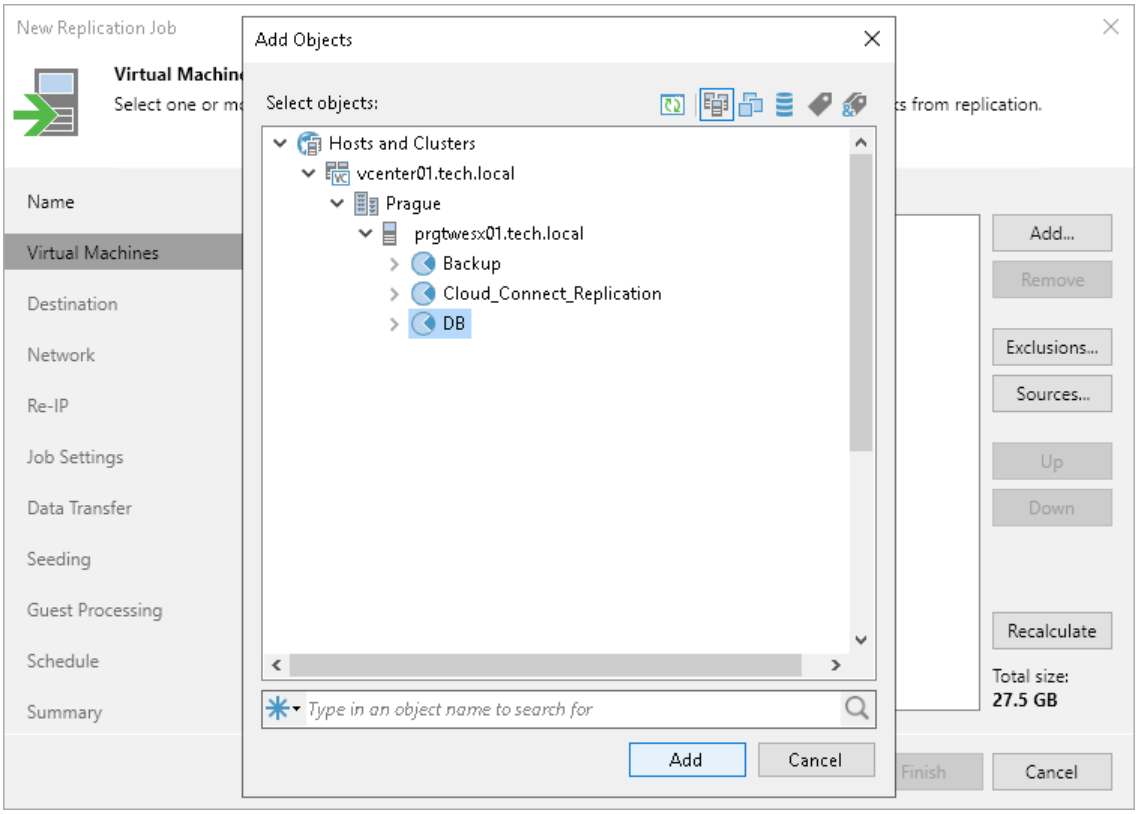

# Step 4. Specify Data Source

You can select a data source from which Veeam Backup & Replication will read VM data:

1. At the **Virtual Machines** step of the wizard, click **Source**.
2. In the **Source Repositories** window, select one of the following options:
   - **From production storage**. In this case, Veeam Backup & Replication will retrieve VM data from datastores connected to the source ESXi host.
   - **From backup files**. In this case, Veeam Backup & Replication will read VM data from the backup chain already existing in the selected backup repository. This option can be used in the replica from backup scenario. For more information, see Replica from Backup.

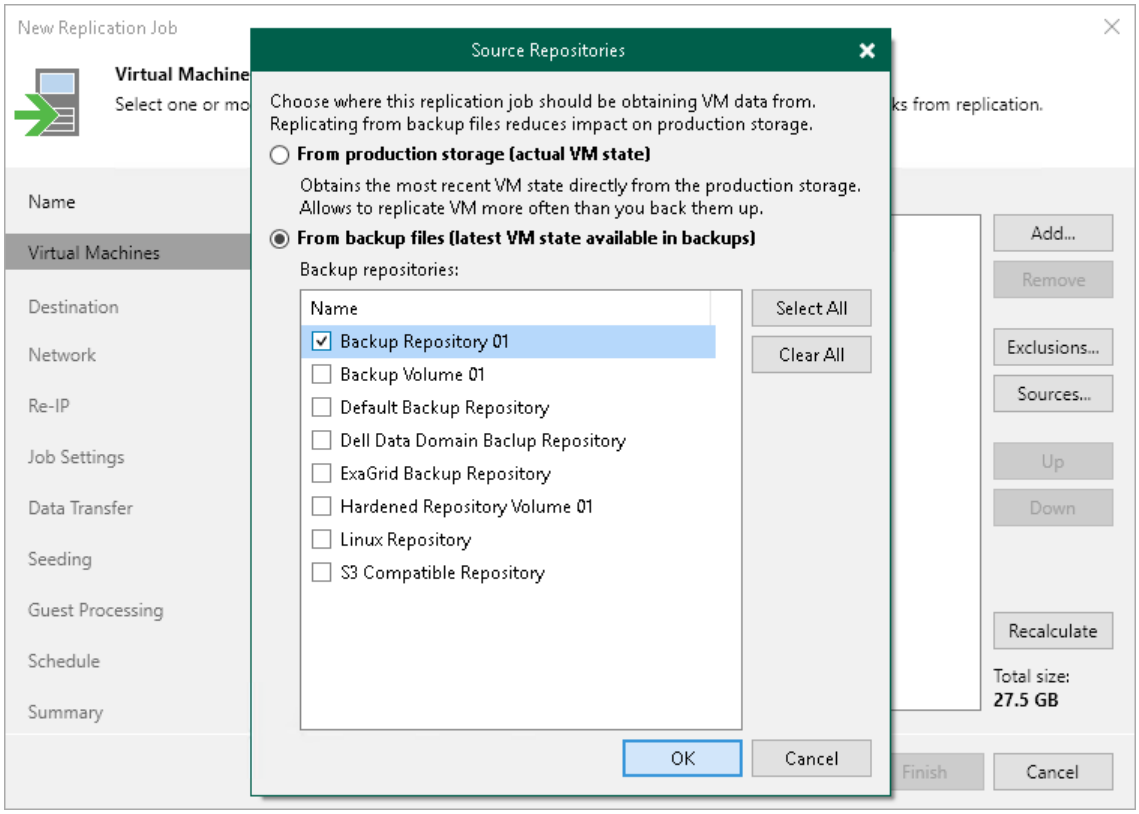

# Step 5. Exclude Objects from Replication Job

After you have added VMs and VM containers to the replication job, you can specify which objects you want to exclude from replicas. You can exclude the following types of objects:

- VMs or VM containers
- VM disks

> **NOTE**
>
> To make the replication process faster and reduce the size of created replicas, Veeam Backup & Replication automatically excludes the following objects from replication:
>
> - VM log files
> - VM templates from VM containers

## Excluding VMs or VM Containers

To exclude VMs from a VM container:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, check that the **VMs** tab is selected and click **Add**.
3. In the **Add Objects** window, select VMs or VM containers that you want to exclude from being replicated and click **Add**.

   You can use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only VMs that have all the selected tags will be excluded from the job.

   You can also use the **Show full hierarchy** check box to display the hierarchy of all VMware Servers added to the Veeam Backup & Replication infrastructure.

4. Click **OK**.

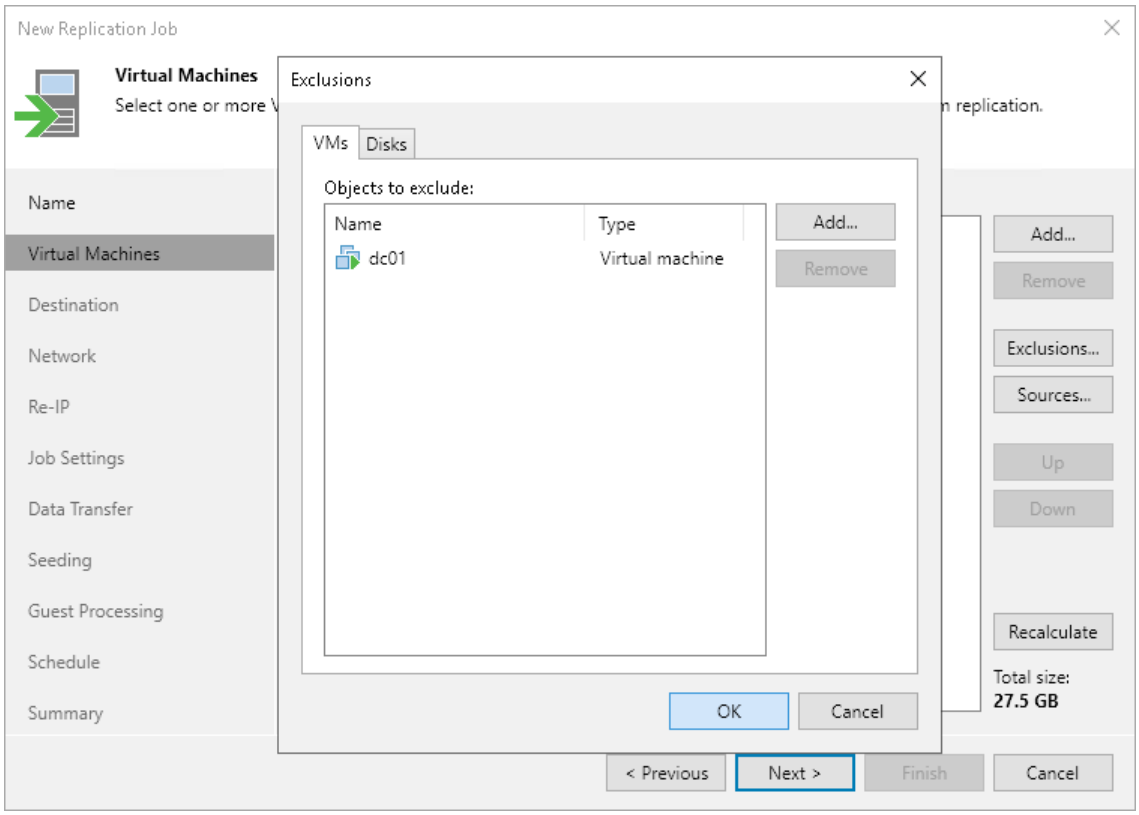


# Excluding Disks

To exclude VM disks:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, do the following:
   a. Switch to the **Disks** tab.
   b. If you want to exclude disks of VMs that are added as a part of containers, click **Add**. In the **Add Objects** window, select the necessary VMs and click **Add**. Veeam Backup & Replication will include these VMs in the list as standalone objects.
   c. In the **Disks to process** list, select VMs or VM containers whose disks you want to exclude.
   d. Click **Edit**.
3. In the **Select Disks** window, select disks that you want to replicate: all disks, 0:0 disks (as a rule, system disks) or specific IDE, SCSI, SATA or NVMe disks. Disks that you do not select will be excluded from processing. Click **OK**.
4. In the **Exclusions** window, click **OK**.

> **NOTE**
>
> If you exclude disks from being replicated and enable application-aware processing, Veeam Backup & Replication will still perform application-aware processing for the excluded disks. This means that VSS will process disk data.

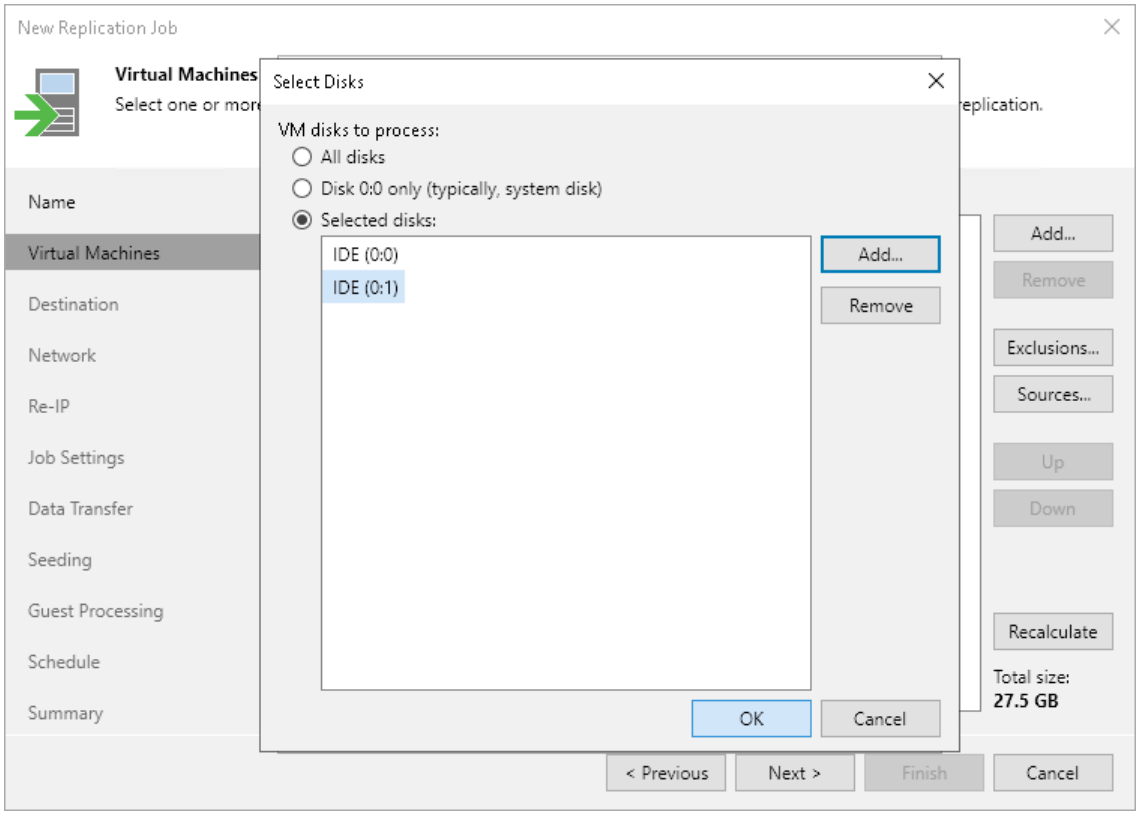
New Replication Job
Virtual Machines
Select one or more
eplication.
Name
Virtual Machines
Destination
Network
Re-IP
Job Settings
Data Transfer
Seeding
Guest Processing
Schedule
Summary
Add...
Remove
Exclusions...
Sources...
Up
Down
Recalculate
Total size:
27.5 GB
< Previous
Next >
Finish
Cancel
Select Disks
VM disks to process:
All disks
Disk 0:0 only (typically, system disk)
Selected disks:
IDE (0:0)
IDE (0:1)
Add...
Remove
OK
Cancel

# Step 6. Specify VM Processing Order

At the **Virtual Machines** step of the wizard, click **Up** and **Down** to change the processing order. VMs at the top of the list have a higher priority and will be processed first.

> **NOTE**
>
> Consider the following:
>
> - VMs inside a VM container are processed at random. To ensure that VMs are processed in the defined order, you must add them as standalone VMs, not as a part of containers.
> - The processing order may differ from the order that you have defined. For example, if resources of a VM that is higher in the priority are not available, and resources of a VM that is lower in the priority are available, Veeam Backup & Replication will process the VM with the lower priority first.

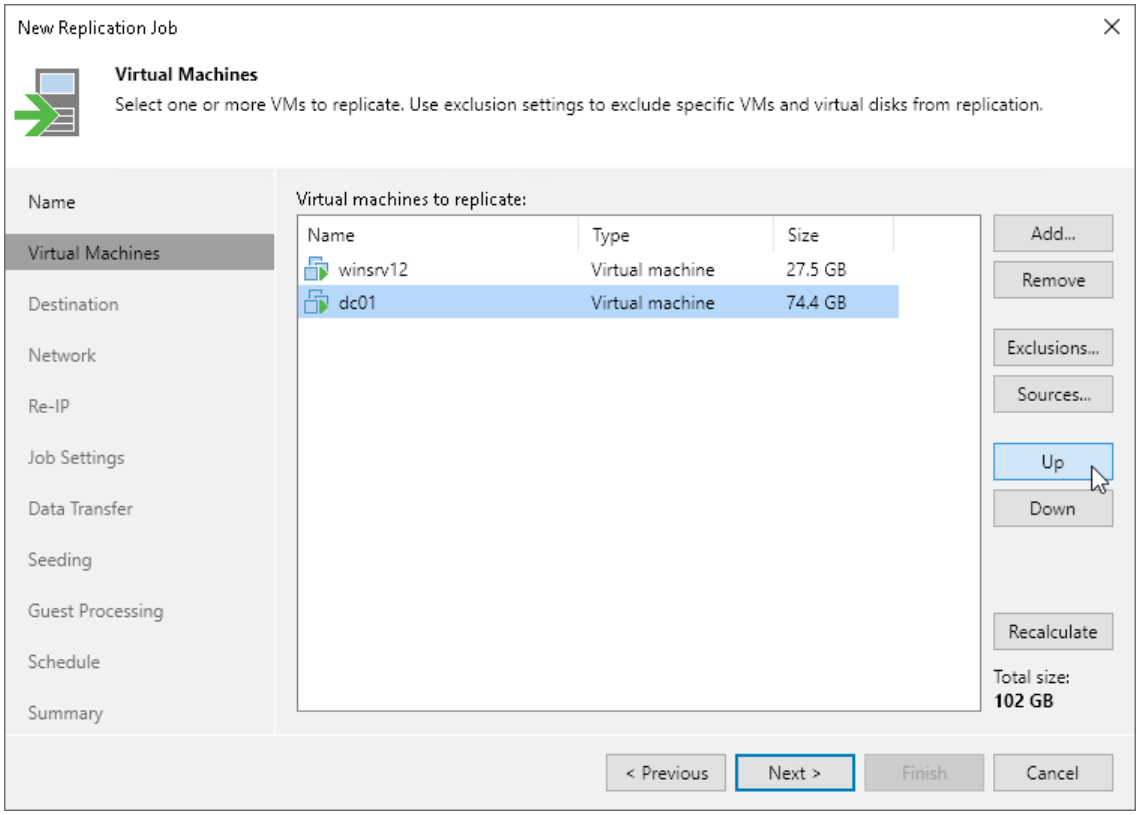

# Step 7. Specify Replica Destination

At the **Destination** step of the wizard, select a target host or cluster, resource pool, folder and datastore for replicas, and types of replica disks:

1. Next to the **Host or cluster** field, click **Choose** and select a host or cluster where replicas must be registered. If you select a cluster or vCenter Server, the replication process will become more sustainable — the replication process will not fail if there is at least one available host in the cluster.

   If you select a cluster as a destination, Veeam Backup & Replication will request VMware to send the list of available hosts, and will select the first host in this list as the destination for the replicas. These replicas will be stored on the datastore with the most free disk space.

2. Next to the **Resource pool** field, click **Choose** and select a resource pool to which replicas will be added.

   If you have selected to replicate multiple VMs and want to add individual replicas to other resource pools:

   a. Click the **Pick resource pool for selected replicas** link.

   b. In the **Choose Resource Pool** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose Resource Pool** window, select the necessary VMs in the **Replica VM resource pool** list. At the bottom of the window, click **Resource Pool**.

   e. In the **Select Resource Pool** window, select the necessary resource pool and click **OK**.

3. Next to the **VM folder** field, click **Choose** and select a folder where all VM files will be stored. Note that the **VM folder** section is disabled if you have selected a standalone ESXi host as the target for replicas.

   If you have selected to replicate multiple VMs and want to place individual replicas to other folders:

   a. Click the **Pick VM folder for selected replicas** link.

   b. In the **Choose Folder** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose Folder** window, select the necessary VMs in the **Replica VM folder** list. At the bottom of the window, click **VM Folder**.

   e. In the **Select Folder** window, select the necessary folder.

4. Next to the **Datastore** field, click **Choose** and select a datastore where replica files will be stored. Note that if you have chosen to replicate VMs to a cluster, Veeam Backup & Replication displays only shared datastores.

   If you have selected to replicate multiple VMs and want to place individual replicas to other datastores:

   a. Click the **Pick datastore for selected virtual disks** link.

   b. In the **Choose VM Files Location** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose VM Files Location** window, select the necessary VMs in the **Files location** list. At the bottom of the window, click **Datastore**.

   e. In the **Select Datastore** window, select the necessary datastore.

5. If you want to store replica configuration files and disk files in different datastores:
   a. Click the **Pick datastore for selected virtual disks** link.
   b. In the **Choose VM Files Location** window, click **Add VM**.
   c. In the **Add Objects** window, select the necessary VMs and click **Add**.
   d. In the **Choose VM Files Location** window, expand the necessary VMs in the **Files location** list, and select the necessary files. At the bottom of the window, click **Datastore**.
   e. In the **Select Datastore** window, select the destination for the selected type of files.

> **NOTE**
>
> After a replication job has finished, you can change the target location for replica files:
>
> - If you specify another host or datastore within the same vCenter Server, the target location will be changed only for the replica files of VMs added to the job after the latest job run. To change the location of "old" replica files, you need to disable the job, change the path in the settings and migrate the replica to the specifies location using native VMware vSphere methods.
> - If you specify a datastore or host within another vCenter Server, Veeam Backup & Replication will start new replication chains in the specified target location for all VMs in the replication job.

6. You can change types of replica disks. By default, Veeam Backup & Replication saves disks in the thin type.

   To change replica disk types:
   a. Click the **Pick datastore for selected virtual disks** link.
   b. In the **Choose VM Files Location** window, click **Add VM**.
   c. In the **Add Objects** window, select VMs whose disk type you want to change and click **Add**.
   d. In the **Choose VM Files Location** window, select the necessary VMs in the **Files location** list. At the bottom of the window, click **Disk type**.
   e. In the **Disk Type Settings** window, select a type that will be used to restore replica disk files: same as source, thin, thick lazy zeroed or thick eager zeroed.

      For more information about disk types, see VMware Docs.

   > **NOTE**
   >
   > Disk type change is available only for VMs that use virtual hardware version 7 or later.

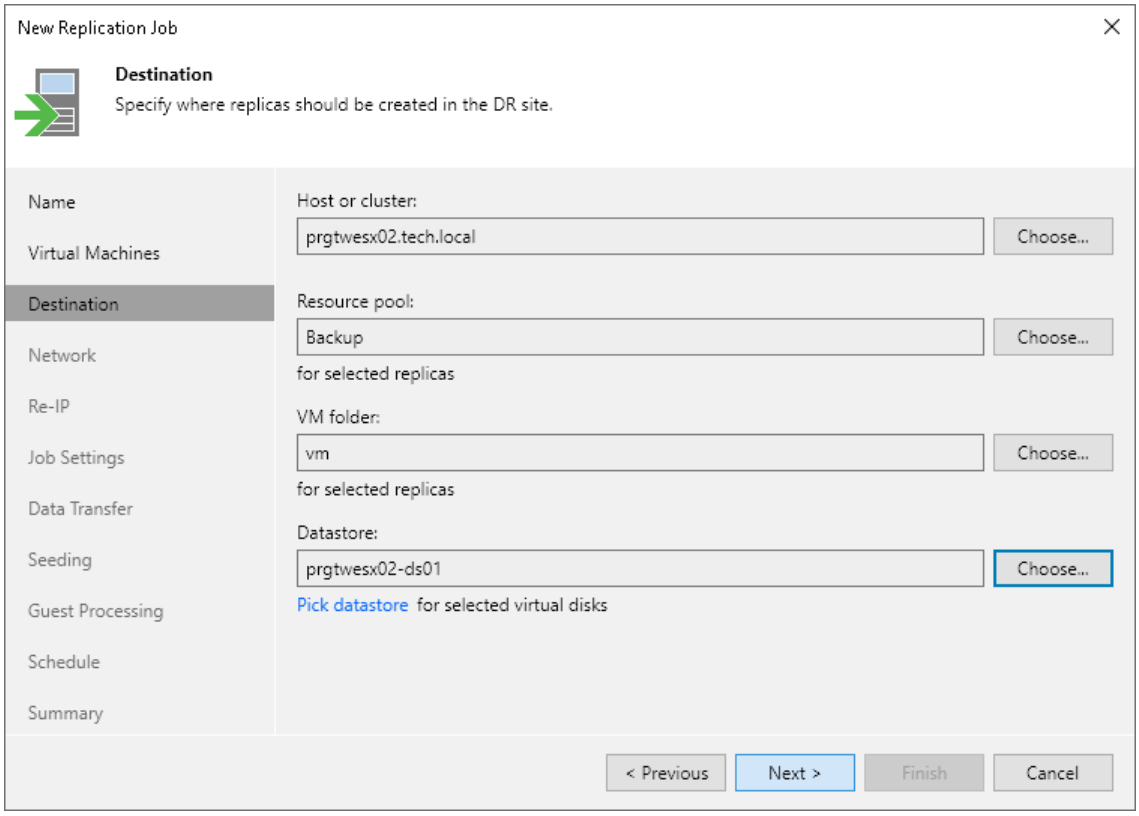
New Replication Job
Destination
Specify where replicas should be created in the DR site.
Name
Virtual Machines
Destination
Network
Re-IP
Job Settings
Data Transfer
Seeding
Guest Processing
Schedule
Summary
Host or cluster:
prgtwesx02.tech.local
Choose...
Resource pool:
Backup
Choose...
for selected replicas
VM folder:
vm
Choose...
for selected replicas
Datastore:
prgtwesx02-ds01
Choose...
Pick datastore for selected virtual disks
< Previous
Next >
Finish
Cancel

# Step 8. Create Network Mapping Table

The **Network** step of the wizard is available if you have selected the **Network remapping** option at the **Name** step of the wizard.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the production site to networks in the disaster recovery (DR) site. When the replication session starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update replica configuration to replace the production networks with the specified networks in the DR site. As a result, you will not have to re-configure network settings manually.

To add a row to the network mapping table:

1. Click **Add**.
2. In the **Network Mapping** window, click **Browse** next to the **Source network** field.
3. In the **Select Network** window, select the production network to which the source workloads are connected and click **OK**.
4. In the **Network Mapping** window, click **Browse** next to the **Target network** field.
5. In the **Select Network** window, select a network in the DR site to which replicas will be connected and click **OK**.
6. In the **Network Mapping** window, click **OK**.

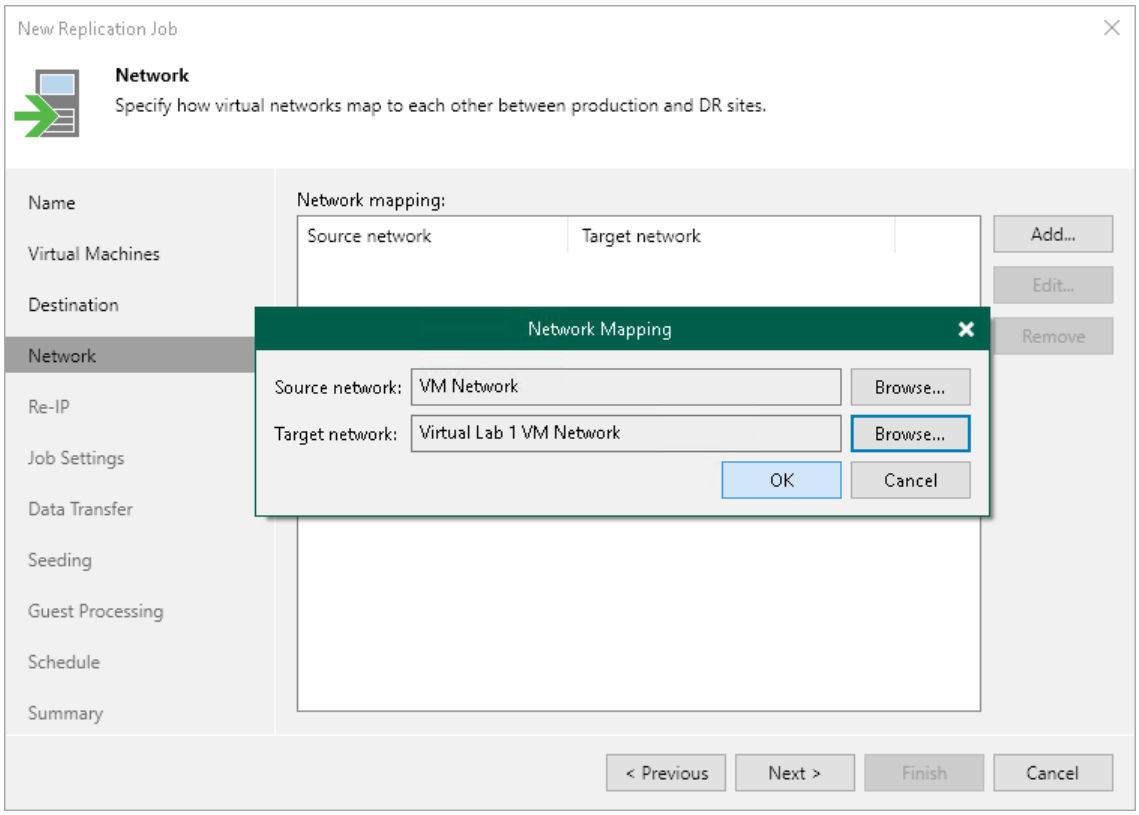

# Step 9. Configure Re-IP Rules

The **Re-IP** step of the wizard is available if you have selected the **Replica re-IP** option at the Name step of the wizard. This step applies only to VMs with Microsoft Windows OSes.

At the **Re-IP** step of the wizard, configure re-IP rules. These rules map IPs in the production site to IPs in the disaster recovery (DR) site. When you perform failover, Veeam Backup & Replication will check the configured re-IP rules and will change replica IPs if the rules apply. VM replicas will get new IP addresses according to the network masks specified in the rules, so that you will be able to reach replicas in the DR site.

To configure a re-IP rule:

1. Click **Add** and select whether you want to configure an IPv4 or IPv6 rule. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
2. In the **Source VM** section, specify an IP numbering scheme used in the production site.

   To facilitate the configuration, Veeam Backup & Replication detects an IP address and subnet mask for the backup server and pre-populates values in the **Source VM** section.
3. In the **Target VM** section, specify an IP address, subnet mask and default gateway that will be used for replicas in the DR site. If required, specify the DNS server addresses. For the IPv4 rules, you can also specify WINS server addresses.
4. In the **Description** field, provide a description.
5. Click **OK**.

NOTE

Consider the following:

- You can specify static IPs or IP ranges. Do not use 0 to specify IP address ranges. In Veeam Backup & Replication, value 172.16.17.0 means a regular IP address 172.16.17.0, not an IP address range. To specify a range, use the asterisk character (*).
- Replica re-IP works only if you perform replica failover using Veeam Backup & Replication. If you power on a VM replica in some other way, for example, manually using vSphere Client, re-IP rules will not be applied to it.
- The backup server OS must support mounting of the system disks of VMs that will be replicated.

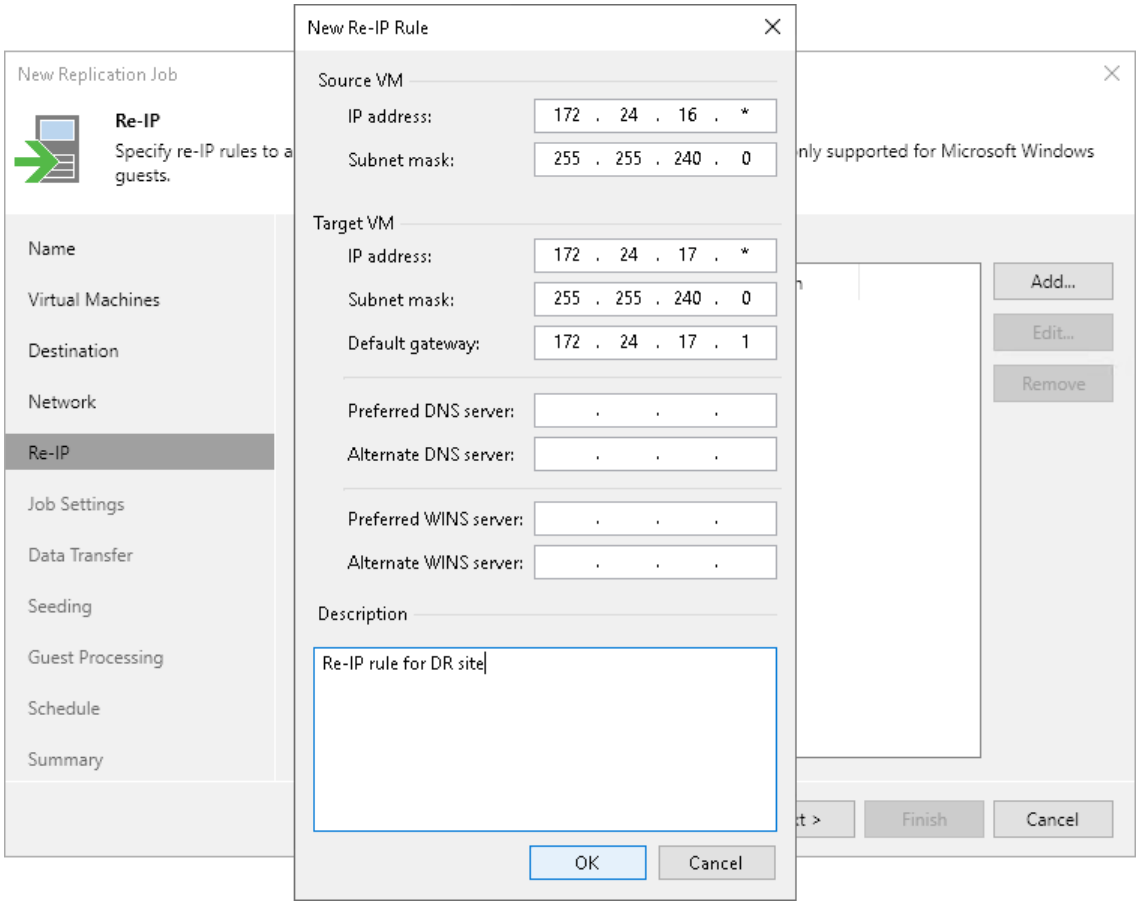
New Replication Job
Re-IP
Specify re-IP rules to a
guests.
nly supported for Microsoft Windows
Name
Virtual Machines
Destination
Network
Re-IP
Job Settings
Data Transfer
Seeding
Guest Processing
Schedule
Summary
Add...
Edit...
Remove
t >
Finish
Cancel
New Re-IP Rule
Source VM
IP address: 172 . 24 . 16 . *
Subnet mask: 255 . 255 . 240 . 0
Target VM
IP address: 172 . 24 . 17 . *
Subnet mask: 255 . 255 . 240 . 0
Default gateway: 172 . 24 . 17 . 1
Preferred DNS server:
Alternate DNS server:
Preferred WINS server:
Alternate WINS server:
Description
Re-IP rule for DR site
OK
Cancel

# Step 10. Specify Replication Job Settings

At the **Job Settings** step of the wizard, specify a backup repository for storing replica metadata, replica name and number of restore points to keep:

1. From the **Repository for replica metadata** list, select a backup repository that will store metadata for VM replicas. For more information on the recommended repository location, see Backup Infrastructure for Replication.

   > **IMPORTANT**
   >
   > Consider the following:
   >
   > - You cannot store VM replica metadata on deduplicating storage appliances. During replication jobs, Veeam Backup & Replication frequently reads and writes small portions of metadata from/to the backup repository. Frequent access to metadata causes low performance of deduplicating storage appliances, which may result in low performance of replication jobs.
   > - You cannot store replica metadata in a scale-out backup repository or object storage repository.

2. In the **Replica name suffix** field, enter a suffix that will be added to the source VM names.

   To register a VM replica on the target host, Veeam Backup & Replication appends the specified suffix to the name of the source VMs. Files of the VM replica are placed to the *VMname_suffix* folder on the selected datastore.

3. In the **Restore points to keep** field, specify the number of restore points that the replication job must maintain. Due to VMware restrictions on the number of VM snapshots, the maximum number of restore points for VM replicas is limited to 28. When the specified number is exceeded, Veeam Backup & Replication removes the earliest restore points.

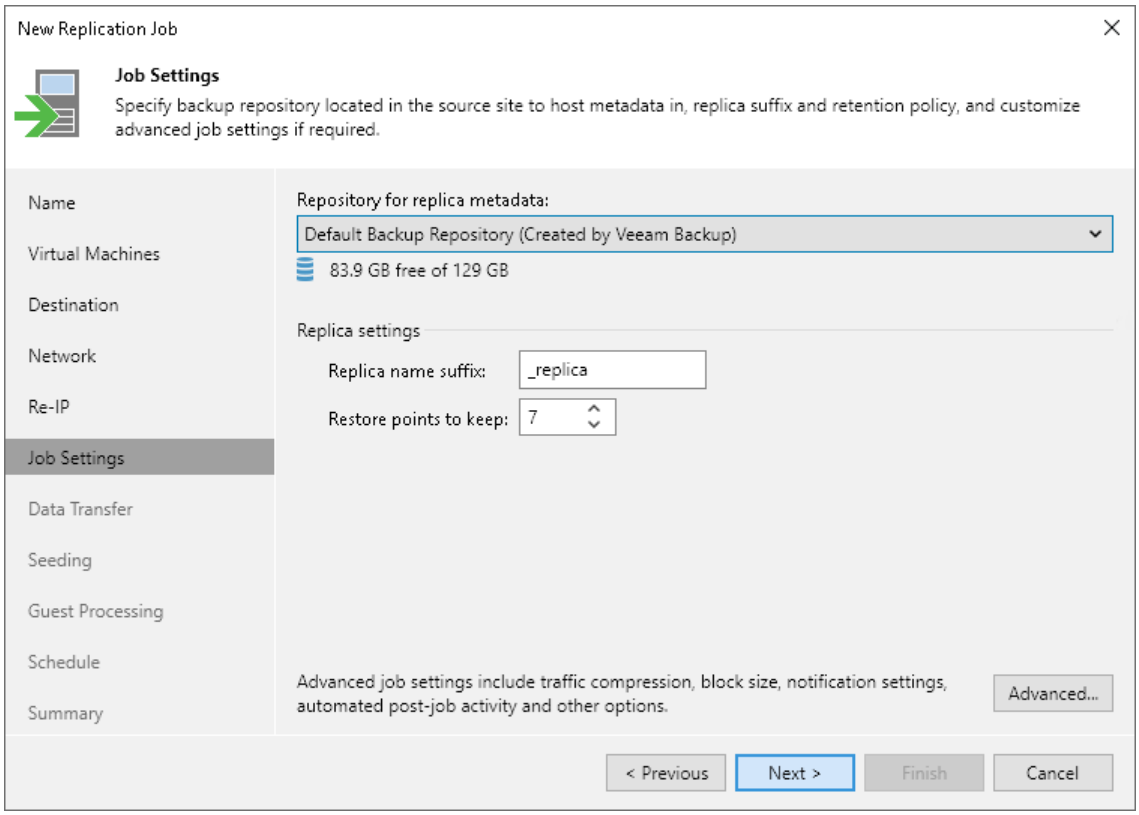

# Step 11. Specify Advanced Replica Settings

At the **Job settings** step of the wizard, you can specify the following settings for the replication job:

- Traffic settings
- Notification settings
- vSphere settings
- Integration settings
- Script settings

> **TIP**
>
> After you specify necessary settings for the replication job, you can save them as default settings. To do this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new replication job, Veeam Backup & Replication will automatically apply the default settings to the new job.

## Traffic Settings

You can optimize data traffic sent over network by specifying which data you want to replicate, data compression level and optimize the job performance and storage usage:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. In the **Advanced Settings** window, check that the **Traffic** tab is selected.
3. [For Microsoft Windows NTFS] By default, Veeam Backup & Replication excludes data blocks of the `hiberfil.sys` and `pagefile.sys` system files from replicas. For more information on how Veeam Backup & Replication excludes data blocks of these system files, see Swap Files.

   If you want to include data blocks of the `hiberfil.sys` and `pagefile.sys` system files into replicas, clear the **Exclude swap file blocks** check box. Note that including these files into replicas will increase their size.
4. [For replication from production storage] By default, Veeam Backup & Replication does not copy deleted file blocks ("dirty" blocks on the VM guest OS) to the target location. For more information, see Deleted File Blocks.

   If you want to include dirty data blocks into VM replicas, clear the **Exclude deleted file blocks** check box. Note that including these files into replicas will increase their size.
5. From the **Compression level** list, select a compression level for VM replicas. For more information on data compression and compressions levels, see Data Compression and Deduplication.

   > **NOTE**
   >
   > Compression is applicable only if VM data is transferred between two backup proxies. If one backup proxy acts as the source and target backup proxy, VM data is not compressed at all.

6. [For replication from production storage] In the **Storage optimization** section, select block size that will be used to process VMs. For more information on the data blocks sizes and how they affect performance, see Storage Optimization.

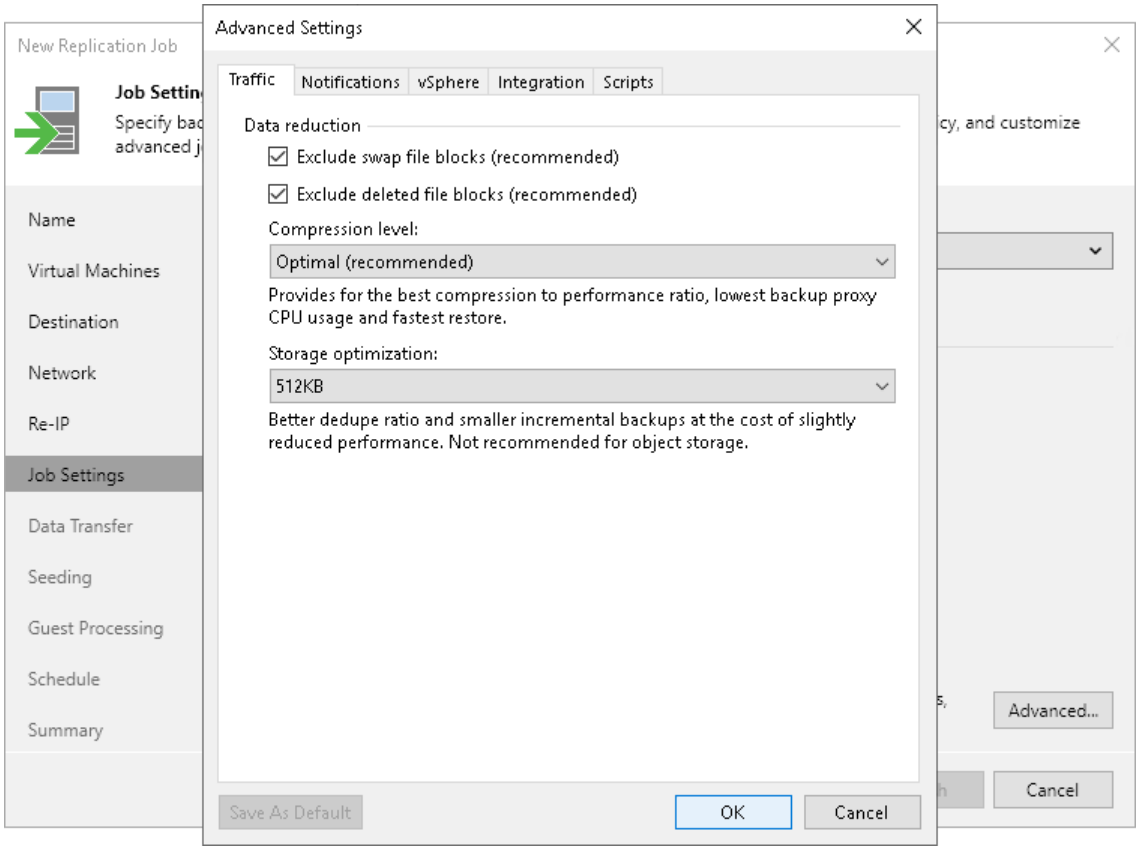


# Notification Settings

To specify notification settings for the replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. In the **Advanced Settings** window, click the **Notifications** tab.
3. To receive SNMP traps when the job completes successfully, select the **Send SNMP notifications for this job** check box.

   SNMP traps will be sent if you configure global SNMP settings in Veeam Backup & Replication and configure software on recipient machine to receive SNMP traps. For more information, see Specifying SNMP Settings.
4. To receive notifications by email in case of job failure, success or warning, select the **Send email notifications to the following recipients** check box. Then configure notification settings:
   a. Check that you have configured global email notification settings as described in section Configuring Global Email Notification Settings.
   b. In the text field, specify a recipient email address. If you want to specify multiple addresses, separate them by a semicolon.
   c. To use global notification settings, select **Use global notification settings**.
   d. To specify a custom notification subject and redefine at which time notifications must be sent, select **Use custom notification settings specified below.** Then specify the following settings:

i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%JobResult%*, *%JobName%*, *%ObjectCount%* (number of VMs in the job) and *%Issues%* (number of VMs in the job that have been processed with the *Warning* or *Failed* status).

ii. Select the **Notify on success**, **Notify on error** or **Notify on warning** check boxes to receive an email notification if the job gets the *Warning*, *Success* or *Error* status.

iii. Select the **Suppress notifications until the last retry** check box to receive the notification about the final job status. If you do not enable this option, Veeam Backup & Replication will send one notification per every job retry.

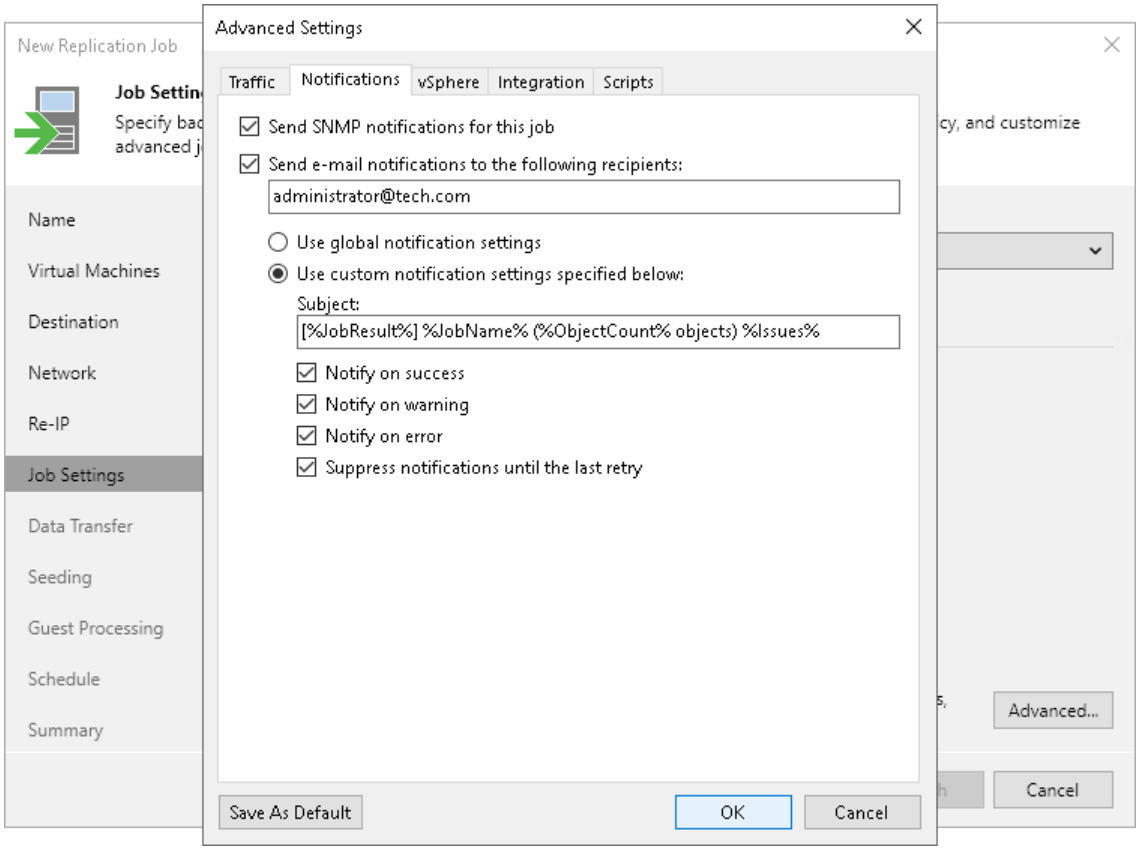


# vSphere Settings

To specify vSphere settings for the replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. Click the **vSphere** tab.
3. Select the **Enable VMware tools quiescence** check box to freeze the file system of processed VMs during replication.

   Depending on the VM version, Veeam Backup & Replication will use the VMware FileSystem Sync Driver (vmsync) driver or VMware VSS component in VMware Tools for VM snapshot creation. These tools are responsible for quiescing the VM file system and bringing the VM to a consistent state suitable for backup. For more information, see VMware Tools Quiescence.
4. In the **Changed block tracking** section, configure VMware vSphere changed block tracking (CBT):

   a. To enable CBT, make sure that the **Use changed block tracking data** check box is selected.

b. To force using CBT even if CBT is disabled in VM configuration, make sure that the **Enable CBT for all processed VMs automatically** check box is selected.

c. To reset CBT after the replication job starts for the first time, make sure that the **Reset CBT on each Active Full backup automatically** check box is selected.

   CBT reset helps avoid issues, for example, when CBT returns incorrect changed data.

For more information on CBT, see Changed Block Tracking.

**IMPORTANT**

You can use CBT for VMs with virtual hardware version 7 or later. These VMs must not have existing snapshots.

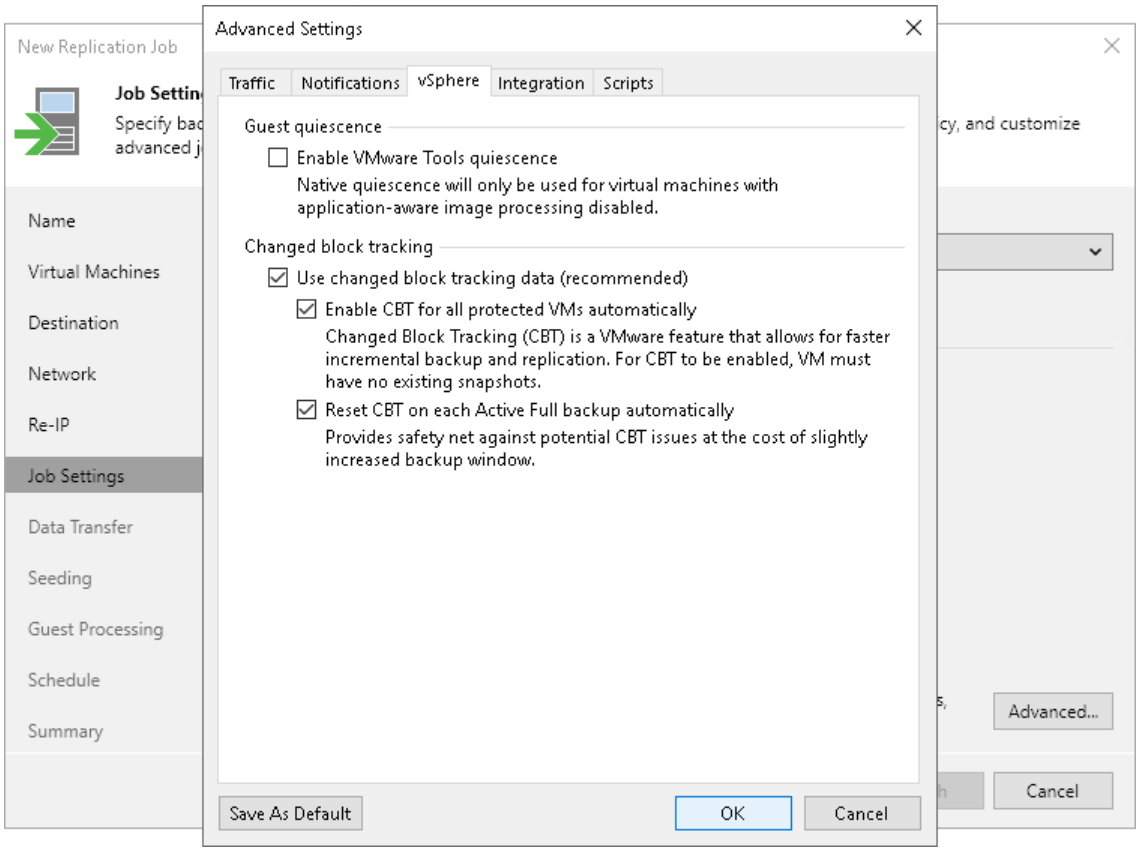


# Integration Settings

On the **Integration** tab, define whether you want to use the Backup from Storage Snapshots technology to create a VM replica. Backup from Storage Snapshots lets you leverage storage snapshots for VM data processing. The technology improves RPOs and reduces the impact of replication activities on the production environment. For more information, see the Backup from Storage Snapshot section in the Storage System Snapshot Integration Guide.

To specify storage integration settings for the replication job:

1. Check.prerequisites. For more information, see the Configuring Backup from Storage Snapshots section in the Storage System Snapshot Integration Guide.
2. At the **Job Settings** step of the wizard, click **Advanced**.
3. Click the **Integration** tab.

4. By default, Backup from Storage Snapshots functionality is enabled. If you do not want to use it, clear the **Enable backup from storage snapshots** check box.
5. If you add to the job many VMs whose disks are located on the same volume or LUN, select the **Limit processed VM count per storage snapshot to** check box and specify the number of VMs for which one storage snapshot must be created.

   In a regular job processing course, Veeam Backup & Replication creates a VMware snapshot for every VM added to the job and then triggers one storage snapshot for all VMs. In some situations, creating VMware snapshots for all VMs may require a lot of time. If you limit the number of VMs per storage snapshot, Veeam Backup & Replication will divide VMs into several groups, trigger a separate storage snapshot for every VM group and read VM data from these snapshots. As a result, the job performance will increase. For more information, see the Limitation on Number of VMs per Snapshot section in the Storage System Snapshot Integration Guide.
6. If Veeam Backup & Replication fails to create a storage snapshot, VMs whose disks are hosted on the storage system will not be processed by the job. To fail over to the regular data processing mode and replicate such VMs, select the **Failover to standard backup** check box.

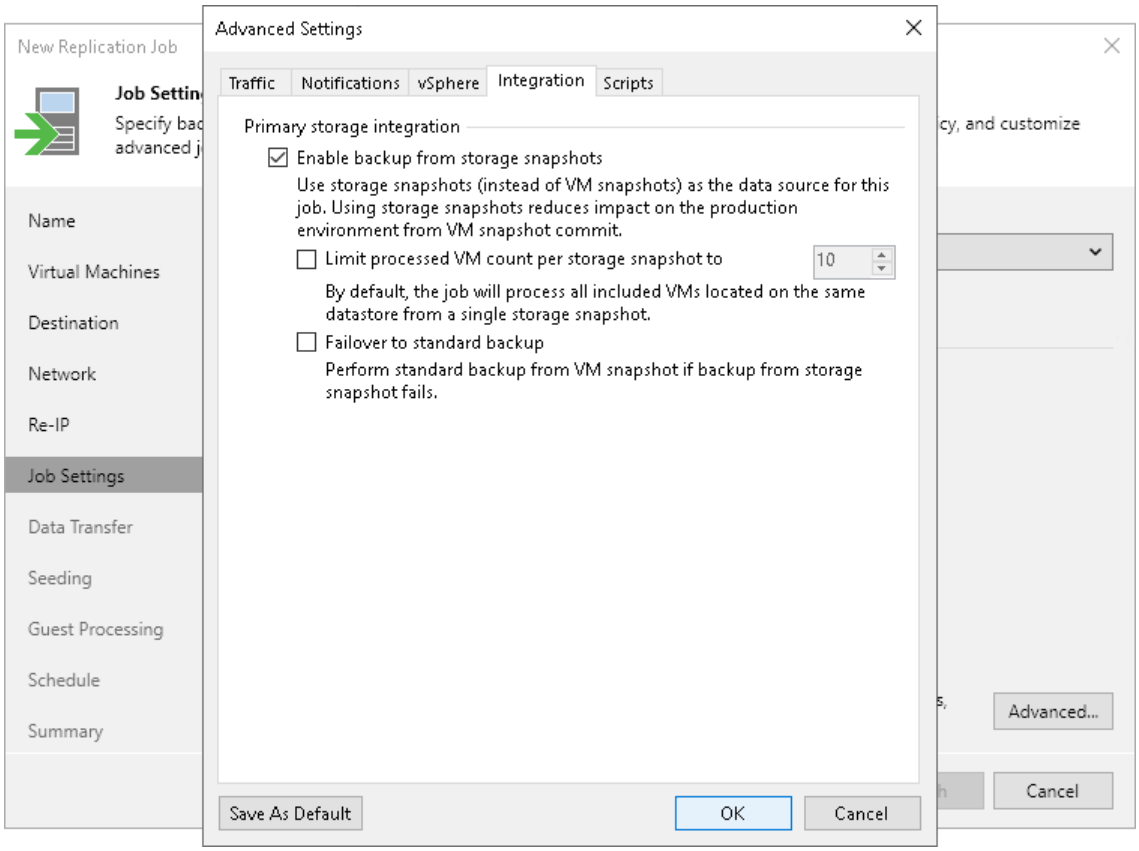


# Script Settings

To specify script settings for the replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. In the **Advanced Settings** window, click the **Scripts** tab.
3. If you want to execute custom scripts before or after the replication job, select the **Run the following script before the job** or **Run the following script after the job** check boxes. Click **Browse** to choose executable files from a local folder on the backup server. The scripts will be executed on the backup server.

4. Configure when to execute pre- and post-replication scripts:
    - To execute scripts after a number of job sessions, select the **Run scripts every... backup session** option and then specify the number of the replication job sessions.
    - To execute scripts on specific week days, select the **Run scripts on selected days only** option. Click **Days** and specify week days on which scripts must be executed.

**NOTE**

Consider the following:

- Custom scripts that you define in the advanced job settings relate to the job itself, not the VM quiescence process. To add pre-freeze and post-thaw scripts for VM image quiescence, use the **Guest Processing** step of the wizard.
- If you select the **Run scripts on the selected days only** option, Veeam Backup & Replication executes scripts only once on each selected day – when the job runs for the first time. During subsequent job runs, scripts are not executed.
- To run the script, Veeam Backup & Replication uses the Service Account under which the Veeam Backup Service is running.

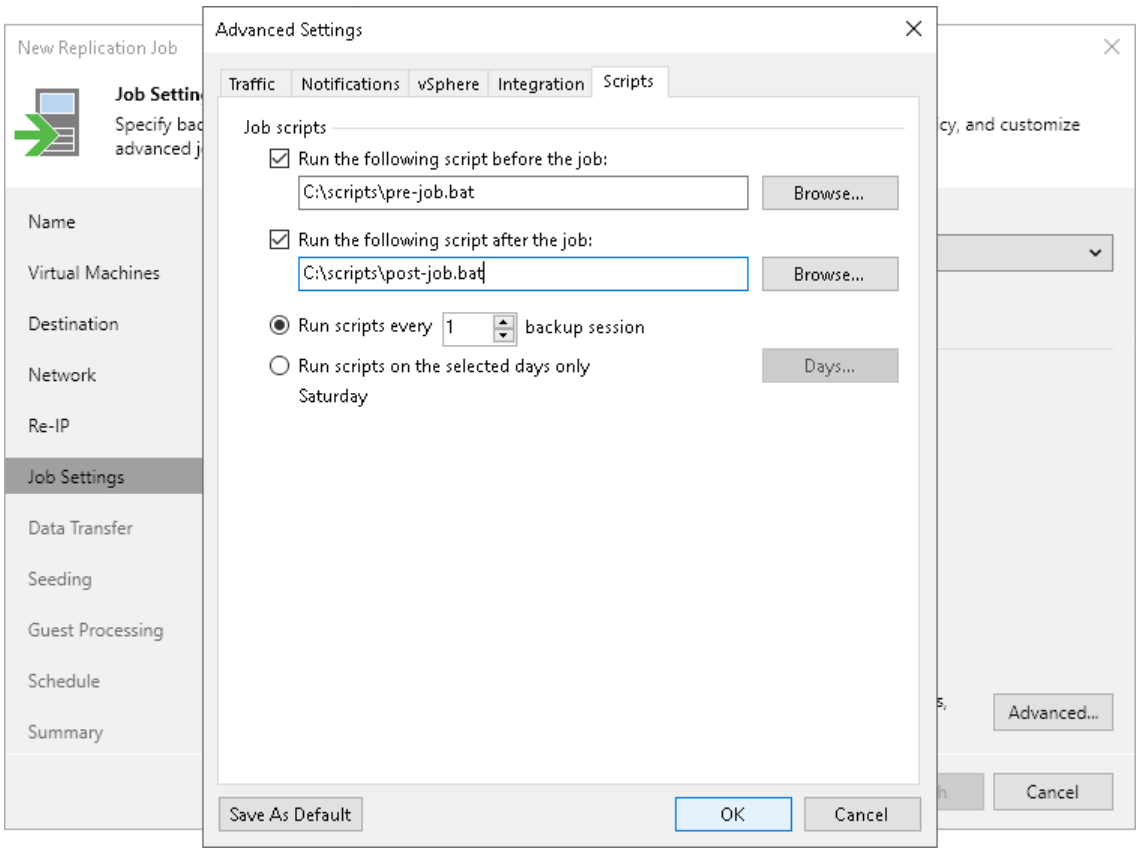

# Step 12. Specify Data Transfer Settings

At the **Data Transfer** step of the wizard, select backup infrastructure components that must be used for the replication process and choose a path for VM data transfer:

1. Specify which backup proxies you want to use:
   - If you want Veeam Backup & Replication to select proxies automatically, leave **Automatic selection** in the **Source proxy** and **Target proxy** fields.

     Veeam Backup & Replication will assign backup proxies for VM processing one by one. Before processing a new VM from the list, Veeam Backup & Replication will check available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use and the current workload on the backup proxies to select the most appropriate backup proxy for VM processing.
   - If you want to select backup proxies manually, do the following:
     i. Click **Choose** next to the **Source proxy** field if you want to select backup proxies in the production site, or next to the **Target proxy** field if you want to select backup proxies in the disaster recovery site.
     ii. In the **Backup Proxy** window, click **Use the selected backup proxy servers only**. Select proxies that you want to use and click **OK**.

> **NOTE**
>
> Consider the following:
>
> - If you plan to replicate VM data within one site, the same backup proxy can act as the source and target backup proxy. For off-site replication, you must deploy at least one backup proxy in each site to establish a stable connection for VM data transfer across sites.
> - We recommend you to select at least two backup proxies to ensure that the job will be performed if one of backup proxies fails or loses its connectivity to the source datastore.

2. Select a path for VM data transfer:
   - To transport VM data directly using backup proxies to the target datastore, select **Direct**.
   - To transport VM data using WAN accelerators, select **Through built-in WAN accelerators**. From the **Source WAN accelerator** list, select the WAN accelerator configured in the source site. From the **Target WAN accelerator** list, select the WAN accelerator configured in the target site.

> **NOTE**
>
> You should not assign one source WAN accelerator to several replication jobs that you plan to run simultaneously. The source WAN accelerator requires a lot of CPU and RAM resources, and does not process multiple replication tasks in parallel. As an alternative, you can create one replication job for all VMs you plan to process over one source WAN accelerator.
>
> The target WAN accelerator, however, can be assigned to several replication jobs. For more information, see Adding WAN Accelerators.

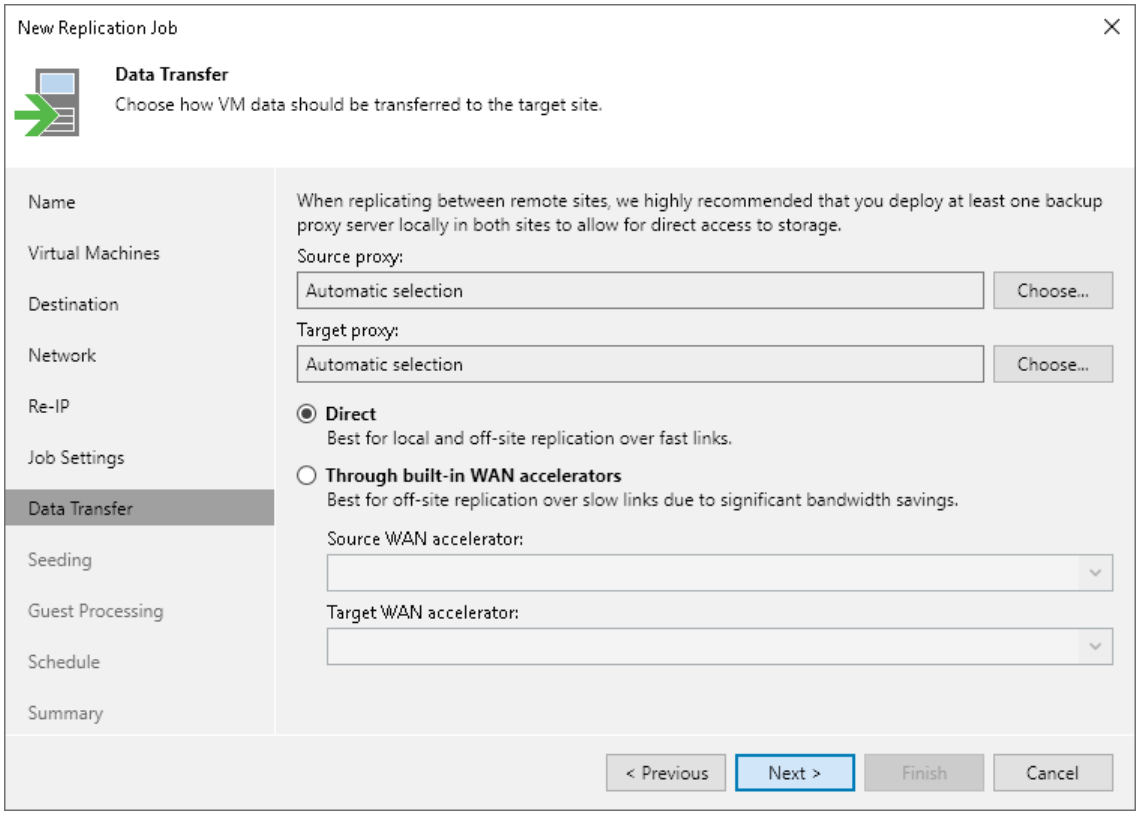
New Replication Job
Data Transfer
Choose how VM data should be transferred to the target site.
Name
Virtual Machines
Destination
Network
Re-IP
Job Settings
Data Transfer
Seeding
Guest Processing
Schedule
Summary
When replicating between remote sites, we highly recommended that you deploy at least one backup proxy server locally in both sites to allow for direct access to storage.
Source proxy:
Automatic selection
Choose...
Target proxy:
Automatic selection
Choose...
Direct
Best for local and off-site replication over fast links.
Through built-in WAN accelerators
Best for off-site replication over slow links due to significant bandwidth savings.
Source WAN accelerator:
Target WAN accelerator:
< Previous
Next >
Finish
Cancel

# Step 13. Define Seeding and Mapping Settings

The **Seeding** step is available if you have selected the **Replica seeding** check box at the Name step of the wizard.

At the **Seeding** step of the wizard, configure replica seeding and mapping. Seeding and mapping help reduce the amount of traffic sent during the initial replica synchronization. For more information on when to use seeding and mapping, see Replica Seeding and Mapping.

If you use replica seeding or mapping, make sure that you select correct backup infrastructure components for the job: source-side backup repository for metadata and backup proxies. It is recommended that you explicitly assign backup proxies in the production site and disaster recovery (DR) site. For more information, see Specify Data Transfer Settings.

> **IMPORTANT**
>
> If the **Replica seeding** check box is enabled in a replication job, all VMs in the job must be covered with seeding or mapping. If a VM is neither has a seed, nor is mapped to an existing VM, it will be skipped from processing.

## Configuring Replica Seeding

To configure replica seeding:

1. Make sure that you have backups of replicated VMs in a backup repository in the DR site. If you do not have the backups, create them as described in section Creating Replica Seeds for CDP.

   > **IMPORTANT**
   >
   > Consider the following:
   >
   > - Backups must be created by Veeam Backup & Replication.
   > - Backups must not reside in a scale-out backup repository.

2. In the **Initial seeding** section, select the **Get seed from the following backup repository** check box.
3. From the list of available backup repositories, select the repository where your replica seeds are stored.

> **NOTE**
>
> If a VM has a seed and is mapped to an existing replica, replication will be performed using replica mapping because mapping has a higher priority.

## Configuring Replica Mapping

To configure replica mapping:

1. Select the **Map replicas to existing VMs** check box.
2. If you want Veeam Backup & Replication to scan the DR site to detect existing copies of VMs that you plan to replicate, click **Detect**.

   If any matches are found, Veeam Backup & Replication will populate the mapping table. If Veeam Backup & Replication does not find a match, you can map a VM to its copy manually.

3. If you want to map a VM manually, select a source VM from the list, click **Edit** and select the copy of this VM on the target host in the DR site.

If there is no existing VM replica in the DR site, you can restore a VM from the backup and map it to the source VM.

To remove a mapping association, select a VM in the list and click **Remove**.

> **NOTE**
>
> The mapping list does not display VMs added to the list of exclusions. For more information, see Exclude Objects from Replication Job.

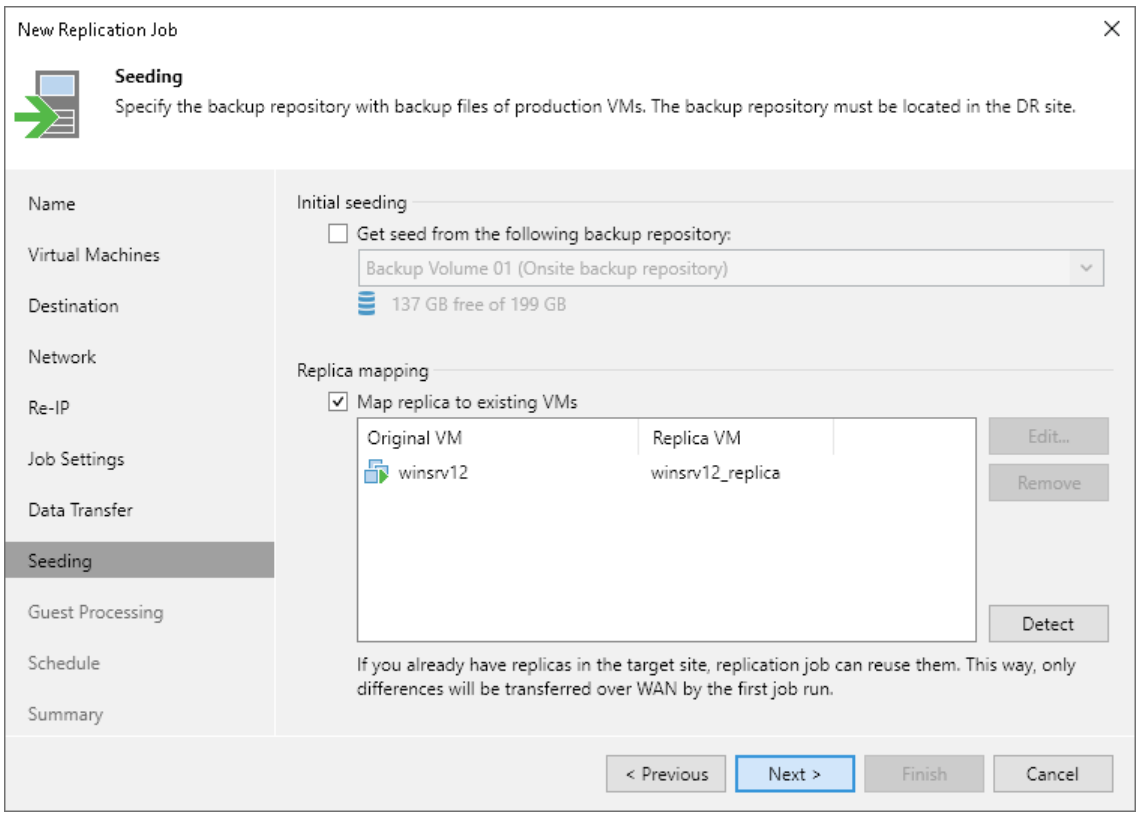

# Step 14. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, enable and configure guest OS processing.

Guest OS processing involves application-aware processing that allows creation of transactionally consistent replicas and guest file system indexing (however, indexing is not available for replicas). In its turn, application-aware processing includes log truncation, execution of custom scripts and guest OS file exclusions. For more information on guest processing, see the Guest Processing section.

To be able to use guest processing, you must also configure user accounts to access guest OSes and guest interaction proxies.

To enable guest OS processing and start configuring it (accounts and guest interaction proxies):

1. Select **Enable application-aware processing**.

   When you select this option, Veeam Backup & Replication enables application-aware processing with the default settings for all VMs. You can further disable application-aware processing for individual VMs and reconfigure the default settings.

2. If you have added Microsoft Windows VMs to be processed, specify which guest interaction proxy Veeam Backup & Replication can use to perform different guest processing tasks:

   - If you want Veeam Backup & Replication to select the guest interaction proxy automatically, leave **Automatic selection** on the **Guest interaction proxy** field.
   - If you want to explicitly specify which servers will perform the guest interaction proxy role, click **Choose**. In the **Guest Interaction Proxy** window, click **Prefer the following guest interaction proxy server**, and select the necessary proxies.

   For more information on the guest interaction proxy, requirements and limitations for it, see Guest Interaction Proxies.

3. From the **Guest OS credentials** list, select a user account that will be used to connect to guest OSes and that has enough permissions. For more information on the permissions and requirements for the user account, see Permissions for Guest Processing.

   [For Microsoft Windows VMs] Veeam Backup & Replication will also use this account to deploy the non-persistent runtime components or use (if necessary, deploy) persistent agent. For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

   [For Linux VMs] If you installed persistent agent components for VMs running Linux or Unix operating systems, select *Use management agent credentials* from the list. For more information, see Persistent Agent Components.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. For more information on adding credentials, see the Credentials Manager section.

   > **NOTE**
   >
   > If you plan to use Kerberos authentication, check limitations and requirements listed in section Guest Processing.

4. To specify credentials for individual workloads, click **Credentials**. Then select the necessary workload and set user credentials for it.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose a VM whose settings you want to customize. Then select the VM in the list and define the necessary settings.

5. To check whether Veeam Backup & Replication can connect to VMs using the specified guest OS credentials and can deploy the non-persistent runtime components or connect to persistent agent components on the guest OSes, click **Test Now**.

After you have enabled application-aware processing for all VMs and configured other settings required for guest processing, you can disable application-aware processing for individual VMs and change the default settings. For more information, see the following sections:

- Application-aware processing general settings
- Microsoft SQL Server transaction log settings
- Oracle archived log settings
- PostgreSQL settings
- VM guest OS file exclusion
- Pre-freeze and post-thaw scripts

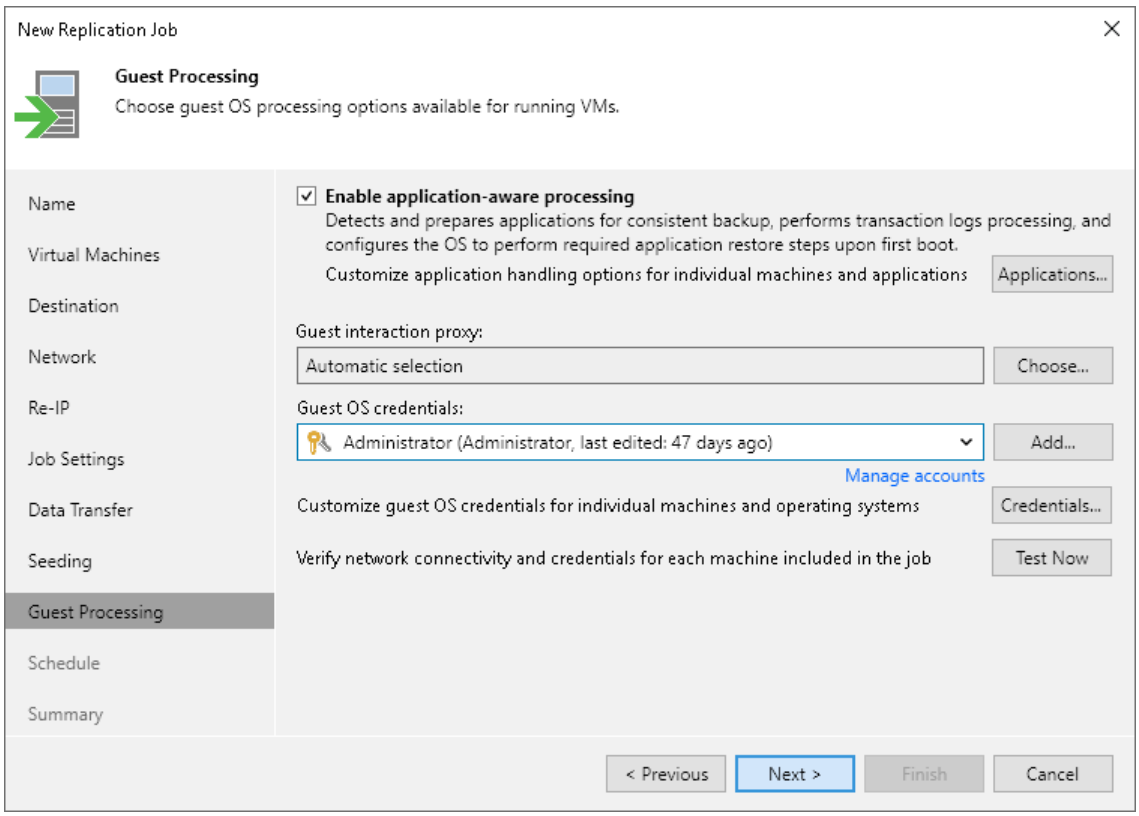


# Application-Aware Processing and Transaction Logs

Application-aware processing helps create transactionally consistent replicas. The transactionally consistent replicas guarantee proper recovery of applications without data loss. For more information on application-aware processing, see Application-Aware Processing.

To configure general application-aware processing settings and specify whether Veeam Backup & Replication processes transaction logs or creates copy-only replicas:

1. At the **Guest Processing** step of the wizard, check that you have selected the **Enable application-aware processing** check box and configured guest OS credentials.
2. At the **Guest Processing** step of the wizard, click **Applications**.

3. In the **Application-Aware Processing Options** window, select workloads for which you want to configure application-aware processing, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.

4. On the **General** tab, in the **Applications** section, specify the behavior scenario for application-aware processing:

   - Select **Require successful processing** if you want Veeam Backup & Replication to stop the replication process if any error occurs during application-aware processing.
   - Select **Try application processing, but ignore failures** if you want to continue the replication process even if an error occurs during application-aware processing. This option guarantees that replication will continue working. However, the resulting replica will be crash consistent, not transactionally consistent.
   - Select **Disable application processing** if you want to disable application-aware processing for the workload.

5. [For Microsoft Exchange and Microsoft SQL Server] In the **VSS Settings** section, specify if Veeam Backup & Replication must process transaction logs or create copy-only replicas:

   a. Select **Process transaction logs with this job** if you want Veeam Backup & Replication to process transaction logs.

      [For Microsoft Exchange VMs] With this option selected, the non-persistent runtime components or persistent components running on the VM guest OS will wait for replication to complete successfully and then trigger truncation of transaction logs. If the replication job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

      [For Microsoft SQL Server] You will need to configure how to process transaction logs.

      > **TIP**
      >
      > To configure log processing for Oracle and PostgreSQL databases, switch to the Oracle and PostrgeSQL tabs.

   b. Select **Perform copy only** if you use another tool to perform guest level processing, and this tool maintains consistency of the database state. Veeam Backup & Replication will create a copy-only replica for the selected VMs. The copy only replica preserves the chain of full and differential files and transaction logs on the VM. For more information, see Microsoft Docs.

6. [For Microsoft Windows VMs] In the **Persistent guest agent** section, select the **Use persistent guest agent** check box to use for application-aware processing persistent guest agents on each protected VM.

   By default, Veeam Backup & Replication uses non-persistent runtime components.
   Veeam Backup & Replication deploys runtime components on each protected VM when the replication job starts, and removes the runtime components as soon as the replication job finishes.

   For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

[For Linux VMs] To use persistent guest agents, you must install Management Agent on protected VMs. For more information, see Persistent Agent Components.

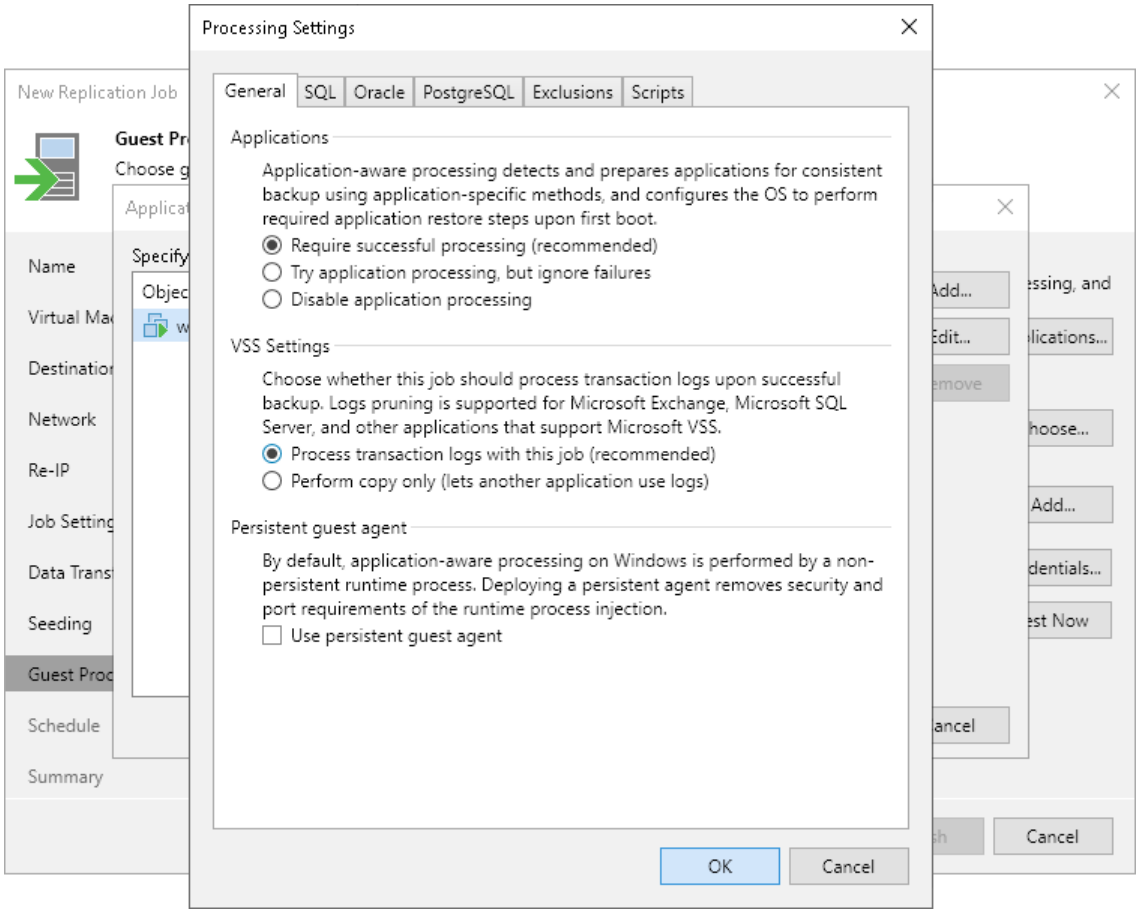


# Microsoft SQL Server Transaction Log Settings

The **SQL** tab is available for VMs that run Microsoft SQL Server and if you have selected **Process transaction logs with this job** when configuring application-aware processing.

To create transactionally consistent backups of an Microsoft SQL Servers, you must check that application-aware processing is enabled and then specify settings of transaction log processing.

# Enabling Application-Aware Processing

Before configuring transaction log processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Microsoft SQL Sever and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

# Specifying Transaction Log Settings

In the **Processing Settings** window, switch to the **SQL** tab and specify how transaction logs must be processed:

- If you want Veeam Backup & Replication to trigger truncation of transaction logs only after the job completes successfully, select **Truncate logs**.

  In this case, the non-persistent runtime components or persistent components will wait for replication to complete and then trigger truncation of transaction logs. If the replication job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

  - If you do not want Veeam Backup & Replication to truncate logs at all, select **Do not truncate logs**.

  This option is recommended if you are using another backup tool to perform VM guest-level backup or replication, and this tool maintains consistency of the database state. In such scenario, Veeam Backup & Replication will not trigger transaction log truncation. After you fail over to the necessary restore point of the VM replica, you will be able to apply transaction logs to get the database system to the necessary point in time between replication job sessions.

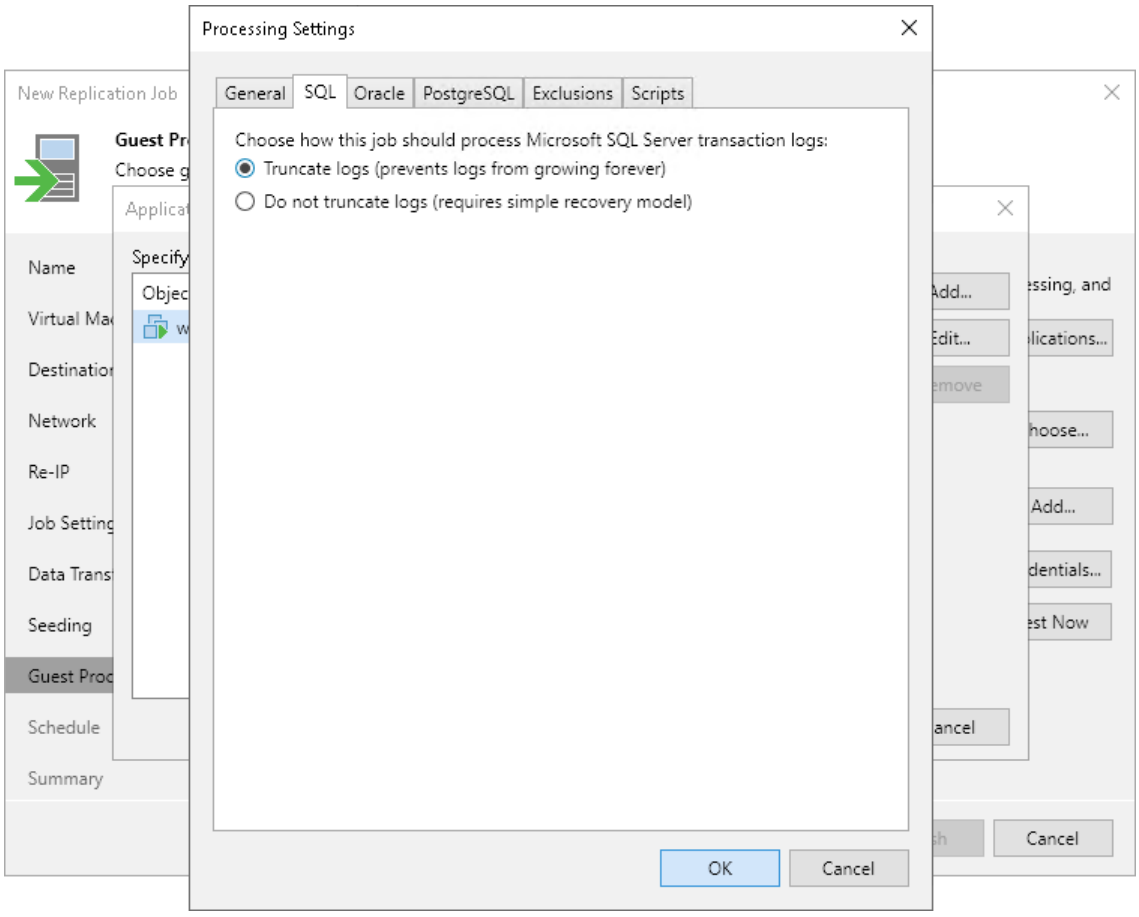


# Oracle Archived Log Settings

The **Oracle** tab applies to Linux VMs that run Oracle.

To create transactionally consistent backups of an Oracle server, you must check that application-aware processing is enabled and then specify settings of archive log processing.

# Enabling Application-Aware Processing

Before configuring archive log processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Oracle server and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

# Specifying Archive Log Settings

To configure how Veeam Backup & Replication must process archive logs of an Oracle server:

1. In the **Processing Settings** window, switch to the **Oracle** tab.
2. From the **Specify Oracle account with SYSDBA privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the Oracle databases. The account that you plan to use must have privileges described in section Permissions.

   You can select **Use guest credentials** from the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle databases.
3. In the **Archived logs** section, specify how to process archived logs:
   - If you want to preserve archived logs on the VM guest OS, select **Do not delete archived logs**. When the replication job completes, the non-persistent runtime components or persistent components will not truncate transaction logs.

     It is recommended that you select this option for databases where the ARCHIVELOG mode is turned off. If the ARCHIVELOG mode is turned on, archived logs may grow large and consume all disk space.
   - If you want to delete archived logs older than <N> hours, select **Delete logs older than <N> hours** and specify the number of hours.

- If you want to delete archived logs larger than <N> GB, select **Delete logs over <N> GB** and specify the size. The specified size refers to the log size of each database, not all databases on the selected Oracle server.

The non-persistent runtime components or persistent components running on the VM guest OS will wait for the replication job to complete successfully and then trigger transaction logs truncation using Oracle Call Interface (OCI). If the job does not manage to replicate the Oracle VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

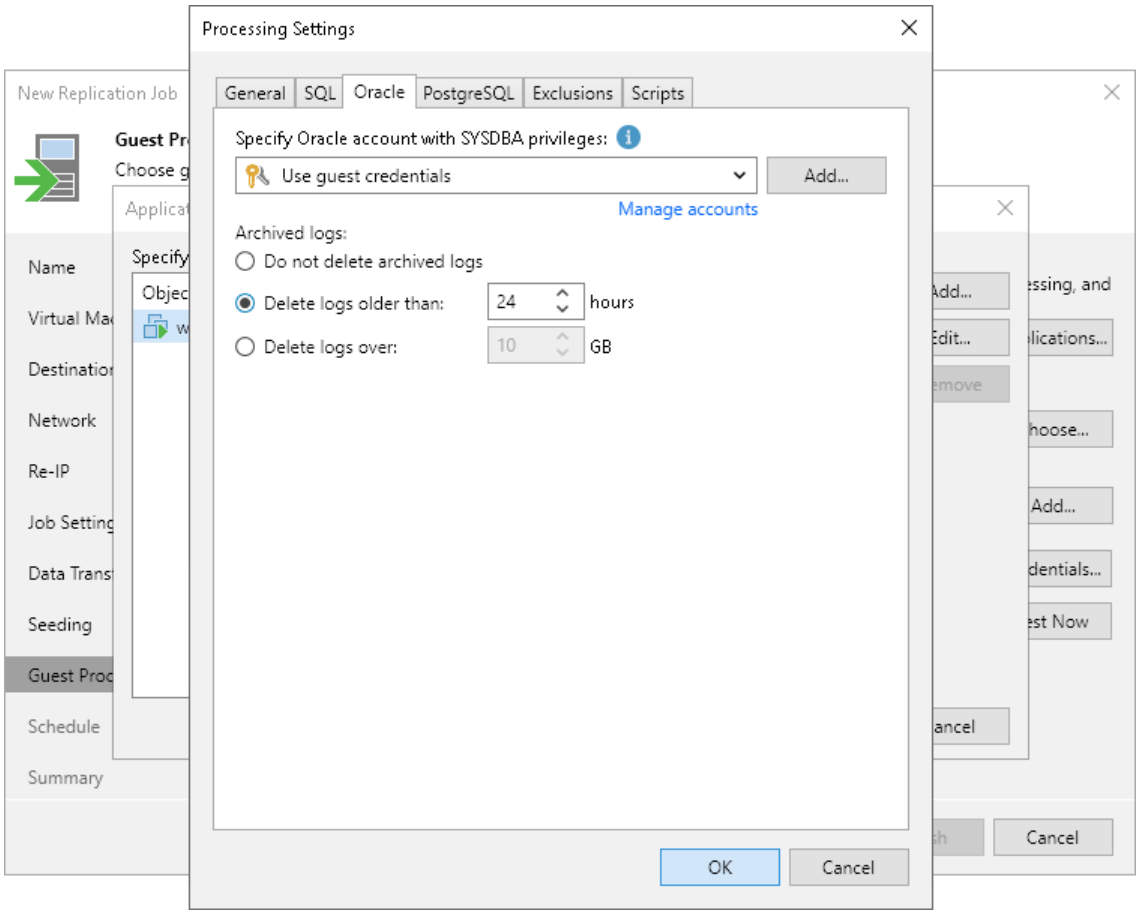


# PostgreSQL Settings

The **PostgreSQL** tab applies to Linux VMs that run PostgreSQL.

To create transactionally consistent backups of a PostgreSQL VM, you must check that application-aware processing is enabled and then specify settings of WAL files processing.

# Enabling Application-Aware Processing

Before configuring WAL files processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.

3. In the displayed list, select the PostgreSQL VM and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.

4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

**NOTE**

By default, Veeam Backup & Replication recursively scans the `/etc/postgresql`, `/var/lib/postgresql` and `/var/lib/pgsql` directories for the configuration files of PostgreSQL instances. If you keep configuration files in custom directories, the *pgsqlagent agent* will use its own `VeeamPostgreSQLAgent.xml` configuration file that is located in the `/etc/veeam/` directory. The pgsqlagent agent configuration file must be a single line XML.

To explicitly include or exclude specific configuration files from rescan, you can add the following commands to the `VeeamPostgreSQLAgent.xml` file:

- `ExcludeConfigDirs` — use this element to exclude configuration files.
- `AddConfigDirs` — use this element to include configuration files.

For example: `<config AddConfigDirs="/opt/psql/" ExcludeConfigDirs="/var/lib/postgresql/13/main45/,/var/lib/postgresql/13/maindd/" />`

# Specifying WAL Files Settings

To define how Veeam Backup & Replication will process WAL files on this VM, do the following:

1. In the **Processing Settings** window, click the **PostgreSQL** tab.
2. From the **Specify PostgreSQL account with superuser privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the PostgreSQL instance. The account must have privileges described in section Permissions. You can select **Use guest credentials** from the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the PostgreSQL instance.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. Note that if you select the **System user without password file (peer)** option in the **The specified user is** area, you can add a user account without specifying a password.

3. In the **The specified user is** section, specify how the user selected in the **Specify PostgreSQL account with superuser privileges** drop-down list will authenticate against the PostgreSQL instance:
   - Select **Database user with password** if the account is a PostgreSQL account, and you entered the password for this account in the Credentials Manager.

> **NOTE**
>
> If you want Veeam Backup & Replication to use the user name map authentication, select **Database user with password** and leave the password field empty. Consider the following:
>
> - **Guest OS credentials** specified at the **Guest Processing** step of the wizard will be used as the System-Username.
> - PostgreSQL account specified at the step 2 will be used as the PG-Username.
>
> For more information about the user name maps, see PostgreSQL documentation.

- Select **Database user with password file (.pgpass)** if the password for the account is defined in the `.pgpass` configuration file on the PostgreSQL server. For more information about the password file, see PostgreSQL documentation.
- Select **System user without password file (peer)** if you want Veeam Backup & Replication to use the peer authentication method. In this case, Veeam Backup & Replication will apply the OS account as the PostgreSQL account.

> **IMPORTANT**
>
> If you have added a new PostgreSQL account and want to use it with the peer authentication method, make sure that you have added this account as a Linux user with sufficient permissions.

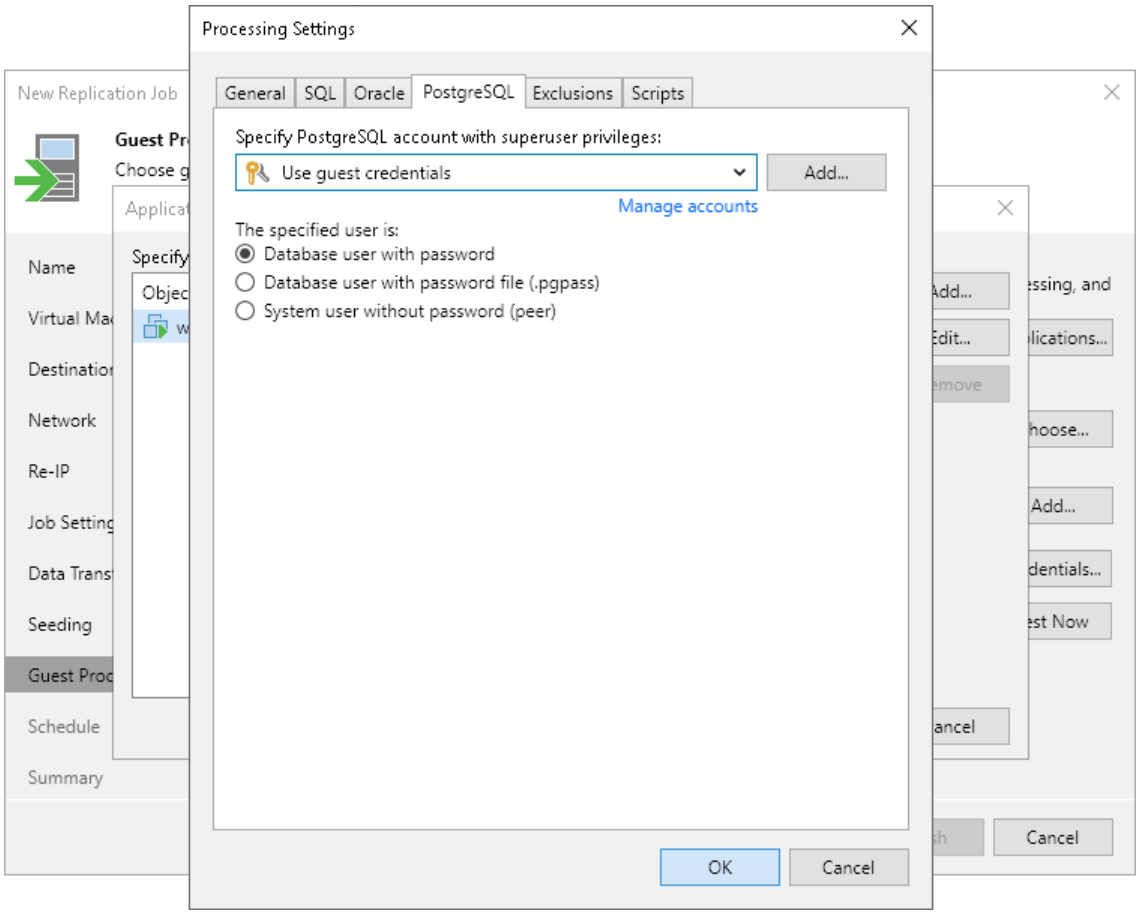


# VM Guest OS File Exclusion Settings

These settings apply only to Microsoft Windows workloads.

To exclude guest OS files and folders from being replicated:

1. At the **Guest Processing** step of the wizard, check that you have selected the **Enable application-aware processing** check box and configured guest OS credentials.
2. In the **Application-Aware Processing Options** list, select workloads for which you want to exclude files and folders and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
3. In the **Processing Settings** window, switch to the **Exclusions** tab and specify whether you want to exclude or include files and folders:
   - To remove individual files and folders from replicas, select **Exclude the following files and folders** and click **Add**.
   - To include only the specified files and folders in replicas, select **Include only the following files and folders** and click **Add**.
4. In the **Specify Folder** window, specify which files and folders you want to include or exclude. For the methods that you can use to specify the list of exclusions or inclusions, see VM Guest OS Files.

> **NOTE**
>
> When you select files to be included or excluded, consider requirements and limitations that are listed in section Requirements and Limitations for VM Guest OS File Exclusion.

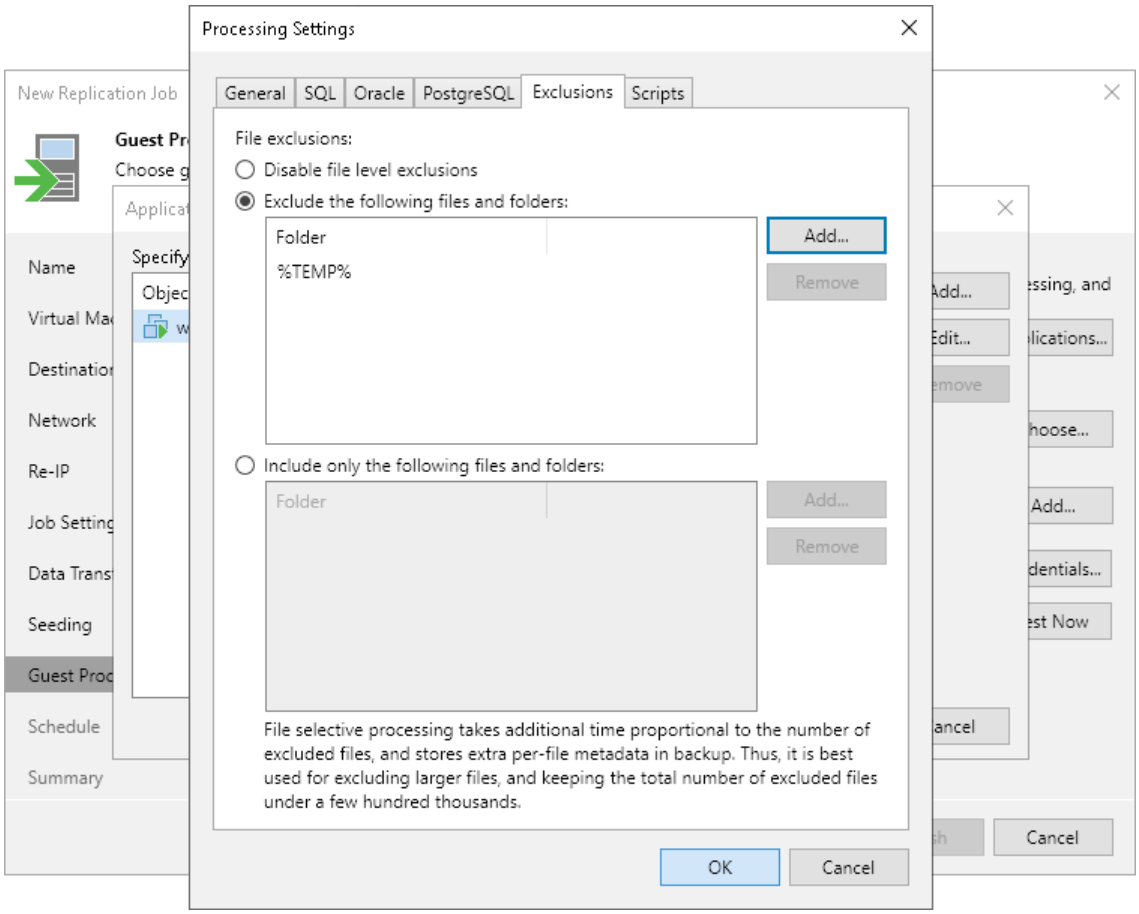

# Pre-Freeze and Post-Thaw Script Settings

If you plan to replicate VMs running applications that do not support VSS, you can instruct Veeam Backup & Replication to run custom pre-freeze and post-thaw scripts for these VMs. The pre-freeze script quiesces the VM file system and application data to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script brings the VM and applications to their initial state.

To specify pre-freeze and post-thaw scripts:

1. At the **Guest Processing** step of the wizard, check that you have selected the **Enable application-aware processing** check box and configured guest OS credentials.
2. At the **Guest Processing** step of the wizard, click **Applications**.
3. In the **Application-Aware Processing Options** list, select workloads for which you want to configure scripts, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. Click the **Scripts** tab.
5. In the **Script processing mode** section, select a scenario for script execution:
   - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the replication process if scripts fail.
   - Select **Ignore script execution failures** if you want to continue the replication process even if script errors occur.
   - Select **Disable script execution** if you do not want to run scripts for the VM.
6. In the **Windows scripts** section, specify paths to scripts for Microsoft Windows VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.
7. In the **Linux scripts** section, specify paths to scripts for Linux VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

   If you plan to replicate a VM container with Microsoft Windows and Linux VMs, you can select to execute both Microsoft Windows and Linux scripts. When replication starts, Veeam Backup & Replication will automatically determine which OS type is installed on the VM and use the correct scripts for this VM.

> **TIP**
>
> Beside pre-freeze and post-thaw scripts for VM quiescence, you can instruct Veeam Backup & Replication to run custom scripts before the job starts and after the job completes. For more information, see Script Settings.

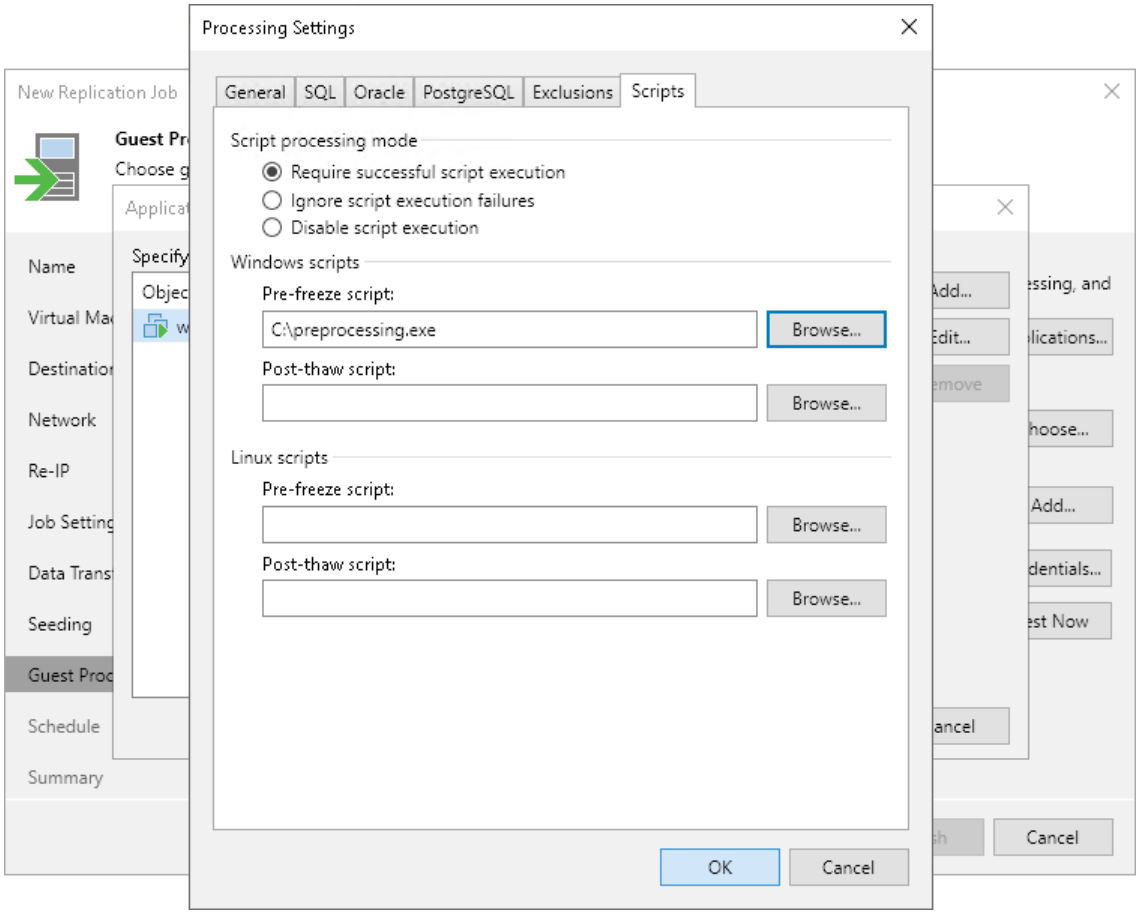
Processing Settings
General
SQL
Oracle
PostgreSQL
Exclusions
Scripts
Script processing mode
Require successful script execution
Ignore script execution failures
Disable script execution
Windows scripts
Pre-freeze script:
C:\preprocessing.exe
Browse...
Post-thaw script:
Browse...
Linux scripts
Pre-freeze script:
Browse...
Post-thaw script:
Browse...
OK
Cancel
New Replication Job
Name
Virtual Ma
Destinatio
Network
Re-IP
Job Setting
Data Trans
Seeding
Guest Proc
Schedule
Summary
Cancel

# Step 15. Define Job Schedule

At the **Schedule** step of the wizard, select to run the replication job manually or schedule the job to run on a regular basis:

1. To run the replication job automatically, select the **Run the job automatically** check box. If you do not select this check box, you will have to start the job manually.
2. Define scheduling settings for the job:
   - To run the job at specific time daily, on weekdays or with specific periodicity, select **Daily at this time**. In the fields on the right of the check box, specify the time and required days.
   - To run the job once a month on specific days, select **Monthly at this time**. In the fields on the right of the check box, specify the necessary days.

     > **NOTE**
     >
     > When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a set time interval, do the following:
     i. Select **Periodically every**. In the field on the right, select the necessary time unit: *Hours* or *Minutes.*
     ii. If you want to specify the permitted time window for the job, click **Schedule**. In the **Time Periods** window, specify the schedule.

     If you want to shift the schedule, specify the offset in the **Start time within an hour** field. For example, you schedule the prohibited hours from 08:00 AM to 10:00 AM, and set the offset value to 25. The schedule will be shifted forward, and the prohibited hours will be from 8:00 AM and to 10:25 AM.

     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run the job by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.

   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.
   - To chain jobs, use the **After this job** field. For the first job in the chain, you must specify an automatic time schedule, otherwise the chained jobs will not be started. For the chained jobs, select the **After this job option** and choose the preceding job from the list. For more information on job chaining and recommendations for it, see Chained Jobs.

3. In the **Automatic retry** section, select the **Retry failed VMs processing** if Veeam Backup & Replication must attempt to run the job again for VMs whose processing failed for some reason. Enter the number of attempts to run the job and define time spans between them.

   If you select continuous schedule for the job, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job sessions.

4. In the **Backup window** section, specify a time interval within which the job must be completed. The backup window prevents the job from overlapping with production hours and ensures the job does not provide unwanted overhead on your production environment.

   To set up a backup window for the job:

   a. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**.

   b. In the **Time Periods** window, define the allowed hours and prohibited hours for VM replication. If the job exceeds the allowed window, the job will be automatically terminated.

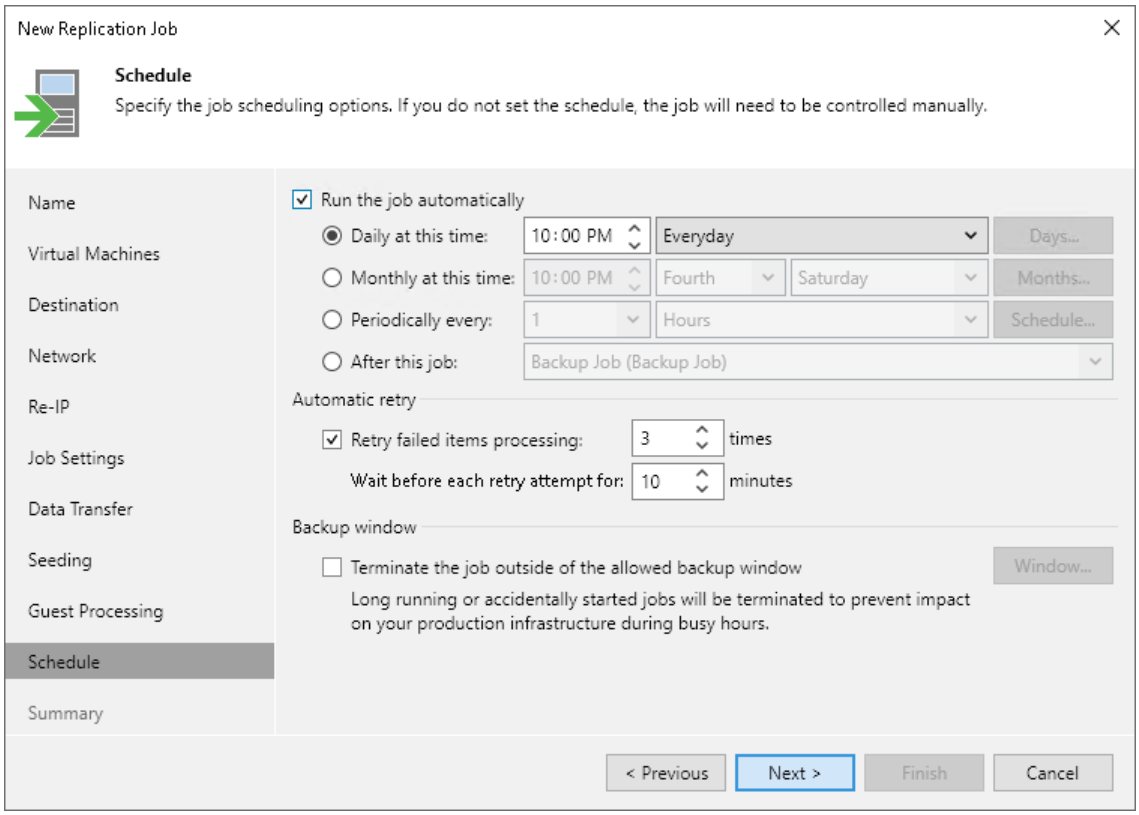

# Step 16. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the replication job. If you want to start the job right after you close the wizard, select the **Run the job when I click Finish** check box, otherwise leave the check box unselected. Then click **Finish** to close the wizard.

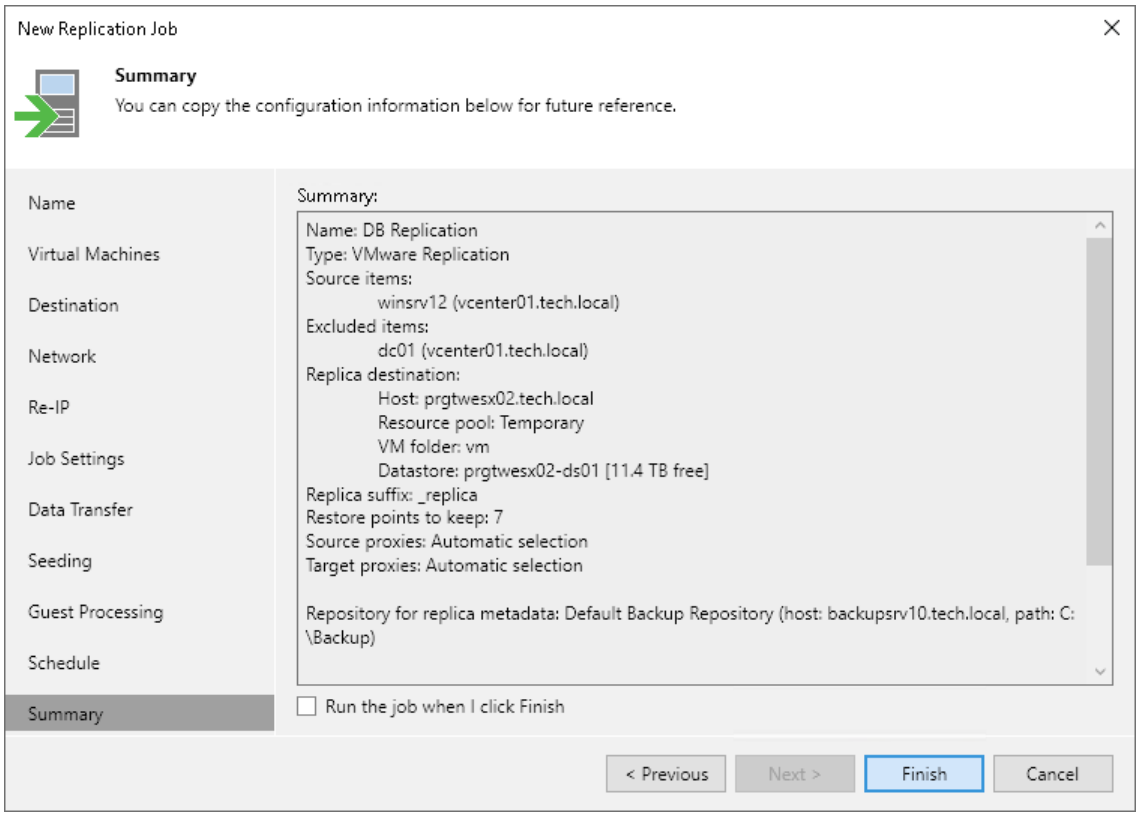

# Creating Replica Seeds

To use replica seeding in a replication job, you must have backups of replicated VMs in a backup repository in the disaster recovery (DR) site. These backups are known as replica seeds. For more information on seeding and when to use it, see Replica Seeding and Mapping.

If you do not have replica seeds in the DR site, do the following:

1. Create a backup of VMs that you plan to replicate as described in section Creating Backup Jobs. As the target repository for this job, select a backup repository in the production site. Then run the job.

   If you already have backups containing the necessary VMs, there is no need to configure and run a new backup job. For seeding, you can use any existing backups created by Veeam Backup & Replication. The backup must include VBK and VBM files. If you have a full backup and a chain of forward increments, you can use VIB files together with the VBK and VBM files. In this case, Veeam Backup & Replication will restore VMs from the seed to the latest available restore point.

2. Copy the backup from the backup repository in the production site to a backup repository in the DR site.

   You can move the backup using a file copy job or any other appropriate method, for example, copy the backup to a removable storage device, ship the device to the DR site and copy backups to the backup repository in the DR site.

   If you do not have a backup repository in the DR site, you need to create the repository as described in section Backup Repositories.

   > **IMPORTANT**
   >
   > You cannot copy backups to a scale-out backup repository in the DR site.

3. After the backup is copied to the backup repository in the DR site, perform rescan of this backup repository as described in section Rescanning Backup Repositories. Otherwise, Veeam Backup & Replication will not be able to detect the copied backup.

# Managing Replicas

To view all created replicas, open the **Home** view and navigate to the **Replicas** node. The working area displays the full list of the created replicas. Here, you can view replica properties and delete replicas from the configuration database or disk.

# Viewing Replica Properties

You can view summary information about created replicas. The summary information provides the following data:

- Available restore points
- Date of restore points creation
- Data size and replica status

To view replica properties:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, right-click the necessary replica and select **Properties**. Alternatively, select **Properties** on the ribbon.

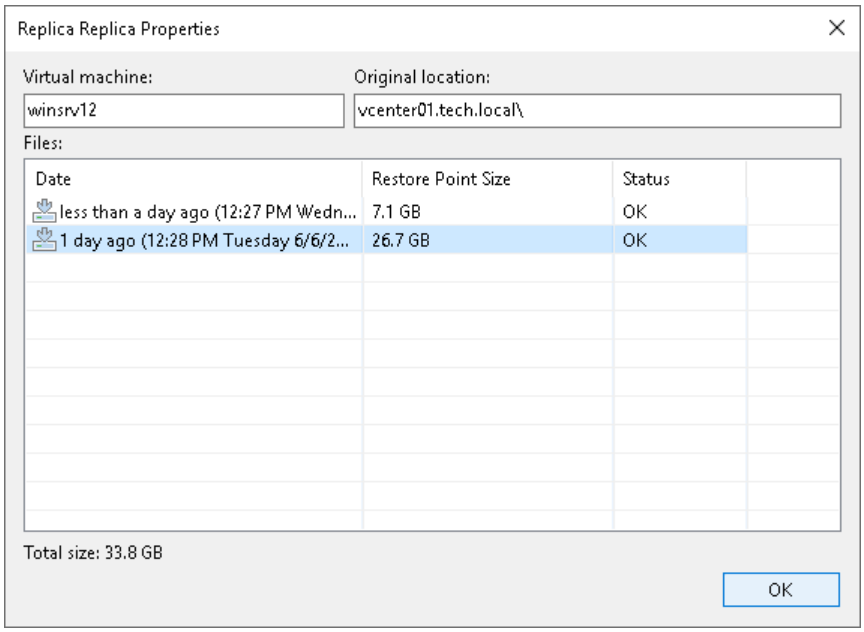

# Rescanning Replicas

You may need to perform replica rescan in the following cases:

- After you delete restore points for one or more replicas. For details, see this Veeam KB article.
- After you restore the configuration database, and the session results show that some hosts used to register replicas were unavailable during the session.

  To check whether any errors occurred during the database restore session, open the **Home** view and select **System** in the inventory pane. In the working area, right-click the **Configuration Database Resynchronize** job and select **Statistics**.

During the rescan process, Veeam Backup & Replication gathers information on replicas that are currently available in backup repositories and updates the list of replicas in the configuration database.

To rescan replicas, do the following:

1. Open the **Home** view.
2. In the inventory pane, right-click the **Replicas** node and select **Rescan Replicas**.

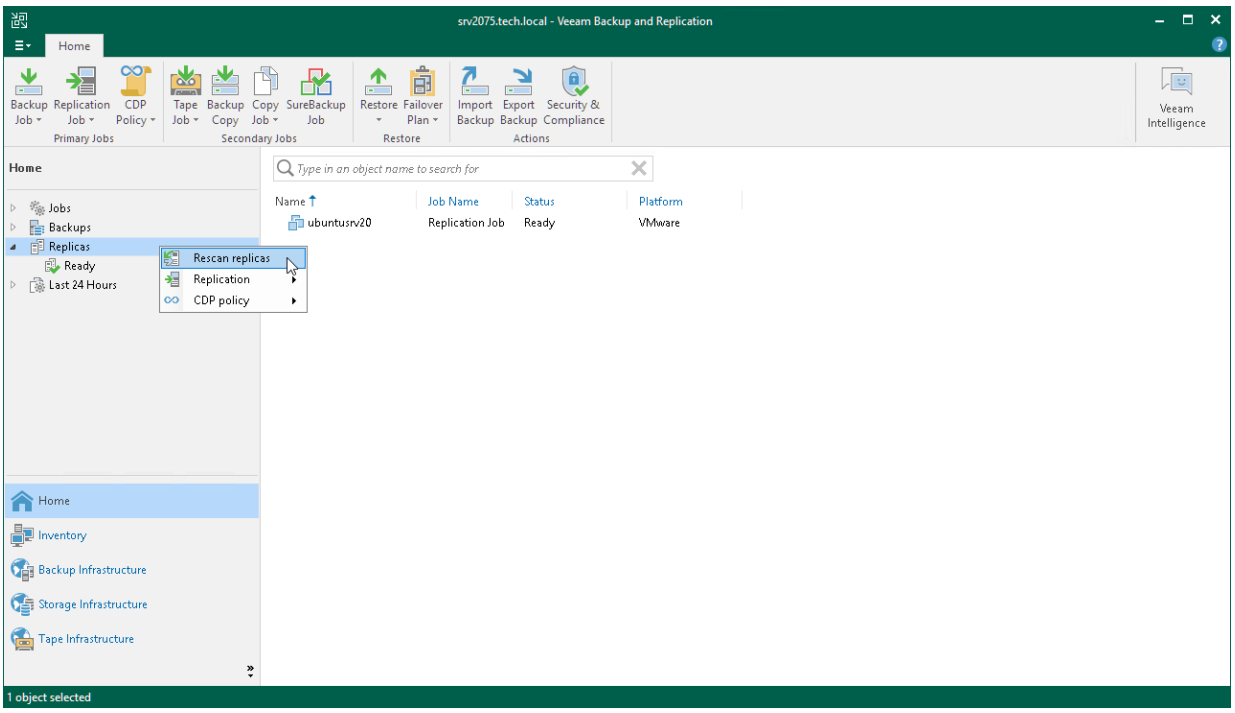

# Removing Replicas from Configuration

When you remove replicas from the configuration, Veeam Backup & Replication removes records about the replicas from the configuration database, stops showing the replicas in Veeam Backup & Replication console and also stops synchronizing their state with the state of the source VMs. However, the actual replicas remain on hosts.

To remove records about replicas from the Veeam Backup & Replication console and configuration database:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.
3. In the working area, select replicas in the *Ready* state and click **Remove from** > **Configuration** on the ribbon. Alternatively, right-click one of the selected replicas and select **Remove from configuration**.

> **NOTE**
>
> Consider the following:
>
> - The **Remove from configuration** operation can be performed only for VM replicas in the *Ready* state. If the VM replica is in the *Failover* or *Failback* state, this option is disabled.
> - When you perform the **Remove from configuration** operation for a VM that is replicated as a standalone object, Veeam Backup & Replication removes this VM from the initial replication job. When you perform the **Remove from configuration** operation for a VM that is replicated as part of a VM container (host, cluster, folder, resource pool, VirtualApp, datastore or tag), Veeam Backup & Replication adds this VM to the list of exclusions in the initial replication job. For more information, see Exclude Objects from Replication Job.

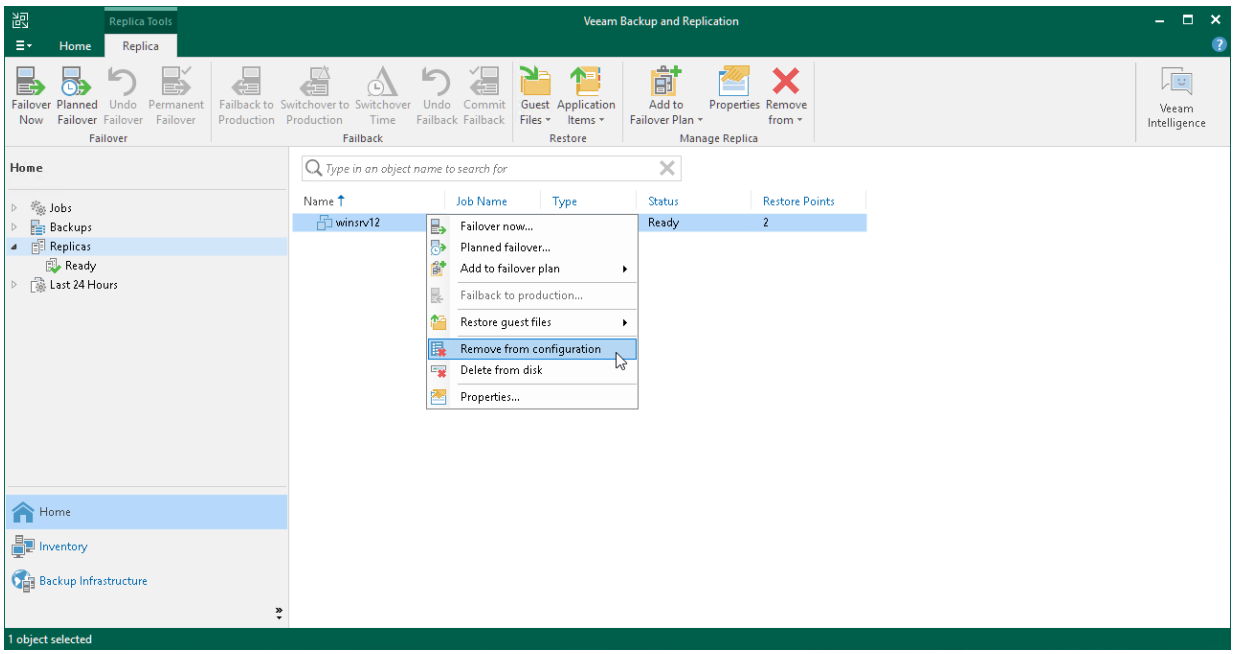

# Deleting Replicas from Disk

When you delete replicas from disks, Veeam Backup & Replication removes the replicas not only from the Veeam Backup & Replication console and configuration database, but also from host storage.

> **NOTE**
>
> Consider the following:
>
> - You can delete records only about replicas that are in the *Ready* state.
> - Do not delete replica files from the destination storage manually, use the **Delete from disk** option instead. If you delete replica files manually, subsequent replication sessions will fail.
> - Unlike the **Remove from configuration** operation, the **Delete from disk** operation does not remove the processed workload from the initial replication job. This means that the replication process will restart for this workload. To avoid this, you can exclude the workload from the replication job or disable the job.

To delete replica files from disks:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.
3. In the working area, select the necessary replica and click **Remove from** > **Disk** on the ribbon. As an alternative, right-click the replica and select **Delete from disk**.

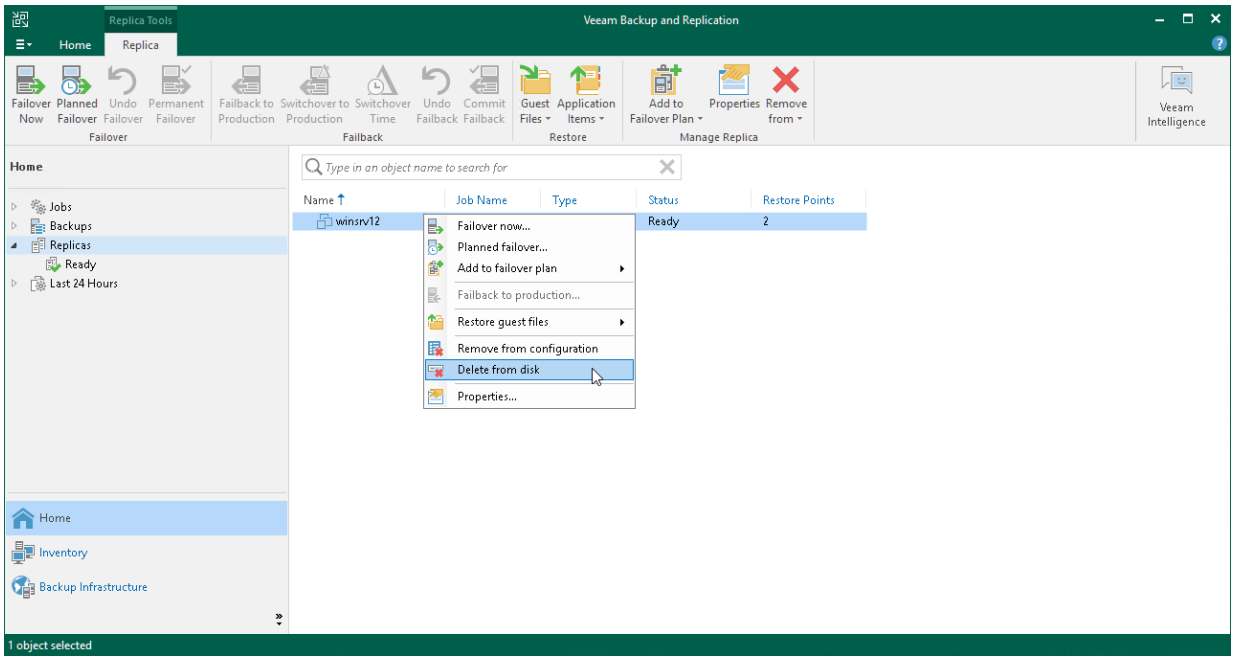

# Managing Replication Jobs

To view all created replication jobs, open the **Home** view and navigate to the **Jobs** > **Replication** node. The working area displays the full list of the created replication jobs. Here, you can manage the jobs: retry, edit, clone, disable and delete jobs.

## In This Section

- Retrying Replication Jobs
- Editing Replication Jobs
- Cloning Replication Jobs
- Disabling and Deleting Replication Jobs

# Retrying Replication Jobs

The retry option is necessary when a replication job fails and you want to retry this operation again. When you perform a retry, Veeam Backup & Replication restarts the operation only for the failed workloads added to the job and does not process VMs that have been processed successfully. As a result, the retry operation takes less time compared to running the job for all workloads.

To retry a failed replication job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary replication job and select **Retry** on the ribbon. Alternatively, you can right-click the necessary replication job and select **Retry**.

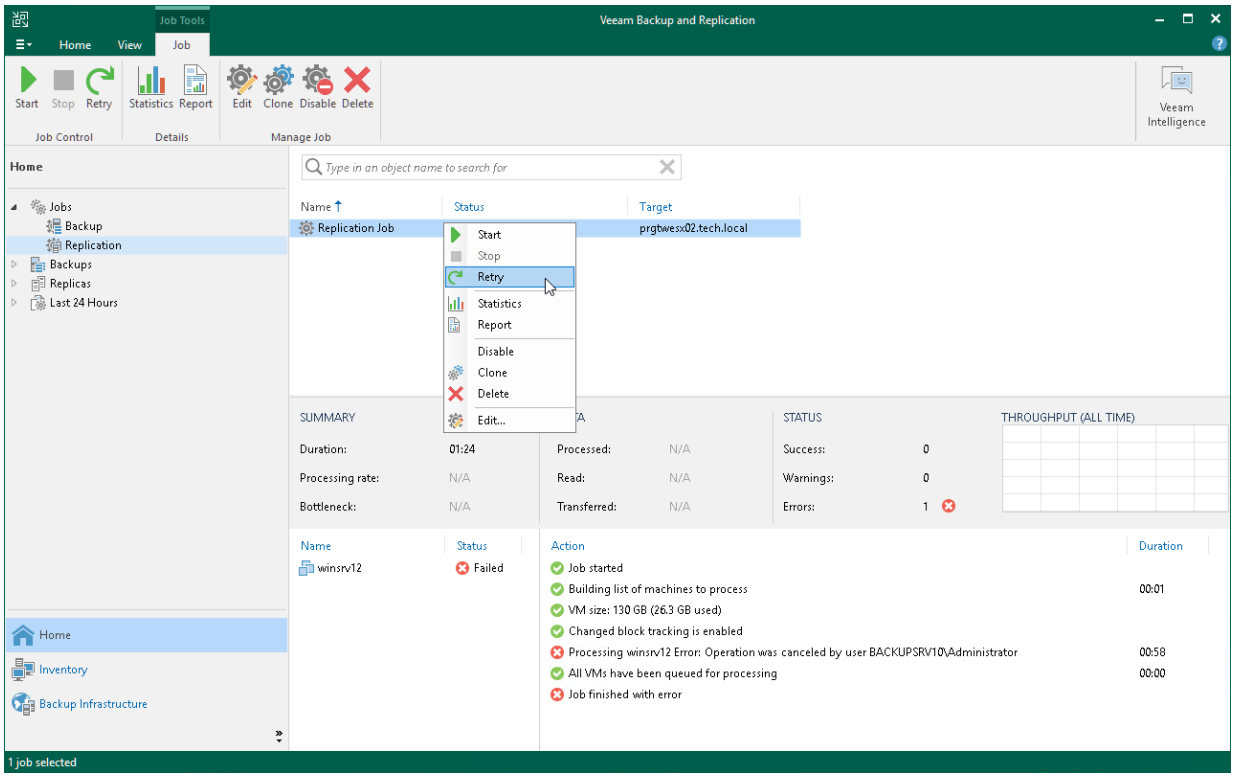

# Editing Replication Jobs

You can edit configured jobs at any moment. For example, you may want to change scheduling settings for the job or add some VMs to the job.

To edit a replication job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary replication job and select **Edit** on the ribbon. Alternatively, you can right-click the necessary replication job and select **Edit**.
4. Follow the instructions provided in Creating Replication Jobs.

> **TIP**
>
> If you are planning to change the repository where replica metadata is stored, first, follow the instructions provided in this Veeam KB article.

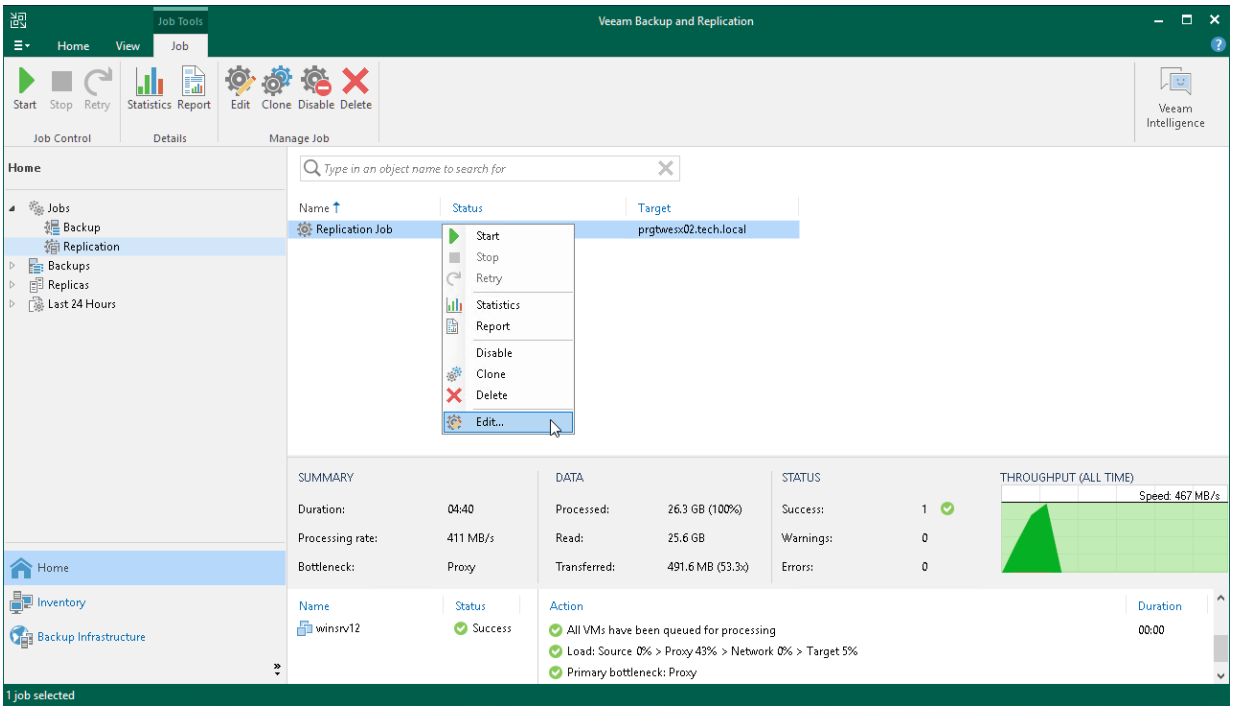

# Cloning Replication Jobs

You can create new jobs by means of job cloning. Job cloning allows you to create an exact copy of any job with the same job settings and edit settings of cloned jobs as required.

The name of the cloned job is formed by the following rule: *<job_name_clone1>*, where job_name is the name of the original job and *clone1* is a suffix added to the original job name. If you clone the same job again, the number in the name will be incremented, for example, *job_name_clone2*, *job_name_clone3* and so on.

To clone a job:

1. Open the **Home** view.
2. In the inventory pane, select the **Jobs** > **Replication** node.
3. In the working area, select the job and click **Clone** on the ribbon or right-click the job and select **Clone**.
4. After a job is cloned, you can edit all its settings, including the job name.

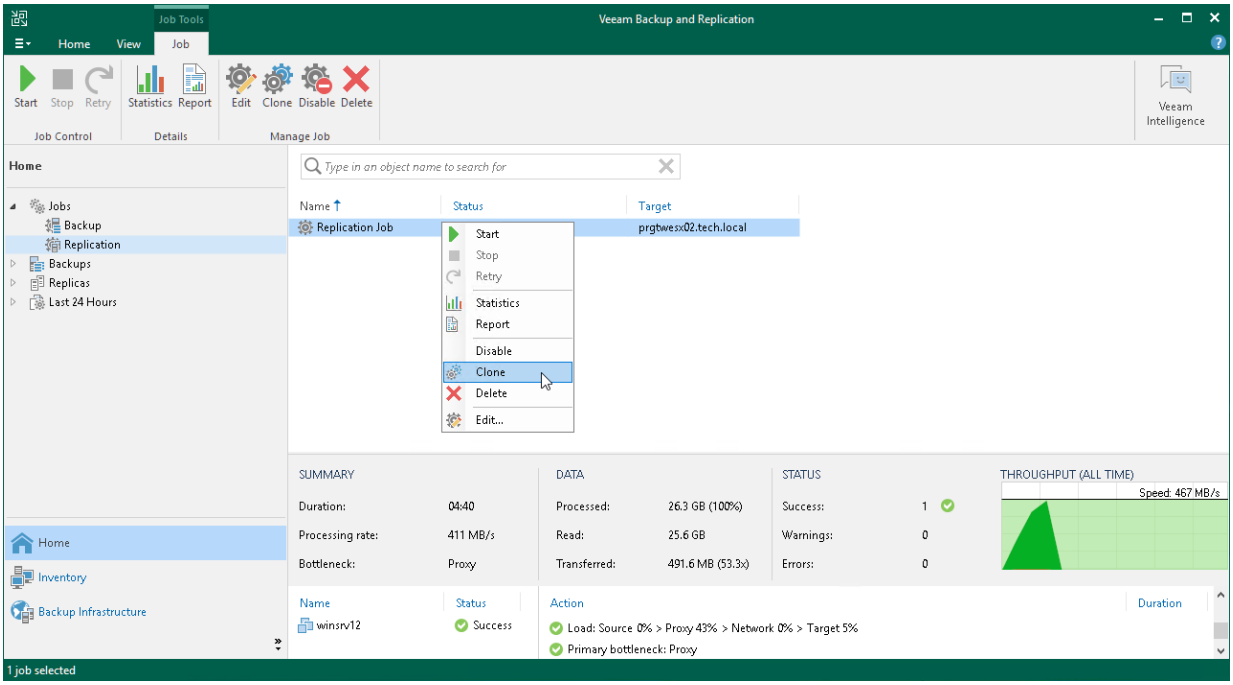

# Disabling and Deleting Replication Jobs

You can temporary disable scheduled jobs. The disabled job is paused for some period of time and is not run by the specified schedule. You can enable a disabled job at any time. You can also permanently delete a job from Veeam Backup & Replication and from the configuration database.

## Disabling Jobs

To disable a replication job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary policy and select **Disable** on the ribbon. Alternatively, you can right-click the necessary policy and select **Disable**.

To enable a disabled replication job, select it and click **Disable** on the ribbon once again.

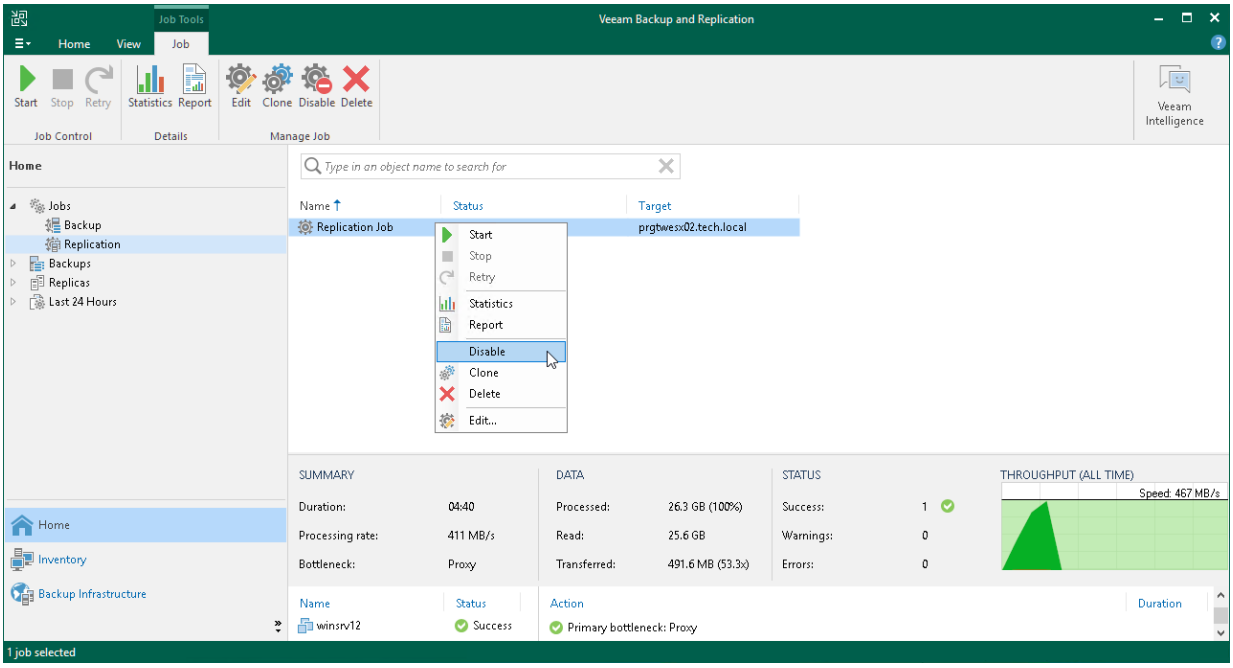


## Deleting Jobs

To delete a replication job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary replication job and select **Delete** on the ribbon. Alternatively, you can right-click the necessary replication job and select **Delete**.

> **NOTE**
>
> Veeam Backup & Replication allows you to delete only stopped replication jobs.

After you delete the job, the replicas created by this job are displayed under the **Replicas** node. If you want to remove replicas from the Veeam Backup & Replication console and configuration database but keep them on hosts, follow the instructions provided in Removing Replicas from Configuration. If you want to remove replicas not only from Veeam Backup & Replication, but also from host storage, follow the steps in Deleting Replicas from Disk.

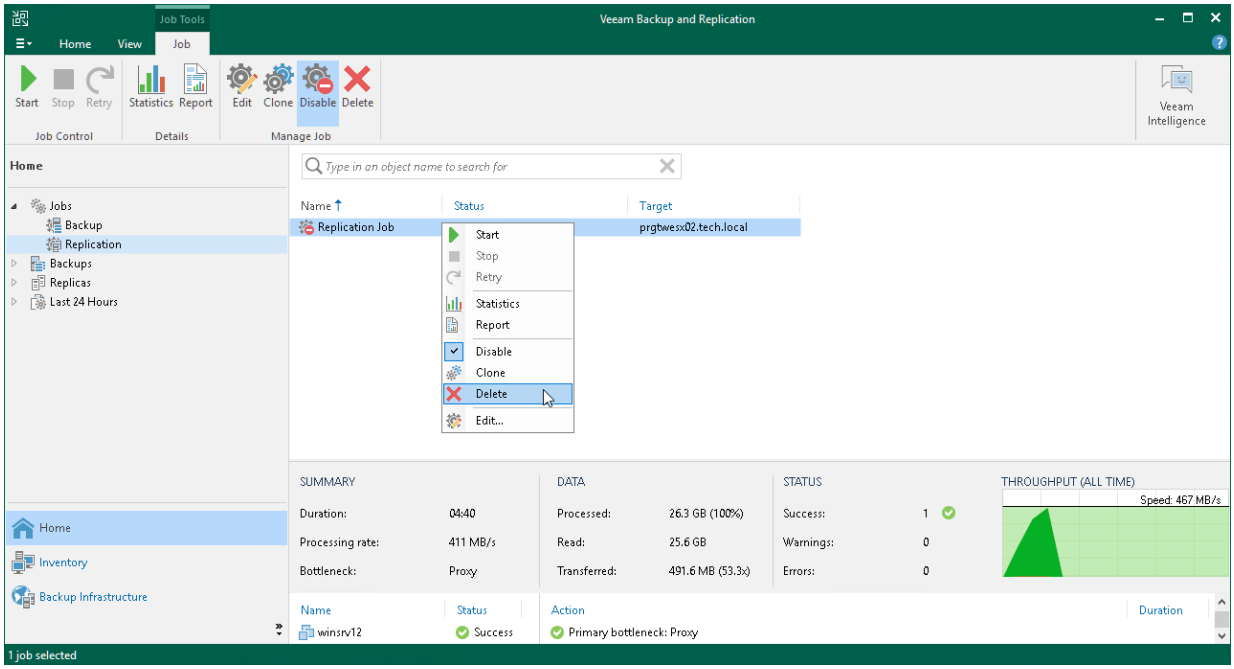

# Failover and Failback for Replication

Failover and failback operations help you ensure that your business will function even if a disaster strikes your production site. Failover is a process of switching from the VM on the source host to its VM replica on a host in the disaster recovery site. Failback is a process of returning from the VM replica to the source VM.

Veeam Backup & Replication provides the following failover and failback operations:

- **Perform failover**

  When you perform failover, you shift all processes from the source VM in the production site to the VM replica in the disaster recovery site. During failover, changes made on the VM replica are not reflected on the source VM.

  Failover is an intermediate step that needs to be finalized: you can undo failover, perform permanent failover or perform failback.

  For more information on how failover is performed, see Failover.

- **Perform planned failover**

  When you perform planned failover, you shift all processes from the source VM to its replica. During failover, changes made on the VM replica are not reflected on the source VM.

  Planned failover is helpful when you know that the source VM is about to go offline, for example, you plan to perform datacenter maintenance, and you want to proactively switch the workload to the replica. The procedure is designed to transfer the current workload, that is why it does not suggest to select a restore point.

  For more information on how planned failover is performed, see Planned Failover.

- **Create failover plan**

  When you create a failover plan, you define the order in which Veeam Backup & Replication must perform failover for VMs, and an interval of time for which Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list.

  For more information on failover plans, see Failover Plans.

- **Perform permanent failover**

  When you perform permanent failover, you permanently switch from the source VM to a VM replica and use this replica as the source VM.

  For more information on how permanent failover is performed, see Permanent Failover.

- **Undo failover**

  When you undo failover, you shift all processes back to the source VM and discard all changes made to the VM replica while it was running.

  You can use the undo failover scenario if you have failed over to the VM replica for testing and troubleshooting purposes, and you do not need to synchronize the source VM state with the current state of the replica.

  For more information on how failover undo is performed, see Undoing Failover.

- **Perform failback**

  When you perform failback, you shift all processes back to the source VM and send to the source VM all changes that took place while the VM replica was running. During failover, changes made on the source VM are not sent to the VM replica.

If the source host is not available, you can recover a VM with the same configuration as the source VM and switch to it. For more information on how failback is performed, see Failback.

When you perform failback, changes are only sent to the source/recovered VM but not published. You must test whether the source/recovered VM works with these changes. Depending on the test results, you can do the following:

- **Commit failback**. When you commit failback, you confirm that changes on the source/recovered VM work as expected and you want to get back to the source VM.

  For more information on how failback commit is performed, see Failback Commit.

- **Undo failback**. When you undo failback, you confirm that changes on the source/recovered VM are not working as expected and you want to discard them, and then to get back to the VM replica.

  For more information on how failback undo is performed, see Failback Undo.

Veeam Backup & Replication supports failover and failback operations for one VM and for several VMs. In case one or several hosts fail, you can use batch processing to restore operations with minimum downtime.

The following scheme can help you decide at which moment which operations are preferable.

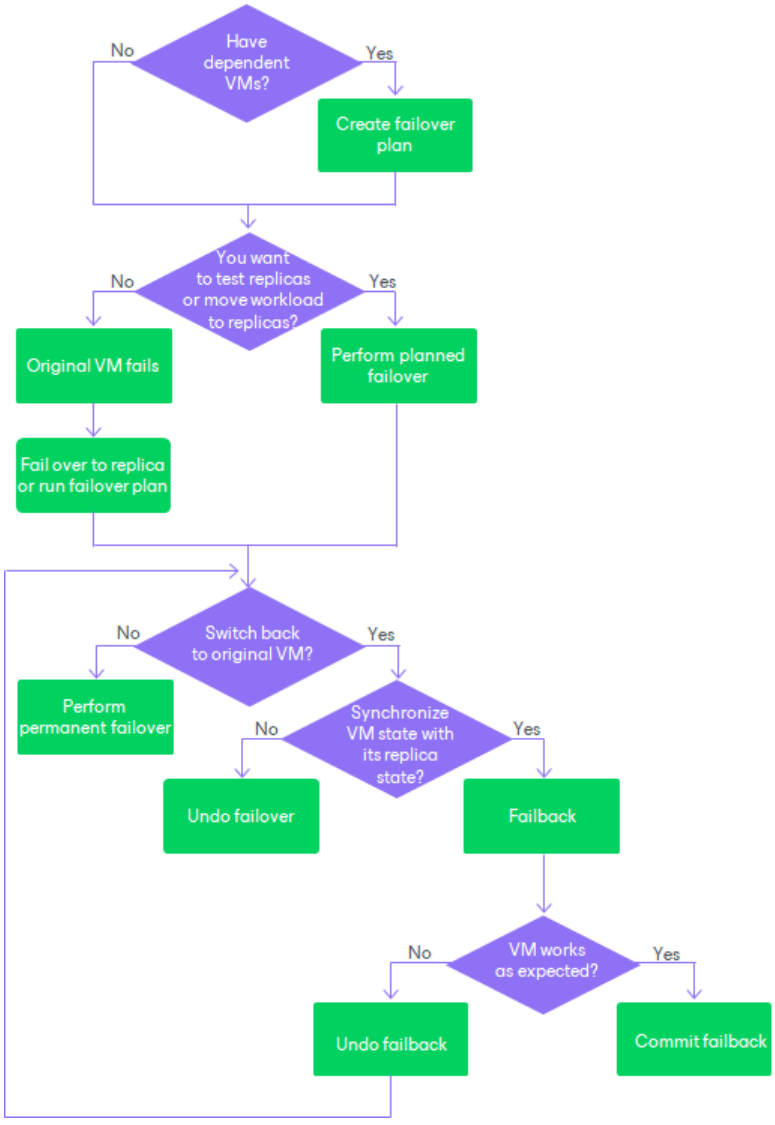

# Failover Plans

A failover plan helps you perform failover for dependent VMs one by one, as a group. To do this automatically, you can prepare a failover plan.

In the failover plan, you define the order in which VMs must be processed and an interval of time for which Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list. The failover plan helps ensure that some VMs, such as a DNS server, are already running at the time the dependent VMs start.

> **IMPORTANT**
>
> The failover plan must be created in advance.

In case the primary VM group goes offline, you can start the failover plan manually. When you start the plan, you can choose to fail over to the latest state or select the point in time to which VM replicas must be started. Veeam Backup & Replication will look for the closest restore points to this point in time and use them to start VM replicas. The source VMs will not be powered off.

The failover process is performed in the following way:

1. For each VM, Veeam Backup & Replication detects its replica. The VMs whose replicas are already in the *Failover* or *Failback* state are skipped from processing.
2. The replica VMs are started in the order they appear in the failover plan within the set time intervals.

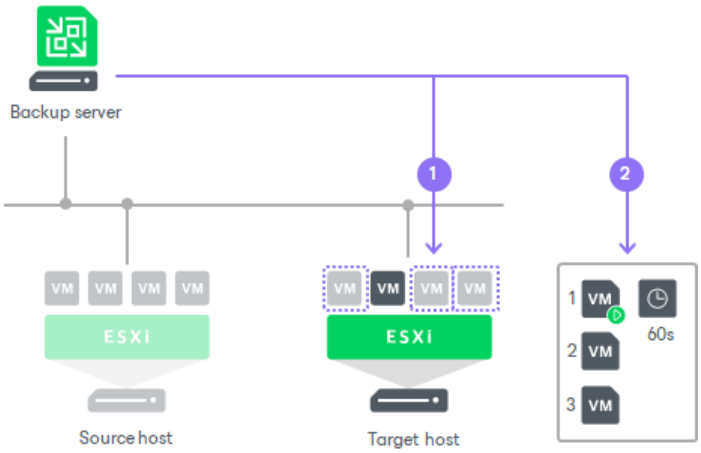


# Limitations for Failover Plans

The maximum number of VMs that can be started simultaneously when you run a failover plan is 10. If you have added more VMs to the failover plan and scheduled them to start simultaneously, Veeam Backup & Replication will wait for the first VMs in the list to fail over and then start the failover operation for subsequent VMs. This limitation helps reduce the workload on the production infrastructure and backup server.

For example, if you have added 14 VMs to the failover plan and scheduled them to start at the same time, Veeam Backup & Replication will start the failover operation for the first 10 VMs in the list. After the 1st VM is processed, Veeam Backup & Replication will start the failover operation for the 11th VM in the list, then for the 12th VM and so on.

# Finalizing Failover Plans

Failover is a temporary intermediate step that needs to be finalized. You can finalize group failover in the same ways as regular failover: undo failover, perform permanent failover or failback.

When you perform failback or permanent failover, you need to process each VM individually. For more information, see Performing Failback and Performing Permanent Failover. When you undo failover, you can process the whole group. For more information, see Undoing Failover by Failover Plans.

# Creating Failover Plans

To create a failover plan, use the **New Failover Plan** wizard.

## Before You Begin

Before you create a failover plan, check the following prerequisites:

- [If you plan to select VMs from replication jobs] VMs that you plan to include in the failover plan must be successfully replicated at least once.
- [If you plan to select VMs from replication jobs] VM replicas must be in the *Ready* state.
- If you plan to use pre-failover and post-failover scripts for the failover plan, you must create scripts before you configure the failover plan.

# Step 1. Launch New Failover Plan Wizard

To launch the **New Failover Plan** wizard, do one of the following:

- On the **Home** tab, click **Failover Plan** and select **VMware vSphere**.
- Open the **Home** view. In the working area, select VMs that you want to add to a failover plan. On the ribbon, click **Add to Failover Plan** > **New failover plan** if you want to create a new failover plan, or **Add to Failover Plan** > *<Plan Name>* if you want to add VMs to an existing failover plan.
- Open the **Home** view. In the working area, select VMs that you want to add to a failover plan and right-click one of them. Select **Add to failover plan** > **New failover plan** if you want to create a new failover plan, or **Add to failover plan** > *<Plan Name>* if you want to add VMs to an existing failover plan.

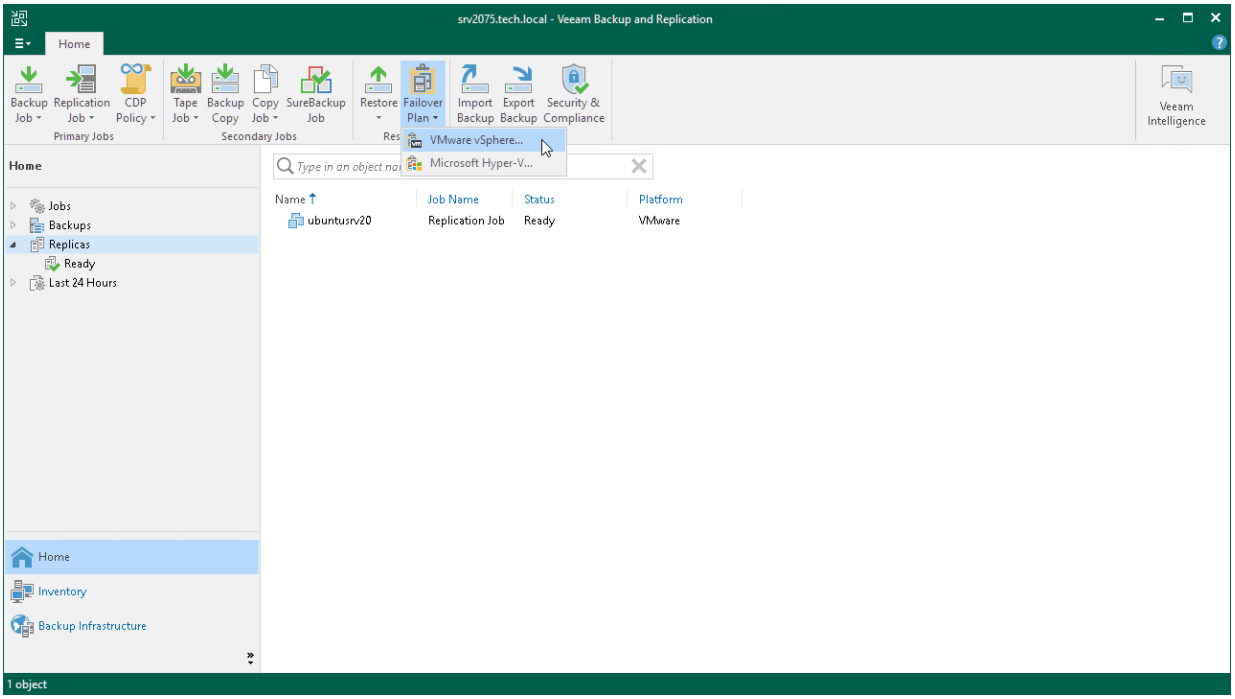

## Step 2. Specify Failover Plan Name and Description

At the **General** step of the wizard, specify a name and description for the failover plan.

If you want to execute custom scripts before or after the failover plan, select the **Pre-failover script** and **Post-failover** script check boxes and click **Browse** to choose executable files. For example, you may want stop some applications on production VMs before the failover plan starts or send an email to backup administrators after the failover plan finishes.

The scripts will be executed on the backup server. Veeam Backup & Replication supports the script files in the following formats: BAT, CMD, EXE and PS1.

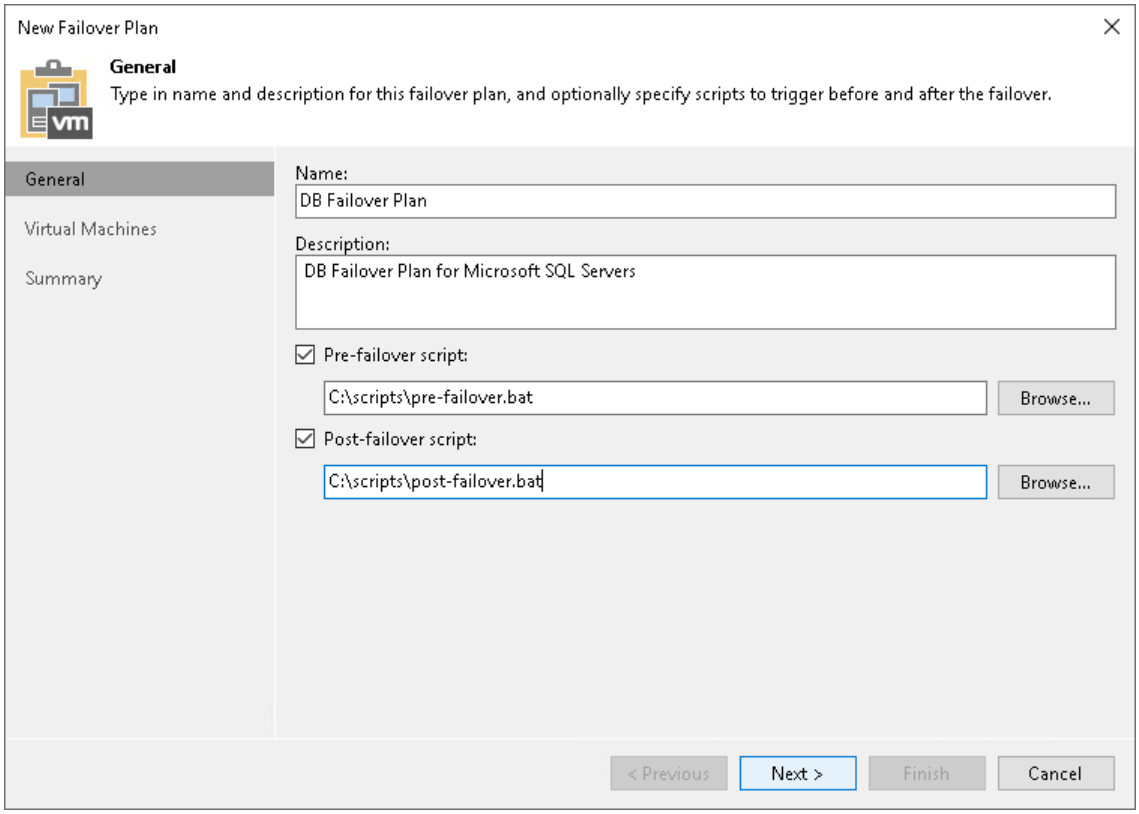

## Step 3. Select VMs

At the **Virtual Machines** step of the wizard, select VMs that you want to add to the failover plan. You can add separate VMs and whole VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags).

To add VMs and VM containers:

1. Click **Add VM**.
2. Select where to browse for VMs and VM containers:
   - **From infrastructure** – browse the virtual environment and select VMs or VM containers. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list.
   - **From replicas** – browse existing replication jobs and select all VMs or specific VMs from replication jobs.

To quickly find the necessary VMs or VM containers, you can use the search fields.

> **NOTE**
>
> A source from which you add a VM to a failover plan does not affect whether you fail over to the latest or specific restore point. It is the command that you select when starting a failover plan that defines the restore point. For more information, see Running Failover Plans.

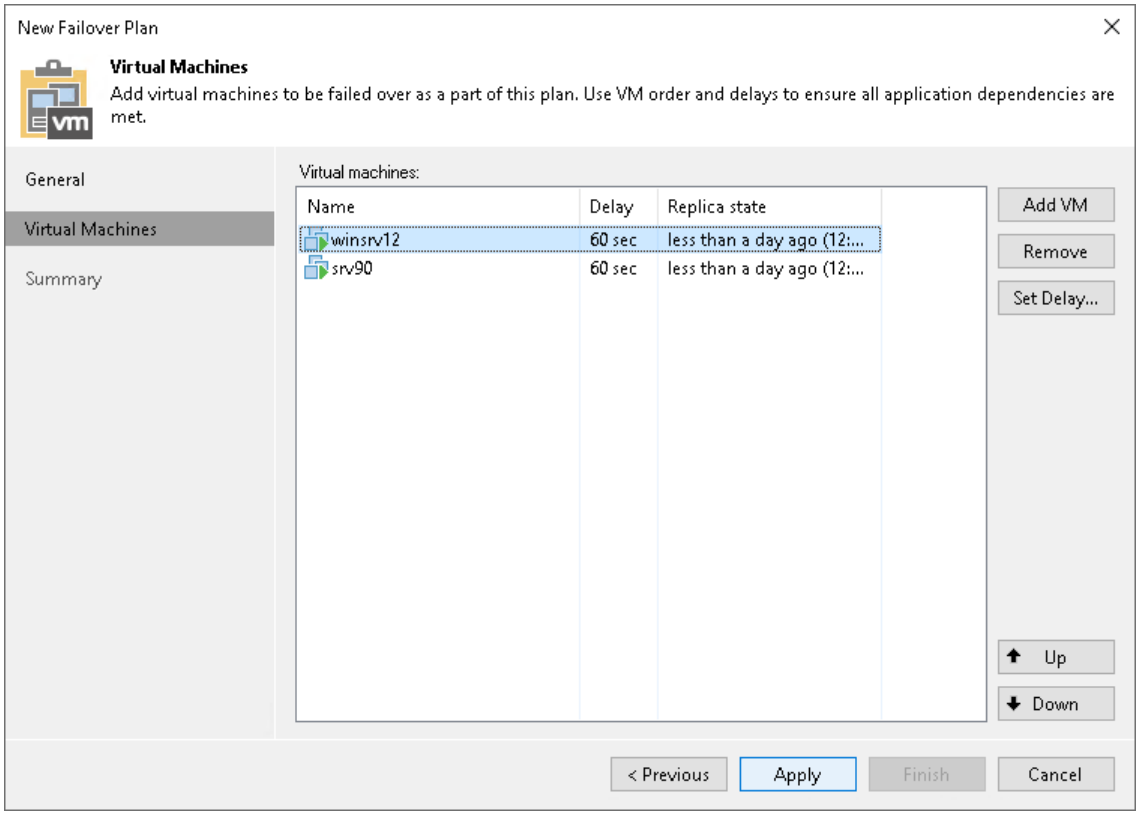

# Step 4. Define VM Failover Order

At the **Virtual Machines** step of the wizard, click **Up** and **Down** to change the processing order. VMs at the top of the list have a higher priority and will be started first. If some VMs provide environment for other dependent VMs, make sure that they are started first.

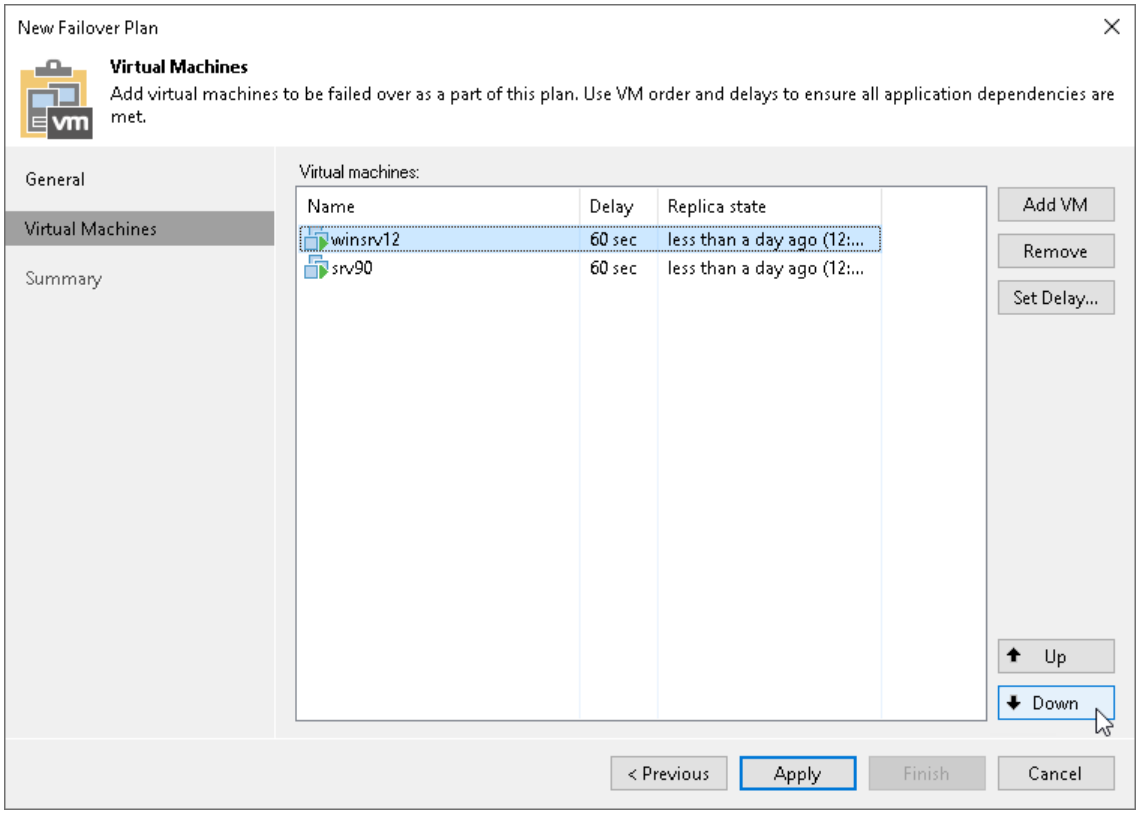

## Step 5. Set Time Delay

After you have set the order for VMs in the failover plan, you need to set a time delay for VMs. The delay time defines for how long Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list. You can use time delays to make sure that some VMs are already running at the moment dependent VMs start.

For example, you have added 2 VMs to the failover plan and set a time delay to 60 seconds for the first VM in the list. Veeam Backup & Replication will perform failover in the following manner: Veeam Backup & Replication will start the failover operation for the first VM in the list, then wait for 60 seconds and start the failover operation for the second VM in the list.

> **NOTE**
>
> Time delays can be specified for all VMs in the list except the last one. If you do not specify time delays, VMs will be started simultaneously.

To set the time delay for a VM:

1. Select it and click **Set Delay** on the right or double-click the VM in the list.
2. Enter the time interval that you consider sufficient for this VM to boot.

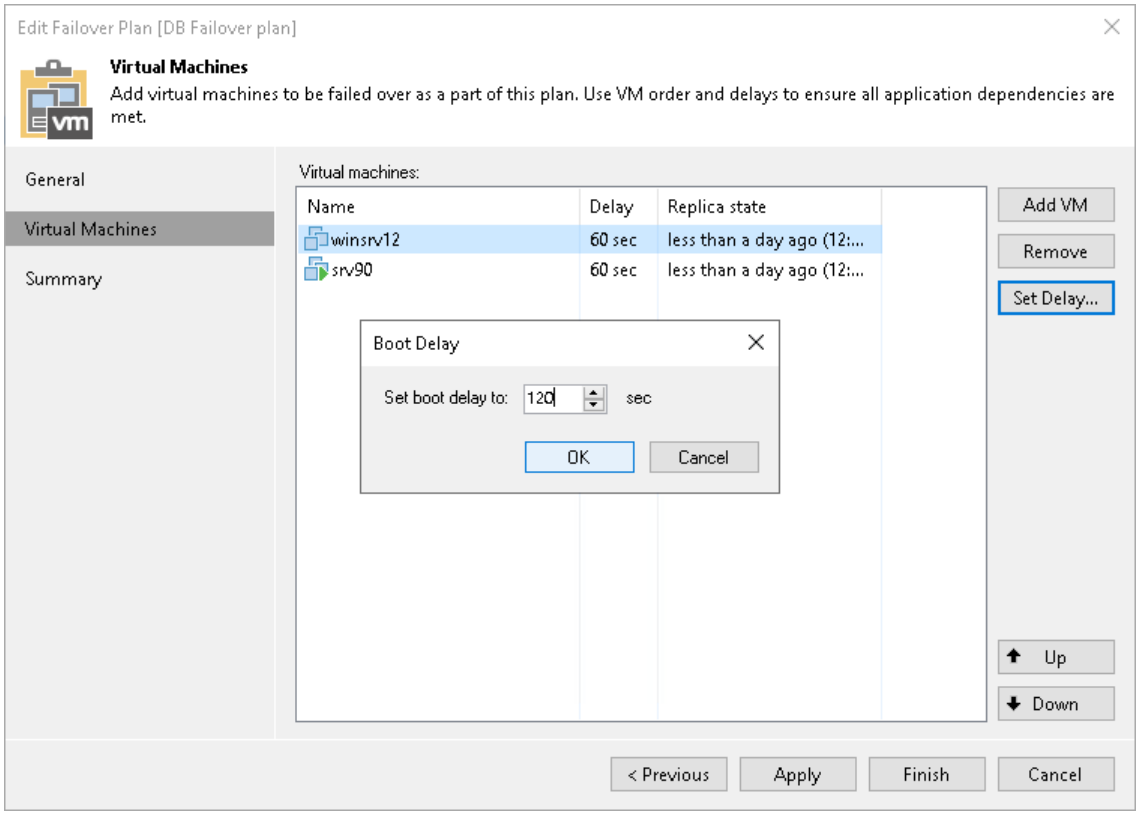

## Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details for the configured failover plan and click **Finish** to create the plan.

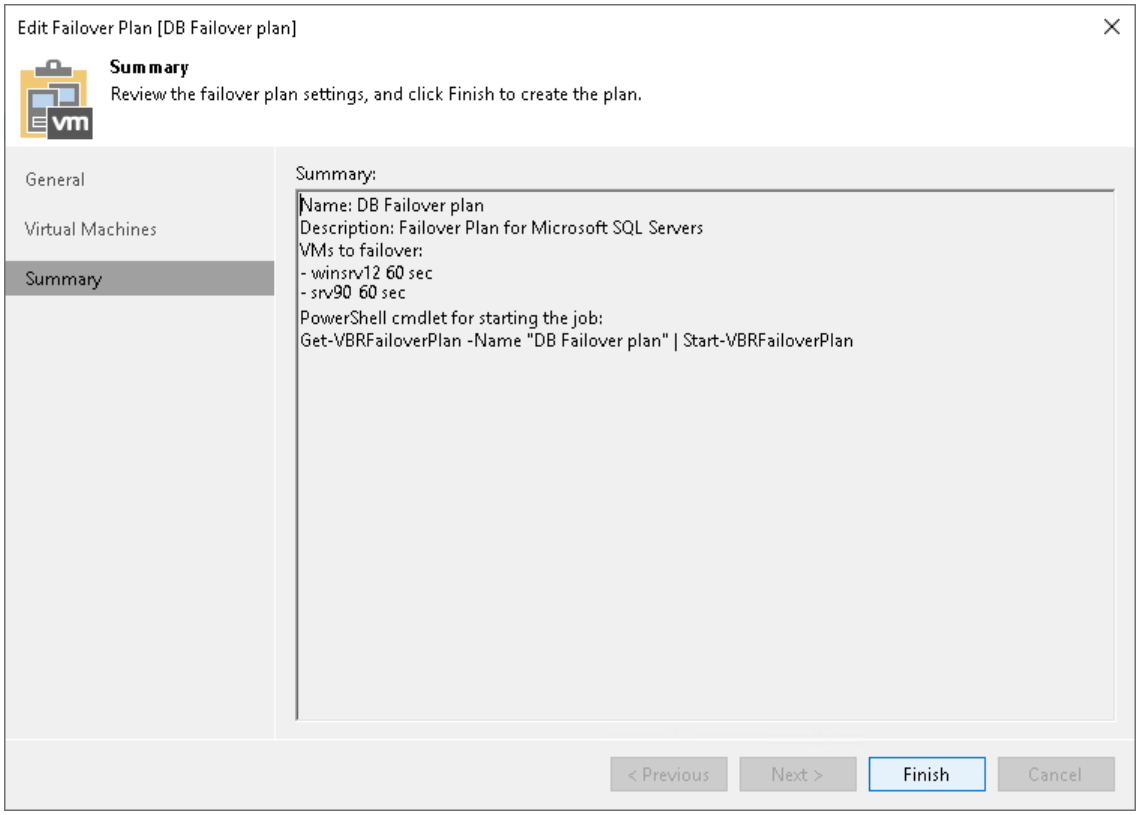


# Running Failover Plans

You have the following options to run a failover plan:

- You can fail over to latest restore points of VM replicas.

  In this case, Veeam Backup & Replication searches for the latest restore point of VM replicas across all replication jobs configured on the backup server. For example, you have 2 jobs that replicate the same VM: *Job* 1 has created the most recent point at 2:00 AM and *Job* 2 has created the most recent restore point at 3:00 AM. When you run the failover plan using the **Start** command, Veeam Backup & Replication will pick the restore point created at 3:00 AM with *Job 2*.

- You can fail over to specific restore points of VM replicas.

# Failing Over to Latest Restore Points

To fail over to the latest restore points of VM replicas:

1. Open the **Home** view.
2. Expand the **Replicas** node.
3. Select **Failover Plans**.
4. In the working area, right-click the failover plan and select **Start**.

# Failing Over to Specific Restore Points

To fail over to specific restore points of VM replicas:

1. Open the **Home** view.
2. Expand the **Replicas** node.
3. Select **Failover Plans**.
4. In the working area, right-click the failover plan and select **Start to**.
5. In the displayed window, select the backup date and time. Veeam Backup & Replication will find the closest restore point prior to the entered value for each VM and will fail over to it.

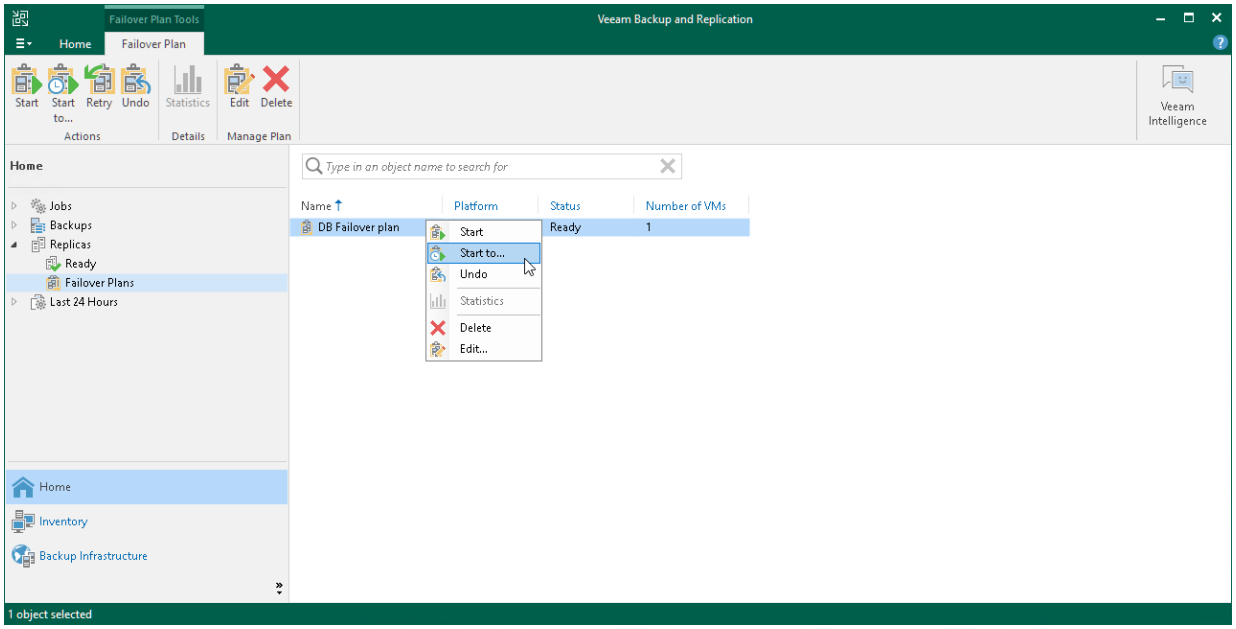


# Undoing Failover by Failover Plans

You can undo failover for all VMs added to the failover plan at once. When you undo failover, you switch the workload back to source VMs and discard all changes that were made to VM replicas during failover. If some of the VMs were already failed back, for example manually by the user, they are skipped from processing.

Veeam Backup & Replication starts the failover undo operation for a group of 5 VMs at the same time. The time interval between the operation starts is 10 seconds. For example, if you have added 10 VMs to the failover plan, Veeam Backup & Replication will undo failover for the first 5 VMs in the list, then will wait for 10 seconds and undo failover for the remaining 5 VMs in the list. Time intervals between the operation starts help Veeam Backup & Replication reduce the workload on the production environment and the backup server.

To undo failover by a failover plan:

1. Open the **Home** view.
2. Expand the **Replicas** node.
3. Select **Failover Plans**.
4. In the working area, right-click the failover plan and select **Undo**.

5. In the displayed dialog box, click **Yes** to confirm the operation.

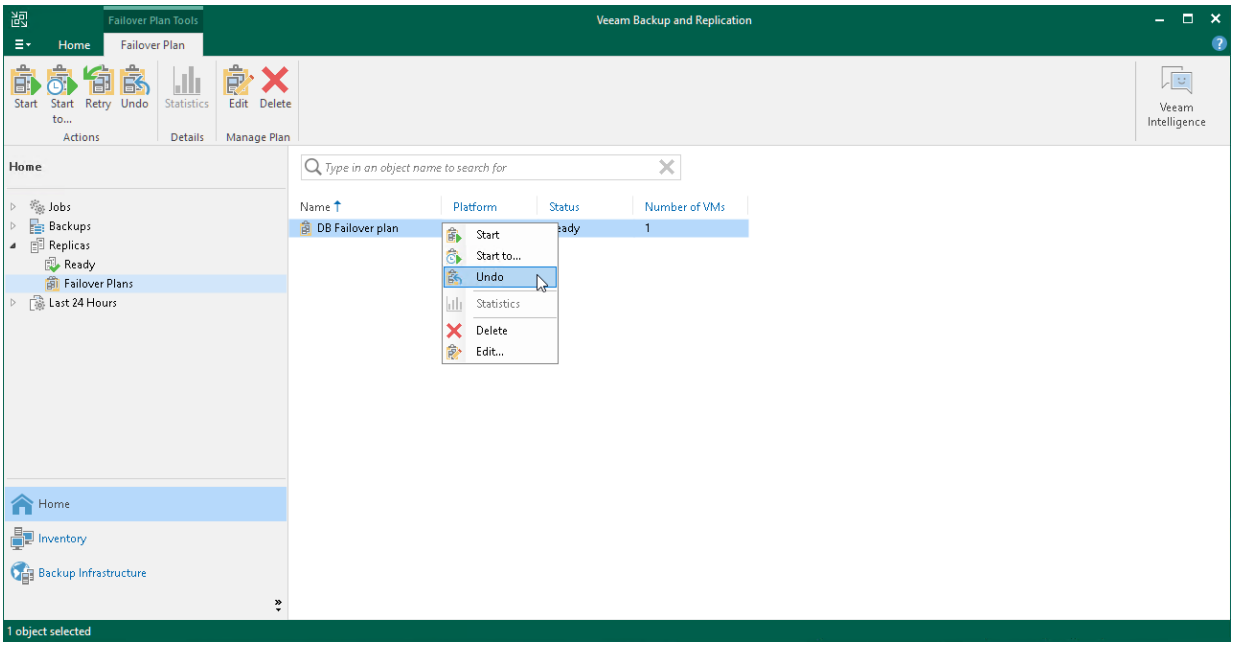

# Failover

Failover is a process when Veeam Backup & Replication switches processes from the source VM in the production site to its VM replica in the disaster recovery site. During failover, Veeam Backup & Replication recovers the VM replica to the required restore point and shifts all I/O processes from the source VM to its replica. As a result, you have a fully functional VM within a couple of seconds, and your users can access services and applications with minimum disruption.

You can fail over to replicas not only when a disaster strikes the production site, but also to test replicas for recoverability. You can perform failover while the source VM is running. After all the necessary tests, you can undo failover and get back to the normal mode of operation. If the source VMs and VM replicas are located in the same network, consider temporary disconnecting the source VMs from the network to avoid IP address or machine name conflicts. As an alternative way of testing, Veeam Backup & Replication also provides the SureReplica technology. For more details, see SureReplica.

> **IMPORTANT**
>
> Use Veeam Backup & Replication to perform failover operations. Avoid powering on a replica manually — this may disrupt further replication operations or cause loss of important data.

The failover operation is performed in the following way:

1. Veeam Backup & Replication rolls back the VM replica to the required restore point. To do this, it reverts the VM replica to the necessary snapshot in the replica chain.
2. Veeam Backup & Replication powers on the VM replica. The state of the VM replica is changed from *Ready* to *Failover*.

   If you perform failover for testing or disaster recovery (DR) simulation purposes, and the source VM still exists and is running, the source VM remains powered on.

   > **NOTE**
   >
   > Veeam Backup & Replication stops all replication activities for the source VM until its replica is returned to the *Ready* state.

3. All changes made to the VM replica while it is running in the *Failover* state are written to the delta file of the snapshot, or restore point, to which you have selected to roll back.

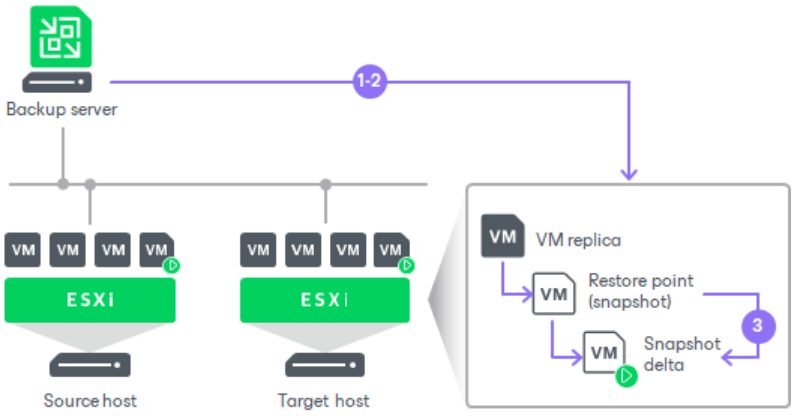

# Finalizing Failover

Failover is an intermediate step that needs to be finalized. You can use one of the following operations:

- Undo failover.
- Perform permanent failover.
- Perform failback.

# Performing Failover

For more information on failover, see Failover and Failover and Failback for Replication.

To perform failover, use the **VMware Failover** wizard.

## Before You Begin

Before you fail over to a VM replica, check the following prerequisites:

- The failover operation can be performed for VMs that have been successfully replicated at least once.
- VM replicas must be in the *Ready* state.

## Step 1. Launch Failover Wizard

To launch the **Failover** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from replica** > **Entire replica** > **Failover to a replica**.
- Open the **Home** view. In the inventory pane select **Replicas** > **Ready**. In the working area, select the necessary replica and click **Failover Now** on the ribbon.
- Open the **Home** view. In the inventory pane select **Replicas** > **Ready**. In the working area, right-click the necessary replica and select **Failover Now**.

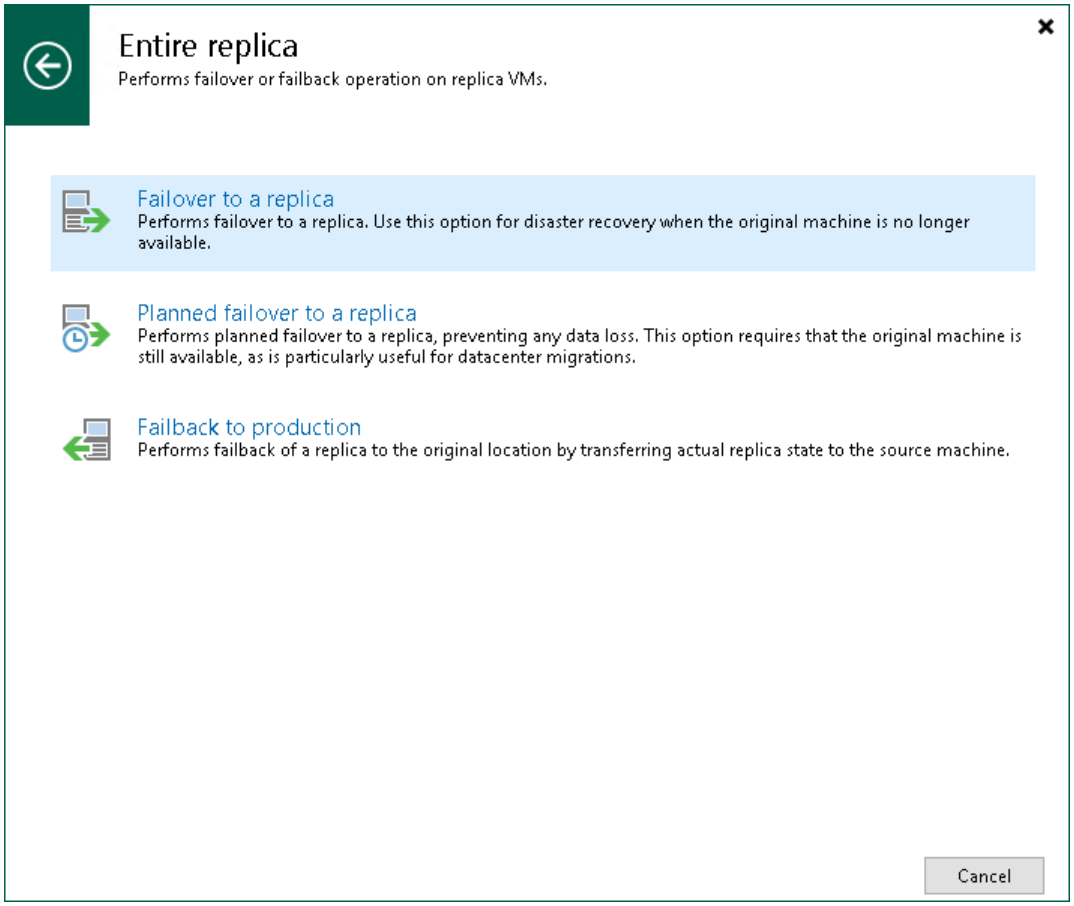

## Step 2. Select VMs

At the **Virtual Machines** step of the wizard, you can modify a list of VMs from which you fail over. To add VMs or VM containers, click **Add** > **From infrastructure** if you want to add VMs from the virtual infrastructure, or **Add** > **From replicas** if you want to add VMs from existing replicas. Then select the necessary VMs or VM containers. If you select VM containers, Veeam Backup & Replication will expand them to a plain VM list.

> **NOTE**
>
> Make sure that VMs you select from the virtual environment have been successfully replicated at least once.

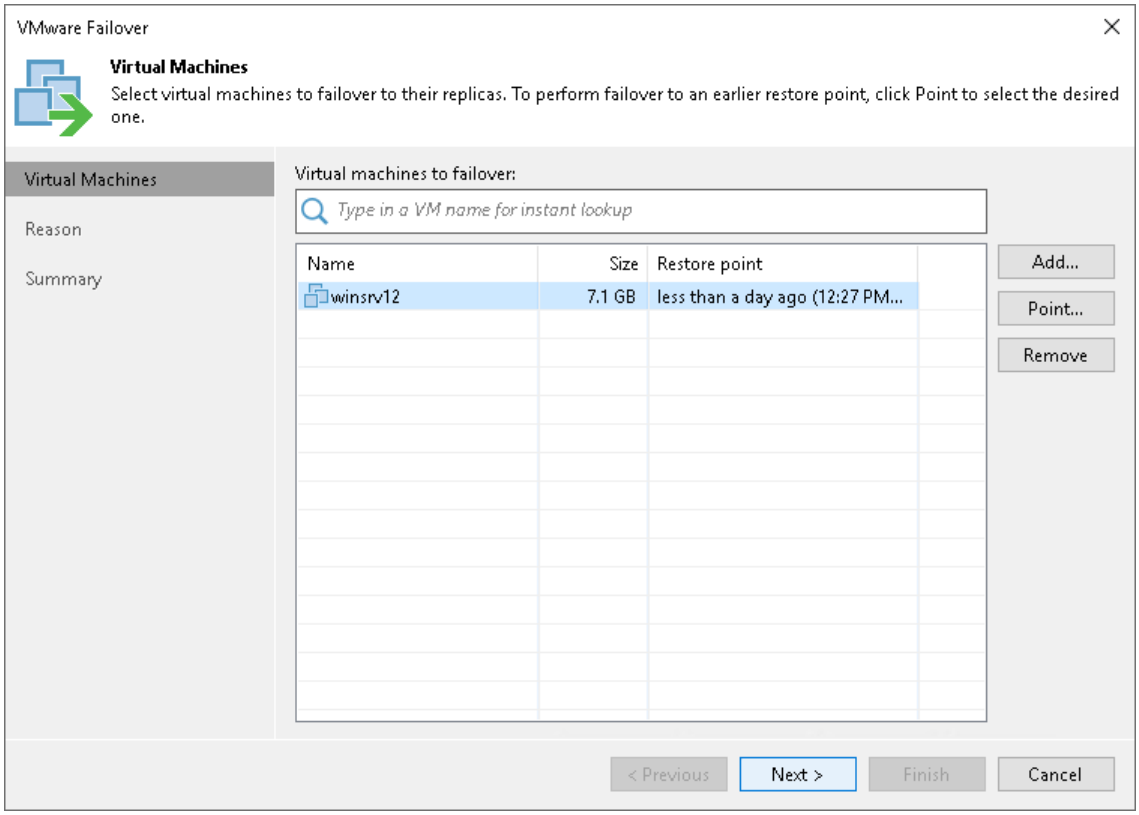

# Step 3. Select Restore Points

By default, Veeam Backup & Replication uses the latest valid restore point of the VM replica. However, you can fail over to an earlier state of the VM. If you have chosen to perform failover for several VMs, you can select the necessary restore point for every VM in the list.

To select a restore point for a VM:

1. In the **Virtual machines to failover** list, select the necessary VM and click **Point**.
2. In the **Restore Points** window, select the necessary restore point.

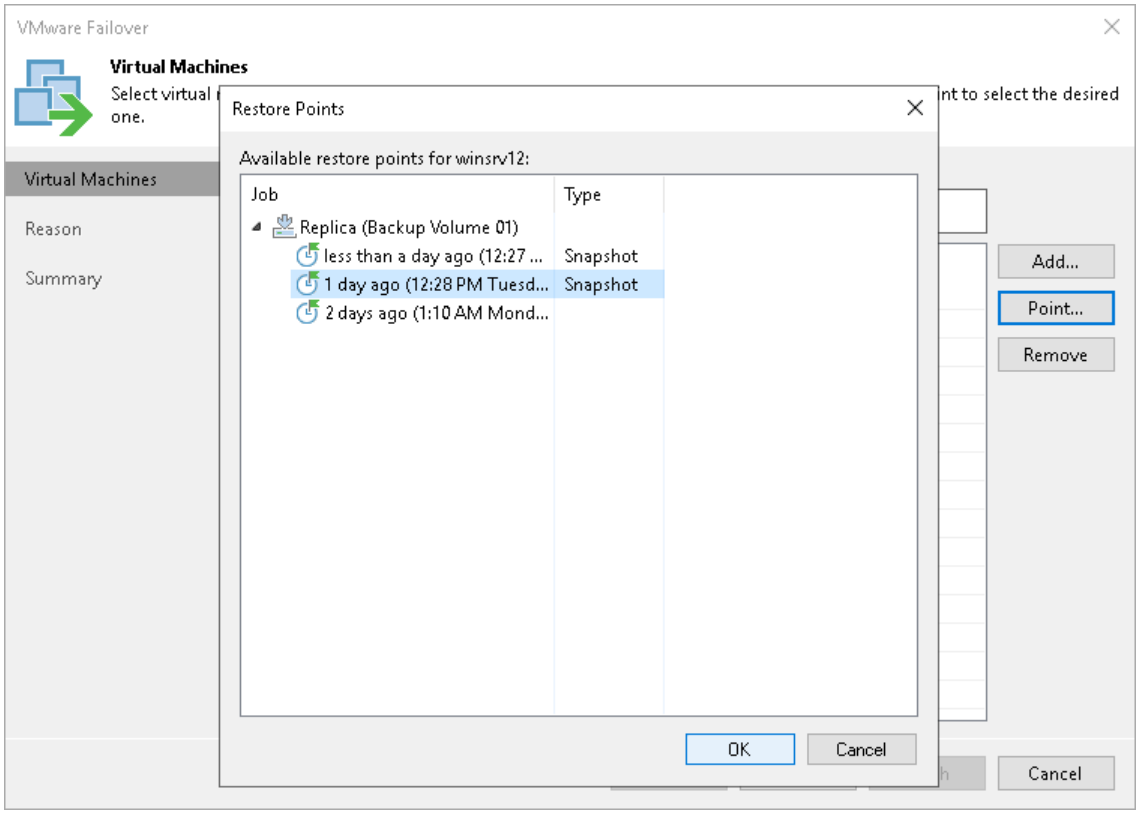

# Step 4. Specify Failover Reason

At the **Reason** step of the wizard, enter a reason for failing over to the replicas. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

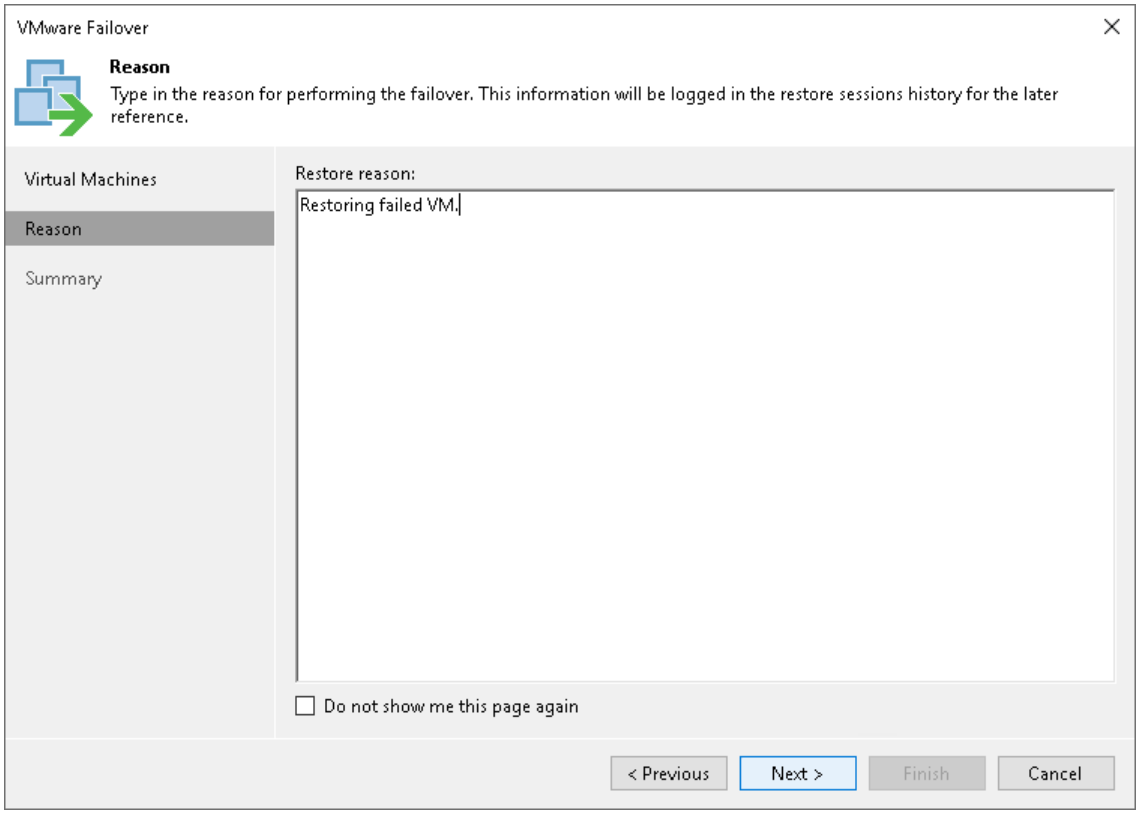

## Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the failover task and click **Finish** to exit the wizard. When the failover process is complete, the VM replicas will be started on the target host.

# What You Do Next

Failover is an intermediate step that needs to be finalized. You can finalize failover in the following ways:

- Perform permanent failover.
- Undo failover.
- Perform failback.

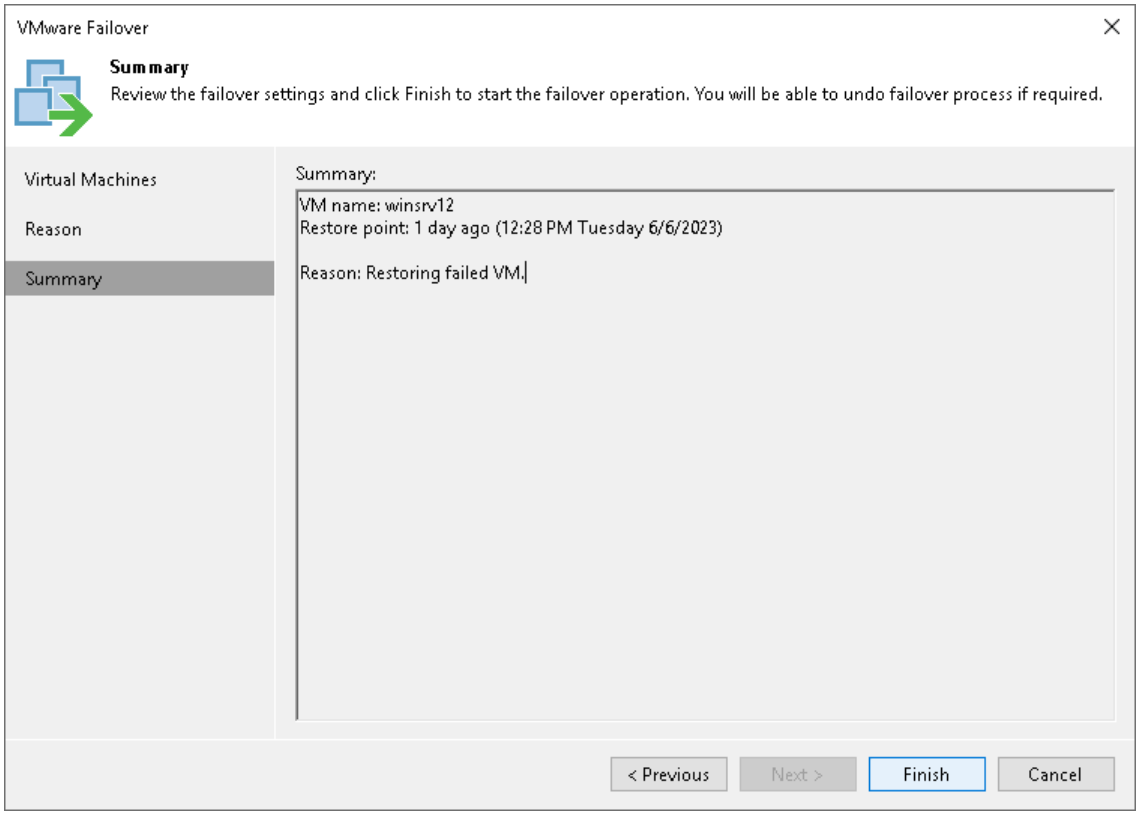

# Permanent Failover

Permanent failover is one of the ways to finalize failover. When you perform permanent failover, you permanently switch from the source VM to its replica. As a result of permanent failover, the VM replica stops acting as a replica and starts acting as the production VM.

The permanent failover operation is performed in the following way:

1. Veeam Backup & Replication removes snapshots (restore points) of the VM replica from the snapshot chain and deletes associated files from the datastore. Changes that were written to the snapshot delta file are committed to the VM replica disk files to bring the VM replica to the most recent state.
2. Veeam Backup & Replication removes the VM replica from the list of replicas in the Veeam Backup & Replication console.
3. To protect the VM replica from corruption after permanent failover is complete, Veeam Backup & Replication reconfigures the current replication job by adding the source VM to the list of exclusions. Note that other jobs are not modified automatically. When the replication job starts, the source VM is skipped from processing. As a result, no data is written to the working VM replica.

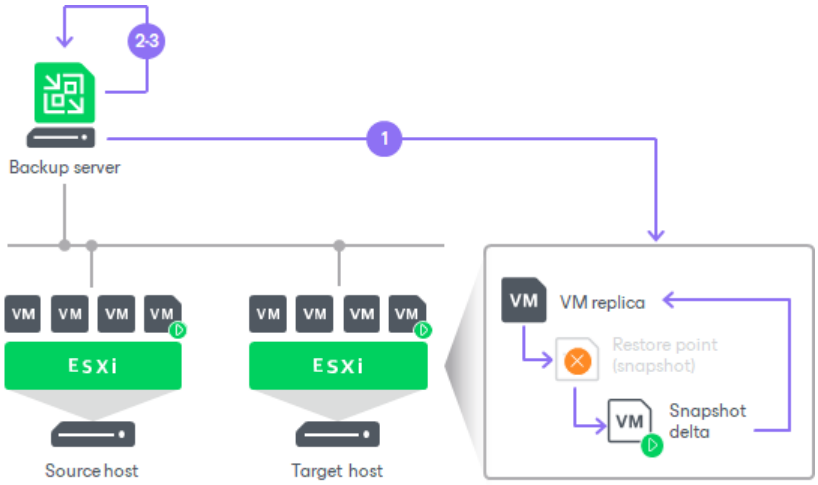


# Performing Permanent Failover

For more information on permanent failover, see Permanent Failover and Failover and Failback for Replication.

To perform permanent failover, do one of the following:

- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, select the necessary replica and click **Permanent Failover** on the ribbon.

- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, right-click the necessary replica and select **Permanent failover**.

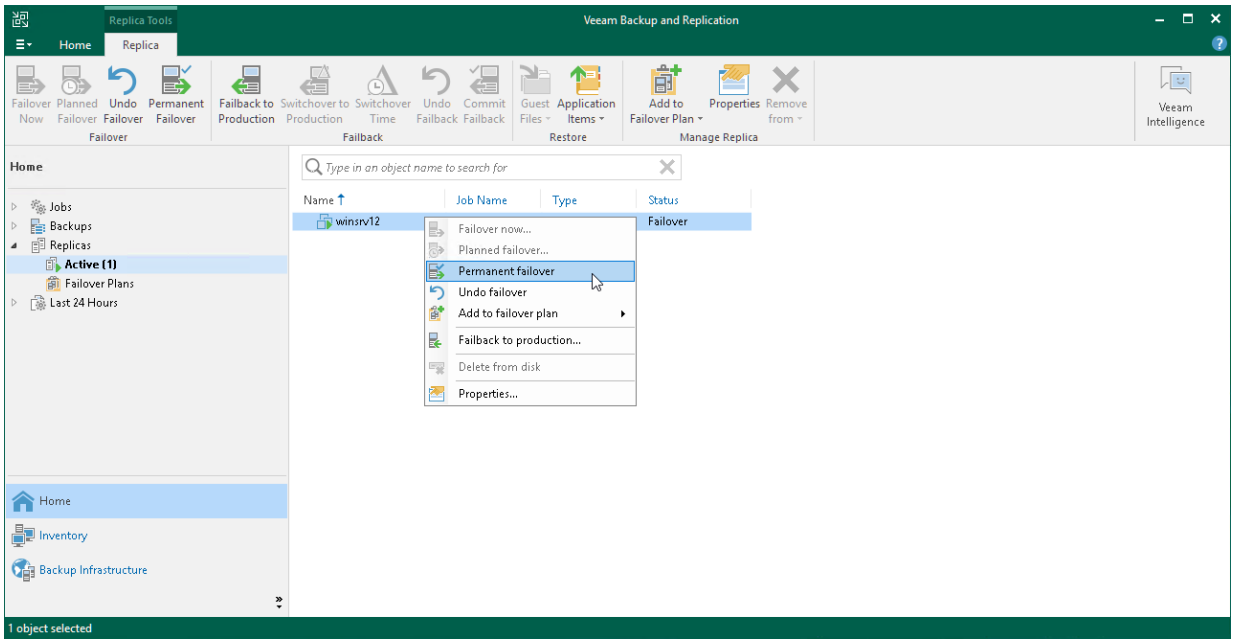

# Planned Failover

Planned failover is a process when you manually launch switching from a primary VM to its replica with minimum interrupting in operation. Planned failover is helpful when you know that your primary VMs are about to go offline and you need to proactively switch the workload from source VMs to their replicas. You can use the planned failover, for example, if you plan to perform datacenter migration, maintenance or software upgrade of the primary VMs. You can also perform planned failover if you have noticed some signs of the approaching disaster.

As the procedure is designed to transfer the current workload to the replica, it does not suggest selecting a restore point to switch.

When you start the planned failover, Veeam Backup & Replication performs the following operations:

1. The failover process triggers the replication job to perform an incremental replication run and copy the un-replicated changes to the replica.
2. The guest OS of the VM is shut down or the VM is powered off.

   If VMware Tools are installed on the VM, Veeam Backup & Replication tries to shut down the VM guest OS. If nothing happens after 15 minutes, Veeam Backup & Replication powers off the VM. If VMware Tools are not installed on the VM or the VM is suspended, Veeam Backup & Replication powers off the VM.
3. The failover process triggers the replication job to perform another incremental replication run and copy the portion of last-minute changes to the replica. The replica becomes fully synchronized with the source VM.
4. The VM is failed over to its replica.
5. The VM replica is powered on.

During the planned failover, Veeam Backup & Replication creates two helper restore points (steps 1 and 3) that are not deleted afterwards. You can see these restore points in the list of restore points for the VM. You can use the restore points later to roll back to the necessary VM replica state.

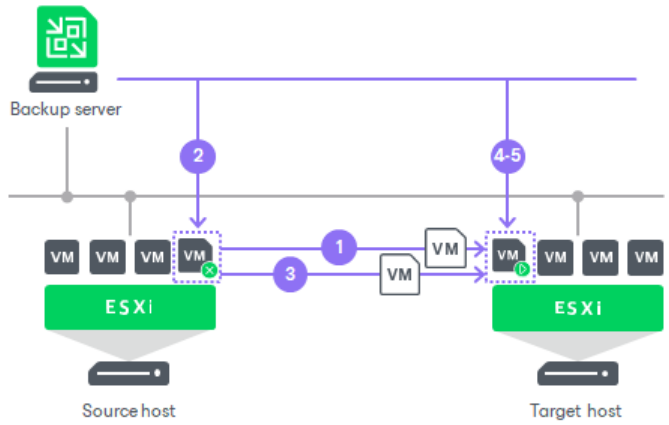


**NOTE**

During planned failover, Veeam Backup & Replication always retrieves VM data from the production infrastructure, even if the replication job uses the backup as a data source. This approach helps Veeam Backup & Replication synchronize the VM replica to the latest state of the production VM.

# Finalizing Planned Failover

When your primary host is online again, you can switch back to it. You can finalize planned failover in the same ways as regular failover: undo failover, perform permanent failover or failback.

# Limitations for Planned Failover

Planned failover has the following limitations:

- If you start planned failover for several VMs that are replicated with one replication job, these VMs will be processed one by one, not in parallel.
- Each planned failover task for each VM is processed as a separate replica job session. If a backup proxy is not available and the session has to wait for resources, job sessions for other VMs in the same task cannot be started before the current session is finished.
- The user account under which you launch the planned failover operation must have the *Veeam Backup Administrator* role or *Veeam Backup Operator* and *Veeam Restore Operator* roles in Veeam Backup & Replication. For more information, see Managing Users and Roles.

# Performing Planned Failover

For more information on planned failover, see Planned Failover and Failover and Failback for Replication.

To perform planned failover, use the **Planned Failover** wizard.

## Before You Begin

Before you perform planned failover, check the following prerequisites:

- VMs for which you plan to perform planned failover must be successfully replicated at least once.
- VM replicas must be in the *Ready* state.

# Step 1. Launch Planned Failover Wizard

To launch the **VMware Planned Failover** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware vSphere** > **Restore from replica** > **Entire replica** > **Planned failover to a replica**.
- Open the **Home** view, expand the **Replicas** node. In the working area, select one or more VMs and click **Planned Failover** on the ribbon. You can also right-click one of the selected VMs and click **Planned Failover**.
- Open the **Inventory** view. In the working area, select one or more VMs and right-click one of the selected VMs and click **Restore** > **Planned Failover**.

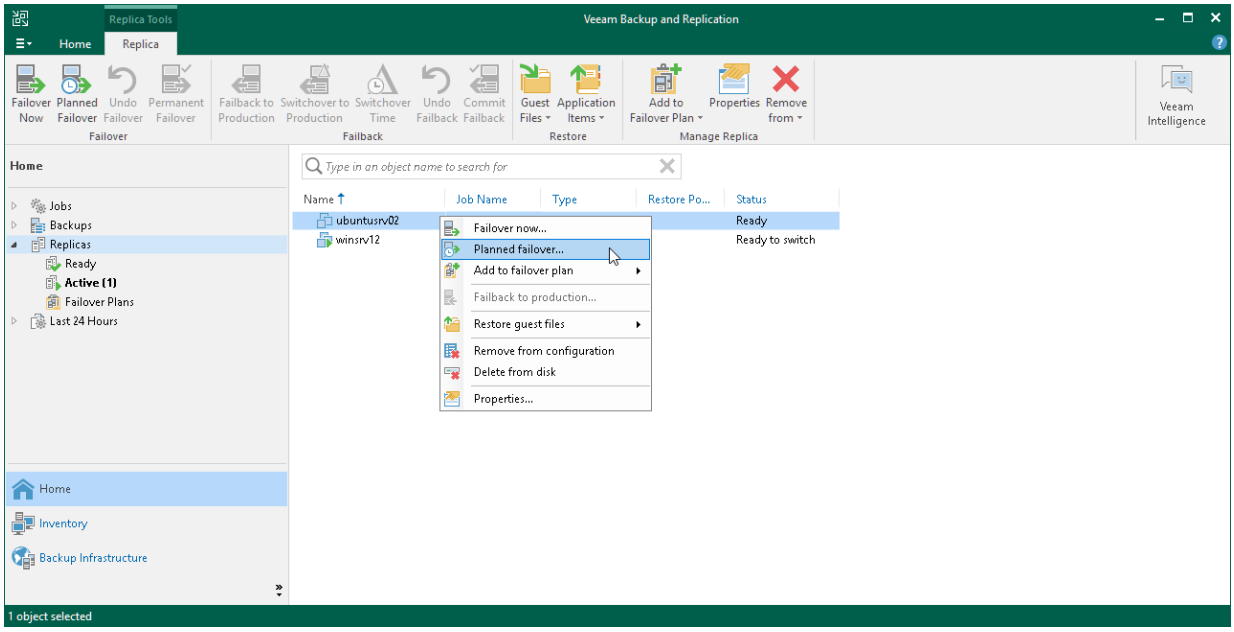

# Step 2. Select VMs

At the **Virtual Machines** step of the wizard, you can modify a list of VMs from which you fail over. To add VMs or VM containers, click **Add** > **From infrastructure** if you want to add VMs from the virtual infrastructure, or **Add** > **From replicas** if you want to add VMs from existing replicas. Then select the necessary VMs or VM containers. If you select VM containers, Veeam Backup & Replication will expand them to a plain VM list.

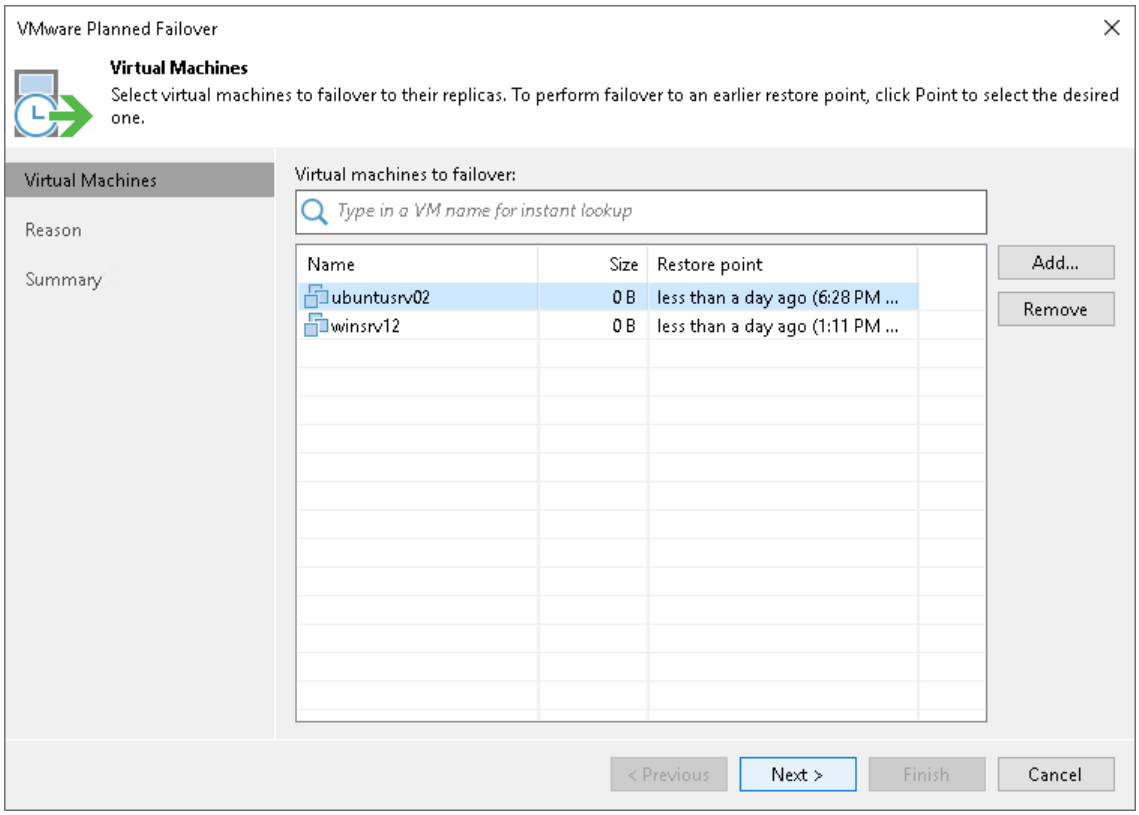

## Step 3. Specify Failover Reason

At the **Reason** step of the wizard, enter a reason for failing over to the replicas. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

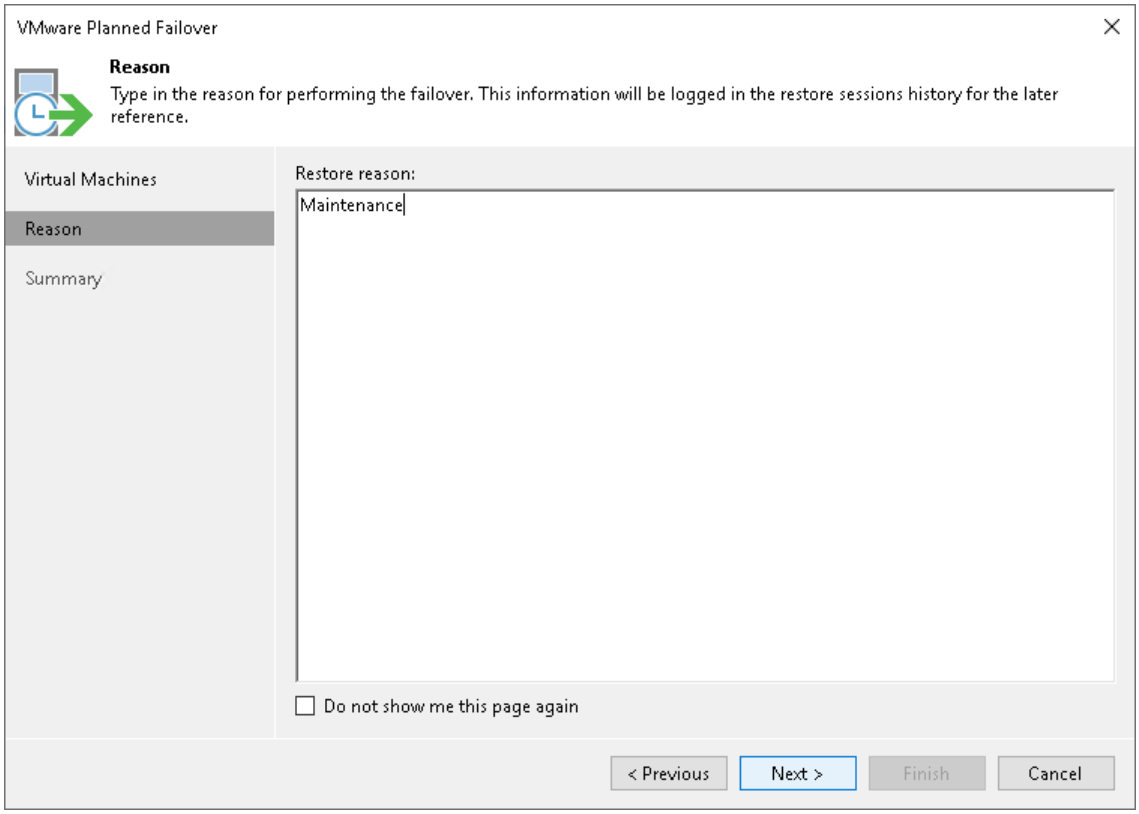

## Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the failover task and click **Finish** to start the failover process. Once planned failover is complete, VM replicas will be started on the target host.

# What You Do Next

Planned failover is an intermediate step that needs to be finalized. You can finalize failover in the following ways:

- Perform permanent failover.
- Undo failover.
- Perform failback.

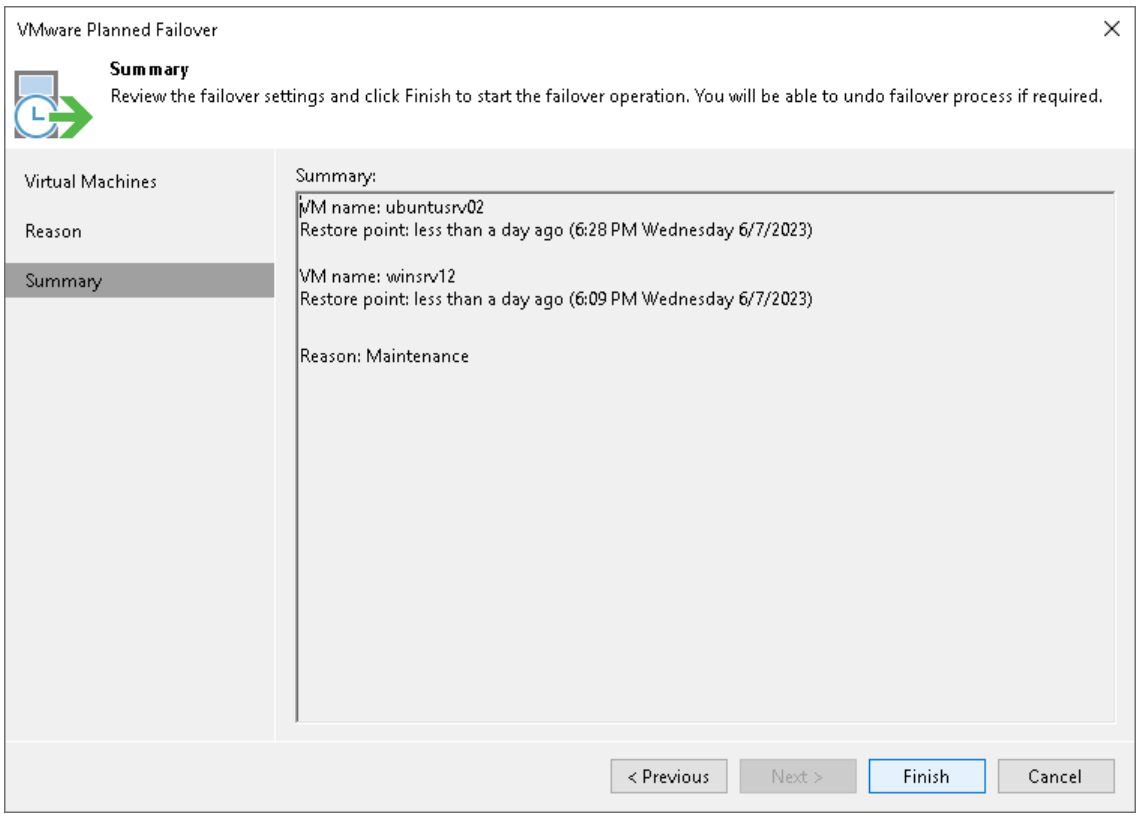

# Failover Undo

Failover undo is one of the ways to finalize failover. When you undo failover, you switch back from a VM replica to the source VM. Veeam Backup & Replication discards all changes made to the VM replica while it was in the *Failover* state. You can use the undo failover scenario if you have failed over to the VM replica for testing and troubleshooting purposes and want to get back to the normal operation mode.

The failover undo operation is performed in the following way:

1. Veeam Backup & Replication reverts the VM replica to its pre-failover state. To do this, Veeam Backup & Replication powers off the VM replica and gets it back to the state of the latest snapshot in the snapshot chain. Changes that were written to the snapshot delta file while the VM replica was in the *Failover* state are discarded.
2. The state of the VM replica gets back to *Ready*, and Veeam Backup & Replication resumes replication activities for the source VM on the source host.

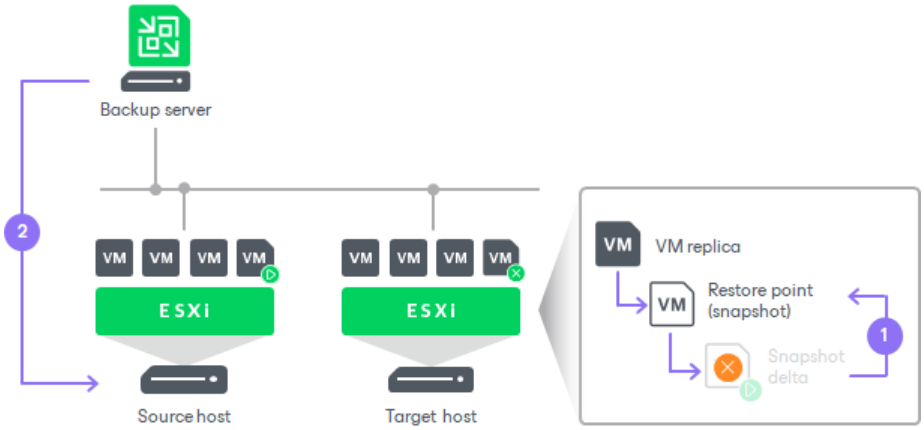


# Undoing Failover

For more information on failover undo, see Failover Undo and Failover and Failback for Replication.

To undo failover:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas > Active** node.
3. In the working area, select the necessary replica and click **Undo Failover** on the ribbon. As an alternative, right-click one of the selected replicas and click **Undo failover**.
4. In the displayed window, do the following:
   a. If you want to force failover undo, select the **Force undo failover** check box.

      When you force failover undo, Veeam Backup & Replication attempts to perform the failover undo operation in a regular way. If the host on which the VM replica resides is unavailable, Veeam Backup & Replication changes the VM replica state to *Ready* in the configuration database and console. This helps avoid failure of the failover undo operation.

b. Click **Yes**.

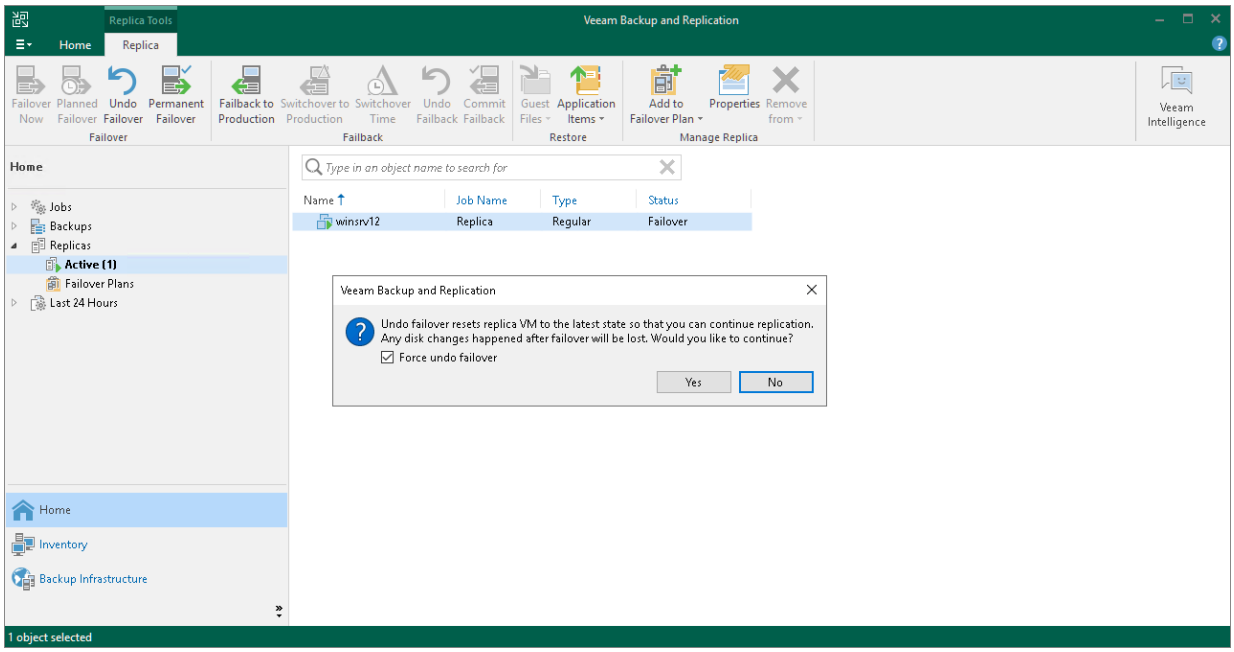
Replica Tools
Veeam Backup and Replication
Home
Replica
Failover Now
Planned Failover
Undo Failover
Permanent Failover
Failover
Failback to Production
Switchover to Production
Switchover Time
Undo Failback
Commit Failback
Failback
Guest Files
Application Items
Restore
Add to Failover Plan
Properties
Remove from
Manage Replica
Veeam Intelligence
Home
Jobs
Backups
Replicas
Active (1)
Failover Plans
Last 24 Hours
Type in an object name to search for
Name
Job Name
Type
Status
winsrv12
Replica
Regular
Failover
Veeam Backup and Replication
Undo failover resets replica VM to the latest state so that you can continue replication. Any disk changes happened after failover will be lost. Would you like to continue?
Force undo failover
Yes
No
Home
Inventory
Backup Infrastructure
1 object selected

# Failback

Failback is one of the ways to finalize failover. When you perform failback, you switch back to the production VM from a VM replica, shift I/O processes from the disaster recovery site to the production site.
Veeam Backup & Replication also sends all change made to the VM replica while it was in the *Failover* state to the production VM. However, note that these changes are only sent to the production VM but not published.

Veeam Backup & Replication provides you the following options to perform failback:

- You can fail back to the source VM in the original location.
- You can fail back to a VM already recovered to a new location. This VM must be recovered before you perform failback. For example, you can recover the VM from a backup.
- You can fail back to a VM recovered from a replica to a new location, or to any location but with different settings. The VM will be recovered from the replica during the failback process.

The first two options help you decrease recovery time and the use of the network traffic because Veeam Backup & Replication needs to transfer only differences between the source/recovered VM and VM replica. For the third option, Veeam Backup & Replication needs to transfer the whole VM data, including its configuration and virtual disk content. Use the third option if there is no way to use the source VM or restore it from a backup.

Veeam Backup & Replication performs failback in two phases:

- **First phase**: Veeam Backup & Replication synchronizes the state of the production VM (the source VM, an already recovered VM or a VM that will be recovered from the replica) with the current state of its replica. This phase may take a lot of time especially if VM is large. While Veeam Backup & Replication performs the first phase of failback, VM replicas are still up and running, users can access these VMs and perform daily routine tasks as normal.
- **Second phase**: Veeam Backup & Replication switches all processes from the VM replica to the production VM, turns off the replica and also sends to the production VM changes made to the VM replica since the end of the first phase.

  The time when the second phase starts depends on how you want to switch from the replica to the production VM. You can switch to the production VM automatically, at the scheduled time or manually. If you select to switch automatically, the second phase will start right after the first phase finishes. If you select to switch at the scheduled time or manually, the second phase will start at the time you want.

The process of failing back to the source VM or an already recovered VM differs from the process of failing back to a VM recovered from a replica:

- How failback to the source VM and already recovered VM works.
- How failback to a VM recovered from a replica works.

## How Failback to Source VM or Already Recovered VM Works

When you fail back to the source VM or an already recovered VM, Veeam Backup & Replication performs the following operations during the first phase:

1. If the production VM is running, Veeam Backup & Replication powers it off.
2. Veeam Backup & Replication creates a working failback snapshot for the production VM.
3. Veeam Backup & Replication creates a failback protective snapshot for the VM replica. You can use this snapshot to return to the pre-failback state of the VM replica afterwards.

4. Veeam Backup & Replication calculates the difference between disks of the production VM and disks of the VM replica in the Failover state. Difference calculation helps Veeam Backup & Replication understand what data needs to be transferred to the production VM to synchronize its state with the state of the VM replica.

   [For VMware vSphere version prior to 7.0] If you fail back to the source VM in the original location and you have enabled the **Quick rollback** option, difference calculation can be performed much faster than without this option enabled. For more information on quick rollback, see Quick Rollback.

5. Veeam Backup & Replication transfers the data that was detected at the previous step to the production VM. The transferred data is written to the delta file of the working failback snapshot on the production VM.

6. Veeam Backup & Replication removes the working failback snapshot from the production VM.

7. Veeam Backup & Replication changes the state of the VM replica from *Failover* to *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication creates a working failback snapshot on the production VM.

2. The guest OS of the VM replica is shut down or the VM replica is powered off.

   If VMware Tools are installed on the VM replica, Veeam Backup & Replication tries to shut down the replica guest OS. If nothing happens after 15 minutes, Veeam Backup & Replication powers off the VM replica. If VMware Tools are not installed on the VM or the VM is suspended, Veeam Backup & Replication powers off the VM. The VM replica remains powered off until you commit failback or undo failback.

3. Veeam Backup & Replication creates a failback protective snapshot for the VM replica. The snapshot acts as a new restore point and saves the pre-failback state of the VM replica. You can use this snapshot to return to the pre-failback state of the VM replica afterwards.

4. Sends data changed on the VM replica while it was in the *Ready to switch* state to the working failback snapshot on the production VM.

5. Veeam Backup & Replication removes the protective snapshot from the VM replica.

6. Veeam Backup & Replication removes the working failback snapshot from the production VM. Changes written to the delta file of this snapshot are committed to the production VM disks.

7. The state of the VM replica is changed from *Ready to switch* to *Failback*. Veeam Backup & Replication temporarily puts replication activities for the production VM on hold.

8. [If you fail back to a VM already recovered to a new location] Veeam Backup & Replication updates the ID of the source VM in the Veeam Backup & Replication configuration database. The ID of the source VM is replaced with the ID of the recovered VM.

9. If you have selected to power on the production VM after failback, Veeam Backup & Replication powers on the production VM on the host.

# How Failback to VM Recovered from Replica Works

When you fail back to a VM recovered from a replica, Veeam Backup & Replication performs the following operations during the first phase:

1. Veeam Backup & Replication requests vCenter Server to create on the target host an empty VM with the same configuration as the VM replica. vCenter Server registers the created production VM.

2. Veeam Backup & Replication creates a working failback snapshot for the production VM.

3. Veeam Backup & Replication creates a failback protective snapshot for the VM replica. You can use this snapshot to return to the pre-failback state of the VM replica afterwards.
4. Veeam Backup & Replication transfers data of the VM replica to the production VM to update the production VM state to the VM replica state.
5. Veeam Backup & Replication removes the working failback snapshot from the production VM.
6. Veeam Backup & Replication changes the state of the VM replica from *Failover* to *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the same operations as described in section How Failback to Source VM or Already Recovered VM Works.

# Finalizing Failback

Failback is an intermediate step that needs to be finalized. If the production VM works as expected and you want to get back to it, commit failback. If the VM does not work as expected, undo failback.

# Failback on VSAN

Due to specifics of VSAN data storage organization, Veeam Backup & Replication cannot get the difference between disks of a VM replica located on VSAN and disks of the source VM in a regular manner.
Veeam Backup & Replication needs to read VM disks data anew in every failback process phase. As a result, failback for VMs replicas on VSAN slightly differs from the regular failback course.

Before Veeam Backup & Replication starts the failback process, it checks the location of VM replica disks. If at least one disk is located on VSAN, Veeam Backup & Replication performs failback in the following way:

1. Veeam Backup & Replication creates a working failback snapshot for the source VM.
2. For every VM disk, Veeam Backup & Replication performs the following actions:
    a. If you fail back to the source VM location, Veeam Backup & Replication calculates the difference between the VM replica disk and the source VM disk. To do this, Veeam Backup & Replication reads the whole amount of disk data from VSAN, and transfers only changed data to the source VM side.
    b. If you fail back to a new location, Veeam Backup & Replication transfers the whole disk without calculating the difference.
3. The VM replica is powered off.
4. Veeam Backup & Replication creates a protective failback snapshot for the VM replica. Using the protective failback snapshot, Veeam Backup & Replication detects what changes took place on the VM replica while VM disk data was being transported. As well as before, Veeam Backup & Replication reads the whole amount of VM disks data but transports only those data blocks that have changed since the VM disks transfer.

The rest of the failback process does not differ from the regular failback process.

# Quick Rollback

If you fail back from a VM replica to the source VM in the original location, you can instruct Veeam Backup & Replication to perform quick rollback. Quick rollback significantly reduces the failback time and has little impact on the production environment.

During failback with the quick rollback option enabled, Veeam Backup & Replication does not calculate digests for entire VM replica disks to get the difference between the source VM and VM replica. Instead, it queries CBT to get information about disk sectors that have changed, and calculates digests only for these disk sectors. As a result, digest calculation is performed much faster. After that, Veeam Backup & Replication performs failback in a regular way: transport changed blocks to the source VM, powers off the VM replica and synchronizes the source VM with the VM replica once again.

It is recommended that you use quick rollback if you fail back to the source VM after a problem that has occurred at the level of the guest OS of the VM replica — for example, there has been an application error or a user has accidentally deleted a file on the VM replica guest OS. Do not use quick rollback if the problem has occurred at the VM hardware level, storage level or due to a power loss.

# Requirements for Quick Rollback

To perform quick rollback, make sure that the following requirements are met:

- You must perform failback to the source VM in the original location.
- CBT must be enabled for the source VM.
- The VM replica must be created with the **Use changed block tracking data** option enabled.

# Limitations for Quick Rollback

The following limitations apply to quick rollback:

- Due to changes in VMware vSphere 7.0 and later, the replica failback operation forces digest recalculation for both source and target VMs. That is why the **Quick rollback** option is ignored for ESXi hosts starting from version 7.0.
- During the first replication job session after failback with quick rollback, CBT on the source VM is reset. Due to that Veeam Backup & Replication will read data of the entire VM.
- Quick rollback can be performed in the Direct NFS access, Virtual appliance, Network transport mode. The Direct SAN access transport mode cannot be used for quick rollback due to VMware limitations.

# Performing Failback

For more information on failback, see Failback and Failover and Failback for Replication.

To switch back to the source VM, use the **Failback** wizard.

## Before You Begin

Before you perform failback, check the following prerequisites:

- VMs for which you plan to perform failback must be successfully replicated at least once.
- VM replicas must be in the *Failover* state.

## Step 1. Launch Failback Wizard

To launch the **Failback** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from replica** > **Entire replica** > **Failback to production**.
- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, select the necessary replica and click **Failback to Production** on the ribbon.
- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, right-click the necessary replica and select **Failback to production**.

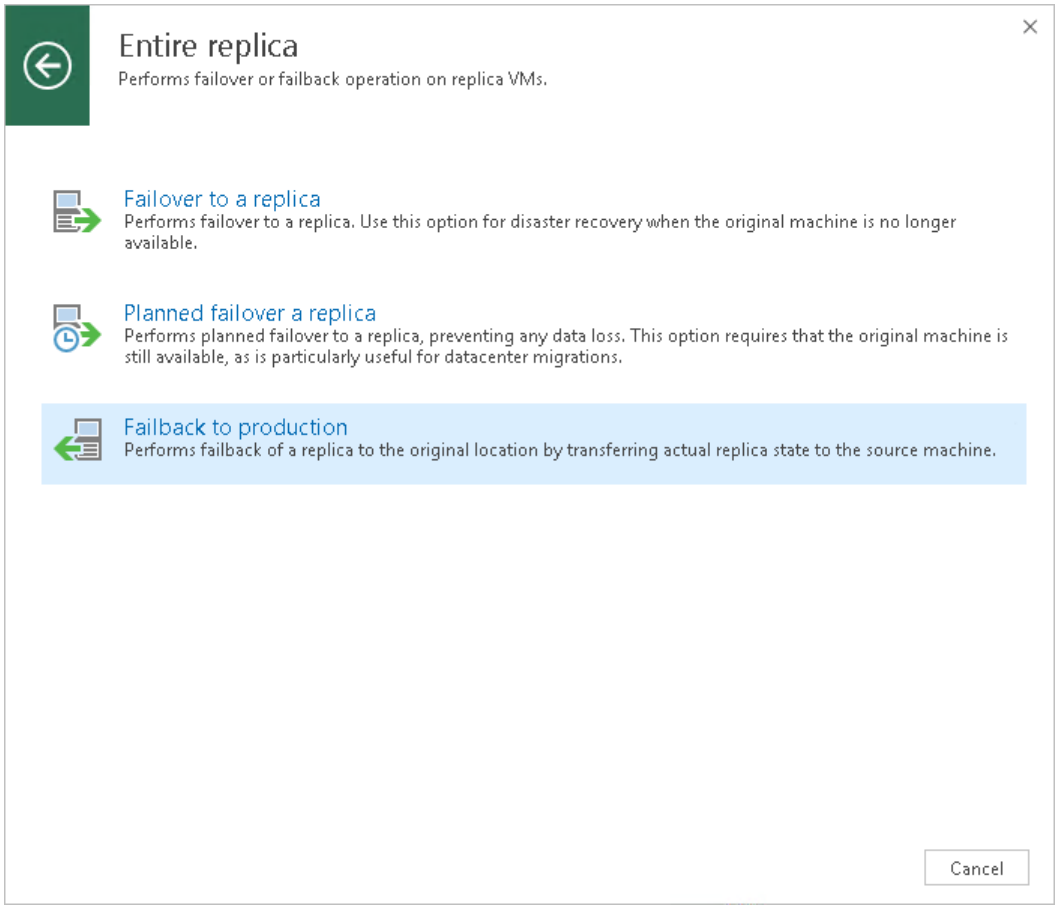

## Step 2. Select VM Replicas to Fail Back

At the **Replica** step of the wizard, select replicas from which you want to fail back.

To update the list of replicas that are ready for failback (replicas in the *Failover* state), click **Populate**.

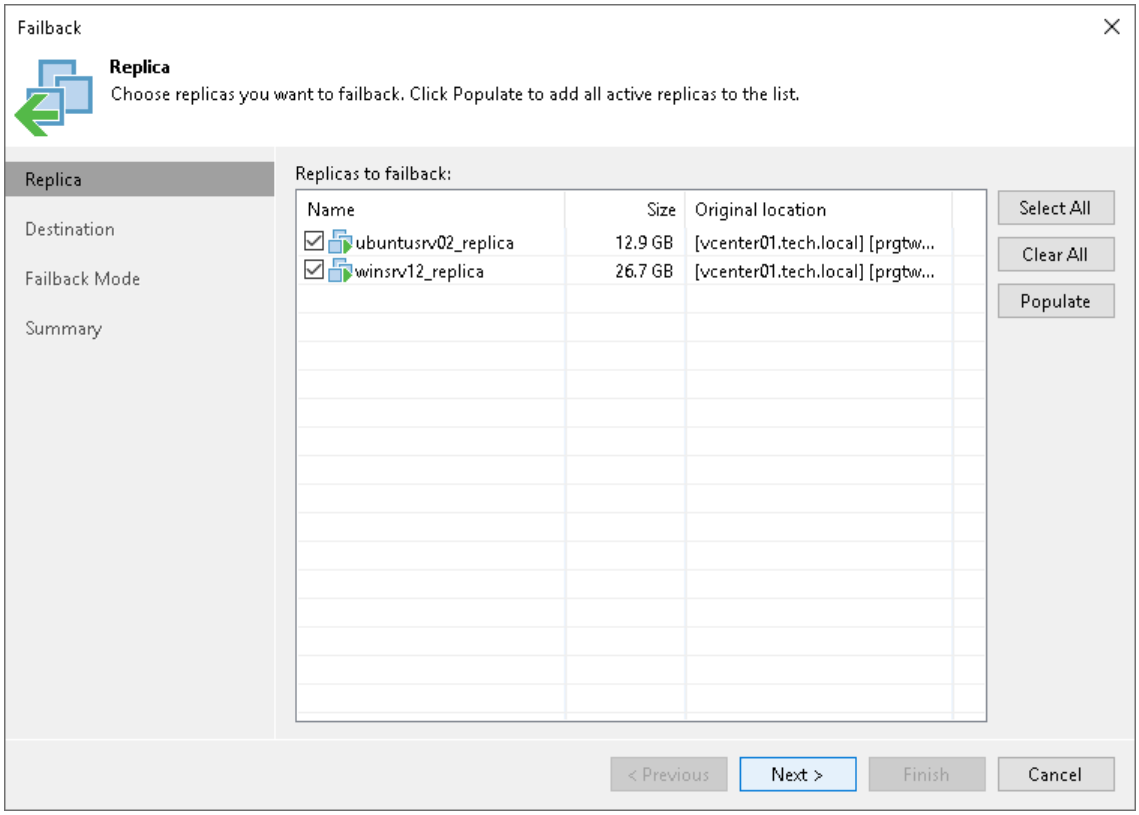

## Step 3. Select Failback Destination

At the **Destination** step of the wizard, select a failback destination and backup proxies for data transfer during failback:

1. Select a destination for failback. Veeam Backup & Replication supports the following options:
   - **Failback to the original VM** – select this option if you want to fail back to the source VMs that reside on the source hosts. Veeam Backup & Replication will synchronize the state of the source VMs with the current state of their replicas to apply any changes that occurred to the VM replicas while running in the DR site.

     If you select this option, you will proceed to the Failback Mode step of the wizard.
   - **Failback to the original VM restored in a different location** – select this option if the source VMs have already been recovered to a new location, and you want to switch to the recovered VMs from their replicas. Veeam Backup & Replication will synchronize the state of the recovered VMs with the current state of the VM replicas to apply any changes that occurred to the replicas while running in the DR site.

     If you select this option, you will proceed to the Target VM step of the wizard.
   - **Failback to the specified location** – select this option if you want to recover VMs from replicas. You can recover VMs to a new location, or to any location but with different settings (such as network settings, virtual disk type, configuration file path and so on). Select this option if there is no way to fail back to the source VM or an already recovered VM.

     If you select this option, the wizard will include additional steps.

   If you select one of the first two options, Veeam Backup & Replication will send to the source/recovered VMs only differences between the existing virtual disks. Veeam Backup & Replication will not send replica configuration changes such as different IP address or network settings (if replica Re-IP and network mapping were applied), new hardware or virtual disks added while the replicas were in the *Failover* state.

   If you select **Failback to the specified location**, Veeam Backup & Replication will send to the specified location whole VM data, including VM configurations and virtual disk content.

2. To select which backup proxies will be used for data transfer, click **Pick backup proxies for data transfer**. By default, Veeam Backup & Replication selects proxies automatically.

   If you leave automatic proxy selection, Veeam Backup & Replication will check available backup proxies before processing each VM from the VM list. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use, the current workload on the backup proxies to select the most appropriate resource for VM processing.

   If you want to assign proxies manually, use the following instructions. If VMs and their replicas reside in different sites, select at least one backup proxy in the production site and one proxy in the disaster recovery site. If VMs and replicas reside in the same site, you can use the same backup proxy as the source and target one. We recommend you to select at least two backup proxies in each site to ensure that failback will be performed in case one proxy fails or looses the network connection.

3. [For VMware vSphere prior to version 7.0; for failback to the original VMs] If you want to fasten failback, and the source VMs had problems at the guest OS level, select the **Quick rollback** check box. For more information on quick rollback, its requirements and limitations, see Quick Rollback.

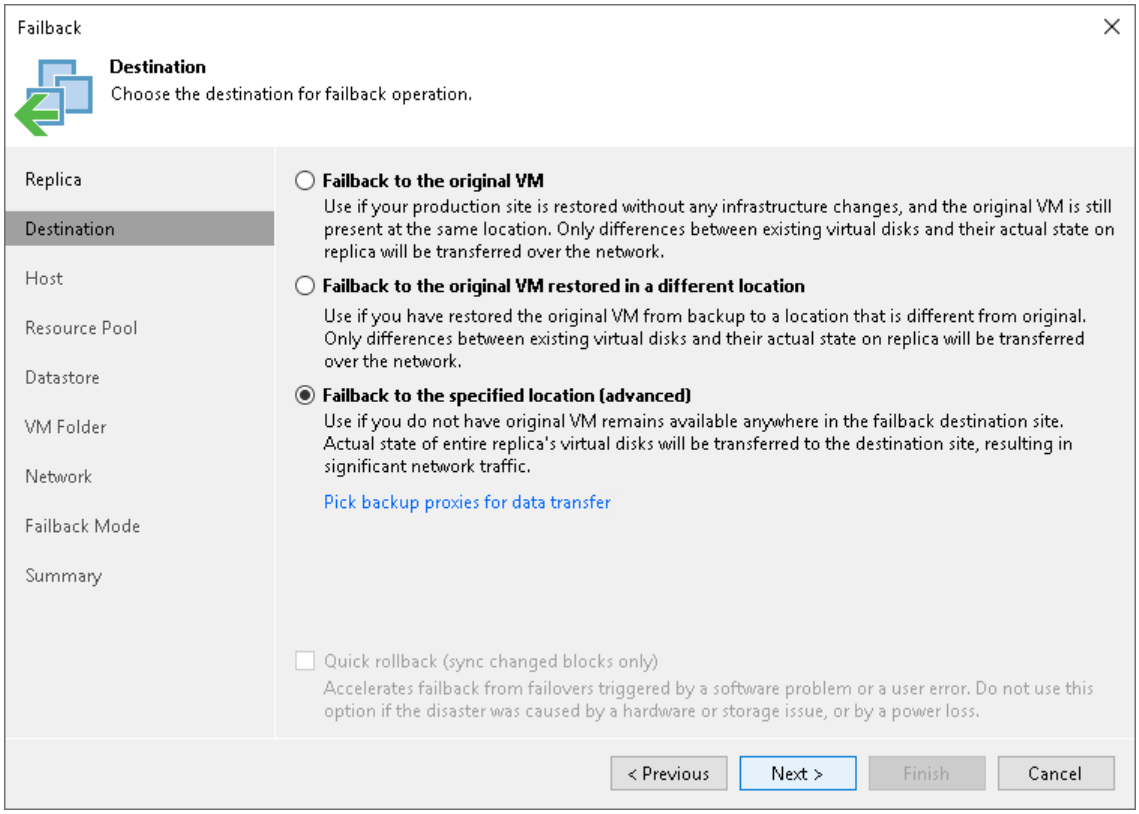


# Restoring Storage Policies

If the replicated VM was associated with the storage policy, in the failback to source location scenario, Veeam Backup & Replication will associate the restored VM with this storage policy.

When you click **Next**, Veeam Backup & Replication will check storage policies in the virtual environment and compare this information with the information about the replica storage policy. If the original storage policy has been changed or deleted, Veeam Backup & Replication will display a warning. You can select one of the following options:

- **Current** – the restored VM will be associated with the profile with which the source VM in the production environment is currently associated.
- **Default** – the restored VM will be associated with the profile that is set as default for the target datastore.
- **Stored** – the restored VM will be associated with the profile that was assigned to the source VM at the moment of replication.

For more information, see Storage Profiles.

# Step 4. Select Target Host or Cluster

The **Host** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Host** step of the wizard, specify names for the recovered VMs and hosts or clusters where the recovered VMs will be registered. To do this, select the necessary VMs and use the **Host** and **Name** buttons.

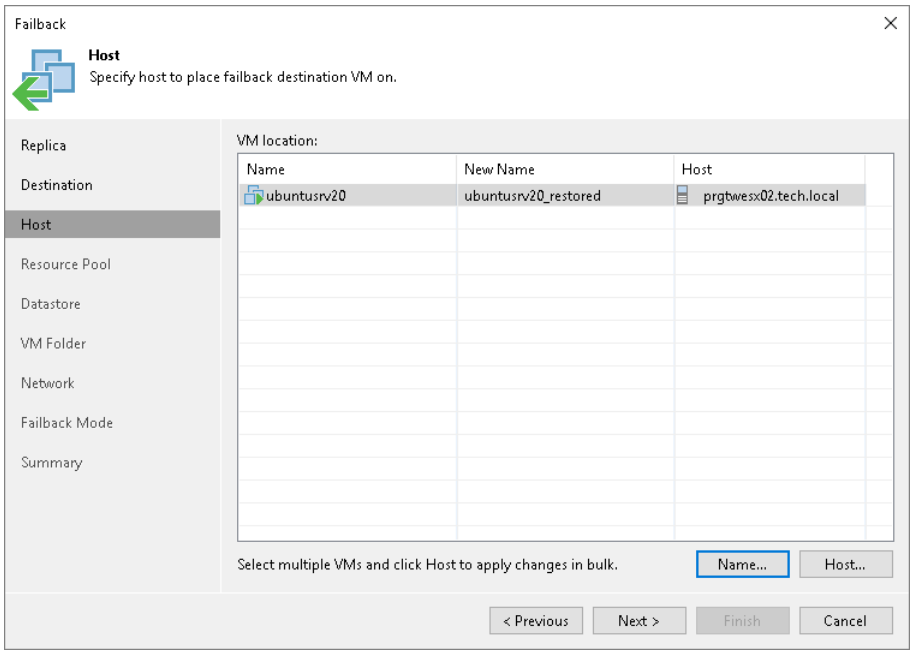

# Step 5. Select Target Resource Pool

The **Resource Pool** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Resource Pool** step of the wizard, select resource pools to which the recovered VMs will be added. To do this, select VMs that you want to add to the same resource pool, click **Pool** and select the necessary resource pool in the **Select Resource Pool** window.

As an alternative, you can select a vApp to which the restored VM will be included. To find the necessary vApp, at the left bottom corner of the **Select Resource Pool** window, click the resource pool icon ( ) and select *VirtualApp*.

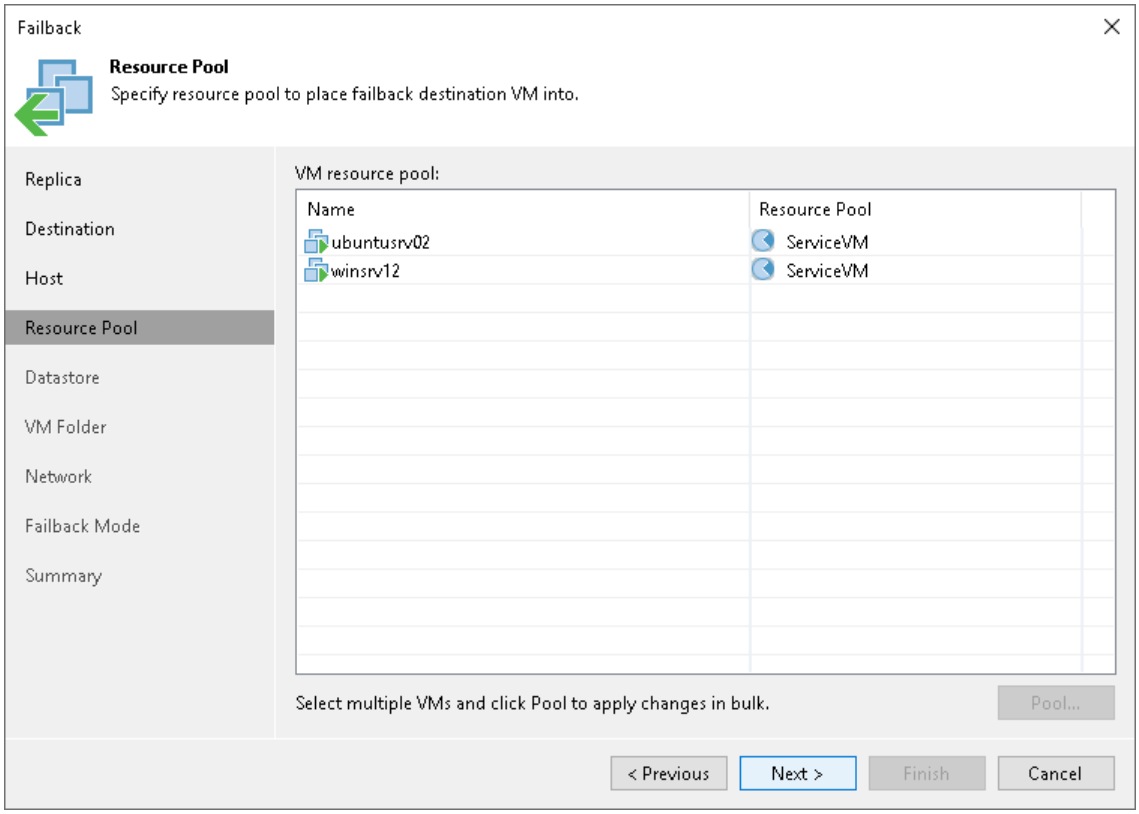

# Step 6. Select Target Datastore

The **Datastore** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Datastore** step of the wizard, specify datastores where you want to store configuration files and disk files of VMs that will be recovered. Also, you can change disk types.

1. To change a datastore where VM files will be stored, select the necessary VMs and click **Datastore**. From the list of available datastores, select the necessary datastore.

   If configuration and disk files of VMs must be placed to different datastores, select files of the necessary type, click **Datastore** and select the necessary datastore.

2. To change disk type settings, select the necessary disk files and click **Disk Type**. In the **Disk Type Settings** window, select the necessary disk type. For more information about disk types, see VMware docs.

   By default, Veeam Backup & Replication preserves disk types of the source VMs.

   > **NOTE:**
   >
   > You can change disk types only for VMs with Virtual Hardware version 7 or later.

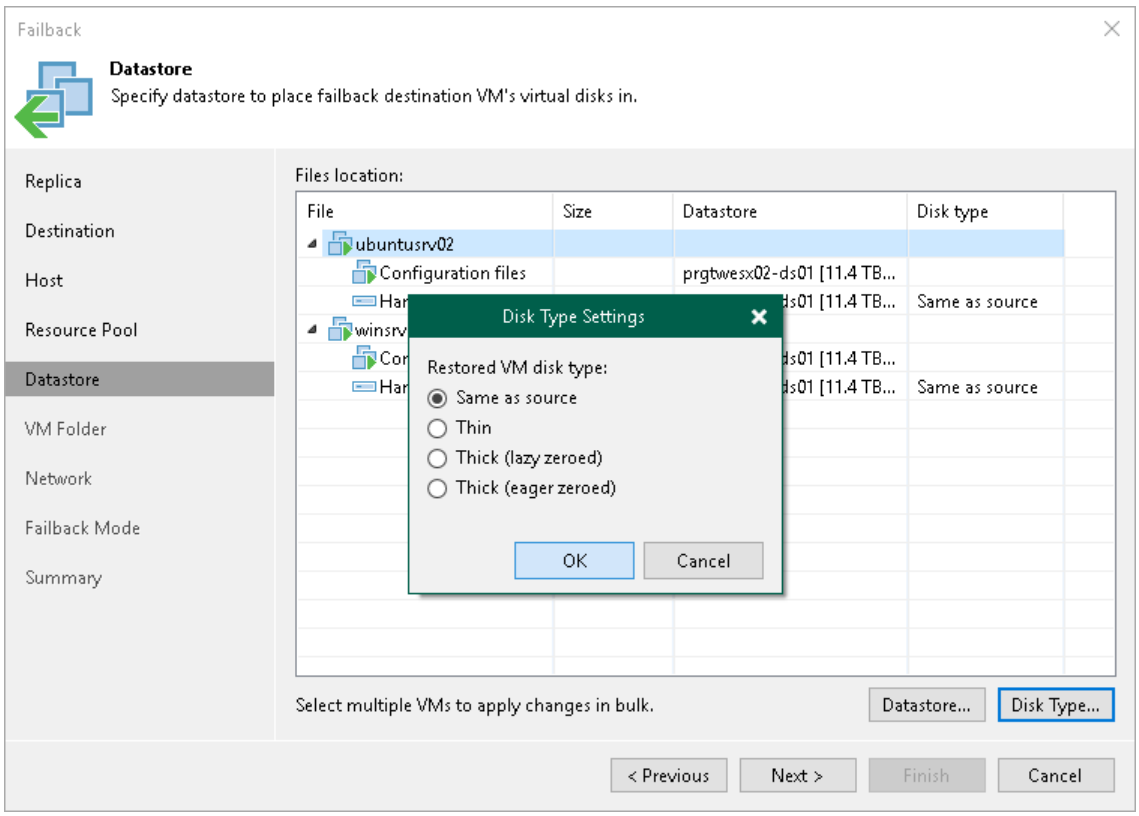

# Step 7. Select Target Folder

The **VM Folder** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **VM Folder** step of the wizard, specify folders in the target datastores where all files of the recovered VMs will be stored.

If you want the recovered VMs to have the same tags as the source VMs, select the **Restore VM Tags** check box.

> **NOTE**
>
> Consider the following:
>
> - You can select destination folders only if you recover VMs to destinations other than standalone hosts.
> - You can recover VM tags only if you recover VMs to their original locations, and the source VM tags are still available on the source vCenter Server.

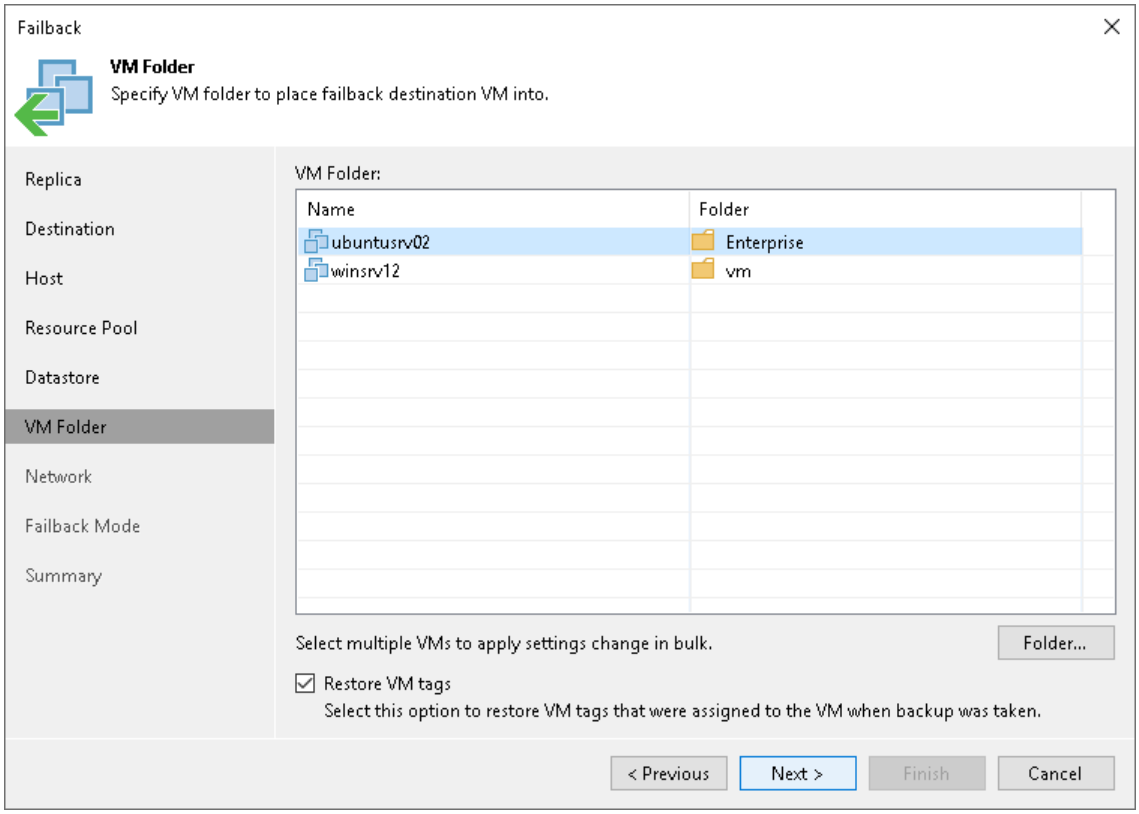

# Step 8. Select Target Network

The **Network** step is available if you have selected the **Failback to the specified location** option at the Destination step. This step applies if you fail back to VMs recovered to new locations, and if networks in those locations differ from networks in the disaster recovery (DR) site.

At the **Network** step of the wizard create a network mapping table. This table maps networks in the DR site to networks in the site where the recovered VMs will reside. Veeam Backup & Replication will use the network mapping table to update configuration files of VMs on the fly, during the failback process.

To change networks to which the restored VMs will be connected:

1. In the **Network connections** list, select the necessary VMs and click **Network**.

   If VMs are connected to multiple networks, select networks which you want to map.

2. In the list of available networks, select a network to which the recovered VMs will be connected.

If you do not want to connect the recovered VMs to any virtual network, select the necessary VMs and click **Disconnect**.

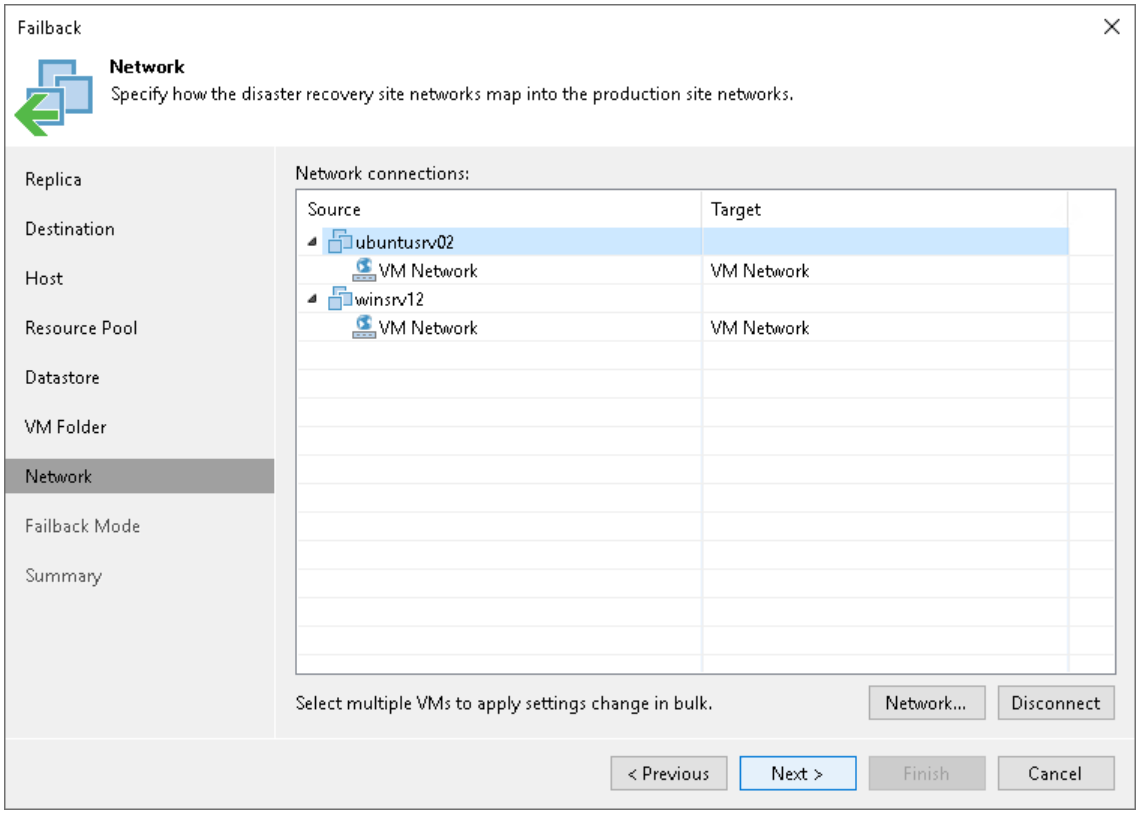

# Step 9. Map Replicas to Restored VMs

The **Target VM** step is available if you have selected the **Failback to the original VM restored in a different location** option at the Destination step.

At the **Target VM** step of the wizard, specify to which VMs you want to fail back from replicas. These VMs must be already restored from backups in the required location.

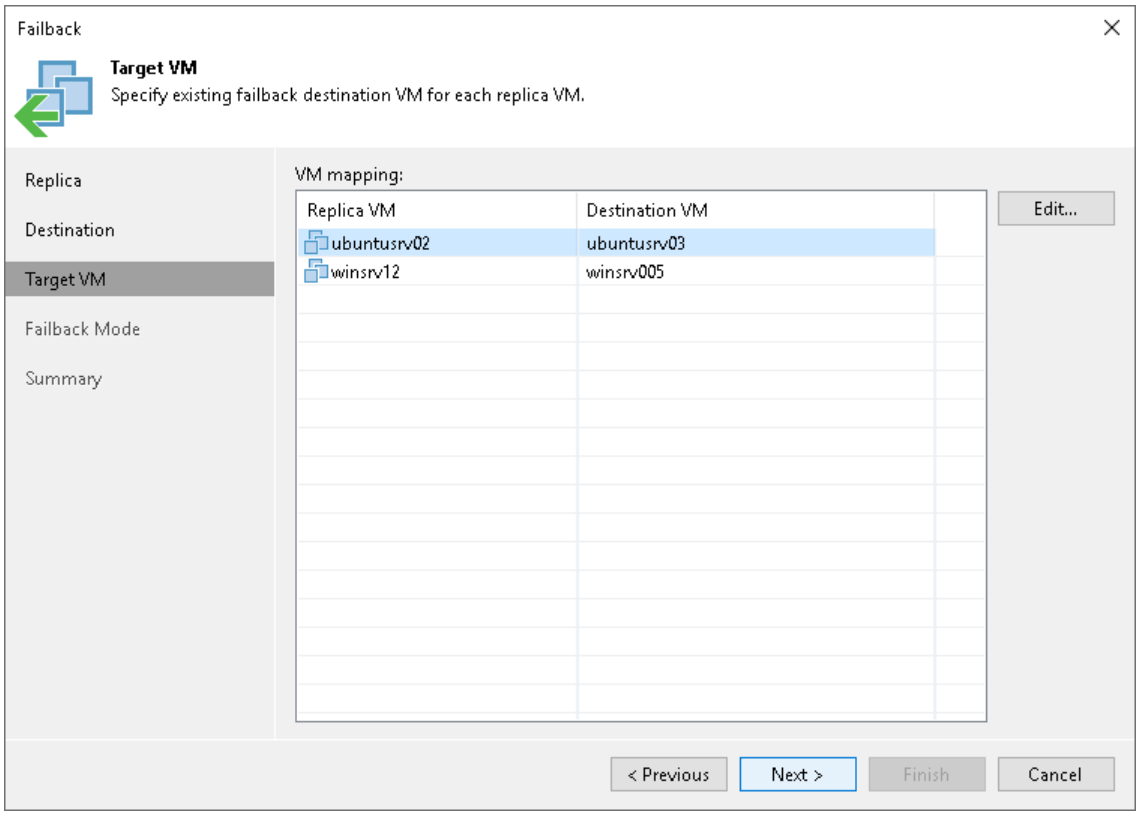

# Step 10. Schedule Switch to Production VMs

At the **Failback Mode** step of the wizard, specify when switch from replicas to production VMs must be performed:

- Select **Auto** if you want Veeam Backup & Replication to perform the switch automatically right after the state of the production VMs is synchronized with the state of their replicas.
- Select **Scheduled** if you want Veeam Backup & Replication to perform the switch at a specific time.
- Select **Manual** if you want to perform the switch manually.

If you select the **Scheduled** or **Manual** option, you can further reset/set the scheduled time or switch to the production VM manually. For more information, see Changing Switching Time and Switching Replicas to Production VMs Manually.

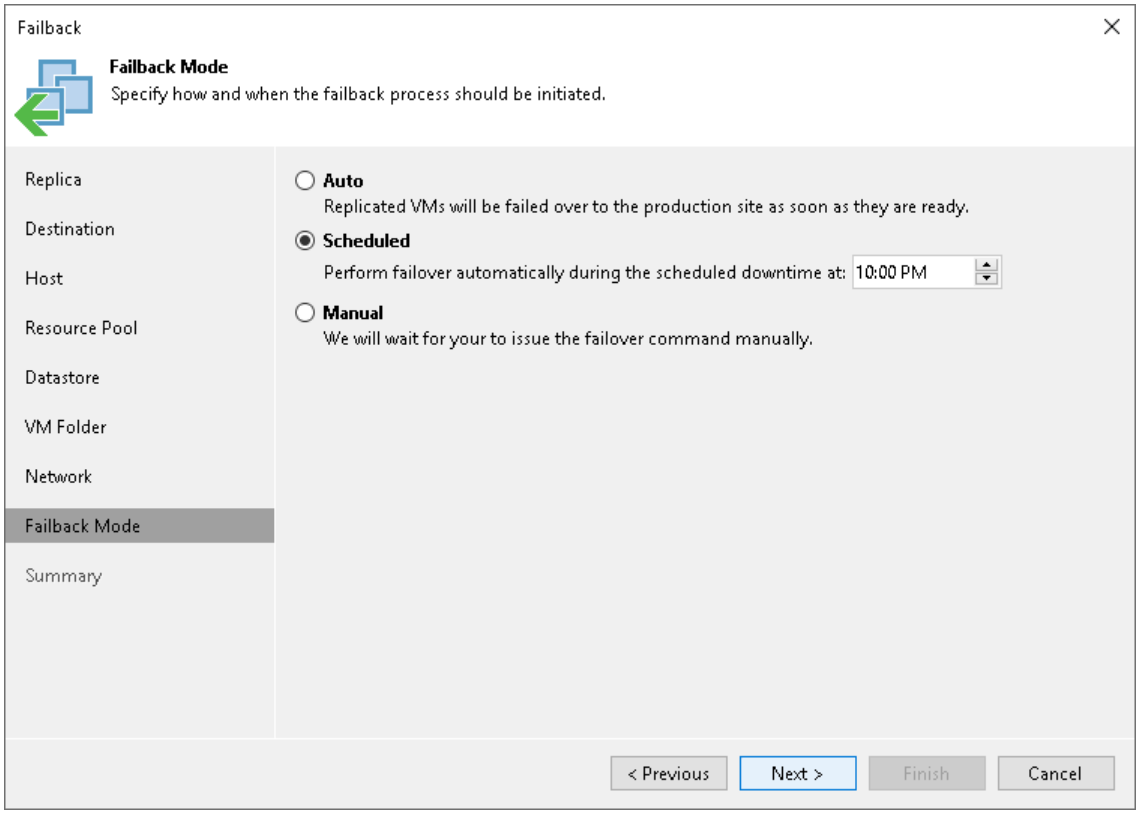

## Step 11. Review Summary and Finish Working with Wizard

At the **Summary** step of the wizard, review the configured failback settings and click **Finish**.

If you want to power on the production VMs right after the switch to production operation is performed, select the **Power on target VM after restoring** check box.

# What You Do Next

Failback is an intermediate step that needs to be finalized. You can finalize failback in the following ways:

- Commit failback.
- Undo failback.

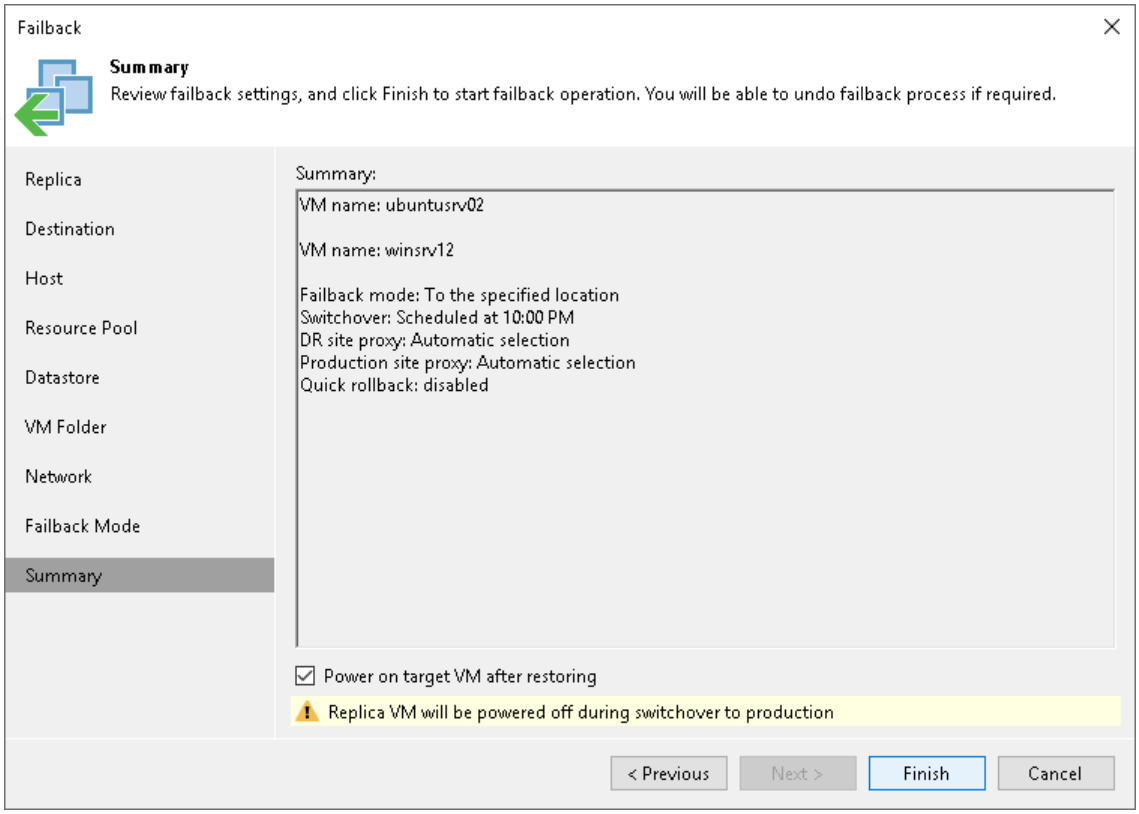


# Switching to Production VMs Manually

The following instructions apply if you have selected to switch from replicas to production VMs manually or at the scheduled time at the **Failback Mode** step of the **Failback** wizard.

To switch to a production VM from its replica, do the following:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. Right-click a replica in the *Ready to switch* state and select **Switchover to production**.

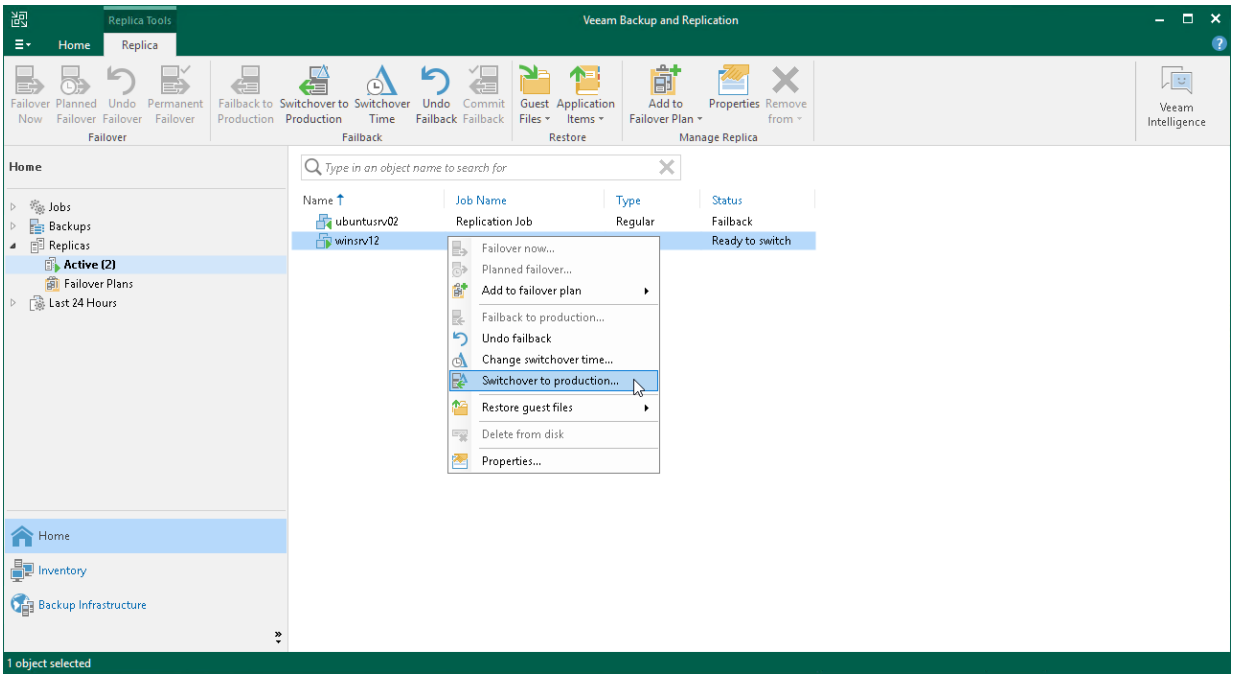


# What You Do Next

After you switch to the production VM, you must finalize failback. You can finalize failback in the following ways:

- Commit failback
- Undo failback

# Changing Switching Time

The following instructions apply if you have selected to switch from replicas to production VMs manually or at the scheduled time at the **Failback Mode** step of the **Failback** wizard.

To change the time when Veeam Backup & Replication will switch from replicas to production VMs:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. Right-click a replica in the *Ready to switch* state and select **Change switching time.**

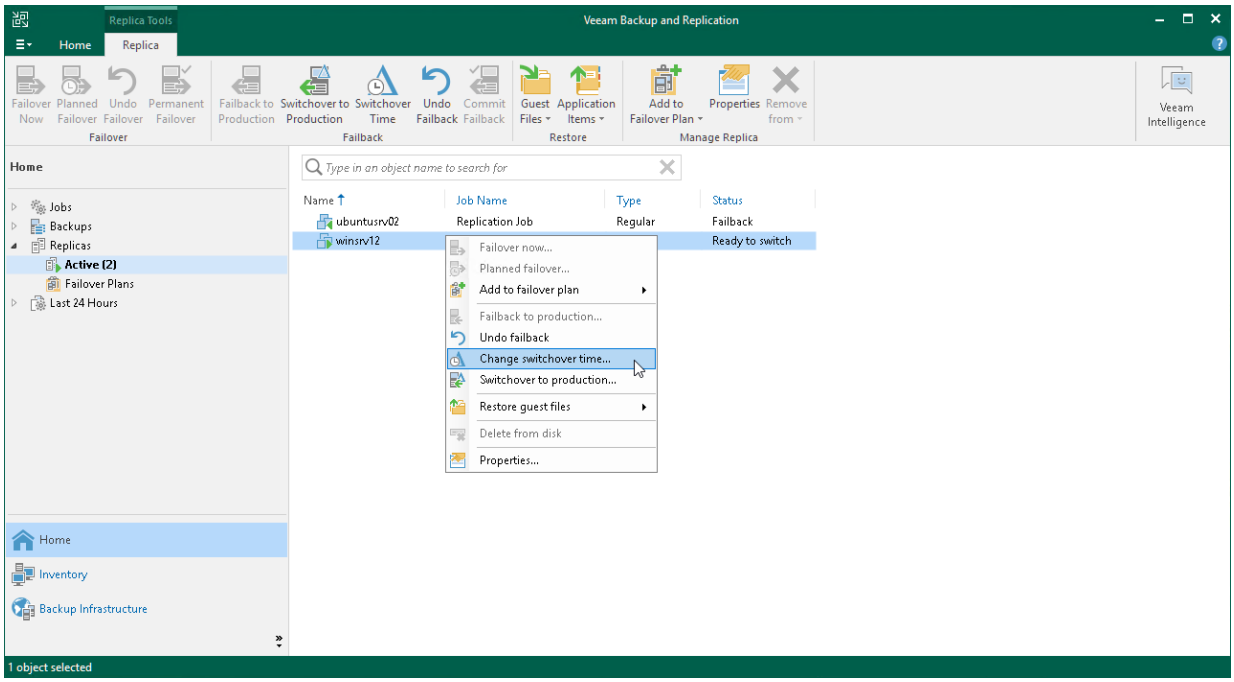

# Failback Commit

Failback commit is one of the ways to finalize failback. When you commit failback, you confirm that the VM to which you failed back (the production VM) and also changes sent to it during failback work as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production VM.

> **NOTE**
>
> If you have selected to switch to the production VM manually, you must first perform the switchover.

The failback commit operation is performed in the following way:

1. Veeam Backup & Replication changes the state of the replica from *Failback* to *Ready*.
2. Further operations depend on whether you have failed back to the source VM or recovered VM:
   - If you have failed back to a VM recovered from a backup or replica, Veeam Backup & Replication reconfigures all existing jobs where the source VM is present and adds the source VM to the list of exclusions. The recovered VM takes the role of the source VM and is included into all jobs instead of the excluded VM. When the replication job starts, Veeam Backup & Replication processes the recovered VM instead of the former source VM.
   - If you have failed back to the source VM, the replication job is not reconfigured. When the replication job starts, Veeam Backup & Replication still processes the source VM.

During failback commit, the failback protective snapshot that saves the pre-failback state of a VM replica is not deleted. Veeam Backup & Replication uses this snapshot as an additional restore point for VM replica. With the pre-failback snapshot, Veeam Backup & Replication needs to transfer fewer changes and therefore puts less load on the network when replication activities are resumed.

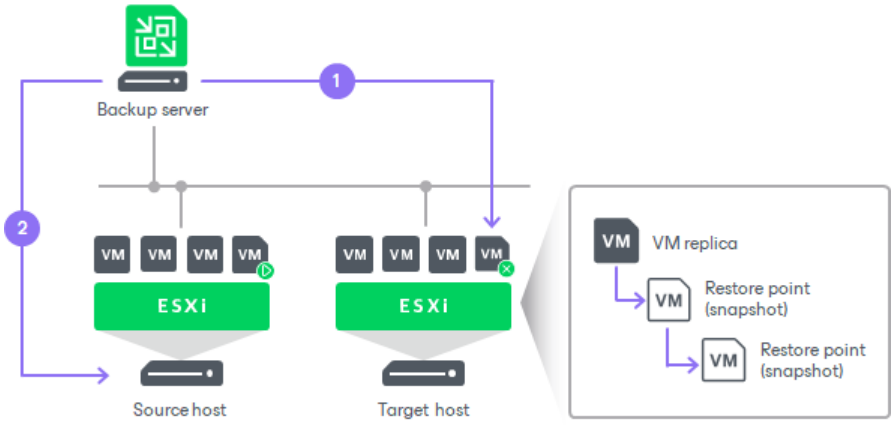


## Committing Failback

For more information on failback commit, see Failback Commit and Failover and Failback for Replication.

To commit failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. In the working area, select the necessary replica and click **Commit Failback** on the ribbon. As an alternative, you can right-click the replica and select **Commit failback**.

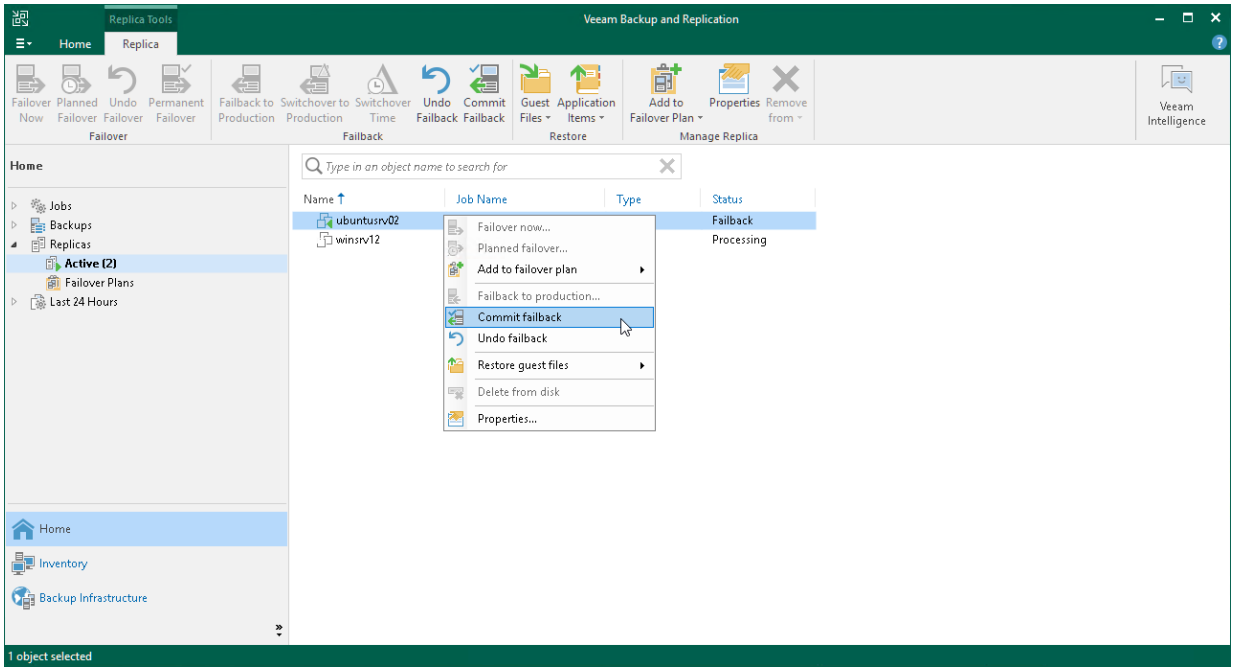

# Failback Undo

Failback undo is one of the ways to finalize failback. When you undo failback, you confirm that the VM to which you failed back (the production VM) and changes sent to it during failback work in a wrong way and you want to get back to the replica.

The failback undo operation is performed in the following way:

1. Veeam Backup & Replication deletes the protective failback snapshot on the VM replica.
2. Veeam Backup & Replication powers on the VM replica and changes the VM replica state from *Failback* to *Failover*.

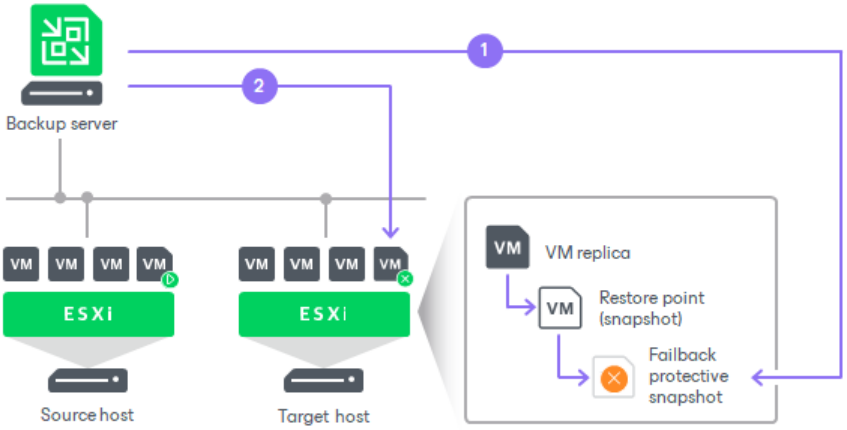


## Undoing Failback

For more information on failback undo, see Failback Undo and Failover and Failback for Replication.

To undo failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. In the working area, select the necessary replica and click **Undo Failback** on the ribbon. Alternatively, you can right-click the necessary replica and select **Undo Failback**.

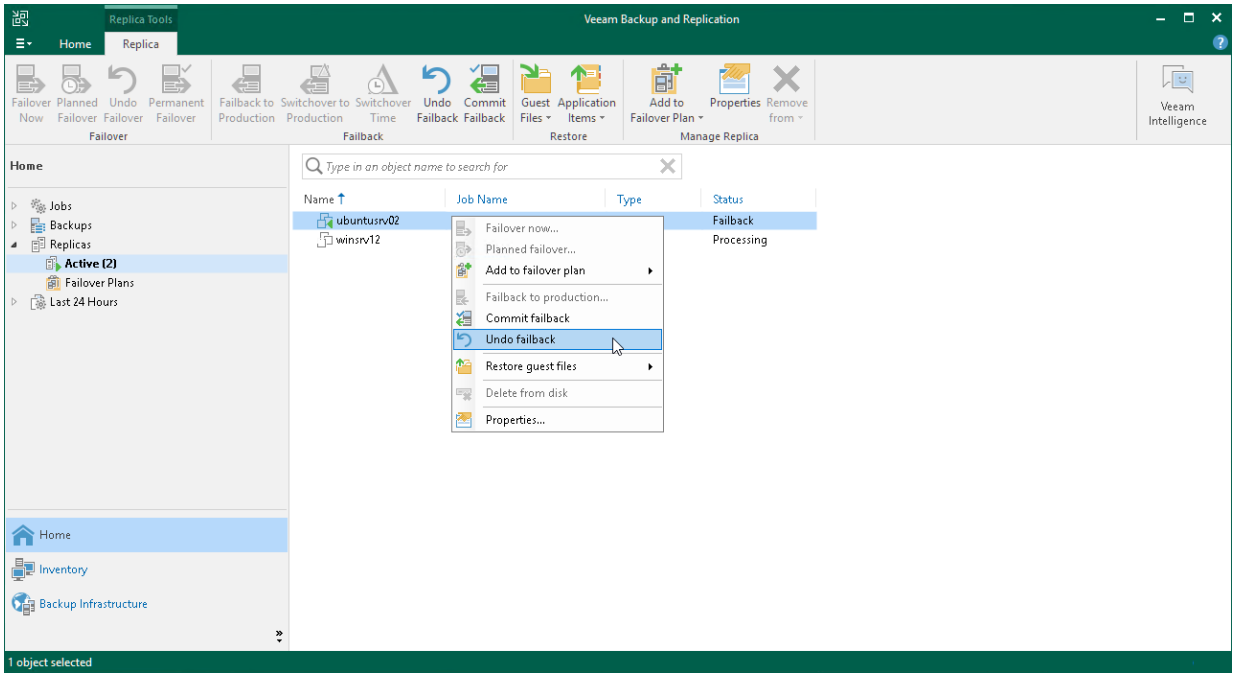

# Continuous Data Protection (CDP)

Continuous data protection (CDP) is a technology that helps you protect mission-critical VMware virtual machines when data loss for seconds or minutes is unacceptable. CDP also provides minimum recovery time objective (RTO) in case a disaster strikes because CDP replicas are in a ready-to-start state.

## Data Replication

First, CDP creates replicas and, then, keeps these replicas up to date.

CDP constantly replicates I/O operations performed on VMs. To read and process I/O operations in transit between the protected VMs and their underlying datastore, CDP uses vSphere APIs for I/O filtering (VAIO) that gives an option not to create snapshots. Because CDP is always on and does not create snapshots, it allows reaching a lower recovery point objective (RPO) compared to the snapshot-based replication — near-zero RPO which means almost no data loss.

Data of I/O operations is stored on the target datastore and relates to short-term restore points. The short-term restore points allow you to recover a VM to a state back to seconds or minutes ago (depending on the RPO that you specify) in case a disaster strikes. Information about short-term restore points is maintained in a special journal. This journal stores records about short-term restore points for a maximum of 168 hours (7 days). If you want to recover a VM to an older state, Veeam Backup & Replication allows you to create additional restore points that contain a VM state back to hours or days ago. Such restore points are called long-term restore points.

To facilitate replication over slow connections, Veeam Backup & Replication optimizes traffic transmission. It filters out unnecessary data blocks such as duplicate data blocks, zero data blocks and compresses replica traffic.

To replicate a VM, you need to configure required backup infrastructure components and create a CDP policy.

## Data Recovery

To recover a VM to a short-term or long-term restore point, you need to fail over to its replica.

When you fail over to a replica, the replica takes over the role of the source VM. After your source VM is repaired, you can fail back to it and transfer all changes that occurred to replica to the source VM. If your source VM cannot be repaired, you can perform permanent failover, that is, permanently switch from the source VM to the VM replica and use this replica as the source VM. For more information, see Failover and Failback for CDP.

# Backup Infrastructure for CDP

The following backup infrastructure components are required for CDP:

- Backup server
- Source and target hosts
- VMware CDP proxies

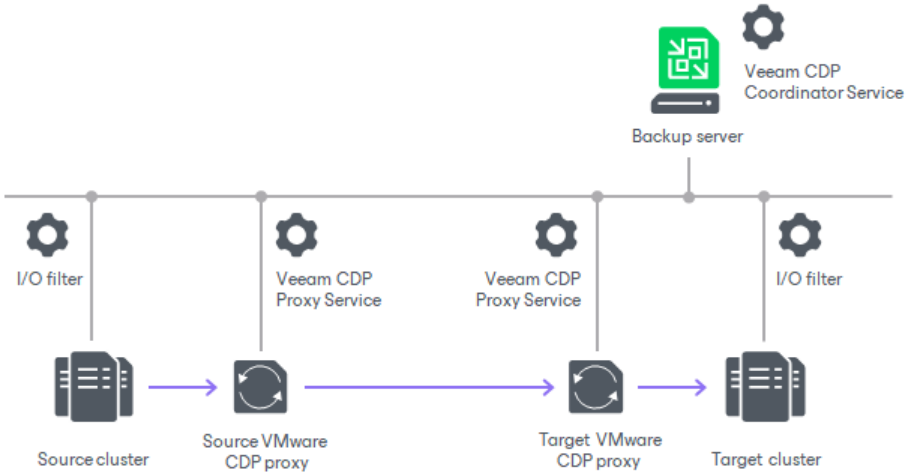


## Backup Server

The backup server is the configuration, administration and management core of the backup infrastructure. The backup server runs the Veeam CDP Coordinator Service. This service coordinates replication and data transfer tasks, and controls resource allocation. We recommend you to place the backup server in the target site or as a separate unit.

For more information on the backup server, see Backup Server.

## Source and Target Hosts

The source and the target hosts are two terminal points between which replicated VM data is moved. The source and target hosts must be a part of the same cluster or two different clusters. In turn, clusters must be managed by the same vCenter Server or two different vCenter Servers connected to the same backup server. For more information on requirements to the hosts and how to add vCenter Servers to the backup infrastructure, see the Requirements and Limitations and Adding VMware vSphere Servers sections.

The source and target hosts perform the following tasks:

- The source host reads VM disk data, reads and processes I/O operations and sends data to source proxies. The data is sent uncompressed.
- The target host receives data from target proxies and saves this data to replicas on the datastore. Also, the target host manages replicas: creates replicas, retains restore points and so on.

# I/O Filter on Hosts

To be able to use hosts for CDP, you must install the I/O filter on each cluster where hosts reside. After you launch the I/O filter installation, Veeam Backup & Replication requests the vCenter Server to install and configure the I/O filter on existing clusters and hosts. For more information on how to install the filter, see Installing I/O Filter.

It is the I/O filter that reads and processes I/O operations in transit between the protected VMs and their underlying datastore and that sends/receives data to/from VMware CDP proxies. Also, the filter communicates with the Veeam CDP Coordinator Service on the backup server and notifies the service that the backup infrastructure must be reconfigured if any proxy becomes unavailable. This I/O filter is built on the basis of vSphere API for I/O filtering (VAIO).

# VMware CDP Proxies

A VMware CDP proxy is a component that operates as a data mover and transfers data between the source and target hosts. We recommend you to configure at least two VMware CDP proxies: one (source proxy) in the production site and one (target proxy) in the disaster recovery site.

The source and target VMware CDP proxies perform the following tasks:

- The source proxy prepares data for short-term restore points from data received from the source host, compresses and encrypts the data (if encryption is enabled in the network traffic rules). Then sends it to the target proxy.
- The target proxy receives the data, decompresses and decrypts it, and then sends to the target host.

For more information on VMware CDP proxies, their requirements, limitations and deployment, see VMware CDP Proxies.

# Considerations and Limitations

If you plan to use CDP to protect your workloads, consider the following requirements and limitations.

## Licensing

CDP is included in the Veeam Universal License. When using a legacy socket-based license, the Enterprise Plus edition is required.

## Platforms and Datastores

- CDP is supported for a limited number of platforms and versions. For more information on platforms and requirements for them, see Supported Platforms and Applications.
- CDP supports specific source and target datastores. For more information, see Veeam CDP Source and Target.

## Infrastructure Components

- The backup server must have at least 16 GB RAM.
- VMware CDP proxy must meet the requirements listed in section VMware CDP Proxy.
- The network between the infrastructure components must suit the load on your infrastructure. We recommend that you use at least 100 Mbps network. Slower networks cannot handle the generated disk traffic without interruption. The better performance is proven on 10 Gbps networks or faster and MTU 9000.
- You must configure CDP for one cluster on one backup server only. If you add a cluster to another backup server and install the I/O filter (filter required for CDP) on it, CDP policies on the first backup server will fail. For more information on the I/O filter, see I/O Filter.
- When you upgrade Veeam Backup & Replication to the current version, you can postpone the upgrade of the I/O filter on the vCenter Servers to a later time. Veeam Backup & Replication supports the postponed upgrade for the minor versions within the 12 version. You must postpone the upgrade or upgrade all clusters on the vCenter Servers. Partially upgraded vCenter Servers or clusters have limited functionality. You cannot add VMs from such vCenter Servers or clusters to CDP policies, commit failback and perform some other operations.
- The maximum number of CDP policies that can be created on one backup server is 100.
- If you upgrade the primary version of hosts in the vCenter Server (for example, from ESXi version 7 to 8), you must upgrade the I/O filter version on this server. Perform this upgrade either alongside the ESXi upgrade or immediately after it. For more information, see Updating and Uninstalling I/O Filter.
- If you add a new cluster or host to a vCenter Server after the I/O filter is installed on the existing clusters and hosts, you must launch the I/O Filter Management wizard. Make sure that check boxes are selected near all clusters where the I/O filter must be present and finish the wizard. After you finish the wizard, Veeam Backup & Replication requests the vCenter Server to install the I/O filter (if required), and to configure it.

## Virtual Machines

- Sizes of all VM disks must be set as integer values.

- VMs that you plan to protect must not have snapshots at the moment the CDP policy starts for the first time. For example, if a VM is already added to a backup job, make sure that the scheduled backup session does not overlap with the CDP policy first run.
- The hardware versions of VMs that you plan to protect must be supported on the target cluster.
- CDP works only for powered on VMs. However, for VMs with IDE hard drives you must power off VMs before the initial synchronization. This is required to apply a new storage policy to VMs. For more information, see How CDP Works.
- CDP is not supported for VMs to which VM storage policies with multiple datastore specific rule sets apply. If you need to define datastores to be used for placement of the VMs, you can add a tag rule for vSAN instead of a tag based placement datastore specific rule. For more information on how to create tag rules for vSAN, see VMware Docs.
- Shared disks, physical RDM and SCSI bus sharing are not supported. Note that vRDM disks are supported.
- One VM can be processed only by one CDP policy.
- During the initial synchronization, Veeam Backup & Replication needs double space allocation for VMs with thick-provisioned disks on the target host. After the initial synchronization finishes, Veeam Backup & Replication frees half of the allocated space for such VMs. This is a part of mechanism required to support protection of disks newly added to the source VMs.
- The maximum number of disks attached to one VM is 50.
- If you change the VM disk size after the replication process starts, Veeam Backup & Replication starts the replication process anew. This action removes restore points created before the size change and starts the initial synchronization.
- The maximum number of long-term restore points per disk is 95.
- For VMs encrypted on VMware side, consider the following:
  - Make sure that the **Allow I/O filters before encryption** parameter is set to *True* for the VM storage policy component. For more details, see VMware vSphere documentation.
  - To create encrypted CDP policy, select the encrypted datastore at the **Destination** step of the wizard. For more information, see Creating CDP Policies.
  - Encryption for transaction log files that contain incremental changes for disks is not supported. For more information, see CDP Replication Chain.

# Replicas

- On the target host, Veeam Backup & Replication does not allow you to migrate replicas using VMware vSphere Storage vMotion. Note that host vMotion is allowed. However, note that during failover, host vMotion is not allowed.
- Replicas can be powered on only using the failover operation; powering on replicas manually is not supported.
- Veeam Backup & Replication supports testing. For more information, see SureReplica.
- SureReplica verification does not prevent CDP policy from running.

# How CDP Works

This section describes how CDP works during replication. To learn how CDP works during data recovery, see Failover and Failback for CDP.

CDP workflow during replication is divided into two parts: backup infrastructure component configuration and data transfer.

During the configuration, Veeam Backup & Replication configures the required backup infrastructure components. Veeam Backup & Replication also reconfigures the components if something changes in the infrastructure. During data transfer, Veeam Backup & Replication creates short-term and long-term restore points by sending disk data blocks and changes made to them. For more information on restore points, see CDP Replication Chain.

Backup infrastructure component configuration and data transfer are constant processes.

## Component Configuration Algorithm

The following steps apply during the initial configuration, that is, when a CDP policy starts for the first time or after the policy is enabled:

1. Veeam CDP Coordinator Service reads policy settings from the configuration database and creates a list of VM tasks to process. For every VM added to the CDP policy, Veeam Backup & Replication creates a new task.
2. Veeam CDP Coordinator Service checks that required backup infrastructure components are available.
3. Veeam CDP Coordinator Service queries information about VMs added to the CDP policy and virtualization hosts from the vCenter Server.
4. Veeam CDP Coordinator Service requests vCenter Server to create on the target host empty replicas with the same configuration as source VMs, but with empty virtual disks. vCenter Server registers the created replicas.
5. Veeam CDP Coordinator Service requests vCenter to apply the Veeam CDP Replication storage policy to virtual disks of VMs on the source ESXi hosts. This storage policy adds the component that is required for CDP and that gets data of all I/O operations. For more information on storage policies, see VMware Docs.

   The Veeam CDP Replication storage policy itself is created on the vCenter Server when you install the I/O filter. For more information on how to install the filter, see Installing I/O Filter.
6. Veeam CDP Coordinator Service selects which VMware CDP proxies will be used for data transfer and sets a number of rules for data transfer, such as network traffic throttling rules and so on.

   If you select automatic proxy selection when configuring the CDP policy, Veeam Backup & Replication analyzes the current workload on CDP proxies and selects a VMware CDP proxy according to the following priority rules (starting from the most preferable one):

   - VMware CDP proxy on a physical machine.
   - VMware CDP proxy on a VM located in the same cluster — that is, the source proxy in the cluster where source VMs are located (on any host), the target proxy in the cluster where VM replicas are located (on any host).
   - Other VMware CDP proxies.

   For more information on how to specify the proxy selection mode, see Specify Data Transfer and Replica Settings.

7. Veeam CDP Coordinator Service sends to the backup infrastructure components configurations required for CDP. This configuration includes such information as RPO, short-term and long-term retention settings.

After the initial configuration finishes, Veeam Backup & Replication starts monitoring the backup infrastructure. If something changes in the infrastructure or CDP policy settings, Veeam Backup & Replication reconfigures the components. Consider the following examples:

- If a VMware CDP proxy becomes unavailable, Veeam CDP Coordinator Service on the backup server gets a notification that this proxy is no longer available. Then, the service selects another proxy.

  > **NOTE**
  >
  > If you add new proxies to the backup infrastructure, Veeam Backup & Replication does not use these proxies for the already created CDP policies — that is, Veeam Backup & Replication does not reconfigure the infrastructure. If you have selected automatic proxy selection for a CDP policy and want to use the newly added proxies, disable and then enable the CDP policy. If you have selected the proxies manually, edit the CDP policy settings and add the required proxies.

- If virtual disks are added to source VMs, Veeam CDP Coordinator Service requests vCenter to create the disks on the target host, applies the storage policy and selects VMware CDP proxies to transfer disk data.
- If new VMs are added to the CDP policy, Veeam CDP Coordinator Service requests vCenter to create replicas with empty disks, applies the storage policy to the VM disks and selects which VMware CDP proxies to use for data transfer.
- If the source VMs were migrated to other host or datastore using VMware vSphere vMotion, Veeam CDP Coordinator Service analyzies how data will be transferred after the migration and selects VMware CDP proxies to optimize data transfer.

Component reconfiguration requires Veeam CDP Coordinator Service to be working. If the coordinator goes down, existing CDP policies still work, create and remove short-term restore points. Long-term restore points are not created and removed because it is the coordinator who manages them. However, if any of the components goes out of service, for example, VMware CDP proxy goes offline or VMware vSphere vMotion changes the infrastructure, CDP policies start failing until Veeam CDP Coordinator Service is repaired.

# Data Transfer Algorithm

Data transfer starts right after Veeam Backup & Replication configures the backup infrastructure components (performs the initial configuration). Data transfer differs during the initial synchronization and during the incremental synchronization.

As a rule, the initial synchronization is performed when disk data is sent to the target host for the first time. During the initial synchronization, Veeam Backup & Replication sends data for full copies of virtual disks and creates the very first restore points.

During the incremental synchronization, Veeam Backup & Replication mainly sends data for incremental changes made to virtual disks and creates short-term and long-term restore points. For more information on restore points and files created for replicas, see CDP Replication Chain.

## Data Transfer Algorithm During Initial Synchronization

The following steps apply to data transfer during the initial synchronization:

1. On the source host, the I/O filter reads all data from VM disks and sends it to the source VMware CDP proxies.

   As the source VMs are still running, data for the already transferred data blocks can change. The I/O filter intercepts these changes and sends them to the proxies. Sending changes instead of whole changed data blocks helps minimize traffic sent over the network.

2. The Veeam CDP Proxy Service on the source proxies compresses the received data and sends it to the target proxy.

3. The Veeam CDP Proxy Service on the target proxies decompresses the received data. Then sends data to the target host.

4. The I/O filter on the target host saves the received data to virtual disks. The saved data relates to the very first long-term restore point. This restore point is crash-consistent.

   > **NOTE**
   >
   > Veeam Backup & Replication starts creating long-term and short-term restore points only after the initial synchronization finishes. For VMware Cloud Director, the initial synchronization must finish for all VMs in a vApp.

If traffic encryption is configured and the IP addresses of components fall under the rules, these components also encrypt data before sending it. They also decrypt the received data before performing any operations with data if it was encrypted. For more information, see Enabling Traffic Encryption.

After the initial synchronization finishes, Veeam Backup & Replication starts the incremental synchronization.

## Data Transfer Algorithm During Incremental Synchronization

During the incremental synchronization, Veeam Backup & Replication creates short-term and long-term restore points. To create short-term restore points, Veeam Backup & Replication intercepts changes made by transactions on VM disks and sends these changes to the target datastore. Changes are constantly transferred and are saved to transaction logs on the target datastore. To create long-term restore points, Veeam Backup & Replication uses data of short-term restore points and saves the restore points to delta disks. Long-term restore points are created by schedule.

The following steps apply when Veeam Backup & Replication transfers data for short-term restore points:

1. On the source host, the I/O filter intercepts data of all I/O operations and sends this data to the source VMware CDP proxies.

2. Once in the RPO, the Veeam CDP Proxy Service on the source proxies prepares data required for a short-term restore point. For this, Veeam CDP Proxy Service gets the latest state of the data that the source VMware CDP proxies have accumulated.

3. The source Veeam CDP Proxy Service compresses data and sends it to the target proxy.

4. The target Veeam CDP Proxy Service decompresses the received data. Then sends data to the target host.

5. The I/O filter on the target host saves the received data to transaction logs.

The following steps apply when long-term restore point creation is scheduled:

1. [If application-aware image processing is enabled for the CDP policy] Veeam Backup & Replication connects to VM guest OSes, deploys non-persistent runtime components or connects/deploys persistent agent components on VM guest OSes and performs in-guest processing tasks such as quiescing applications on the VM and creating a consistent view of application data. For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.
2. On the target datastore, the I/O filter forms a long-term restore point using data of short-term restore points created since the creation of the previous long-term restore point. Data for the new long-term restore point is saved to a delta disk.

> **NOTE**
>
> If traffic encryption is configured and the IP addresses of components fall under the rules, these components also encrypt data before sending it. They also decrypt the received data before performing any operations with data if it was encrypted. For more information, see Enabling Traffic Encryption.

# CDP Policy Statuses

Based on the workflow, CDP policies can have the following statuses:

- *Initial sync* — the initial synchronization is in process.
- *Syncing* — the incremental synchronization is in process.
- *CBT mode* — the replica data on the target host is not actual. This status can be shown, for example, if a VMware CDP proxy is overloaded and cannot receive or send data. The status can change for *Syncing* after the workload decreases and replica data on the target host is updated to the current VM state on the source host. For more information on how Veeam Backup & Replication behaves in case of data delivery issues, see Guaranteed Delivery.
- *Success, Warning or Error* — the CDP process was successful, had warnings or failed. These statuses are shown for the disabled CDP policies.

# CDP Replication Chain

A replication chain is a sequence of files that allows you to roll back a replica to a specific point in time during failover. Veeam Backup & Replication creates replication chains for each VM added to a CDP policy.

The replication chain contains short-term and long-term restore points. Short-term restore points allow you to roll back replica data to a state created seconds or minutes ago, while long-term restore points — to a state created hours or days ago. Short-term restore points are always crash-consistent, long-term restore points can be crash-consistent or application-consistent depending on how long-term restore points are configured in CDP policy settings.

The replication chain is stored on the target datastore in the <replica_VM_name> folder. The replication chain consists of the files of the following types (only the key file types are listed):

- VMX — the configuration file of the replica.

  The replication chain contains one .vmx file, other files from the list are created per virtual disk.

- VMDK — virtual disk files that store contents of replica hard disk drives.

  On the datastore, you can see files under the following names:

  - *<disk_name>*.vmdk — files that store full copies of virtual disks, that is, store base disk data. These files are created during the initial synchronization and relate to the very first long-term restore point. This restore point is crash-consistent.
  - *<disk_name>-<index>*.vmdk — files that store incremental changes made to virtual disks, that is, store delta disk data. Delta disk files relate to the following:
    - Long-term restore points. In this case, files are created according to the long-term schedule configured in CDP policy settings. These files remain unchanged till they are retained by long-term retention.
    - Technical points. These technical points are required during short-term retention. The technical points creation is connected to the transaction log creation. The technical point is created when the transaction log file reaches its maximum size or when transaction log is created after the short-term retention is reached. For more information on transaction log files, see the TLOG description.
  - *<disk_name>-<index>*.tlog.vmdk — transaction log files that store incremental changes made to virtual disks during RPO set in CDP policy settings. These files relate to short-term restore points that are created once in RPO set in CDP policy settings. One transaction log file stores data for multiple short-term restore points.

    A new transaction log file is created in the following cases:

    - When a transaction log file reaches 512 GB on VMFS, VSAN and vVol datastores. The total size of the transaction log files per disk is maximum 2 TB.
    - When a long-term restore point is created.
    - When the short-term retention is reached. For example, if you set short-term retention to 1 hour, a new transaction log file will be created every hour.
    - After failback commit.
  - *<disk_name>*-interim.vmdk — files for protective virtual disks. Changes made to virtual disks will be written in these files when you perform failover.

> **NOTE**
>
> Although VMDK files look like VMware snapshot files, they are not real snapshots. These files are created by the I/O filter installed on the target host.

- VMFD — files that contains metadata for disks.
- META — files that contain metadata for transaction log files.

To roll back a replica to a specific point in time, the chain of files created for the replica must contain files with data for the base disk (*<disk_name>*-flat.vmdk) and a set of files that contain incremental changes for disks (*<disk-name>*-*<index>*.vmdk + .tlog.vmdk). If any file in the replication chain is missing, you will not be able to roll back to the necessary state. For this reason, you must not delete files from the datastore manually. Instead, you must specify retention policy settings that will let you maintain the desired number of files. For more information on retention policies and how to configure them, see Retention Policies.

# Retention Policies

A retention policy defines for how long Veeam Backup & Replication must store restore points for replicas. Veeam Backup & Replication offers two retention policy schemes:

- Short-term retention
- Long-term retention

## Short-term Retention

Veeam Backup & Replication retains short-term restore points for the number of hours or minutes specified in CDP policy settings. When the retention period is exceeded, Veeam Backup & Replication transforms the replication chain in the following way. The example shows how short-term retention works for a replica with one virtual disk.

1. Veeam Backup & Replication checks whether the replication chain contains outdated short-term restore points.
2. If an outdated restore point exists, Veeam Backup & Replication commits data for the short-term restore point from the transaction log file into the nearest technical point (TP). For more information on technical points, short-term restore points, long-term restore points and when they are created, see CDP Replication Chain.

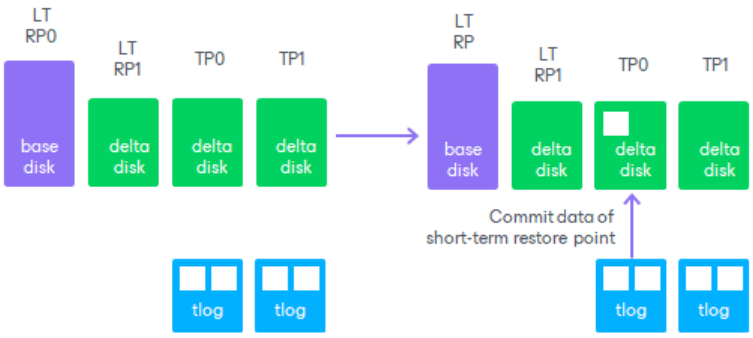


3. If the transaction log file does not contain data for further short-term restore points, Veeam Backup & Replication deletes the transaction log file as redundant — its data has already been committed into the technical delta disk file.

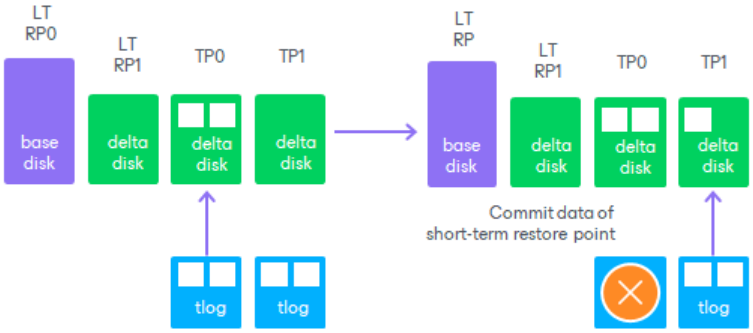


4. After a technical point remains without the related transaction log file, Veeam Backup & Replication considers this technical point outdated and commits data of a newer technical point or a long-term restore point into the outdated technical point. Consider the following:
   - The newer technical restore point must not have the related transaction log.

- The data of the long-term restore point can be committed into a technical point but not vice versa. Long-term restore points remain unchanged till they are retained by the long-term retention.

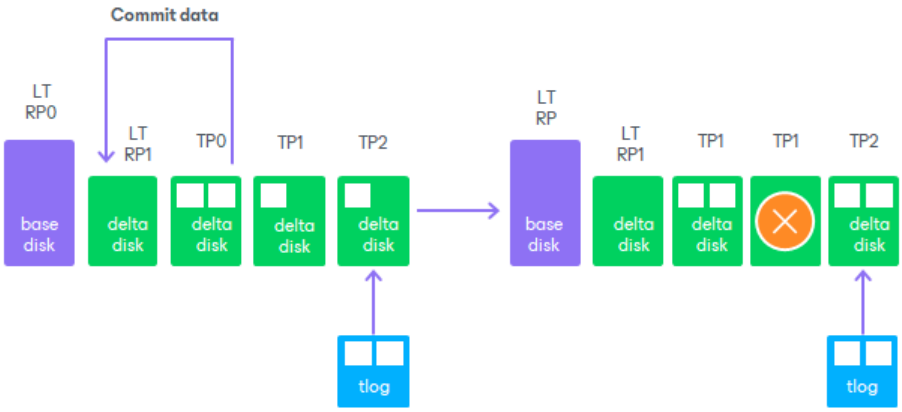


**NOTE**

Due to the chain transformations, Veeam Backup & Replication can actually store short-term restore points for a longer period than specified in the short-term retention policy. This period can be up to 25% longer than the specified short-term retention. Therefore, ensure there is sufficient free space on the target host.

# Long-term Retention

Veeam Backup & Replication retains long-term restore points for the number of days specified in CDP policy settings. When the retention period is exceeded, Veeam Backup & Replication transforms the replication chain in the following way. The example shows how long-term retention works for a replica with one virtual disk.

1. Veeam Backup & Replication checks whether the replication chain contains outdated long-term restore points.
2. If an outdated restore point exists, Veeam Backup & Replication rebuilds the file of an outdated long-term restore point (LT RP) to include data of a newer long-term restore point. To do that, Veeam Backup & Replication commits into the base disk file data from the earliest delta disk file that relates to a long-term restore point. This way, the base disk file 'moves' forward in the replication chain.

   For more information on technical points, short-term restore points, long-term restore points and when they are created, see CDP Replication Chain.

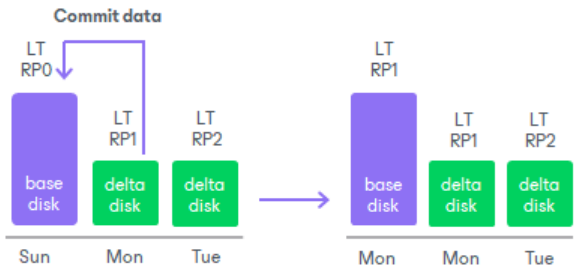

3. Veeam Backup & Replication removes the earliest delta disk file from the chain as redundant — this data has already been committed into the base disk file.

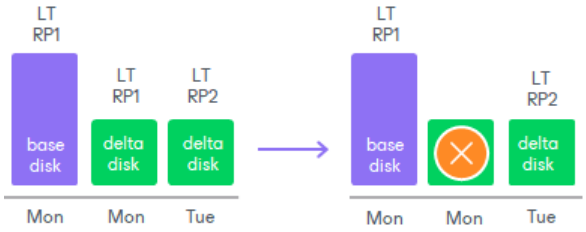

# Guaranteed Delivery

In CDP, data delivery between any two backup infrastructure components is guaranteed by means of the TCP protocol. However, there can be situations when Veeam Backup & Replication is not able to send all data changes generated during CDP in time. For example, if RPO is small, a VM generates a lot of changes, but the infrastructure performance is not enough; or a VMware CDP proxy gets out of work because of a power outage, and all data changes stored on it are lost. This makes data inconsistent and leads to the loss of restore points.

CDP has a special mechanism and tools that return data into the consistent state. The I/O filter on the source host has a mechanism that tracks data blocks that have changed — change tracking (CT). The I/O filter deletes data block addresses from the list of changed blocks only after a confirmation message that data blocks were successfully saved to the replica is received from the target host. In addition, VMware CDP proxies also store data changes until a confirmation message is received from the target host.

How Veeam Backup & Replication uses these tools depend on whether issues occur on the source or target site:

- **Source site**. If the source host does not send the data changes because of the high system load, Veeam Backup & Replication waits till the load decreases. And only then, it gets the list of changed data blocks, reads data blocks from the disk and sends the data to the target host. The target host saves the received data to the datastore. Until this moment, the data on the target host remains inconsistent. As a result, there are "gaps" in the restore point journal and you are not able to restore a VM to some points in time. However, after the target host saves the data, the CDP policy resumes the creation of consistent restore points.

  If the source proxy gets out of work, Veeam Backup & Replication selects another VMware CDP proxy and then behaves as described in the previous paragraph, except for waiting for the load decrease.

- **Target site**. If the target proxy gets out of work, Veeam Backup & Replication selects another VMware CDP proxy and requests data changes from the source VMware CDP proxy. The proxy sends the changes to the target host once again, and the data on the target host becomes consistent almost immediately, so that you are able to restore a VM to any point in time.

  If the connection between the target VMware CDP proxy and the target host is lost, Veeam Backup & Replication checks the host state. If the host is in the Maintenance mode or has disappeared from the cluster, Veeam Backup & Replication starts writing data changes to replicas using another target ESXi host (provided that the host exists and is connected to the datastore where replicas are stored). If the host is not in the Maintenance mode or has not disappeared, Veeam Backup & Replication considers that these are temporary problems with the network and sends data changes again after some time.

# Replica Seeding and Mapping

Replica seeding and mapping are technologies that help reduce the amount of traffic sent over a network. With these technologies, Veeam Backup & Replication does not have to transfer all of VM data from the source host to the target host across the sites during the initial synchronization. For more information on the initial synchronization, see How CDP Works.

You can use seeding and mapping in the following scenarios:

- **Seeding**

  Configure replica seeding if, in a backup repository located in the disaster recovery (DR) site, you have backups of VMs that you plan to replicate. During replication, Veeam Backup & Replication will restore VMs from these backups and will synchronize the state of the restored VMs with the latest state of the source VMs. Then Veeam Backup & Replication will use these restored VMs as replicas.

  For more information on how to create backups that can be used as "seeds" for replica, see Creating Replica Seeds for CDP.

- **Mapping**

  Configure replica mapping if, on the host in the DR site, you have ready-to-use copies of the source VMs. These can be regular or restored VMs, or replicas created by other CDP policies.
  Veeam Backup & Replication will synchronize the state of these ready-to-use VMs with the latest state of the source VMs and will use these VMs as replicas.

You can also configure both replica seeding and replica mapping in the same CDP policy. For example, if a policy includes 2 VMs, you can use seeding for one VM and map the other VM to an existing VM.

> **IMPORTANT**
>
> If seeding or mapping is enabled in a policy, all VMs in the policy must be covered with seeding or mapping. If a VM neither has a seed, nor is mapped to an existing VM, it will be skipped from processing.

# Algorithm for Seeding

Replica seeding includes the following steps:

1. As a preparatory step for replica seeding, you need to create a backup of a VM that you plan to replicate. For more information on how to create a backup that will be used as a "seed" for replica, see Creating Replica Seeds for CDP.
2. When you create a CDP policy, you should point it to a backup repository in the DR site. During the initial synchronization, Veeam Backup & Replication accesses the backup repository where the replica seed is located, and restores the VM from the backup. The restored VM is registered on the target host in the DR site. Files of the restored VM are placed to the location you specify as the replica destination datastore.

   Virtual disks of a replica restored from the backup preserve their format (that is, if the source VM used thin provisioned disks, virtual disks of the replica are restored as thin provisioned).
3. Veeam Backup & Replication synchronizes the restored VM with the latest state of the source VM.

   After successful synchronization, in the **Home** view in the Veeam Backup & Replication console, under **Replicas** node you will see a replica with two restore points. One point will contain the state of the VM from the backup file; the other point will contain the latest state of the source VM you want to replicate.

4. During incremental synchronization, Veeam Backup & Replication transfers only incremental changes in a regular manner.

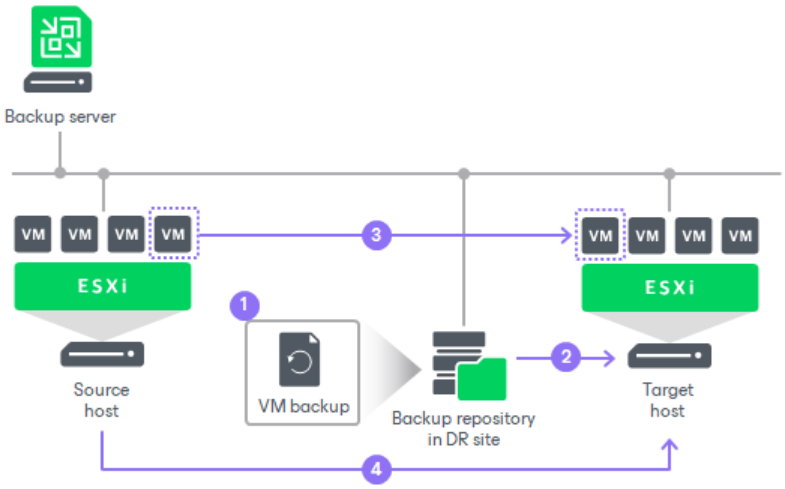


Replica seeding dramatically reduces traffic sent over WAN or slow connections because Veeam Backup & Replication does not send the full contents of the VM image. Instead, it transmits only differential data blocks.

# Algorithm for Mapping

Replication to a mapped VM is performed in the following way:

1. The first step differs depending on which VM you have selected for mapping:
   - If you have selected a regular or restored VM, Veeam Backup & Replication calculates the differences between the source and mapped VM.

     Note that you cannot select a regular or restored VM for CDP with Veeam Cloud Connect.
   - If you have selected a CDP replica, Veeam Backup & Replication imports all restore points of this replica and then calculates the differences between the source and mapped VM. Note that if disk sizes of the source and mapped VM differ, Veeam Backup & Replication will delete all restore points of the mapped VM.
2. To synchronize the state of the mapped VM with the state of the source VM, Veeam Backup & Replication sends the calculated changes to the mapped VM.

   The first and second steps take place during the initial synchronization.
3. During the incremental synchronization, Veeam Backup & Replication transfers only incremental changes in a regular manner.

After the successful initial synchronization, in the **Home** view of Veeam Backup & Replication, under **Replicas** node you will see a replica with restore points. If you have selected for mapping a regular VM or snapshot replica, you will see two restore points: one restore point will contain the latest state of the mapped VM, the other will contain the state of the source VM. If you have selected a CDP replica, you will see all restore points of the mapped VM plus one restore point that will contain the state of the source VM.

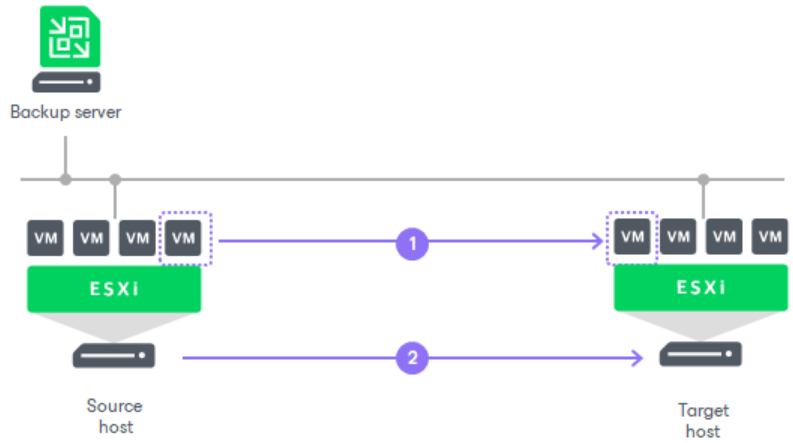

# Installing I/O Filter

To be able to protect VMs with CDP, you must install the I/O filter on each cluster where the VMs that you plan to protect reside and where replicas will reside. For more information on the filter, see Source and Target Hosts.

To install the filter on a specific cluster, open the **Inventory** view. In the inventory pane, navigate to the **Virtual Infrastructure** > **VMware vSphere** > **vCenter Servers** > *<vCenter Server Name>* > *<Cluster Name>* node and right-click it. In the menu, select **Install I/O filter**.

> **IMPORTANT**
>
> If vSphere LifeCycle Management is enabled for your cluster, you need first follow the instructions from this section and then follow the instructions from this Veeam KB article.

To install the I/O filter on multiple clusters in a vCenter Server, use the **VeeamCDP Filter Management** wizard.

# Step 1. Launch I/O Filter Management Wizard

To launch the **VeeamCDP Filter Management** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane, navigate to the **Managed Servers** > **VMware vSphere** > **vCenter Servers** > *<vCenter Server Name>* node. Right-click the node and select **Manage I/O filters**. Alternatively, click **Manage I/O Filters** on the ribbon.
- Open the **Inventory** view. In the inventory pane, navigate to the **Virtual Infrastructure** > **VMware vSphere** > **vCenter Servers** > *<vCenter Server Name>* node. Right-click the node and select **Manage I/O filters**. Alternatively, click **Manage I/O Filters** on the ribbon.

Alternatively, you can install the I/O filter on an individual cluster. To do this, use one of the paths described in the list, open the *<vCenter Server Name>* node and right-click the *<Cluster Name>* node. Select **Install I/O filter**.

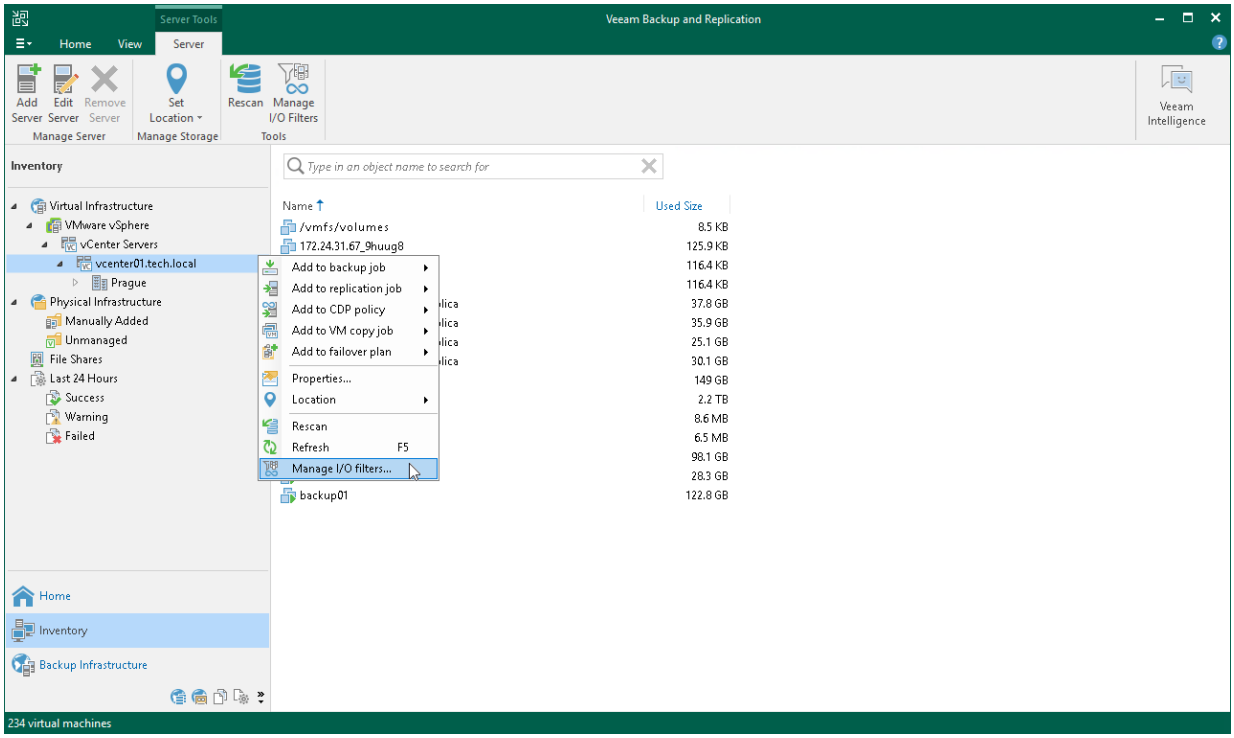

# Step 2. Select Clusters

At the **Clusters** step of the wizard, select check boxes near clusters on which you want to install the I/O filter.

If you select check boxes near clusters where the filters are already installed, Veeam Backup & Replication will update the filters. If you clear check boxes, Veeam Backup & Replication will delete the I/O filter from these clusters.

> **NOTE**
>
> Consider the following:
>
> - If another user has already installed the I/O filter on a cluster, you will be prompted whether to take ownership. For more information, see Taking I/O Filter Ownership.
> - If you add a new cluster to the vCenter server after the I/O filter is installed on the existing clusters, you need to install the I/O filter manually on the newly added cluster. To do that, open the **I/O Filter Management** wizard, make sure that check boxes are selected near all clusters where the I/O filter must be present and finish the wizard.

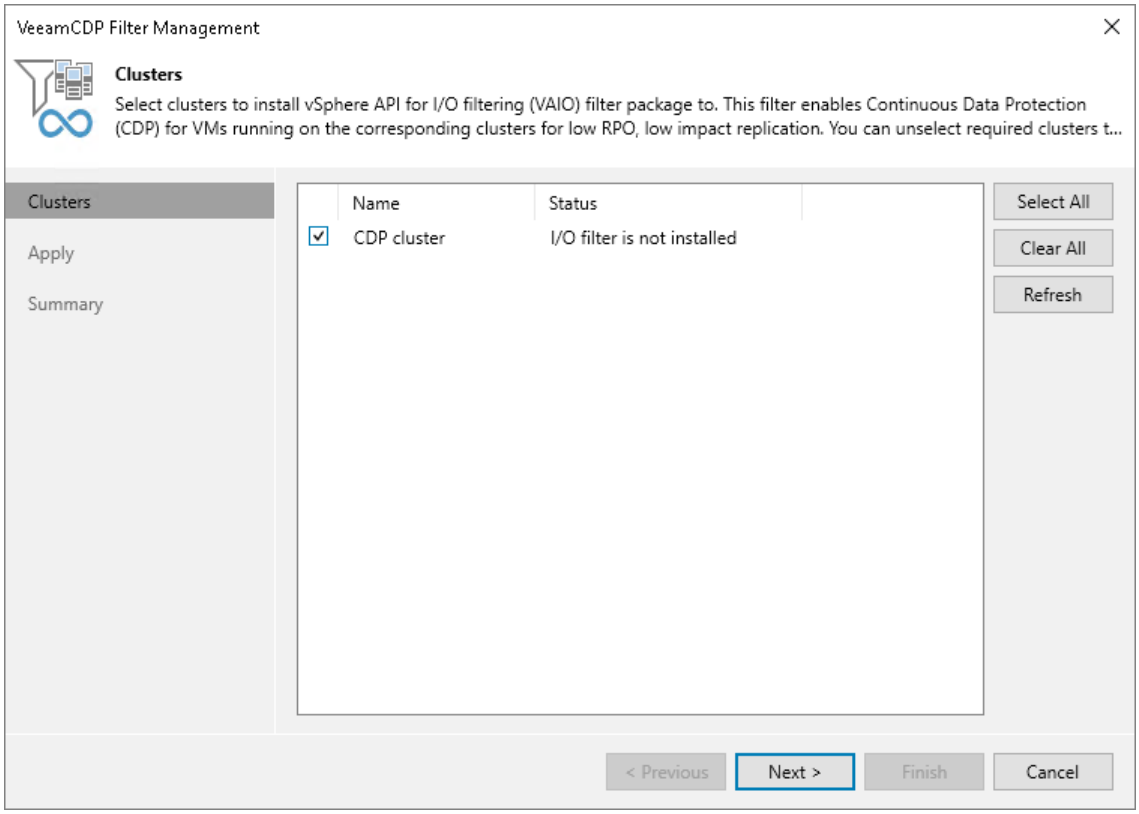

# Step 3. Apply Filter Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs I/O filter. Click **Next**.

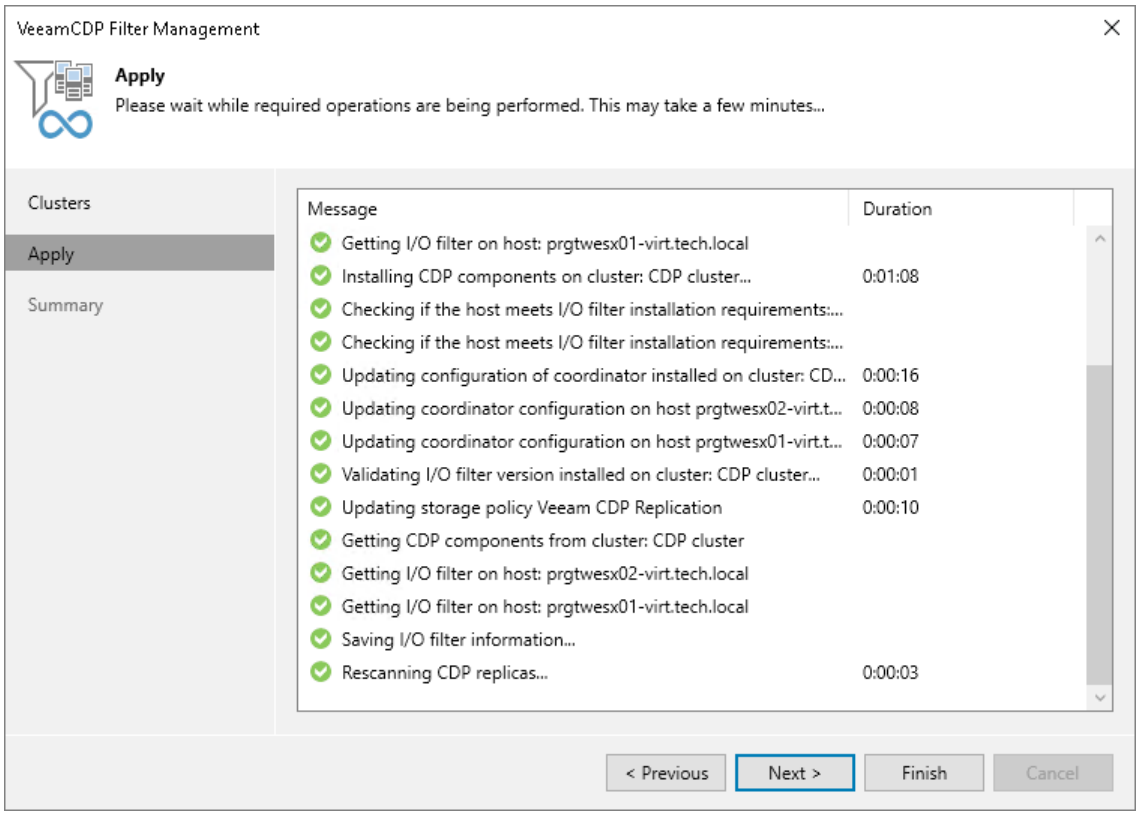

# Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review on which clusters the I/O filter is installed and click **Finish** to exit the wizard.

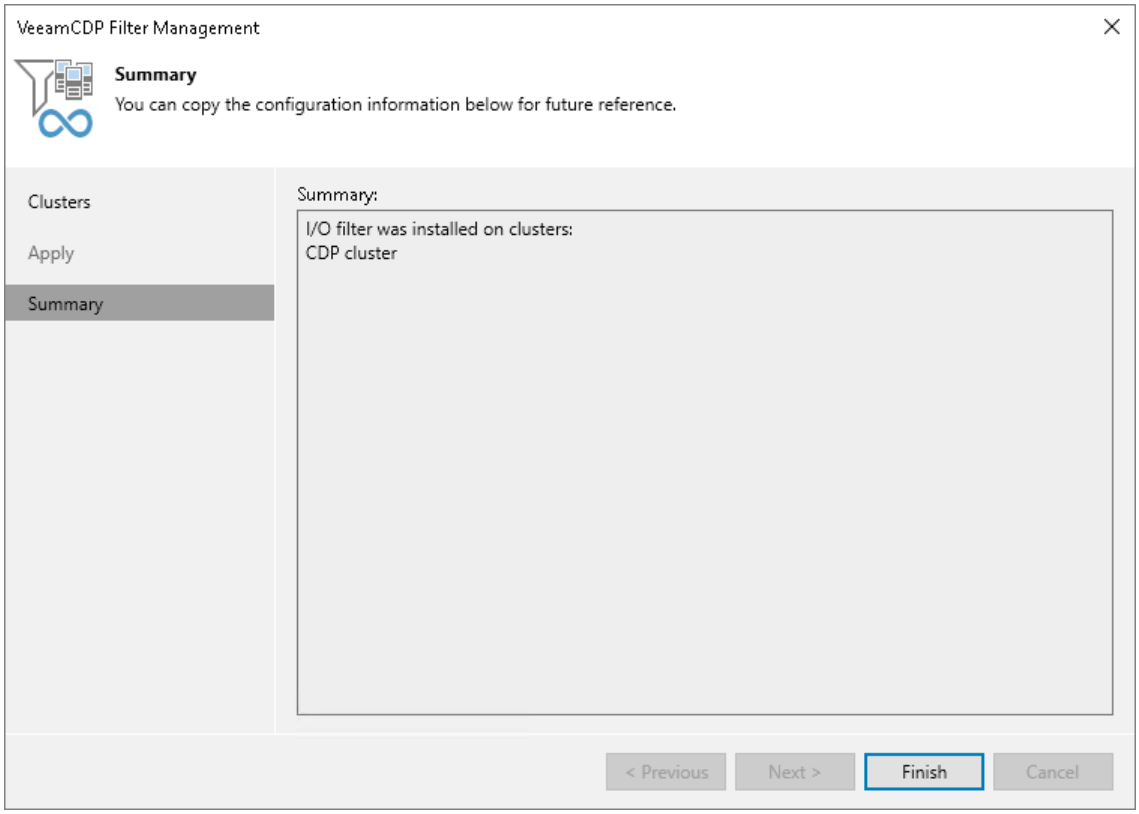

# Updating and Uninstalling I/O Filter

Veeam Backup & Replication allows you to update or uninstall the I/O filter from organization VDCs using the Veeam Backup & Replication console.

## Requirements

Consider the following:

- Make sure that you have disabled or deleted all CDP policies as described in section Disabling and Deleting Policies.
- [When uninstalling the filter] If you have manually assigned the Veeam CDP Replication storage policy to VMs that are parts of the clusters, you must change the storage policy for these VMs. For more information on how to change storage policies, see VMware Docs.

  If replicas are still present in the Veeam Backup & Replication configuration database, you can remove them from the configuration. To see the list of replicas, open the **Home** view. In the inventory pane, click the **Replicas** node.

- When you upgrade Veeam Backup & Replication to the current version, you can postpone the upgrade of the I/O filter on the vCenter Servers to a later time. Veeam Backup & Replication supports the postponed upgrade for the minor versions within the 12 version. You must postpone the upgrade or upgrade all clusters on the vCenter Servers. Partially upgraded vCenter Servers or clusters have limited functionality. You cannot add VMs from such vCenter Servers or clusters to CDP policies, commit failback and perform some other operations.
- If VMware vSphere Distributed Resource Schedule (DRS) is disabled for the clusters, place the hosts of the clusters in the Maintenance mode.

  You can simultaneously place all hosts in the Maintenance mode or place them one by one for each cluster. Veeam Backup & Replication will uninstall/update the filter only on the hosts in the Maintenance mode. That is why you must repeat the procedure (place a host in the Maintenance mode and launch installation/update) for each host in a cluster. When Veeam Backup & Replication uninstalls/updates the filter on the last host in the cluster, it also uninstalls/updates the filter on the cluster.

  If DRS is enabled and the hosts use the shared storage, you do not have to put hosts in the Maintenance mode manually. VMware vSphere will automatically place hosts in the Maintenance mode according to your settings and migrate VMs from one host to another.

## Updating or Uninstalling I/O Filter

To update or uninstall the I/O filter, do the following:

1. Launch the **VeeamCDP Filter Management** wizard as described in section Launch I/O Filter Management Wizard.
2. At the **Clusters** step of the wizard, do the following:
   - To update the filter, make sure that check boxes are selected near the necessary clusters.
   - To uninstall the filter, clear the check boxes near the necessary clusters.
3. Proceed to the last step of the wizard and close the wizard.

As an alternative, you can uninstall the I/O filter from a specific cluster. To do this, open the **Inventory** view. In the inventory pane, navigate to the **Virtual Infrastructure** > **VMware vSphere** > **vCenter Servers** > *<vCenter Server Name>* > *<Cluster Name>* node and right-click it. Select **Uninstall I/O filter**.

**NOTE**

If the upgrade attempt has failed, fix the problems listed in the wizard and launch the upgrade again. During the next upgrade operation, you may see the following error: *"The specified key, name or identifier 'vibUrl' already exists. The VIB contains the same filter as the one to be upgraded"*. If this is the only error, it can be ignored. The error is shown because during the previous upgrade filters were already upgraded on some components. The remaining components are upgraded on the current run. To make sure that upgrade completed successfully, launch the upgrade wizard once again.

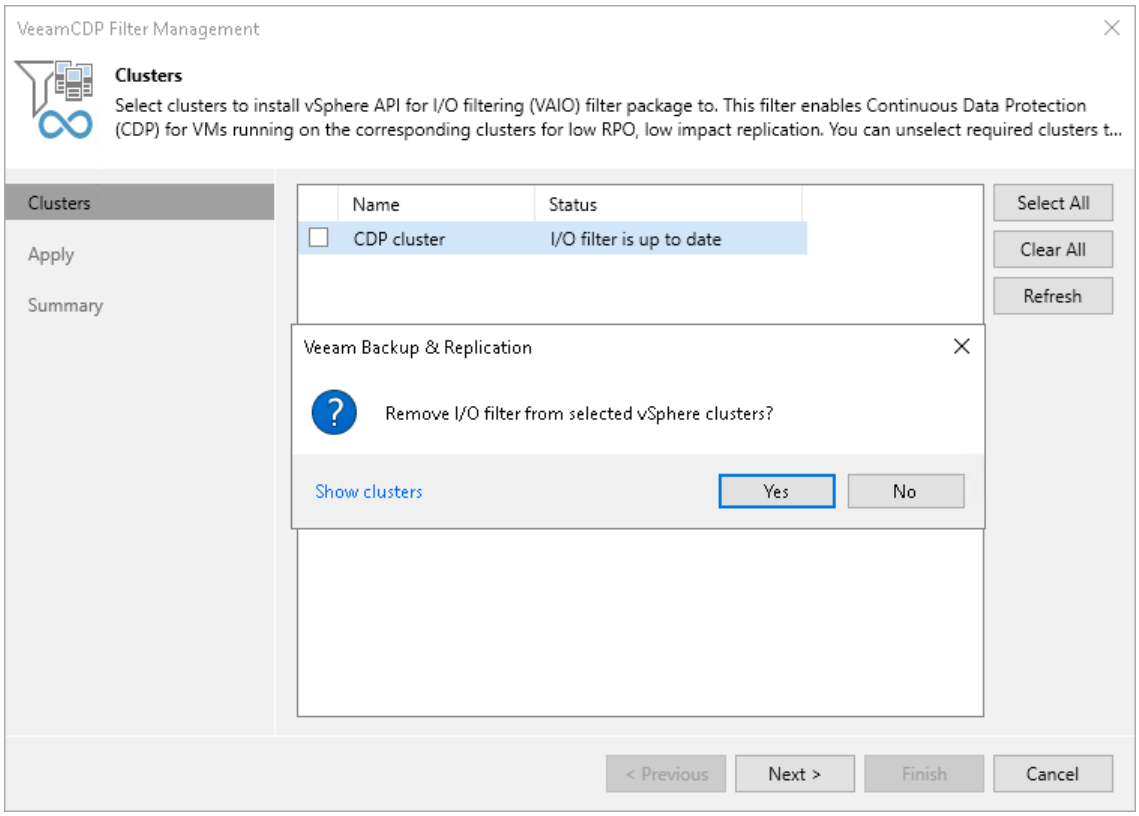

# Taking I/O Filter Ownership

After the backup server finishes installation of the I/O filter on a cluster, the backup server becomes the owner of the I/O filter and its components.

The I/O filter can belong only to one backup server at a time. If a client *B* re-installs the I/O filter that was previously installed by a client *A*, the client *A* loses the ownership privileges over the filter. Losing privileges means that the client *A* will no longer be able to manage CDP policies and replicas. All policies created by the client *A* will fail.

If the I/O filter is managed by another client, you will see a warning when trying to install the filter. If you are sure that the client who installed the I/O filter first does not need it anymore, you can take the ownership. To do that, click **Yes** in the Veeam Backup & Replication window and finish the **I/O Filter Management** wizard as described in section Installing I/O Filter.

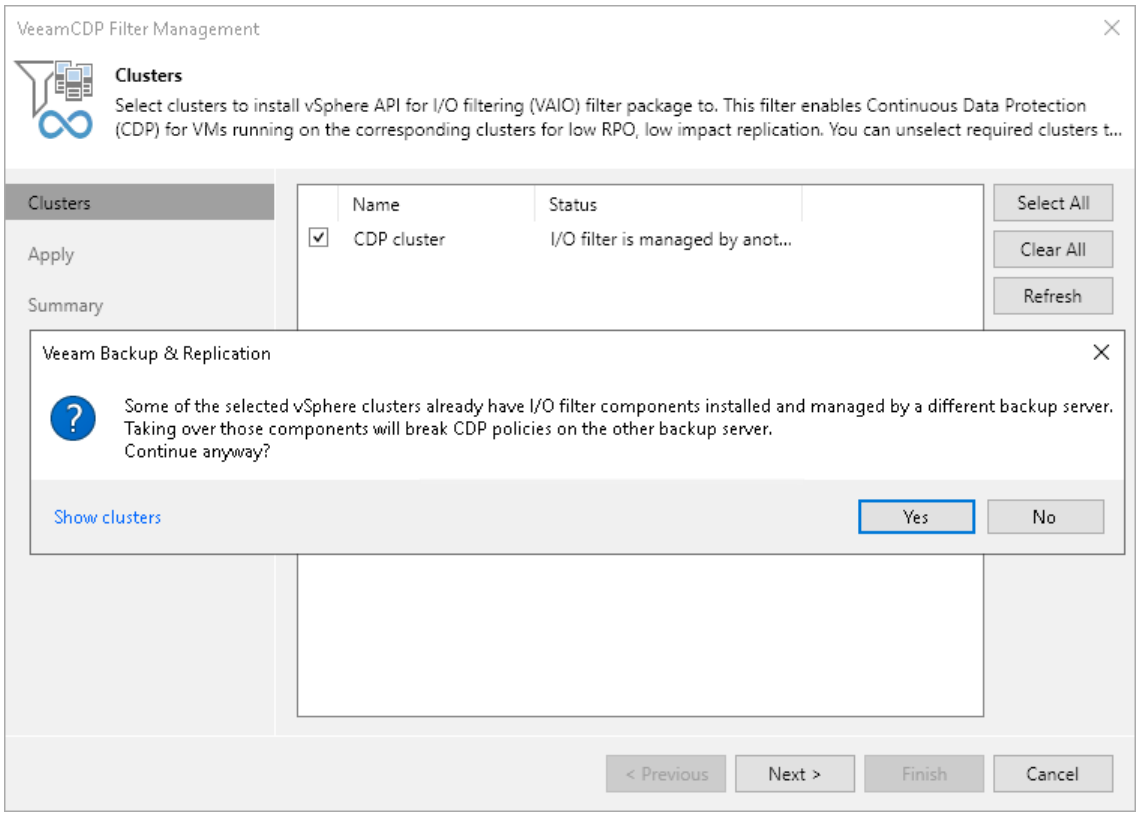

# Creating CDP Policies

To protect VMs with CDP, you must configure a CDP policy. The CDP policy defines which VMs to protect, where to store replicas, how often create short-term and long-term restore points, and so on. One CDP policy can process one or multiple VMs.

To create a CDP policy use the **New CDP Policy** wizard.

# Before You Begin

Before you create a CDP policy, check the following prerequisites:

- Check that all the required components are added to the backup infrastructure. For more information on the required components, see Backup Infrastructure for CDP.
- The I/O filter must be installed on each cluster where VMs that you plan to protect reside. For more information on how to install the filter, see Installing I/O Filter.
- If you plan to use replica seeding, you must create a seed as described in section Creating Replica Seeds for CDP.
- If you plan to protect VMware Cloud Director objects, consider using the dedicated policy. For more information, see CDP for VMware Cloud Director.

# Step 1. Launch New CDP Policy Wizard

To launch the **New CDP Policy** wizard, do one of the following:

- Open the **Home** view. On the ribbon, click **CDP Policy** > **VMware vSphere**.
- Open the **Home** view. In the inventory pane, right-click **Jobs** and select **CDP Policy** > **VMware vSphere**.
- Open the **Inventory** view. In the working area, right-click VMs that you want to replicate. Select **Add to CDP policy** > **New job** if you want to create a new CDP policy, or **Add to CDP policy** > *<Policy Name>* if you want to add VMs to an existing CDP policy.

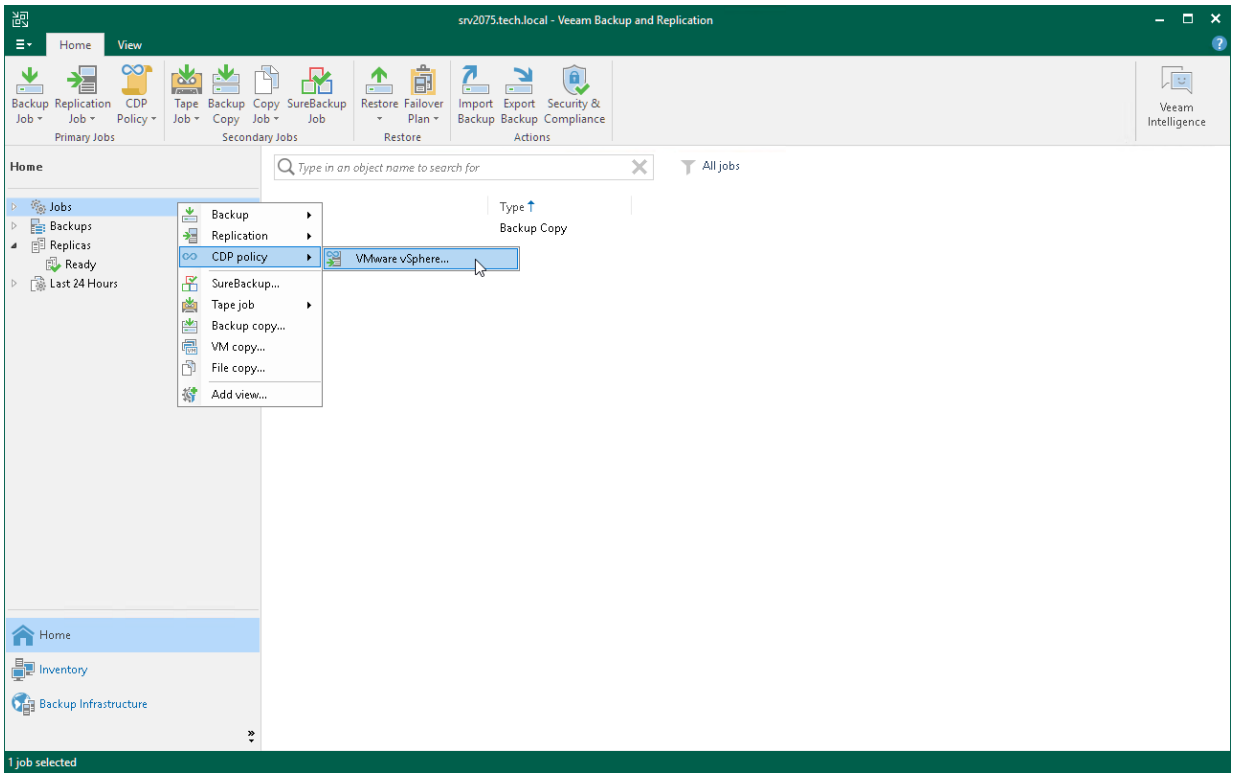

# Step 2. Specify Policy Name and Advanced Settings

At the **Name** step of the wizard, specify a name and description for the CDP policy, and choose whether you want to use replica seeding or network mapping:

1. In the **Name** field, enter a name for the CDP policy.
2. In the **Description** field, provide a description for future reference.
3. If a network between your production and disaster recovery (DR) sites has low bandwidth, and you want to reduce the amount of traffic sent during the initial synchronization of the CDP policy, select the **Replica seeding (for low bandwidth DR sites)** check box.

   When selected, this check box enables the **Seeding** step where you will have to configure replica seeding and mapping.
4. If your DR site networks do not match your production site networks, select the **Network remapping (for DR sites with different virtual networks)** check box.

   When selected, this check box enables the **Network** step where you will have to configure a network mapping table.
5. If the IP addressing scheme in your production site differs from the scheme in the DR site, select the **Replica re-IP (for DR sites with different IP addressing scheme)** check box.

   When selected, this check box enables the **Re-IP** step where you will have to configure replica re-IP rules.

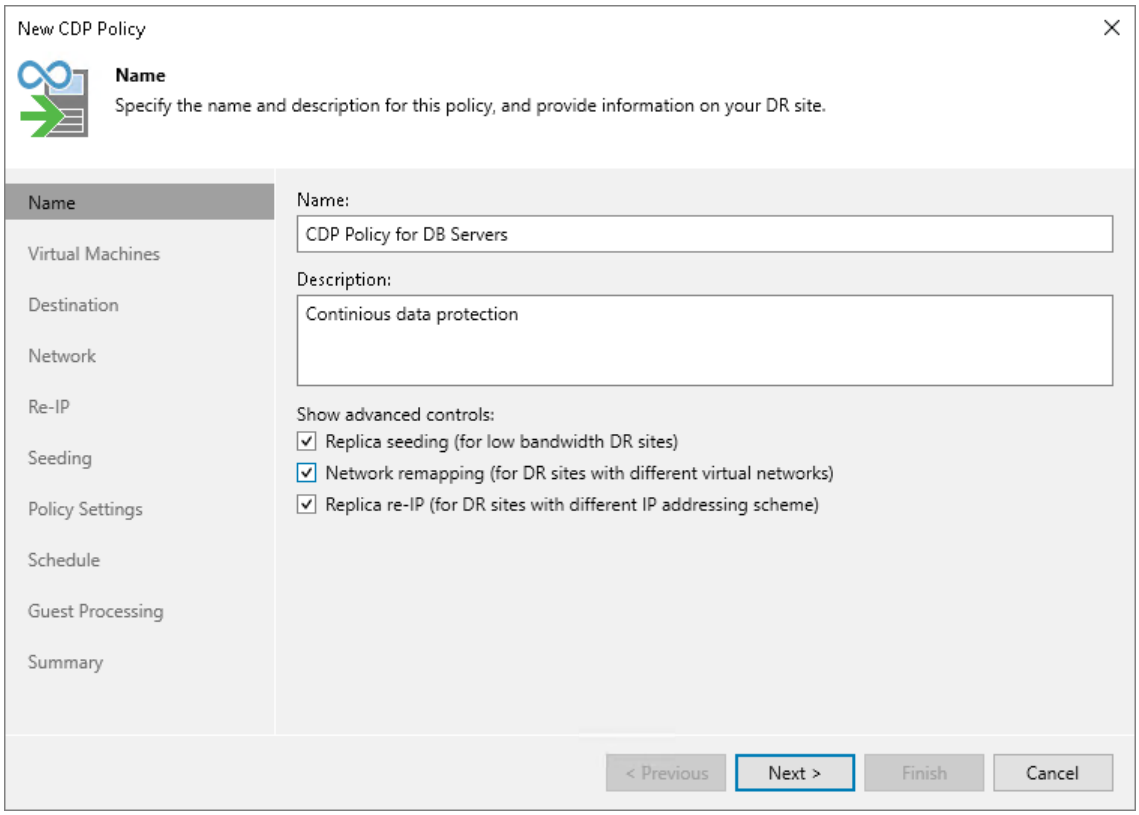

# Step 3. Select VMs to Replicate

At the **Virtual Machines** step of the wizard, select VMs or VM containers that you want to replicate:

1. Click **Add**.
2. In the **Add Object** window, select the necessary VMs or VM containers and click **Add**.

   If you select VM containers (hosts, clusters, folders, resource pools, VirtualApps or datastores) and add new VMs to this container in future, Veeam Backup & Replication will update CDP policy settings automatically to include these VMs.

   > **NOTE**
   >
   > Even if you add VM containers to a CDP policy, only VMs are replicated. VM templates, VM logs, folders and so on are not replicated.

   You can use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only those VMs that have all the selected tags will be processed by the policy.

> **IMPORTANT**
>
> Consider the following:
>
> - You can replicate only VMs that are turned on, the turned off VMs will be skipped from processing.
> - You cannot add to a CDP policy VMs that were already added to other CDP policies created on the same backup server.

To quickly find the necessary VMs, you can use the search field at the bottom of the **Add Object** window. If you want to switch between types of VMs you want to search through, use the button to the left of the search field.

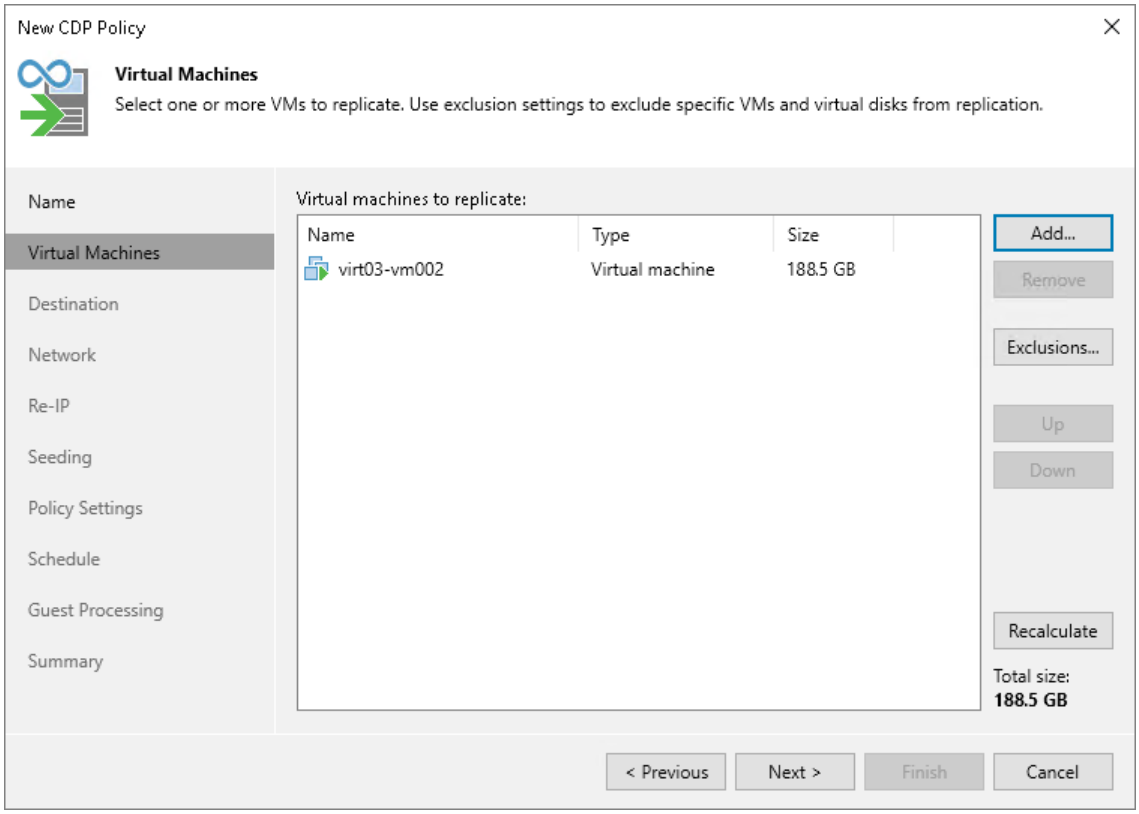

# Step 4. Exclude Objects

After you have added VMs or VM containers to the CDP policy, you can specify which objects you want to exclude from being replicated.

## Excluding VMs and VM Containers

To exclude VMs and VM containers:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, check that the **VMs** tab is selected and click **Add**.
3. In the **Add Objects** window, select VMs or VM containers that you want to exclude from being replicated and click **Add**.

   You can use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only those VMs that have all the selected tags will be excluded from the policy.
4. In the **Exclusions** window, click **OK**.

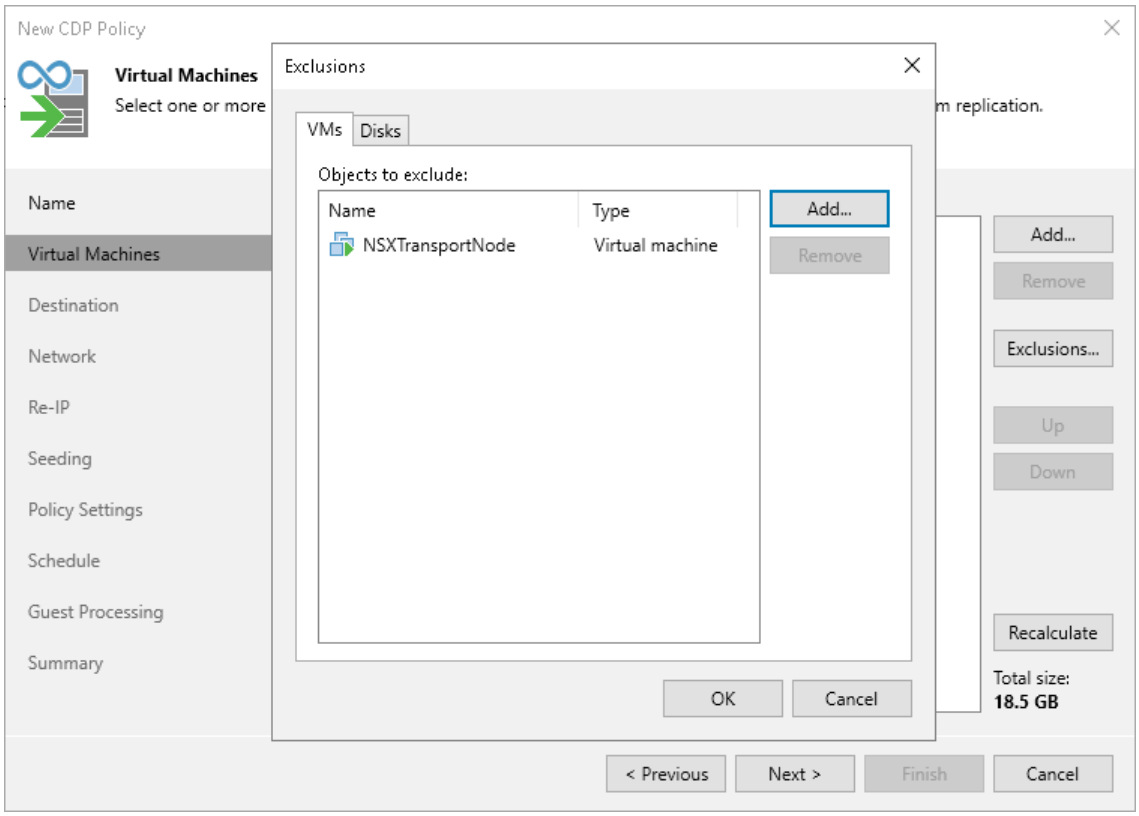


## Excluding Disks

To exclude disks:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.

2. In the **Exclusions** window, do the following:
   a. Switch to the **Disks** tab.
   b. If you want to exclude disks of VMs that are added as a part of containers, click **Add**. In the **Add Objects** window, select the necessary VMs and click **Add**. Veeam Backup & Replication will include these VMs in the list as standalone objects.
   c. In the **Disks to process** list, select the necessary VMs or VM containers.
   d. Click **Edit**.
3. In the **Select Disks** window, select disks that you want to replicate: all disks, 0:0 disks (as a rule, system disks) or specific IDE, SCSI, SATA or NVMe disks. Disks that you do not select will be excluded from processing. Click **OK**.
4. In the **Exclusions** window, click **OK**.

> **NOTE**
>
> If you exclude disks from a backup and enable application-aware processing, Veeam Backup & Replication will still perform application-aware processing for the excluded disks. This means that VSS will process disk data.

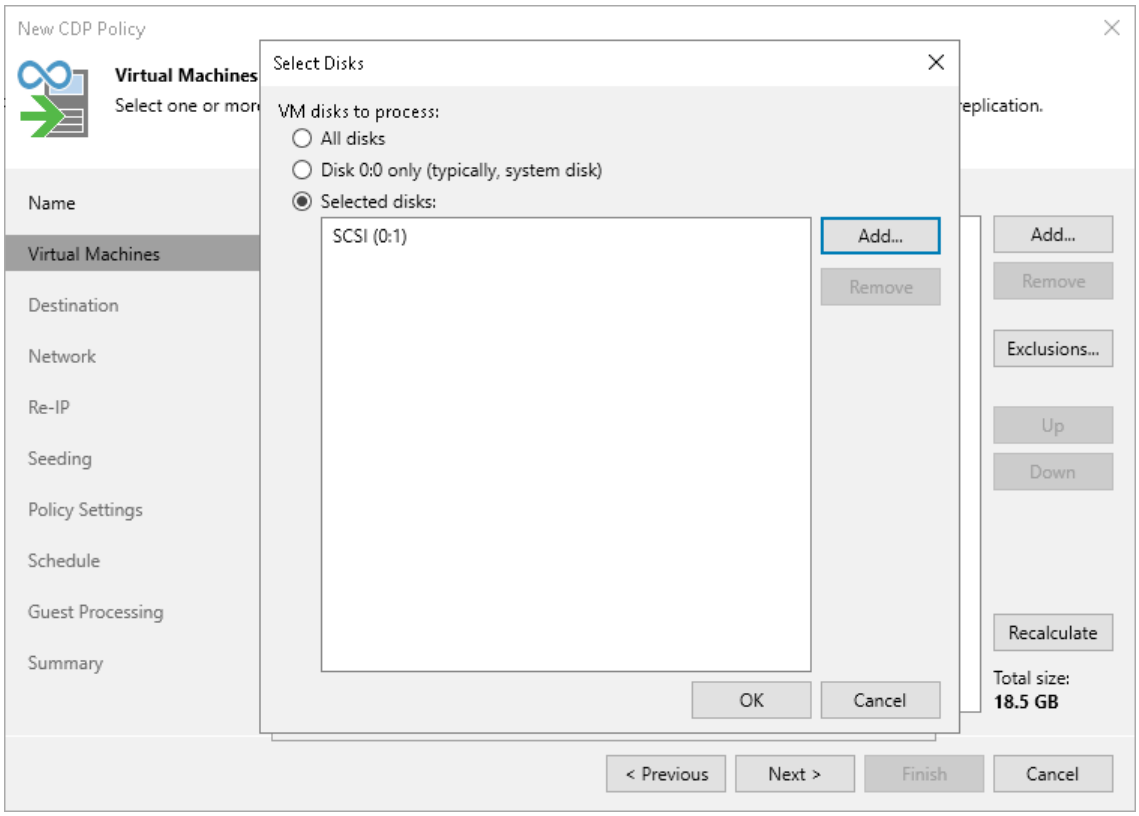

# Step 5. Specify VM Processing Order

At the **Virtual Machines** step of the wizard, click **Up** and **Down** to change the processing order. VMs at the top of the list have a higher priority and will be processed first.

> **NOTE**
>
> Consider the following:
>
> - VMs inside a VM container are processed at random. To ensure that VMs are processed in the defined order, you must add them as standalone VMs, not as a part of containers.
> - The processing order may differ from the order that you have defined. For example, if resources of a VM that is higher in the priority are not available, and resources of a VM that is lower in the priority are available, Veeam Backup & Replication will process the VM with the lower priority first.

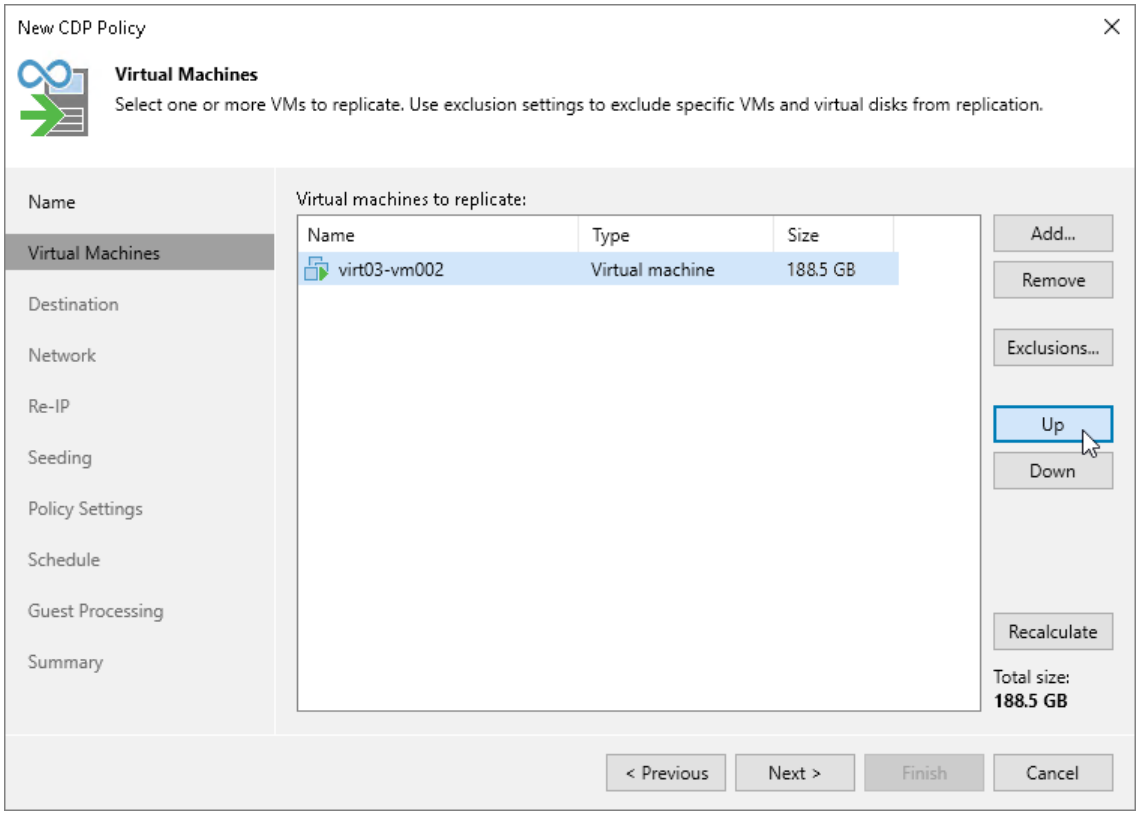

# Step 6. Select Replica Destination

At the **Destination** step of the wizard, select a target host or cluster, resource pool, folder and datastore for replicas, and types of replica disks:

1. Next to the **Host or cluster** field, click **Choose** > **VMware vSphere** and select a host or cluster where replicas must be registered. If you select a cluster or vCenter Server, the replication process will become more sustainable — the replication process will not fail if there is at least one available host in the cluster.

   If you select a cluster as a destination, Veeam Backup & Replication will request VMware to send the list of available hosts, and will select the first host in this list as the destination for the replicas. These replicas will be stored on the datastore with the most free disk space.

2. Next to the **Resource pool** field, click **Choose** and select a resource pool to which replicas will be added.

   If you have selected to replicate multiple VMs and want to add individual replicas to other resource pools:

   a. Click the **Pick resource pool for selected replicas** link.

   b. In the **Choose Resource Pool** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose Resource Pool** window, select the necessary VMs in the **Replica VM resource pool** list. At the bottom of the window, click **Resource Pool**.

   e. In the **Select Resource Pool** window, select the necessary resource pool and click **OK**.

3. Next to the **VM folder** field, click **Choose** and select a folder where all VM files will be stored. Note that the **VM folder** section is disabled if you have selected a standalone ESXi host as the target for replicas.

   If you have selected to replicate multiple VMs and want to place individual replicas to other folders:

   a. Click the **Pick VM folder for selected replicas** link.

   b. In the **Choose Folder** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose Folder** window, select the necessary VMs in the **Replica VM folder** list. At the bottom of the window, click **VM Folder**.

   e. In the **Select Folder** window, select the necessary folder.

4. Next to the **Datastore** field, click **Choose** and select a datastore where replica files will be stored. Note that if you have chosen to replicate VMs to a cluster, Veeam Backup & Replication displays only shared datastores.

   If you have selected to replicate multiple VMs and want to place individual replicas to other datastores:

   a. Click the **Pick datastore for selected virtual disks** link.

   b. In the **Choose VM Files Location** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose VM Files Location** window, select the necessary VMs in the **Files location** list. At the bottom of the window, click **Datastore**.

   e. In the **Select Datastore** window, select the necessary datastore.

5. If you want to store replica configuration files and disk files in different datastores:
   a. Click the **Pick datastore for selected virtual disks** link.
   b. In the **Choose VM Files Location** window, click **Add VM**.
   c. In the **Add Objects** window, select the necessary VMs and click **Add**.
   d. In the **Choose VM Files Location** window, expand the necessary VMs in the **Files location** list, and select the necessary files. At the bottom of the window, click **Datastore**.
   e. In the **Select Datastore** window, select the destination for the selected type of files.
6. You can change types of replica disks. By default, Veeam Backup & Replication saves disks in the thin type.

   To change replica disk types:
   a. Click the **Pick datastore for selected virtual disks** link.
   b. In the **Choose VM Files Location** window, click **Add VM**.
   c. In the **Add Objects** window, select VMs whose disk type you want to change and click **Add**.
   d. In the **Choose VM Files Location** window, select the necessary VMs in the **Files location** list. At the bottom of the window, click **Disk type**.
   e. In the **Disk Type Settings** window, select a type that will be used to restore replica disk files: same as source, thin, thick lazy zeroed or thick eager zeroed.

      For more information about disk types, see VMware Docs.

> **NOTE**
>
> Disk type change is available only for VMs that use virtual hardware version 7 or later.

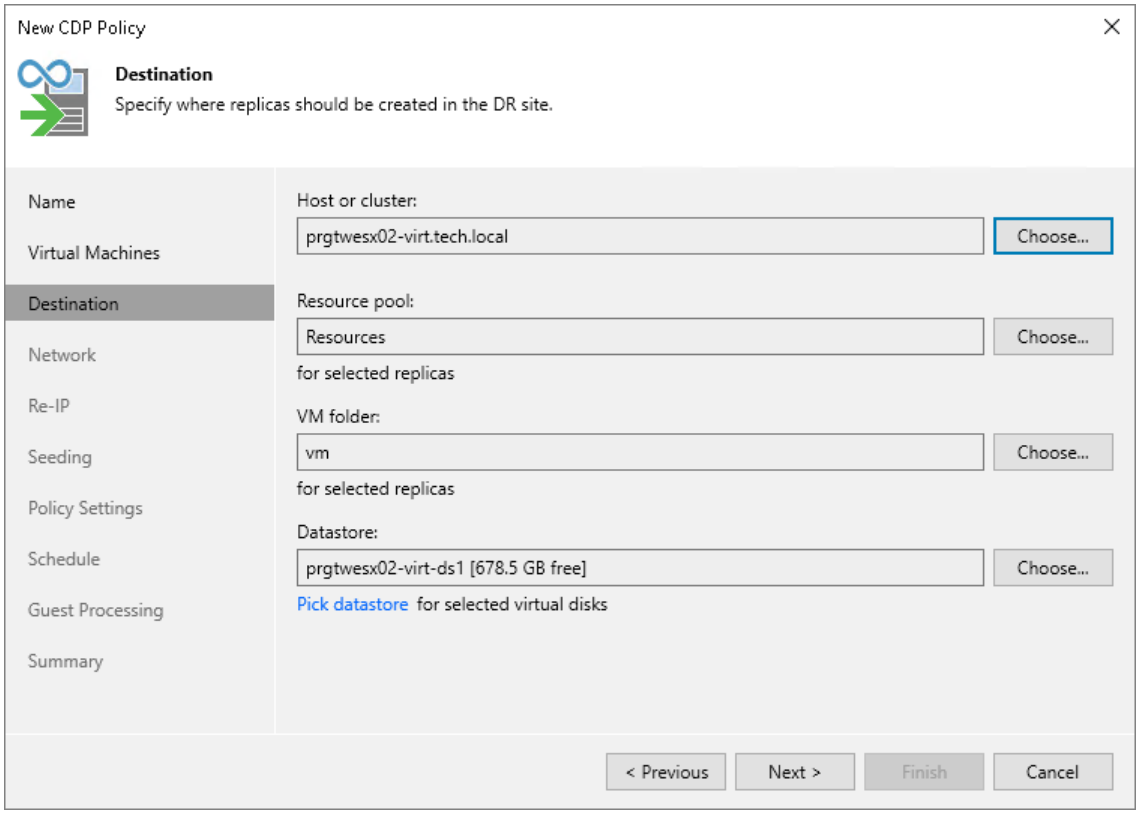

# Step 7. Configure Network Mapping

The **Network** step of the wizard is available if you have selected the **Network remapping** option at the **Name** step of the wizard.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the production site to networks in the disaster recovery (DR) site. When the replication session starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update replica configuration to replace the production networks with the specified networks in the DR site. As a result, you will not have to re-configure network settings manually.

To add a row to the network mapping table:

1. Click **Add**.
2. In the **Network Mapping** window, click **Browse** next to the **Source network** field.
3. In the **Select Network** window, select the production network to which the source workloads are connected and click **OK**.
4. In the **Network Mapping** window, click **Browse** next to the **Target network** field.
5. In the **Select Network** window, select a network in the DR site to which replicas will be connected and click **OK**.
6. In the **Network Mapping** window, click **OK**.

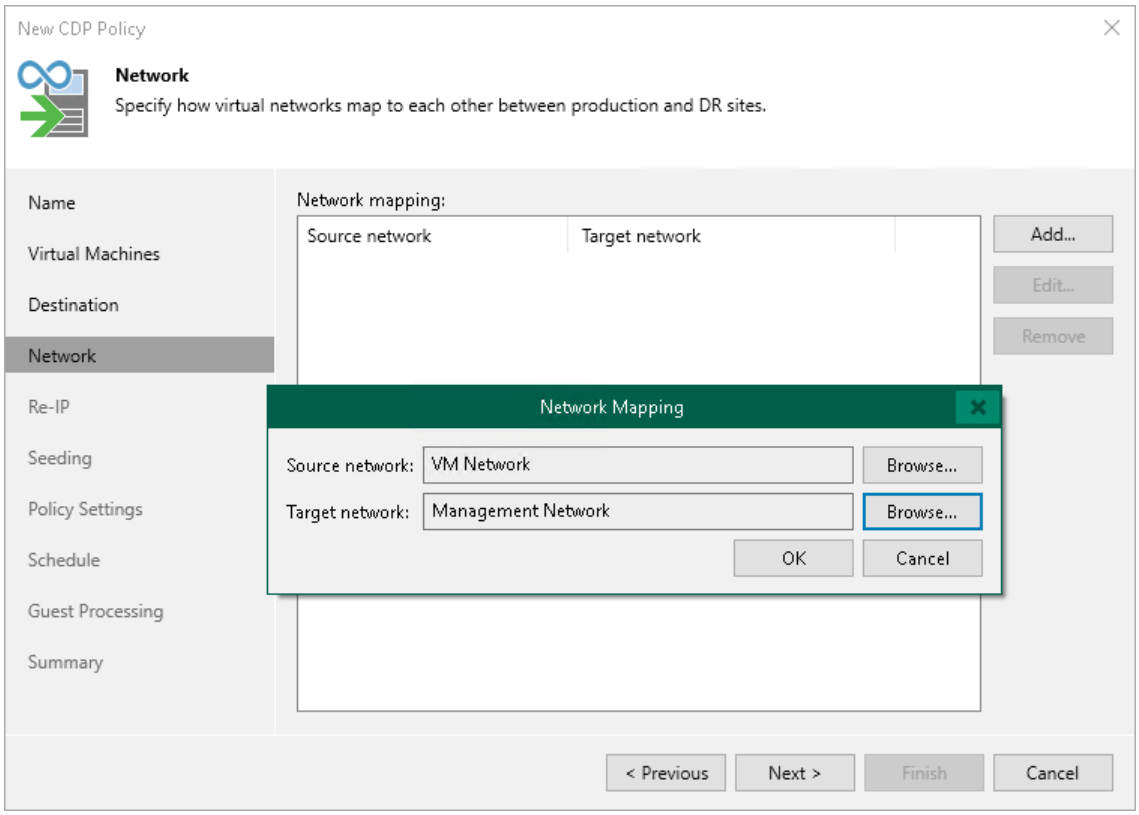

# Step 8. Configure Re-IP Rules

The **Re-IP** step is available if you have selected the **Replica re-IP** check box at the Name step of the wizard. This step applies only to VMs with Microsoft Windows OSes.

At the **Re-IP** step of the wizard, configure re-IP rules. These rules map IPs in the production site to IPs in the disaster recovery (DR) site. When you perform failover, Veeam Backup & Replication will check the configured re-IP rules and will change replica IPs if the rules apply. Replicas will get new IP addresses according to the network masks specified in the rules, so that you will be able to reach replicas in the DR site.

To configure a re-IP rule:

1. Click **Add** and select whether you want to configure an IPv4 or IPv6 rule. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
2. In the **Source VM** section, specify an IP numbering scheme used in the production site.

   To facilitate the configuration, Veeam Backup & Replication detects an IP address and subnet mask for the backup server and pre-populates values in the **Source VM** section.
3. In the **Target VM** section, specify an IP address, subnet mask and default gateway that will be used for replicas in the DR site. If required, specify the DNS server addresses. For the IPv4 rules, you can also specify WINS server addresses.
4. In the **Description** field, provide a description.
5. Click **OK**.

NOTE

Consider the following:

- You can specify static IPs or IP ranges. Do not use 0 to specify IP address ranges. In Veeam Backup & Replication, value 172.16.17.0 means a regular IP address 172.16.17.0, not an IP address range. To specify a range, use the asterisk character (*).
- Replica re-IP works only if you perform replica failover using Veeam Backup & Replication. If you power on a replica in some other way, for example, manually using vSphere Client, re-IP rules will not be applied to it.
- The backup server OS must support mounting of the system disks of VMs that will be replicated.

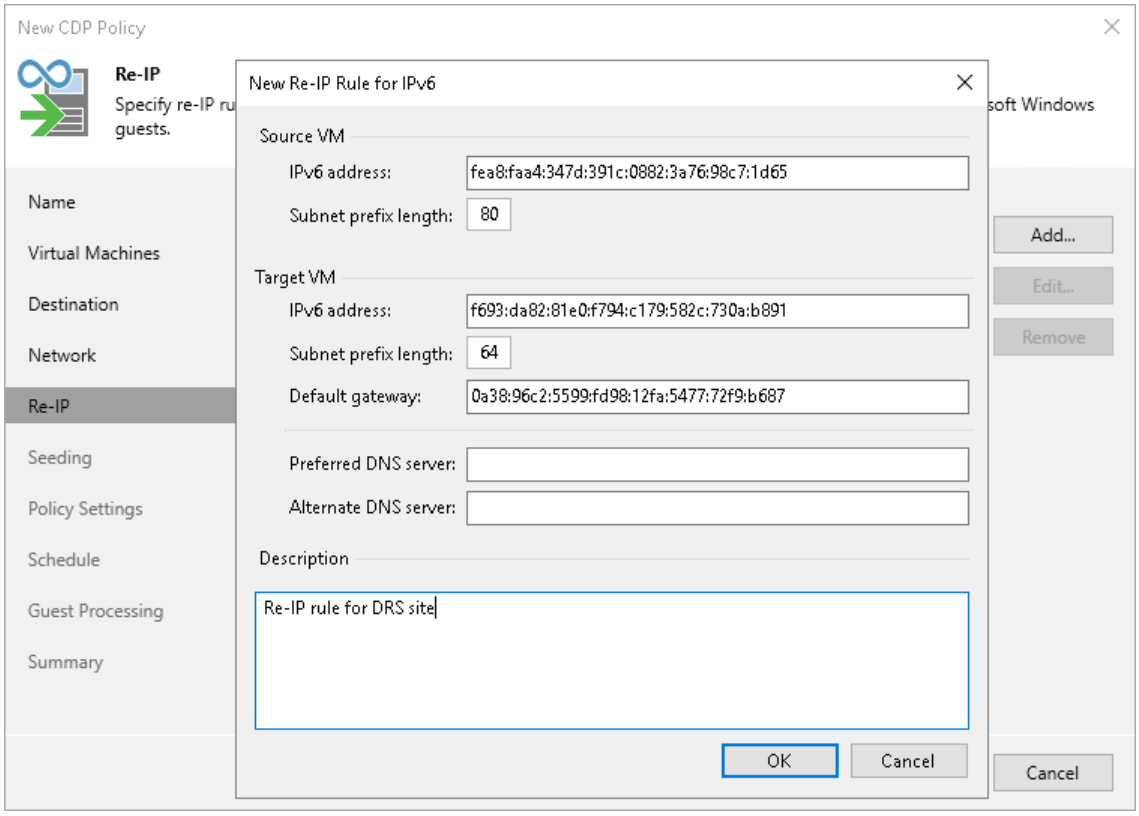
New CDP Policy
Re-IP
Specify re-IP ru
guests.
soft Windows
Name
Virtual Machines
Destination
Network
Re-IP
Seeding
Policy Settings
Schedule
Guest Processing
Summary
Add...
Edit...
Remove
Cancel
New Re-IP Rule for IPv6
Source VM
IPv6 address:
fea8:faa4:347d:391c:0882:3a76:98c7:1d65
Subnet prefix length:
80
Target VM
IPv6 address:
f693:da82:81e0:f794:c179:582c:730a:b891
Subnet prefix length:
64
Default gateway:
0a38:96c2:5599:fd98:12fa:5477:72f9:b687
Preferred DNS server:
Alternate DNS server:
Description
Re-IP rule for DRS site
OK
Cancel

# Step 9. Configure Seeding and Mapping

The **Seeding** step is available if you have selected the **Replica seeding** check box at the Name step of the wizard.

At the **Seeding** step of the wizard, configure replica seeding and mapping. Seeding and mapping help reduce the amount of traffic sent during the initial replica synchronization. For more information on when to use seeding and mapping, see Replica Seeding and Mapping.

> **IMPORTANT**
>
> If the **Replica seeding** check box is enabled in a policy, all VMs in the policy must be covered with seeding or mapping. If a VM is neither has a seed, nor is mapped to an existing VM, it will be skipped from processing.

## Configuring Replica Seeding

To configure replica seeding:

1. Make sure that you have backups of replicated VMs in a backup repository in the DR site. If you do not have the backups, create them as described in section Creating Replica Seeds for CDP.

   > **IMPORTANT**
   >
   > Consider the following:
   >
   > - Backups must be created by Veeam Backup & Replication.
   > - Backups must not reside in a scale-out backup repository.

2. Select the **Get seed from the following backup repository** check box.
3. From the list of available backup repositories, select the repository where your replica seeds are stored.

> **NOTE**
>
> If a VM has a seed and is mapped to an existing replica, replication will be performed using replica mapping because mapping has a higher priority.

## Configuring Replica Mapping

To configure replica mapping:

1. Select the **Map replicas to existing VMs** check box.
2. If you want Veeam Backup & Replication to scan the DR site to detect existing copies of VMs that you plan to replicate, click **Detect**.

   If any matches are found, Veeam Backup & Replication will populate the mapping table. If Veeam Backup & Replication does not find a match, you can map a VM to its copy manually.

3. If you want to map a VM manually, select a source VM from the list, click **Edit** and select the copy of this VM on the target host in the DR site.

To remove a mapping association, select a VM in the list and click **Remove**.

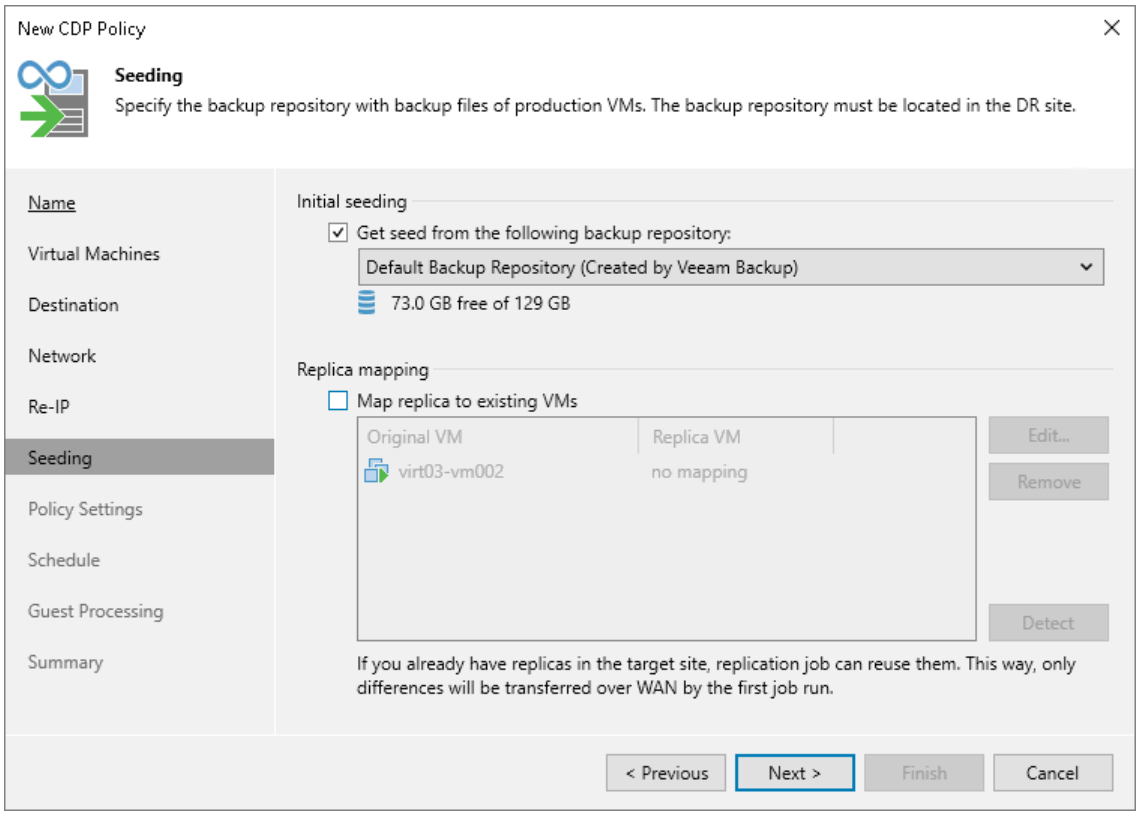

# Step 10. Specify Data Transfer and Replica Settings

At the **Policy Settings** step of the wizard, select VMware CDP proxies that must be used for the CDP policy, specify which suffix to add to replica names and VMware CDP proxies availability:

1. Specify which VMware CDP proxies you want to use:
   - If you want Veeam Backup & Replication to select proxies automatically, leave **Automatic selection** in the **Source proxy** and **Target proxy** fields.

     Veeam Backup & Replication will assign VMware CDP proxies for VM processing one by one. Before processing a new VM from the list, Veeam Backup & Replication will check available VMware CDP proxies.
   - If you want to select VMware CDP proxies manually, do the following:
     i. Click **Choose** next to the **Source proxy** field if you want to select VMware CDP proxies in the production site, or next to the **Target proxy** field if you want to select VMware CDP proxies in the disaster recovery site.
     ii. In the **Backup Proxy** window, click **Use the selected backup proxy servers only**. Select proxies that you want to use and click **OK**.

   > **NOTE**
   >
   > We recommend that you deploy at least two VMware CDP proxies: one CDP proxy in the production site and one CDP proxy in the disaster recovery site.

2. To test whether VMware CDP proxies available in the backup infrastructure can handle replication, click **Test**.

   Veeam Backup & Replication will analyze available CPU on all source and all target VMware CDP proxies, the maximum VM disk write speed during the last hour, and will calculate approximate requirements for VMware CDP proxies. In the **CDP Infrastructure Assessment** window, you will see the calculated values:
   - The **CPU** rows show CPU cores available on all proxies (source or target).
   - The **Proxy RAM** rows show RAM required for CDP and, in parenthesis, RAM available on all proxies (source or target). If values in the parentheses and near the paranthesis are the same, you need to upgrade proxies for which values coincide to provide more resources. For example, you can double up the amount of RAM.
   - The **Proxy Bandwidth** rows show the maximum disk write speed during the last hour and, in parenthesis, available bandwidth based on available cores of source or target proxies.

3. In the **Replica name suffix** field, specify a suffix that will be added to names of replicas.

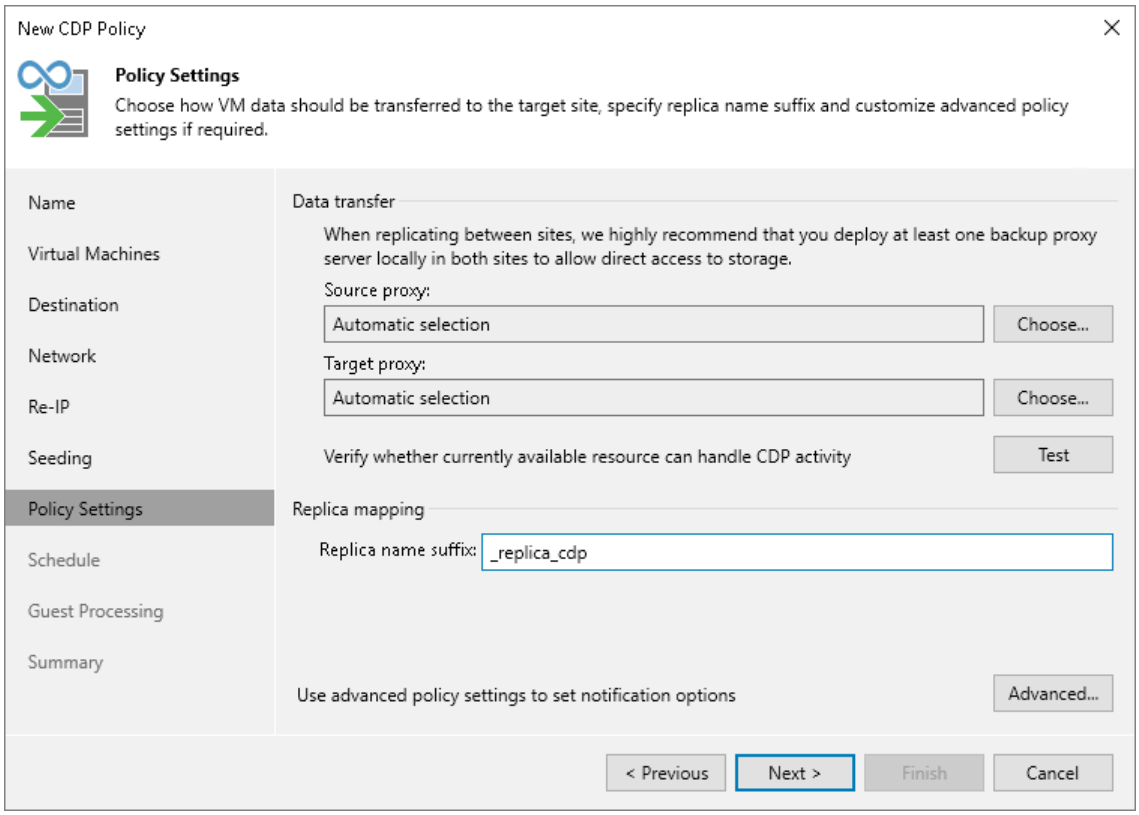


# Related Topics

How CDP Works

# Step 11. Specify Notification Settings

Veeam Backup & Replication can send by email two types of notifications for CDP policies: session reports and RPO reports.

Veeam Backup & Replication sends session reports after a CDP policy session stops. This report includes information on the policy during the session, for example read and transferred data. The session report is configured using the Global Email Notification Settings. Veeam Backup & Replication sends RPO reports after the configured RPO period ends. This report contains information on the maximum delay, SLA and other information. The RPO report is configured at the **Policy Settings** step of the wizard.

At the **Policy Settings** step of the wizard, specify RPO notification settings:

1. At the lower right corner, click **Advanced**.
2. To receive SNMP traps on the CDP policy, select the **Send SNMP notifications for this job** check box.

   SNMP traps will be sent if you configure global SNMP settings in Veeam Backup & Replication and configure software on recipient machine to receive SNMP traps. For more information, see Specifying SNMP Settings.
3. To receive notifications by email in case of policy failure, success or warning, select the **Send email notifications to the following recipients** check box. Then configure notification settings:
   a. Check that you have configured global email notification settings as described in section Configuring Global Email Notification Settings.
   b. In the text field, specify a recipient email address. If you want to specify multiple addresses, separate them by a semicolon.
   c. To use global notification settings, select **Use global notification settings**.
   d. To specify a custom notification subject and redefine at which time notifications must be sent, select **Use custom notification settings specified below.** Then specify the following settings:
      i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%JobResult%*, *%JobName%*, *%ObjectCount%* (number of VMs in the policy) and *%Issues%* (number of VMs in the policy that have been processed with the *Warning* or *Failed* status).
      ii. Select the **Notify on success**, **Notify on error** or **Notify on warning** check boxes to receive email notification if the policy gets the *Warning*, *Success* of *Error* status.

> **NOTE**
>
> A CDP policy will get the *Warning* or *Error* status according to the reporting settings configured at the **Schedule** step of the wizard. The policy will get the *Success* status after the initial configuration succeeds and every day at 8 A.M. if no error or warning occurs.

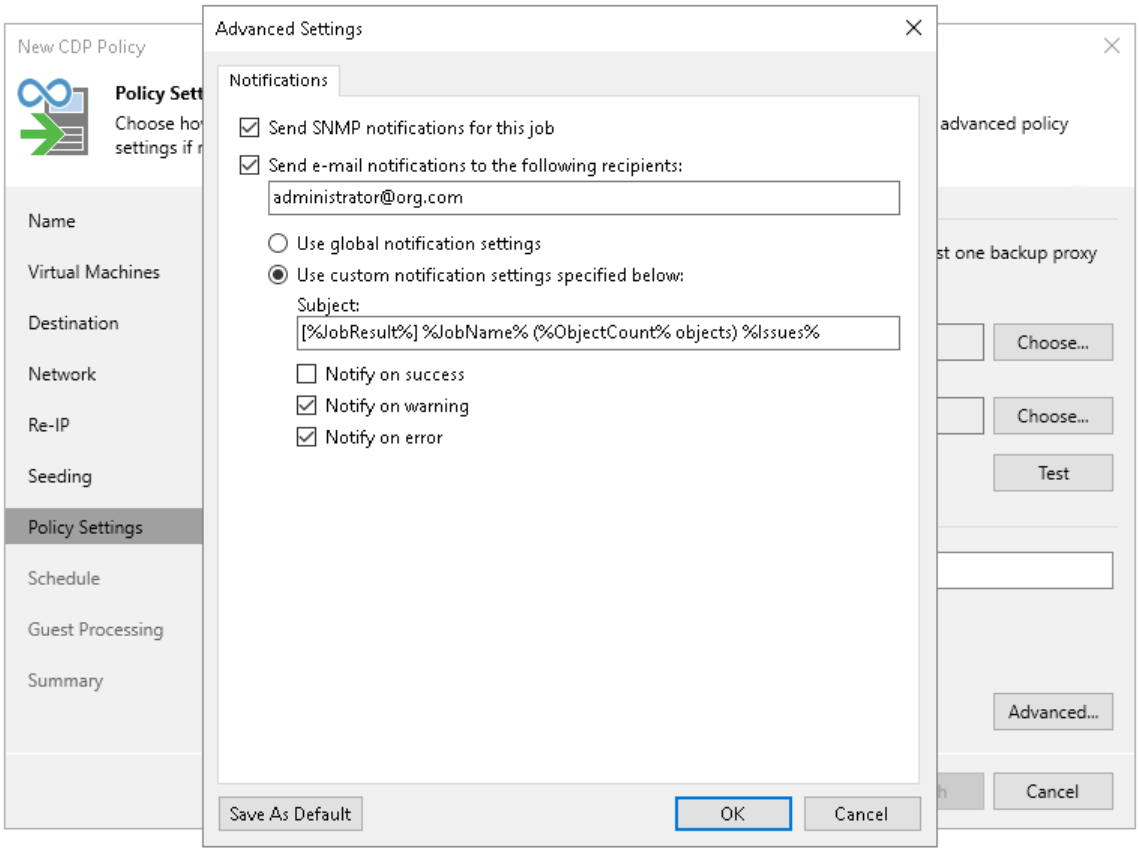
New CDP Policy
Policy Sett
Choose ho
settings if
Name
Virtual Machines
Destination
Network
Re-IP
Seeding
Policy Settings
Schedule
Guest Processing
Summary
advanced policy
st one backup proxy
Choose...
Choose...
Test
Advanced...
Cancel
Advanced Settings
Notifications
Send SNMP notifications for this job
Send e-mail notifications to the following recipients:
administrator@org.com
Use global notification settings
Use custom notification settings specified below:
Subject:
[%JobResult%] %JobName% (%ObjectCount% objects) %Issues%
Notify on success
Notify on warning
Notify on error
Save As Default
OK
Cancel

# Step 12. Specify Replication Schedule

At the **Schedule** step of the wizard, configure the schedule and retention policies:

1. Specify scheduling options:

   a. In the **Recovery point objective** field, specify the necessary RPO in seconds or minutes, that is, how often to create short-term restore points.

      The minimum RPO is 2 seconds, however it can be not optimal if your CDP policy contains many VMs with high workload. The optimal RPO is not less than 15 seconds. The maximum RPO is 60 minutes.

      During every specified period, Veeam Backup & Replication will prepare data for short-term restore points for VM replicas and send this data to the target destination. Note that short-term restore points are crash-consistent.

   b. If you want to prohibit the policy to run at specific time intervals, click **Schedule**. In the schedule box, click **Denied** and select the necessary time area.

2. To instruct the CDP policy to display a warning or error if a newly created restore points are not transferred to the target within the set RPO, click **Reporting**. Then specify when the policy must display error and warning.

   If you have configured email notification settings, Veeam Backup & Replication will mark the policy with the *Warning* or *Error* status and will also send email notifications.

3. In the **Short-term retention** section, configure the short-term retention policy, that is, specify for how long to store short-term restore points.

4. In the **Long-term retention** section, specify when to create long-term restore points and for how long to store them:

   a. In the **Create additional restore points every** field, specify how often you want to create long-term restore points.

   b. In the **Keep restore points for** field, specify for how long to retain these long-term restore points.

c. To specify time periods when Veeam Backup & Replication must create application-consistent and crash-consistent long-term restore points, click **Schedule**. In the schedule box, click **Crash-consistent** or **Application-consistent** and select the necessary time area. By default, Veeam Backup & Replication creates application-consistent backups if you enable application-aware processing at the **Guest Processing** step of the wizard. If you do not enable application-aware processing, Veeam Backup & Replication will create crash-consistent long-term restore points.

   If you want to shift the schedule, specify the offset in the **Start time within an hour** field. For example, you schedule creation of crash-consistent restore points from 00:00 to 01:00, and set the offset value to 25. The schedule will be shifted forward, and the crash-consistent restore points will be created from 0:25 and to 01:25.

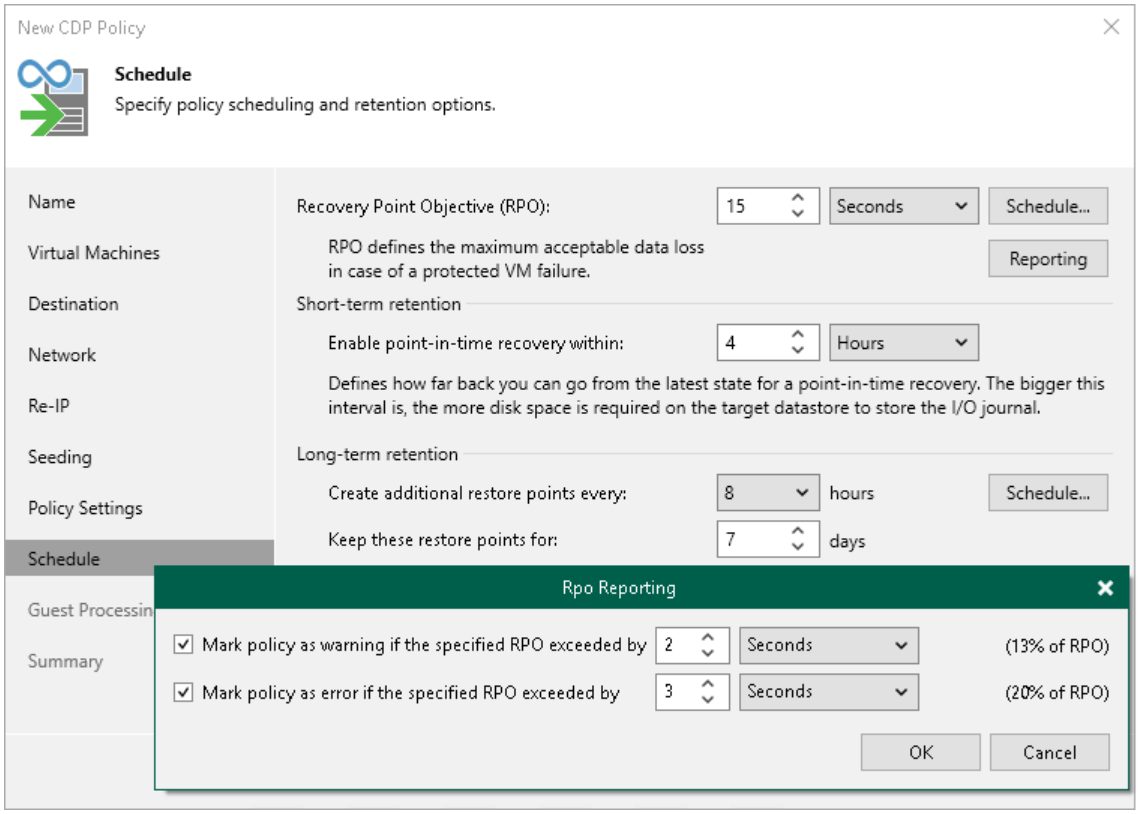

# Step 13. Specify Guest Processing Settings

Settings configured at this step apply only to long-term restore points.

At the **Guest Processing** step of the wizard, enable and configure guest OS processing.

Guest OS processing involves application-aware processing that allows creation of transactionally consistent replicas and guest file system indexing (however, indexing is not available for replicas). In its turn, application-aware processing includes log truncation, execution of custom scripts and guest OS file exclusions. For more information on guest processing, see the Guest Processing section.

To be able to use guest processing, you must also configure user accounts to access guest OSes and guest interaction proxies.

To enable guest OS processing and start configuring it (accounts and guest interaction proxies):

1. Select **Enable application-aware processing**.

   When you select this option, Veeam Backup & Replication enables application-aware processing with the default settings for all VMs. You can further disable application-aware processing for individual VMs and reconfigure the default settings.

2. If you have added Microsoft Windows VMs to be processed, specify which guest interaction proxy Veeam Backup & Replication can use to perform different guest processing tasks:

   - If you want Veeam Backup & Replication to select the guest interaction proxy automatically, leave **Automatic selection** on the **Guest interaction proxy** field.
   - If you want to explicitly specify which servers will perform the guest interaction proxy role, click **Choose**. In the **Guest Interaction Proxy** window, click **Prefer the following guest interaction proxy server**, and select the necessary proxies.

   For more information on the guest interaction proxy, requirements and limitations for it, see Guest Interaction Proxies.

3. From the **Guest OS credentials** list, select a user account that will be used to connect to guest OSes and that has enough permissions. For more information on the permissions and requirements for the user account, see Permissions for Guest Processing.

   [For Microsoft Windows VMs] Veeam Backup & Replication will also use this account to deploy the non-persistent runtime components or use (if necessary, deploy) persistent agent. For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

   [For Linux VMs] If you installed persistent agent components for VMs running Linux or Unix operating systems, select *Use management agent credentials* from the list. For more information, see Persistent Agent Components.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. For more information on adding credentials, see the Credentials Manager section.

   > **NOTE**
   >
   > If you plan to use Kerberos authentication, check limitations and requirements listed in section Guest Processing.

4. To specify credentials for individual workloads, click **Credentials**. Then select the necessary workload and set user credentials for it.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose a VM whose settings you want to customize. Then select the VM in the list and define the necessary settings.

5. To check whether Veeam Backup & Replication can connect to VMs using the specified guest OS credentials and can deploy the non-persistent runtime components or connect to persistent agent components on the guest OSes, click **Test Now**.

After you have enabled application-aware processing for all VMs and configured other settings required for guest processing, you can disable application-aware processing for individual VMs and change the default settings. For more information, see the following sections:

- Application-aware processing and transaction logs
- Microsoft SQL Server transaction log settings
- Oracle archived log settings
- PostgreSQL Settings
- Pre-freeze and post-thaw scripts

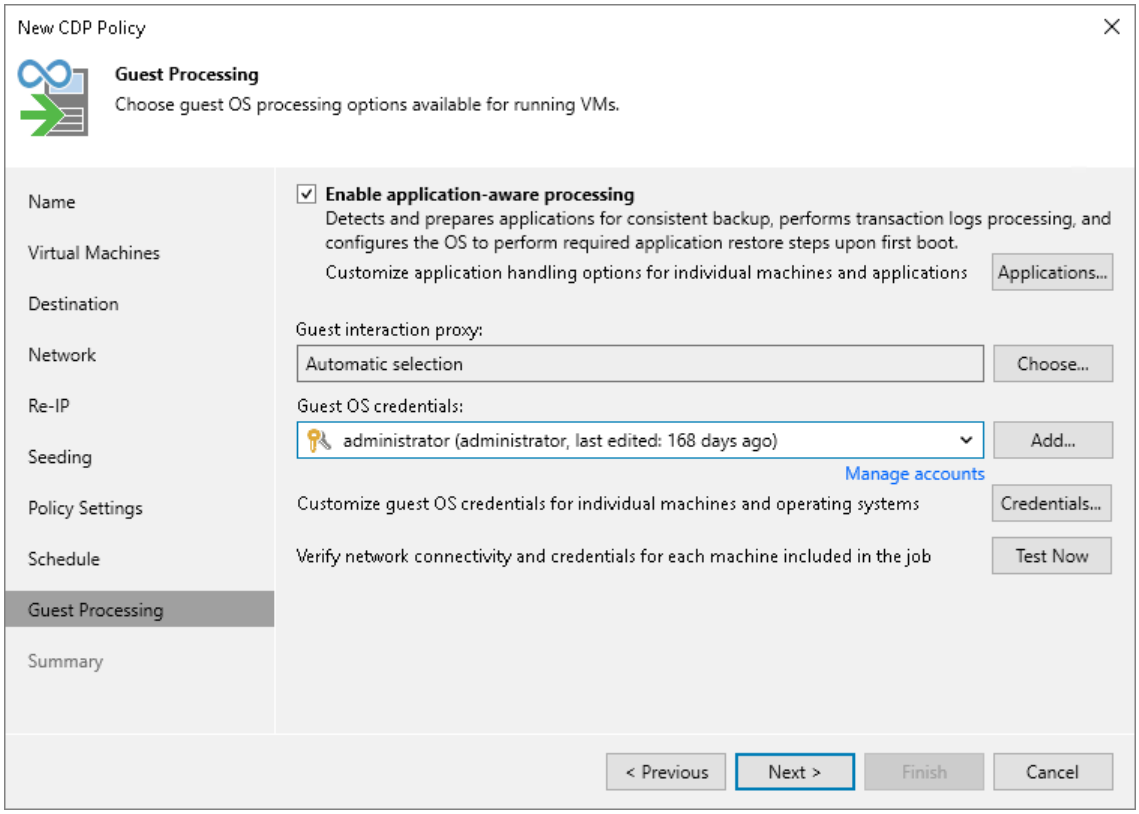


# Application-Aware Processing and Transaction Logs

Application-aware processing helps create transactionally consistent replicas. The transactionally consistent replicas guarantee proper recovery of applications without data loss. For more information on application-aware processing, see Application-Aware Processing.

To configure general application-aware processing settings and specify whether Veeam Backup & Replication processes transaction logs or creates copy-only replicas:

1. At the **Guest Processing** step of the wizard, check that you have selected the **Enable application-aware processing** check box and configured guest OS credentials.
2. At the **Guest Processing** step of the wizard, click **Applications**.
3. In the **Application-Aware Processing Options** window, select workloads for which you want to configure application-aware processing, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. On the **General** tab, in the **Applications** section, specify the behavior scenario for application-aware processing:
   - Select **Require successful processing** if you want Veeam Backup & Replication to stop the replication process if any error occurs during application-aware processing.
   - Select **Try application processing, but ignore failures** if you want to continue the replication process even if an error occurs during application-aware processing. This option guarantees that replication will continue working. However, the resulting replica will be crash consistent, not transactionally consistent.
   - Select **Disable application processing** if you want to disable application-aware processing for the workload.
5. [For Microsoft Exchange and Microsoft SQL Server] In the **VSS Settings** section, specify if Veeam Backup & Replication must process transaction logs or create copy-only replicas:
   a. Select **Process transaction logs with this job** if you want Veeam Backup & Replication to process transaction logs.

      [For Microsoft Exchange VMs] With this option selected, the non-persistent runtime components or persistent components running on the VM guest OS will wait for replication to complete successfully and then trigger truncation of transaction logs. If the replication job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

      [For Microsoft SQL Server] You will need to configure how to process transaction logs.

      > **TIP**
      >
      > To configure log processing for Oracle and PostgreSQL databases, switch to the Oracle and PostrgeSQL tabs.

   b. Select **Perform copy only** if you use another tool to perform guest level processing, and this tool maintains consistency of the database state. Veeam Backup & Replication will create a copy-only replica for the selected VMs. The copy only replica preserves the chain of full and differential files and transaction logs on the VM. For more information, see Microsoft Docs.
6. [For Microsoft Windows VMs] In the **Persistent guest agent** section, select the **Use persistent guest agent** check box to use for application-aware processing persistent guest agents on each protected VM.

   By default, Veeam Backup & Replication uses non-persistent runtime components. Veeam Backup & Replication deploys runtime components on each protected VM when the replication job starts, and removes the runtime components as soon as the replication job finishes.

   For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

[For Linux VMs] To use persistent guest agents, you must install Management Agent on protected VMs. For more information, see Persistent Agent Components.

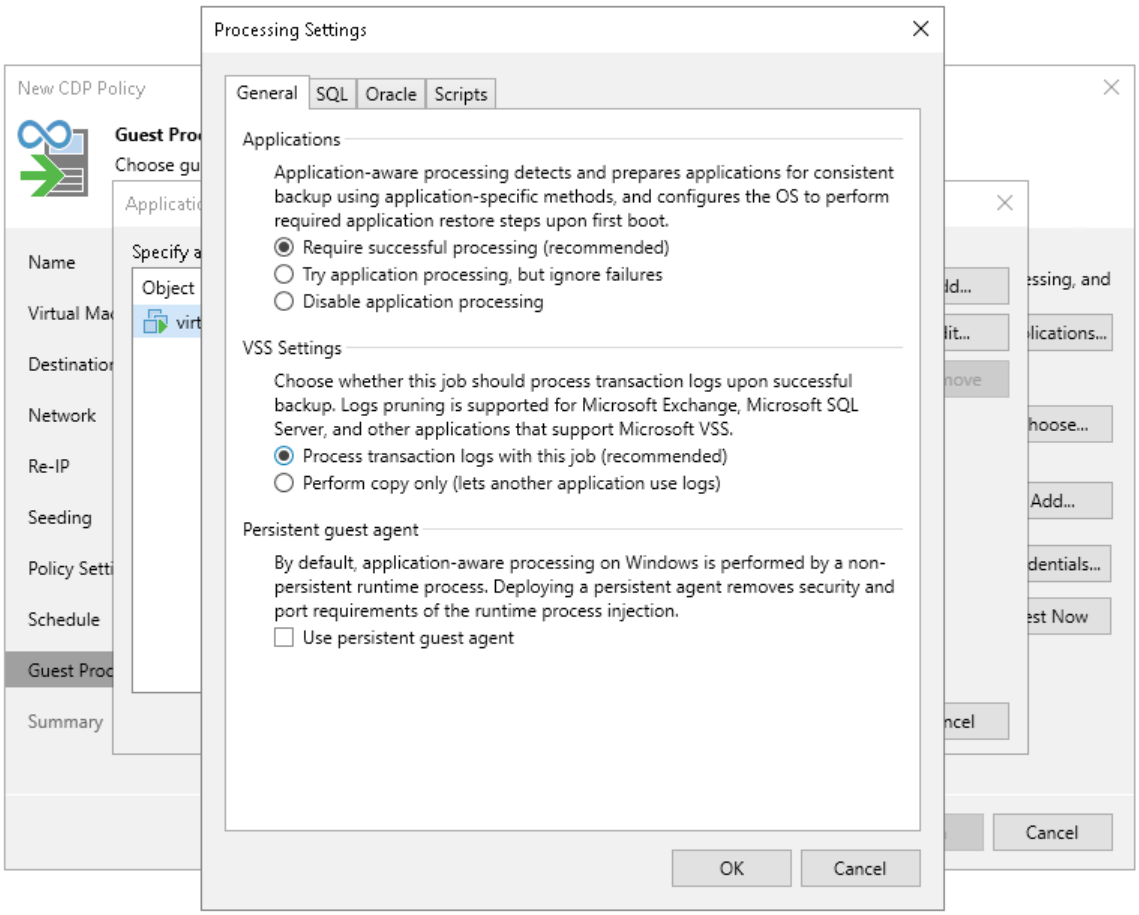


# Microsoft SQL Server Transaction Log Settings

The **SQL** tab is available for VMs that run Microsoft SQL Server and if you have selected **Process transaction logs with this job** when configuring application-aware processing.

To create transactionally consistent backups of an Microsoft SQL Servers, you must check that application-aware processing is enabled and then specify settings of transaction log processing.

# Enabling Application-Aware Processing

Before configuring transaction log processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Microsoft SQL Server and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

# Specifying Transaction Log Settings

In the **Processing Settings** window, switch to the **SQL** tab and specify how transaction logs must be processed:

- If you want Veeam Backup & Replication to trigger truncation of transaction logs after the CDP policy creates a long-term restore point, select **Truncate logs**.

  In this case, transaction logs will be truncated after the CDP policy creates a long-term restore point. If the creation fails, the logs will remain untouched until the next start of the long-term restore point creation.

- If you do not want Veeam Backup & Replication to truncate logs at all, select **Do not truncate logs**.

  This option is recommended if you use another tool to perform VM guest-level replication, and this tool maintains consistency of the database state.

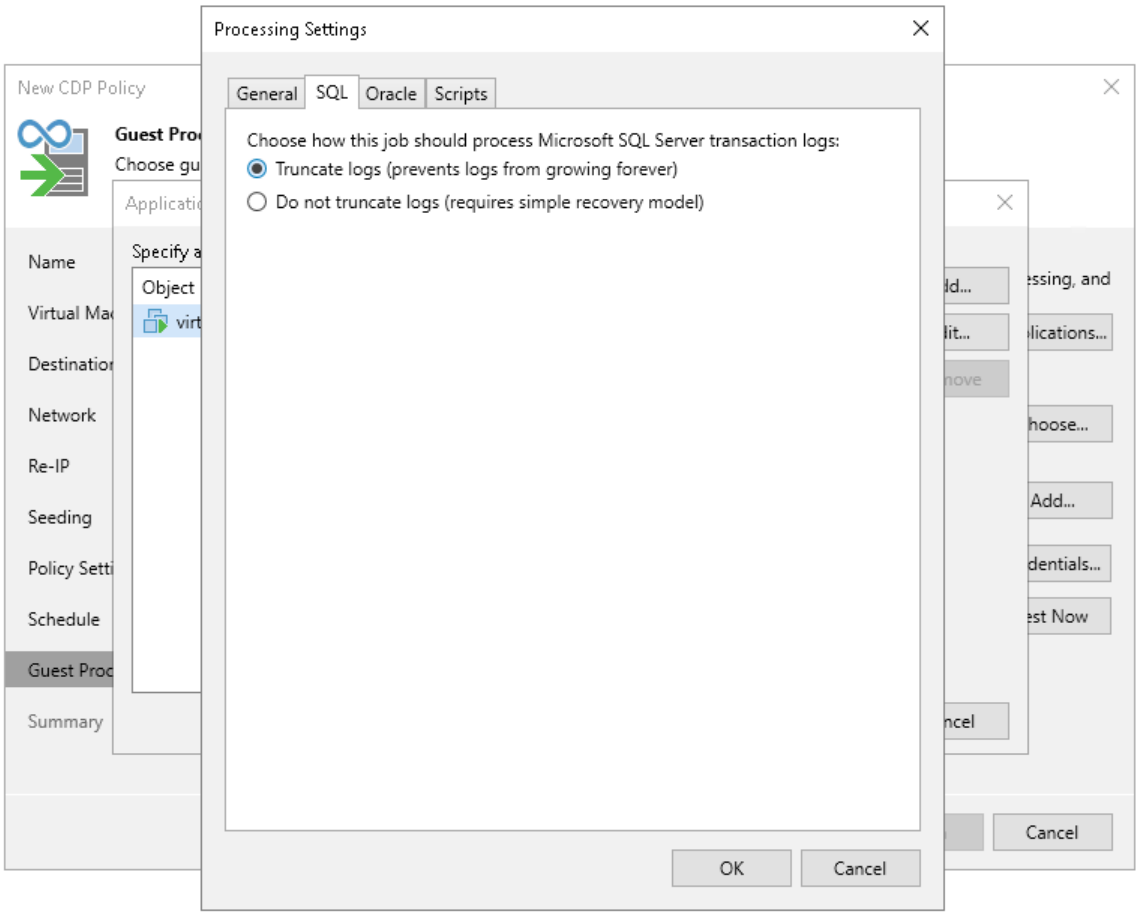


# Oracle Archived Log Settings

The **Oracle** tab applies to Linux VMs that run Oracle.

To create transactionally consistent backups of an Oracle server, you must check that application-aware processing is enabled and then specify settings of archive log processing.

# Enabling Application-Aware Processing

Before configuring archive log processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.

2. Click **Applications**.
3. In the displayed list, select the Oracle server and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

# Specifying Archive Log Settings

To configure how Veeam Backup & Replication must process archive logs of an Oracle server:

1. In the **Processing Settings** window, switch to the **Oracle** tab.
2. From the **Specify Oracle account with SYSDBA privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the Oracle databases. The account that you plan to use must have privileges described in section Permissions.

   You can select **Use guest credentials** from the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle databases.
3. In the **Archived logs** section, specify how to process archived logs:
   - If you want to preserve archived logs on the VM guest OS, select **Do not delete archived logs**. When the replication job completes, the non-persistent runtime components or persistent components will not truncate transaction logs.

     It is recommended that you select this option for databases where the ARCHIVELOG mode is turned off. If the ARCHIVELOG mode is turned on, archived logs may grow large and consume all disk space.
   - If you want to delete archived logs older than <N> hours, select **Delete logs older than <N> hours** and specify the number of hours.

- If you want to delete archived logs larger than <N> GB, select **Delete logs over <N> GB** and specify the size. The specified size refers to the log size of each database, not all databases on the selected Oracle server.

Transaction logs will be deleted using Oracle Call Interface after the CDP policy creates a long-term restore point. If the creation fails, the logs will remain untouched until the next start of the long-term restore point creation.

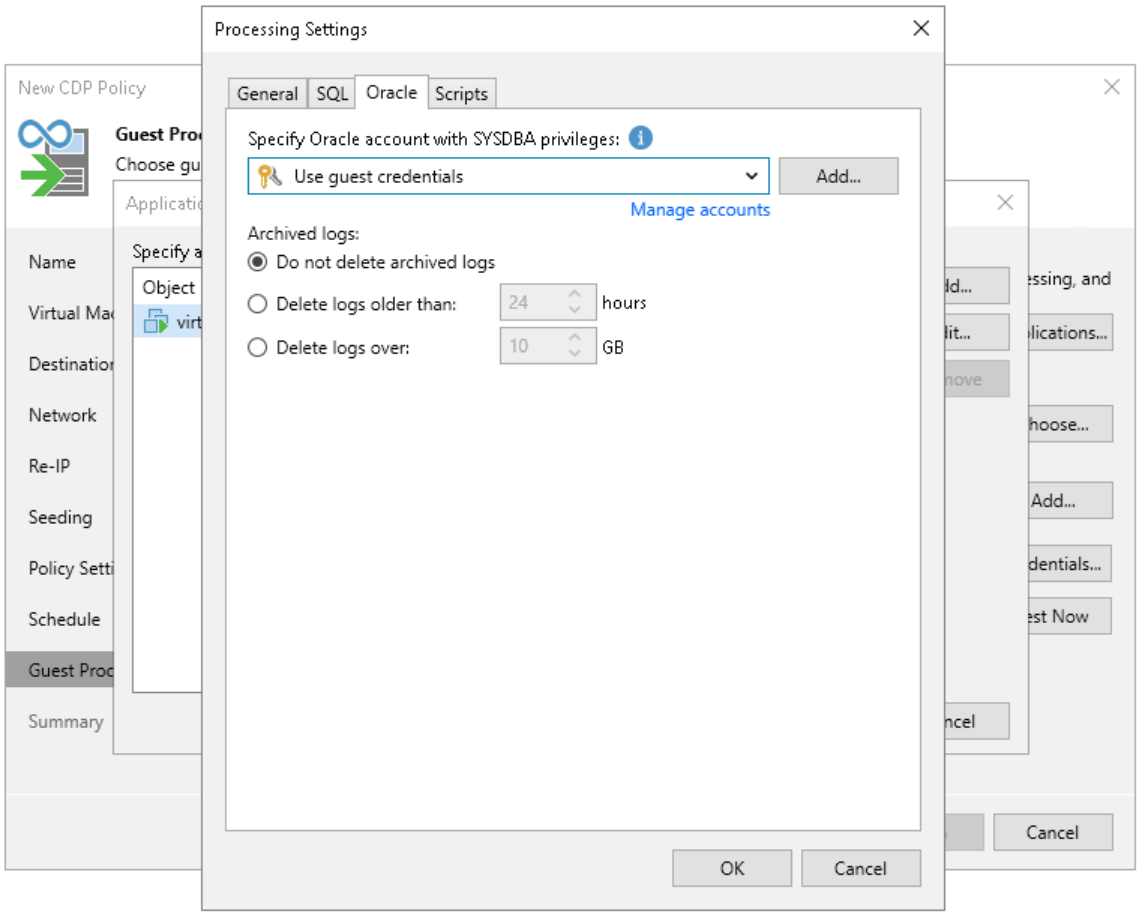


# PostgreSQL Settings

The **PostgreSQL** tab applies to Linux VMs that run PostgreSQL.

To create transactionally consistent backups of a PostgreSQL VM, you must check that application-aware processing is enabled and then specify settings of WAL files processing.

# Enabling Application-Aware Processing

Before configuring WAL files processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the PostgreSQL VM and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.

4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

> **NOTE**
>
> By default, Veeam Backup & Replication recursively scans the `/etc/postgresql`, `/var/lib/postgresql` and `/var/lib/pgsql` directories for the configuration files of PostgreSQL instances. If you keep configuration files in custom directories, the *pgsqlagent agent* will use its own `VeeamPostgreSQLAgent.xml` configuration file that is located in the `/etc/veeam/` directory. The pgsqlagent agent configuration file must be a single line XML.
>
> To explicitly include or exclude specific configuration files from rescan, you can add the following commands to the `VeeamPostgreSQLAgent.xml` file:
>
> - `ExcludeConfigDirs` — use this element to exclude configuration files.
> - `AddConfigDirs` — use this element to include configuration files.
>
> For example: `<config AddConfigDirs="/opt/psql/" ExcludeConfigDirs="/var/lib/postgresql/13/main45/,/var/lib/postgresql/13/maindd/" />`

# Specifying WAL Files Settings

To define how Veeam Backup & Replication will process WAL files on this VM, do the following:

1. In the **Processing Settings** window, click the **PostgreSQL** tab.
2. From the **Specify PostgreSQL account with superuser privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the PostgreSQL instance. The account must have privileges described in section Permissions. You can select **Use guest credentials** from the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the PostgreSQL instance.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. Note that if you select the **System user without password file (peer)** option in the **The specified user is** area, you can add a user account without specifying a password.
3. In the **The specified user is** section, specify how the user selected in the **Specify PostgreSQL account with superuser privileges** drop-down list will authenticate against the PostgreSQL instance:
   - Select **Database user with password** if the account is a PostgreSQL account, and you entered the password for this account in the Credentials Manager.

     > **NOTE**
     >
     > If you want Veeam Backup & Replication to use the user name map authentication, select **Database user with password** and leave the password field empty. Consider the following:
     >
     > - **Guest OS credentials** specified at the **Guest Processing** step of the wizard will be used as the System-Username.
     > - PostgreSQL account specified at the step 2 will be used as the PG-Username.
     >
     > For more information about the user name maps, see PostgreSQL documentation.
   - Select **Database user with password file (.pgpass)** if the password for the account is defined in the `.pgpass` configuration file on the PostgreSQL server. For more information about the password file, see PostgreSQL documentation.

- Select **System user without password file (peer)** if you want Veeam Backup & Replication to use the peer authentication method. In this case, Veeam Backup & Replication will apply the OS account as the PostgreSQL account.

> **IMPORTANT**
>
> If you have added a new PostgreSQL account and want to use it with the peer authentication method, make sure that you have added this account as a Linux user with sufficient permissions.

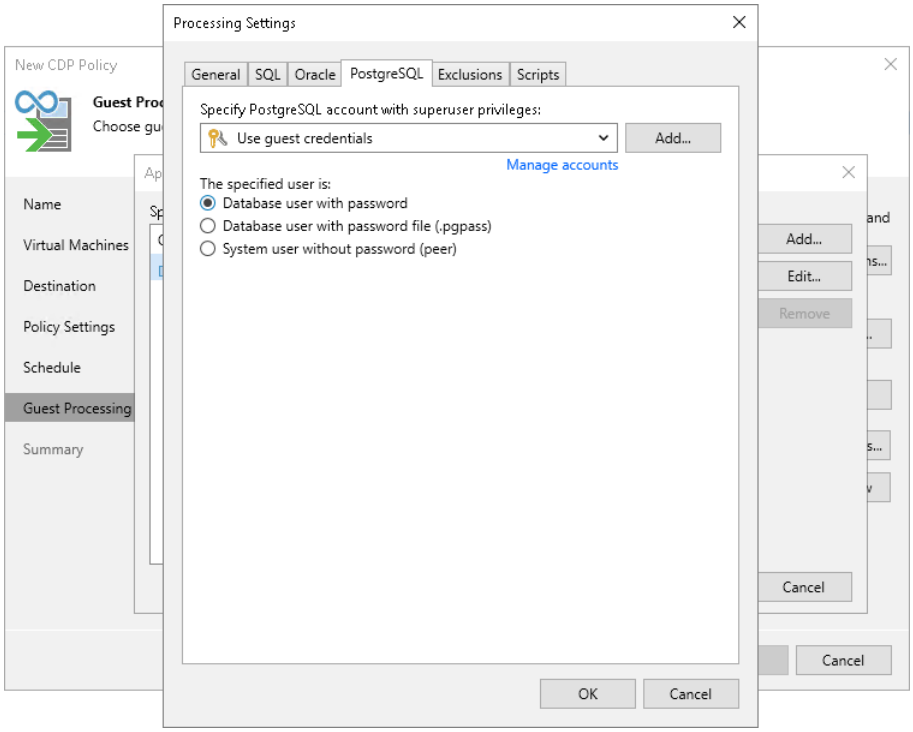


# Script Settings

You can instruct Veeam Backup & Replication to run custom scripts before the CDP policy starts the creation of a long-term restore point and after the policy finishes the creation. For example, these can be pre-freeze and post-thaw scripts for a VM that does not support VSS. The scripts will quiesce the VM file system and application data to bring the VM to a consistent state before the creation of a restore point, and bring the VM and applications to their initial state after the creation finishes.

To specify pre-freeze and post-thaw scripts:

1. At the **Guest Processing** step of the wizard, check that you have selected the **Enable application-aware processing** check box and configured guest OS credentials.
2. At the **Guest Processing** step of the wizard, click **Applications**.
3. In the **Application-Aware Processing Options** list, select workloads for which you want to configure scripts, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. Click the **Scripts** tab.

5. In the **Script processing mode** section, select a scenario for script execution:
   - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the replication process if scripts fail.
   - Select **Ignore script execution failures** if you want to continue the replication process even if script errors occur.
   - Select **Disable script execution** if you do not want to run scripts for the VM.
6. In the **Windows scripts** section, specify paths to scripts for Microsoft Windows VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.
7. In the **Linux scripts** section, specify paths to scripts for Linux VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

   If you plan to replicate a VM container with Microsoft Windows and Linux VMs, you can select to execute both Microsoft Windows and Linux scripts. When replication starts, Veeam Backup & Replication will automatically determine which OS type is installed on the VM and use the correct scripts for this VM.

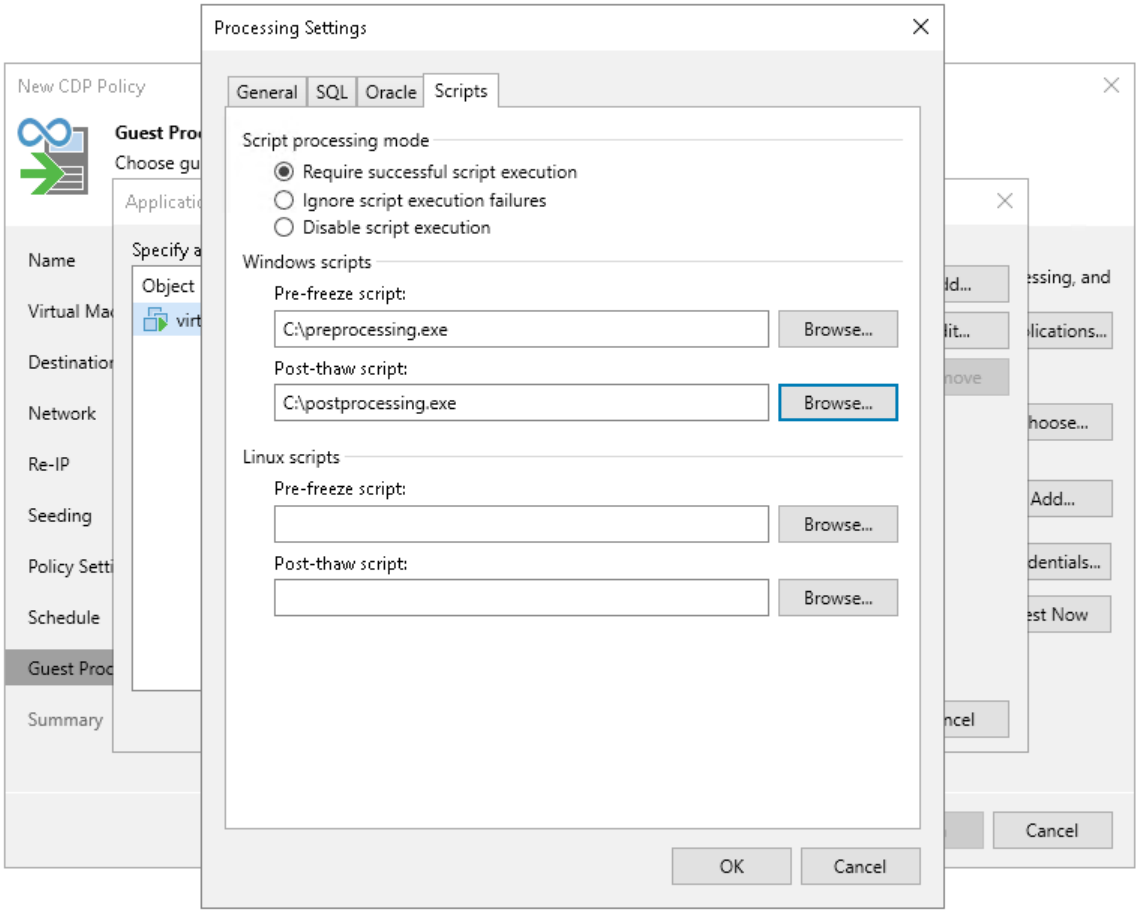

# Step 14. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured settings. If you want to start the policy right after you close the wizard, leave the **Enable the policy when I click Finish** check box selected, otherwise clear the check box. Then click **Finish** to close the wizard.

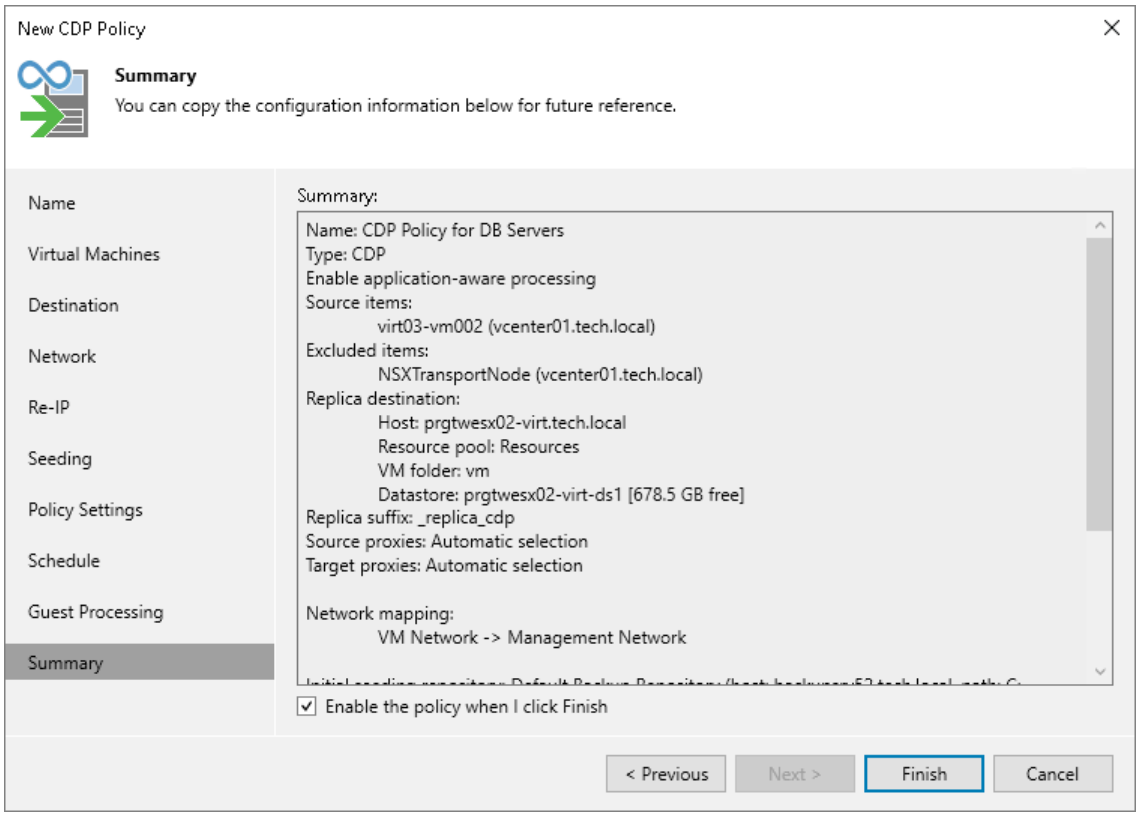

# Creating Replica Seeds for CDP

To use replica seeding in a CDP policy, you must have backups of replicated VMs in a backup repository in the disaster recovery (DR) site. These backups are known as replica seeds. For more information on seeding and when to use it, see Replica Seeding and Mapping.

If you do not have replica seeds in the DR site, do the following:

1. Create a backup of VMs that you plan to replicate as described in section Creating Backup Jobs. As the target repository for this job, select a backup repository in the production site. Then run the job.

   If you already have backups containing the necessary VMs, there is no need to configure and run a new backup job. For seeding, you can use any existing backups created by Veeam Backup & Replication. The backup must include VBK and VBM files. If you have a full backup and a chain of forward increments, you can use VIB files together with the VBK and VBM files. In this case, Veeam Backup & Replication will restore VMs from the seed to the latest available restore point.

2. Copy the backup from the backup repository in the production site to a backup repository in the DR site.

   You can move the backup using a file copy job or any other appropriate method, for example, copy the backup to a removable storage device, ship the device to the DR site and copy backups to the backup repository in the DR site.

   If you do not have a backup repository in the DR site, you need to create the repository as described in section Backup Repositories.

   > **IMPORTANT**
   >
   > You cannot copy backups to a scale-out backup repository in the DR site.

3. After the backup is copied to the backup repository in the DR site, perform rescan of this backup repository as described in section Rescanning Backup Repositories. Otherwise, Veeam Backup & Replication will not be able to detect the copied backup.

# Managing CDP Policies

After you create CDP policies, you can edit, disable and delete them.

To view all created CDP policies, open the Home view and navigate to the **Jobs** > **CDP** node. The working area displays the full list of the created policies. Here, you can manage the policies.

# Viewing Session Statistics and Results

Veeam Backup & Replication allows you to view real-time statistics and session results for a VM added to a CDP policy:

- To view real-time statistics for the policy, open the **Home** view. In the inventory pane select **Jobs** > **CDP**. In the working area, double-click the necessary policy. Alternatively, right-click the policy and select **Statistics**. In the opened window, you will be able to switch between statistics of individual VMs.
- To view real-time statistics for an individual VM, open the **Home** view. In the inventory pane select **Jobs, Last 24 hours or Running**. In the working area, double-click the necessary VM. Alternatively, right-click the VM and select **Statistics**.
- To view statistics on the finished policy sessions, open the **History** view. In the inventory pane select **CDP**. In the working area, double-click the necessary policy session.

The statistics provides detailed data on policy sessions: duration, performance bottlenecks, amount of processed data, read and transferred data.

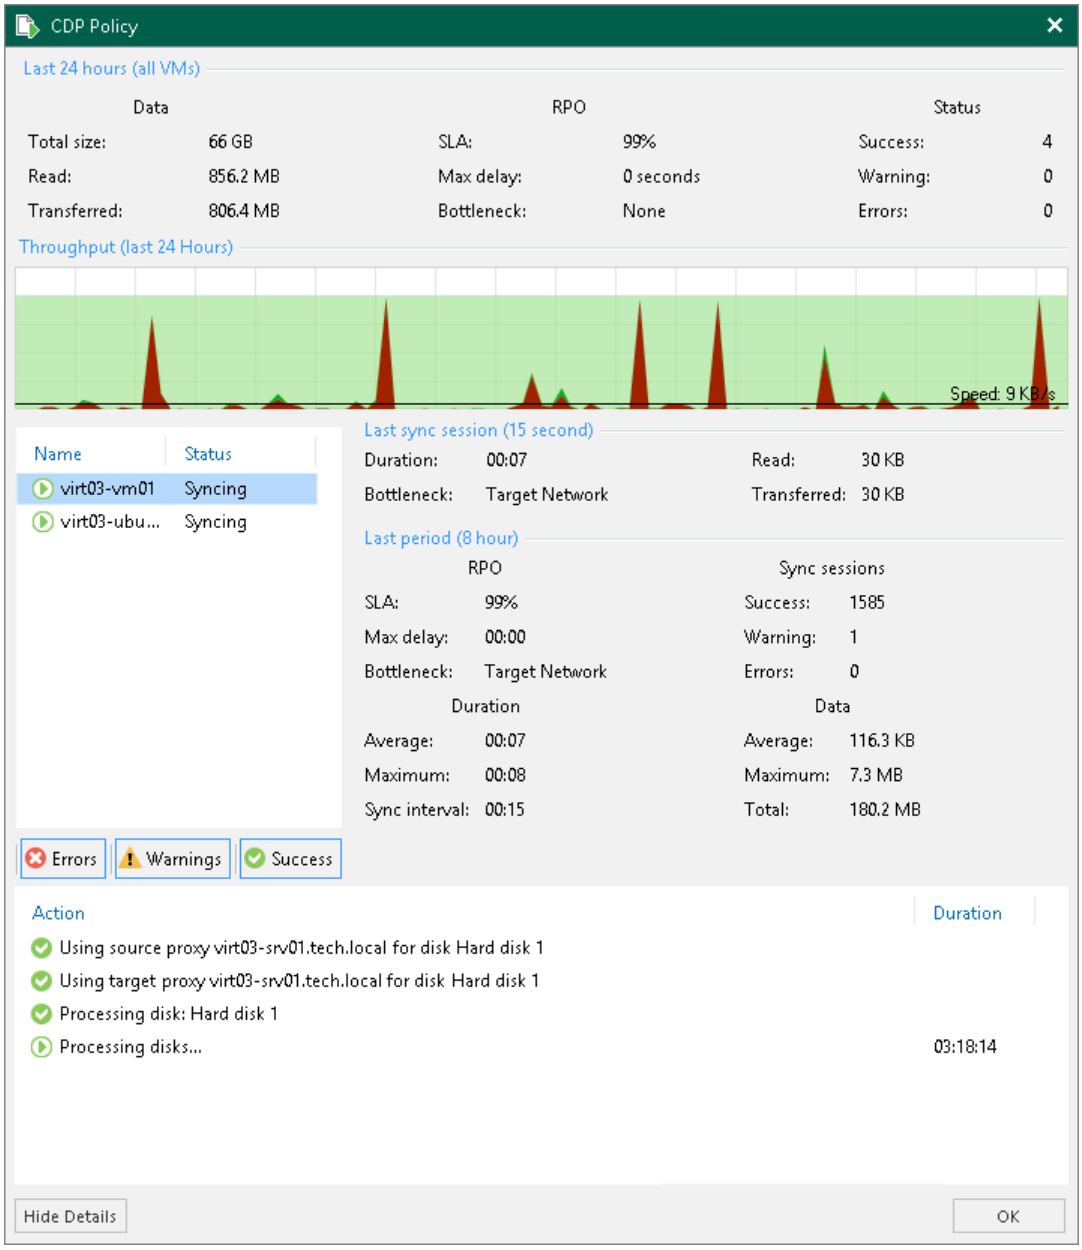

# Statistics Counters

Veeam Backup & Replication displays policy statistics for the following counters:

- The **Last 24 hours (all VMs)** section shows the general information for all VMs for which you have opened the statistics. This information is collected during 24 hours.
  - The **Data** box shows information about processed VM data:
    - **Total size** – total size of the processed data.
    - **Read** – amount of data read from the datastore prior to applying compression and deduplication. The value of this counter is typically lower than the value of the **Total size** counter. Veeam Backup & Replication reads only data blocks that have changed since the last policy session, processes and copies these data blocks to the target.
    - **Transferred** – amount of data transferred from the source VMware CDP proxy to the target VMware CDP proxy. The data is transferred after compression and deduplication.
  - The **RPO** box shows information related to RPO:
    - **SLA** – percentage of sessions completed within the desired RPO.
    - **Max delay** – difference between the configured RPO and time required to transfer and save data.
    - **Bottleneck** – bottleneck in the data transmission process.
  - The **Status box** shows information about task session results. This box informs how many task sessions have completed with the *Success*, *Warning* and *Error* statuses during the 24-hour session. One task session lasts for the period between the creation of two long-term restore points.
- The **Last sync session** section shows general information collected during the last synchronization session of the 24-hour session. One synchronization session lasts for the period between the creation of the two short-term restore points.
  - **Duration** – time period during which data was collected and sent to the target host.
  - **Bottleneck** – bottleneck in the data transmission process.
  - **Read** – amount of data read from the datastore prior to applying compression and deduplication.
  - **Transferred** – amount of data transferred from the source VMware CDP proxy to the target VMware CDP proxy. The data is sent after compression and deduplication.
- The **Last period** section shows general information collected during the last task session of the 24-hour session. One task session lasts for the period between the creation of two long-term restore points.
  - The **RPO** box shows information related to RPO:
    - **SLA** – percentage of sessions completed within the desired RPO.
    - **Max delay** – difference between the configured RPO and time required to transfer and save data.
    - **Bottleneck** – bottleneck in the data transmission process.
  - The **Sync session** box shows information about synchronization session results. This box informs how many synchronization sessions have completed with the *Success*, *Warning* and *Error* statuses during the last task session. One synchronization session lasts for the period between the creation of the two short-term restore points.

- The **Duration** box shows information about duration of synchronization sessions:
  - **Average** — average duration of a synchronization session.
  - **Maximum** — maximum duration of a synchronization session.
  - **Sync interval** — duration of a synchronization session configured in the policy, that is, the specified RPO.
- The **Data** box shows information about processed VM data:
  - **Average** — average amount of data processed within one synchronization session.
  - **Maximum** — maximum amount of data processed within one synchronization session.
  - **Total** — total size of data sent during the task session.

- The pane at the left shows the list of VMs included into the session statistics.
- The pane at the bottom shows the list of operations performed during the session. If you open statistics for a policy, you can see the list of operations for the whole policy or an individual VM. To see the list of operations for an individual VM, click the VM in the pane on the left. To see the list of operations for the whole policy, click anywhere on the blank area in the left pane.

# Colored Graph

To visualize the data transfer process, Veeam Backup & Replication displays a colored graph in the statistics window:

- The green color defines the amount of data read from the source.
- The brown color defines the amount of data transferred to the target.
- The horizontal line defines the current data processing speed.

If the session is still being performed, you can click the graph to view data rate for the last 5 minutes or the last 24 hours. If the session has already ended, the graph will display information for the last 24 hours only.

The colored graph is displayed only for the currently running session or the latest finished session. If you open statistics for past sessions other than the latest one, the colored graph will not be displayed.

# Editing Policies

To edit a CDP policy:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **CDP** node.
3. In the working area, select the necessary policy and select **Edit** on the ribbon. As an alternative, right-click the necessary policy and select **Edit**.
4. Follow the instructions provided in the Creating CDP Policies section.

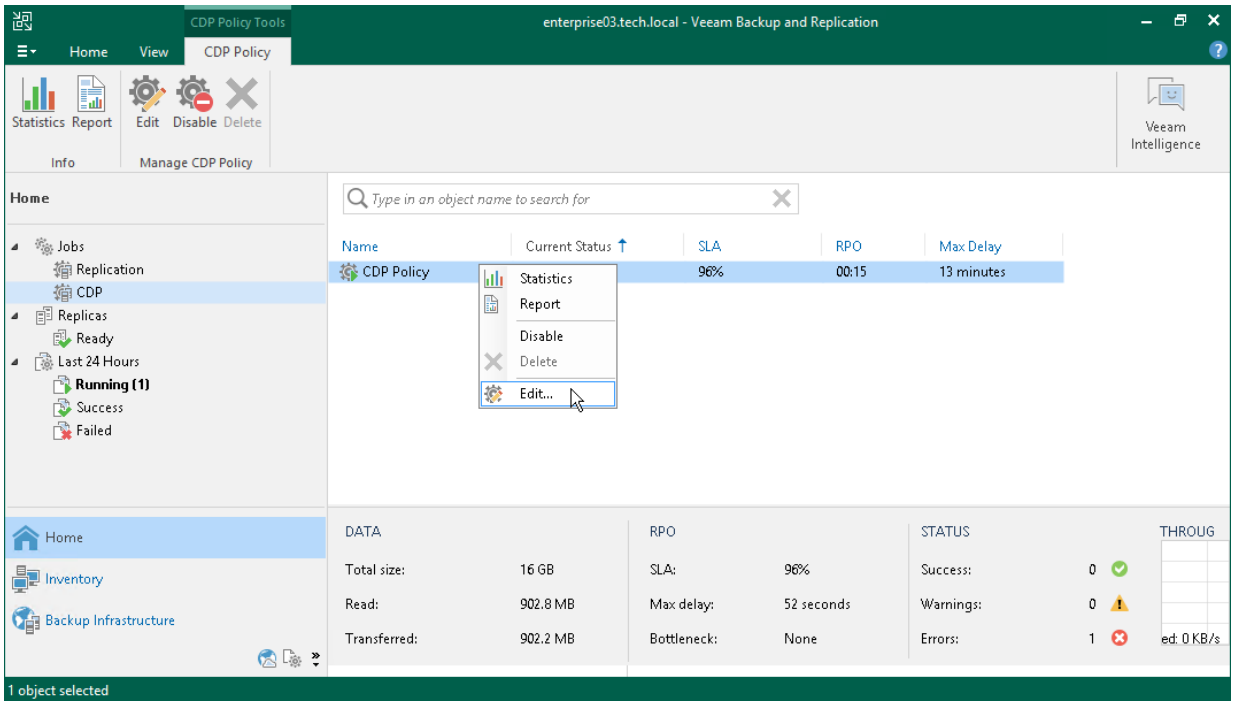

# Cloning Policies

You can create a new CDP policy by cloning an existing one. Policy cloning allows you to create an exact copy of any policy with the same policy settings.

The name of the cloned policy is formed by the following rule: *<policy_name_clone1>*, where *policy_name* is the name of the original policy and *clone1* is a suffix added to the original policy name. If you clone the same policy again, the number in the name will be incremented, for example, *policy_name_clone2*, *policy_name_clone3*, and so on.

## Considerations and Limitations

When cloning a policy, Veeam Backup & Replication can change some policy settings so that cloned policies do not hinder original policies.

- Veeam Backup & Replication disables the cloned policy. To enable the cloned policy, select it in the policy list and click **Disable** on the ribbon. Alternatively, right-click the policy and select **Disable**.
- Before enabling the policy, you must edit the list of protected workloads. One workload can be protected only by one CDP policy.
- In the cloned policy, you cannot change the target destination.

## Cloning Policy

To clone a CDP policy:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.

3. In the working area, select the policy and click **Clone** on the ribbon. Alternatively, right-click the policy and select **Clone**.

   After the policy is cloned, you can edit its settings, including the policy name.

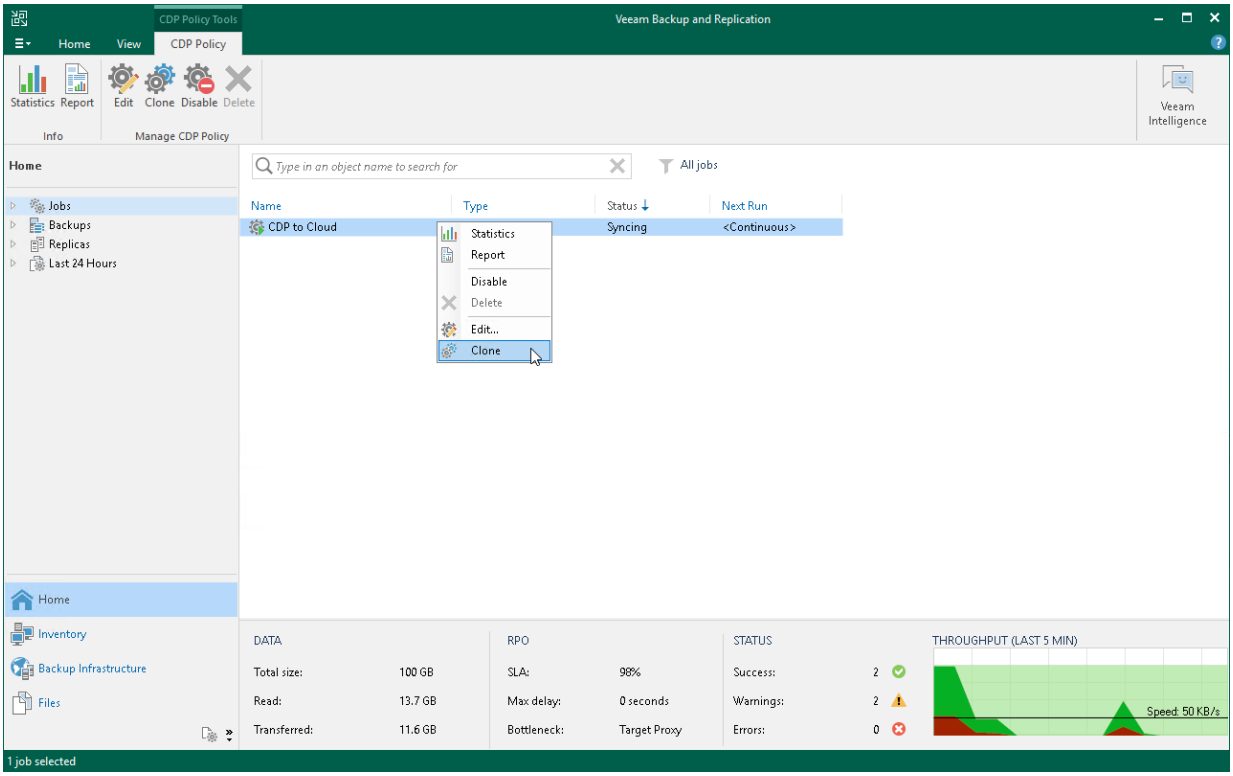

# Disabling and Deleting Policies

Veeam Backup & Replication allows you to temporarily disable or permanently delete created CDP policies.

## Disabling CDP Policies

To disable a CDP policy:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **CDP** node.
3. In the working area, select the necessary policy and select **Disable** on the ribbon. Alternatively, right-click the necessary policy and select **Disable**.

> **TIP**
>
> To enable a disabled policy, select it and click **Disable** once again.

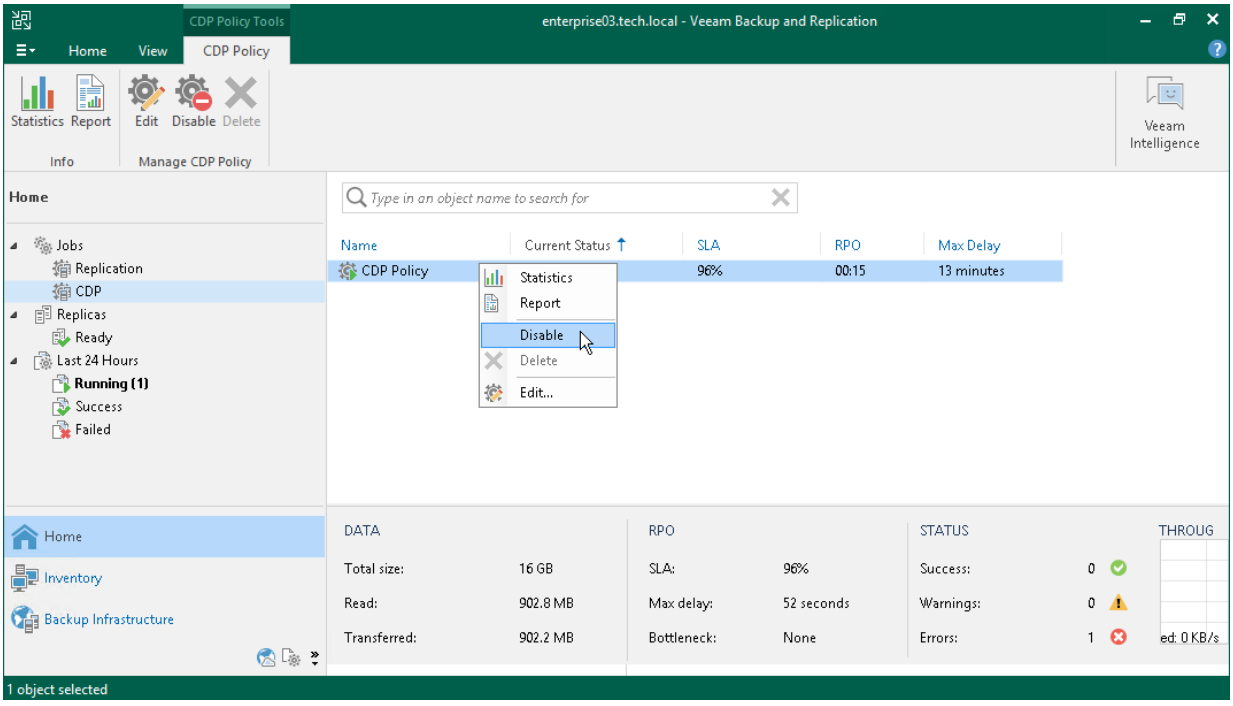


## Deleting CDP Policies

Veeam Backup & Replication allows you to delete only disabled policies. To delete a CDP policy:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **CDP** node.
3. Check that you have disabled the policy that you want to delete.

4. In the working area, select the necessary policy and select **Delete** on the ribbon. Alternatively, right-click the necessary policy and select **Delete**.

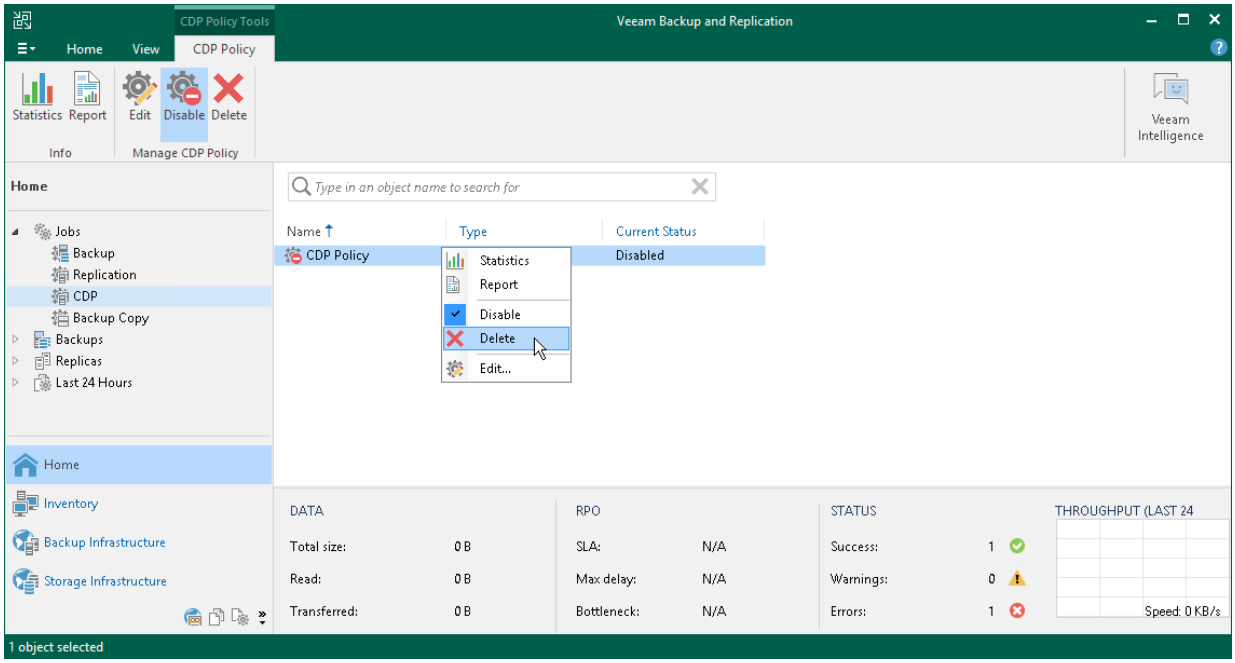

# Managing CDP Replicas

To view all created replicas, open the **Home** view and navigate to the **Replicas** node. The working area displays the full list of the created replicas. Here, you can view replica properties and delete replicas from the configuration database or disk.

# Viewing Replica Properties

To view replica properties:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, right-click the necessary replica and select **Properties**. Alternatively, select **Properties** on the ribbon.

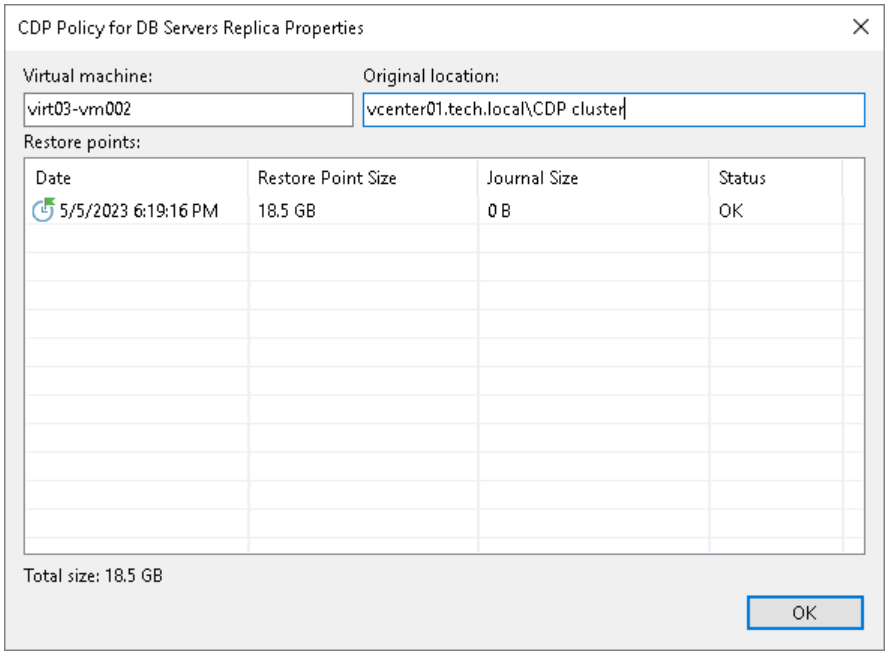

# Rescanning CDP Replicas

You may need to perform replica rescan, for example, if the actual states of replicas differ from the states in the configuration database, if changes were made in the backup infrastructure and so on.

During the rescan process, Veeam Backup & Replication gathers information on replicas that are currently available and updates the list of replicas in the configuration database.

To rescan replicas, do the following:

1. Open the **Home** view.
2. In the inventory pane, right-click the **Replicas** node and select **Rescan Replicas**.

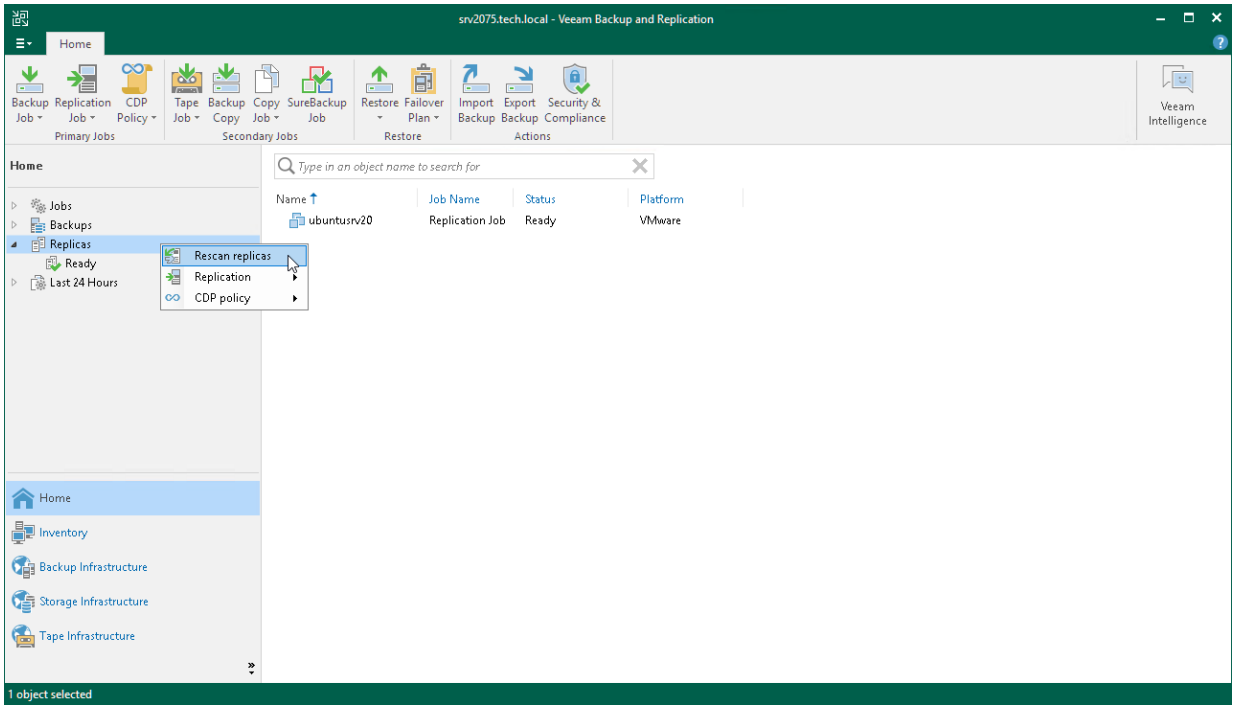

# Removing Replicas from Configuration

When you remove replicas from the configuration, Veeam Backup & Replication removes records about the replicas from the configuration database, stops showing the replicas in Veeam Backup & Replication console and also stops synchronizing their state with the state of the source VMs. However, the actual replicas remain on hosts.

To remove records about replicas from the Veeam Backup & Replication console and configuration database:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.
3. In the working area, select replicas in the *Ready* state and click **Remove from** > **Configuration** on the ribbon. Alternatively, right-click one of the selected replicas and select **Remove from configuration**.

**NOTE**

Consider the following:

- You can remove records only about replicas in the *Ready* state.
- When you remove from the configuration a VM that is replicated as a standalone object, Veeam Backup & Replication removes this VM from the initial replication job. When you remove from the configuration a VM that is replicated as part of a VM container, Veeam Backup & Replication adds this VM to the list of exclusions in the CDP policy.
- When you remove from the configuration a replica, Veeam Backup & Replication consolidates the replica into a regular VM. Veeam Backup & Replication removes technical files, such as transaction log and delta disk files, and commits data from them into the VM.

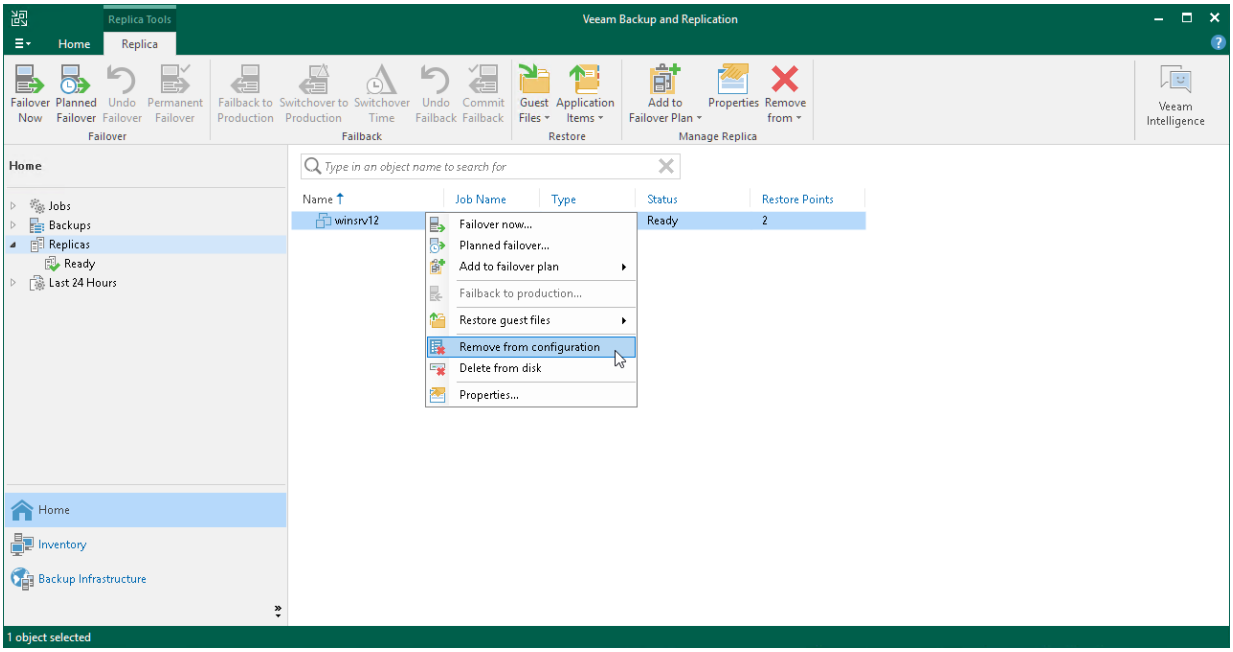

# Deleting Replicas from Disk

When you delete replicas from disks, Veeam Backup & Replication removes the replicas not only from the Veeam Backup & Replication console and configuration database, but also from host storage.

> **NOTE**
>
> Consider the following:
>
> - You can delete records only about replicas that are in the *Ready* state.
> - Do not delete replica files from the destination storage manually, use the **Delete from disk** option instead. If you delete replica files manually, subsequent replication sessions will fail.
> - Unlike the **Remove from configuration** operation, the **Delete from disk** operation does not remove the processed workload from the initial replication job. This means that the replication process will restart for this workload. To avoid this, you can exclude the workload from the replication job or disable the job.

To delete replica files from disks:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.
3. In the working area, select the necessary replica and click **Remove from** > **Disk** on the ribbon. As an alternative, right-click the replica and select **Delete from disk**.

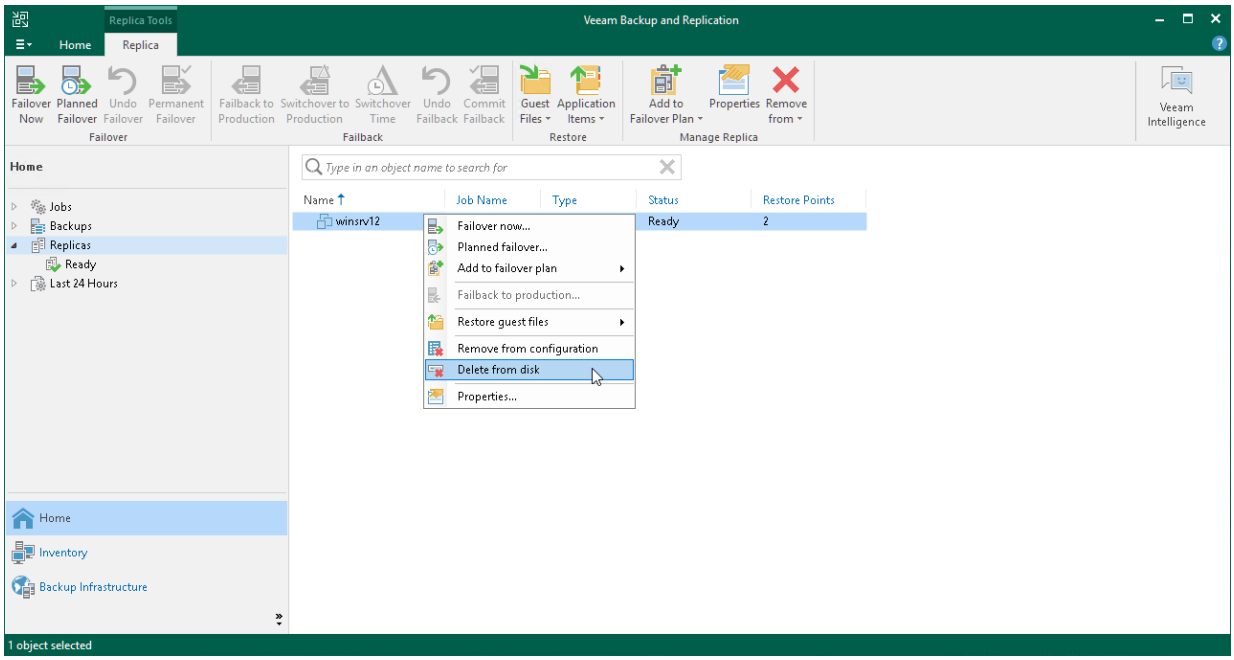

# Failover and Failback for CDP

Failover and failback operations help you ensure that your business will function even if a disaster strikes your production site. Failover is a process of switching from the VM on the source host to its replica on a host in the disaster recovery site. Failback is a process of returning from the replica to the source VM.

Veeam Backup & Replication provides the following failover and failback operations:

- **Perform failover**

  When you perform failover, you shift all processes from the source VM in the production site to the replica in the disaster recovery site. During failover, changes made on the replica are not reflected on the source VM.

  Failover is an intermediate step that needs to be finalized: you can undo failover, perform permanent failover or perform failback.

  For more information on how failover is performed, see Failover.

- **Create failover plan**

  When you create a failover plan, you define the order in which Veeam Backup & Replication must perform failover for VMs, and an interval of time for which Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list.

  For more information on failover plans, see Failover Plans.

- **Perform permanent failover**

  When you perform permanent failover, you permanently switch from the source VM to a replica and use this replica as the source VM.

  For more information on how permanent failover is performed, see Permanent Failover.

- **Perform planned failover**

  When you perform planned failover, you shift all processes from the source VM to its replica. During failover, changes made on the VM replica are not reflected on the source VM.

  Planned failover is helpful when you know that the source VM is about to go offline, for example, you plan to perform datacenter maintenance, and you want to proactively switch the workload to the replica. The procedure is designed to transfer the current workload, that is why it does not suggest to select a restore point.

  For more information on how planned failover is performed, see Planned Failover.

- **Undo failover**

  When you undo failover, you switch back to the source VM and discard all changes made to the replica while it was running.

  You can use the undo failover scenario if you have failed over to the replica for testing and troubleshooting purposes, and you do not need to synchronize the source VM state with the current state of the replica.

  For more information on how failover undo is performed, see Failover Undo.

- **Perform failback**

  When you perform failback, you shift all processes back to the source VM and send to the source VM all changes that took place while the replica was running. During failover, changes made on the source VM are not sent to the replica.

If the source host is not available, you can recover a VM with the same configuration as the source VM and switch to it. For more information on how failback is performed, see Failback.

When you perform failback, changes are only sent to the source/recovered VM but not published. You must test whether the source/recovered VM works with these changes. Depending on the test results, you can do the following:

- **Commit failback**. When you commit failback, you confirm that the source/recovered VM works as expected and you want to get back to it.

  For more information on how failback commit is performed, see Failback Commit.

- **Undo failback**. When you undo failback, you confirm that the source/recovered VM is not working as expected and you want to get back to the replica.

  For more information on how failback undo is performed, see Failback Undo.

Veeam Backup & Replication supports failover and failback operations for one VM and for several VMs. In case one or several hosts fail, you can use batch processing to restore operations with minimum downtime.

The following scheme can help you decide which steps are preferable when you fail over to a replica.

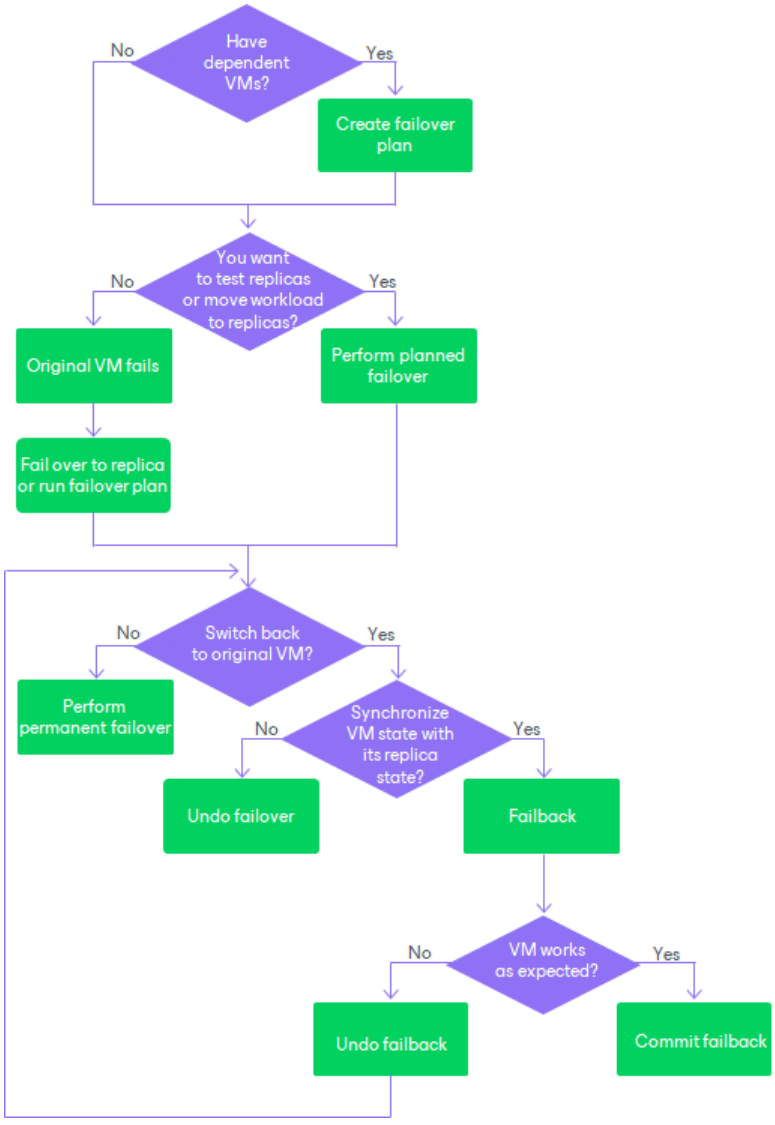

## Related Topics

- Performing Failover
- Failover Plans
- Performing Permanent Failover
- Undoing Failover
- Performing Failback
- Committing Failback
- Undoing Failback

# Failover

Failover is when Veeam Backup & Replication switches processes from the source VM in the production site to its replica in the disaster recovery site. During failover, Veeam Backup & Replication recovers the replica to the required restore point and shifts all I/O processes from the source VM to its replica. As a result, you have a fully functional VM within a couple of seconds, and your users can access services and applications with minimum disruption.

You can fail over to replicas not only when a disaster strikes the production site, but also to test replicas for recoverability. If the source VMs and VM replicas are located in the same network, consider temporary disconnecting the source VMs from the network to avoid IP address or machine name conflicts.

## How Failover Works

The failover operation is performed in the following way:

1. Veeam Backup & Replication puts all replication activities on hold.
2. The state of the replica is changed from *Ready* to *Failover*.
3. Veeam Backup & Replication recovers a replica to the required restore point.
4. Veeam Backup & Replication powers on the replica.

   If you perform failover for testing purposes, and the source VM still exists and is running, the source VM remains powered on.
5. All changes made to the replica disks while the replica is running in the *Failover* state are written to the protective virtual disks (*<disk_name>-interim.vmdk* files).

## Finalizing Failover

Failover is an intermediate step that needs to be finalized. You can perform the following operations:

- Undo failover.
- Perform permanent failover.
- Perform failback.

## Performing Failover

For more information on failover, see Failover and Failover and Failback for CDP.

To perform failover, use the **VMware Failover** wizard.

### Before You Begin

Before you fail over to a replica, check the following prerequisites:

- You can perform failover for VMs that have been successfully replicated at least once.
- Replicas must be in the *Ready* state.
- Host vMotion is not allowed during failover.

# Step 1. Launch VMware Failover Wizard

To launch the **VMware Failover** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from replica** > **Entire replica** > **Failover to a replica**.
- Open the **Home** view. In the inventory pane select the **Replicas** > **Ready** node. In the working area, select the necessary replica and click **Failover Now** on the ribbon.
- Open the **Home** view. In the inventory pane select the **Replicas** > **Ready** node. In the working area, right-click the necessary replica and select **Failover now**.

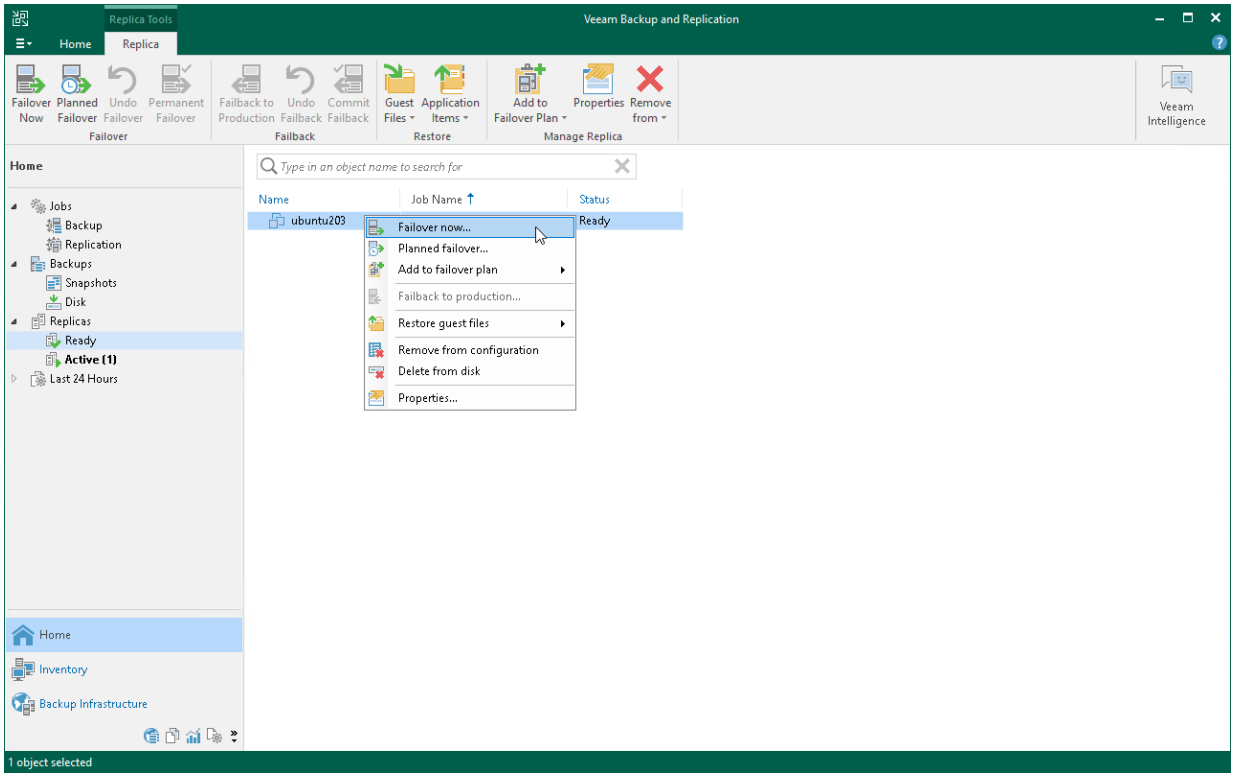

# Step 2. Select VMs

At the **Virtual Machines** step of the wizard, you can modify a list of VMs from which you fail over. To add VMs or VM containers, click **Add** > **From infrastructure** if you want to add VMs from the virtual infrastructure, or **Add** > **From replicas** if you want to add VMs from existing replicas. Then select the necessary VMs or VM containers. If you select VM containers, Veeam Backup & Replication will expand them to a plain VM list.

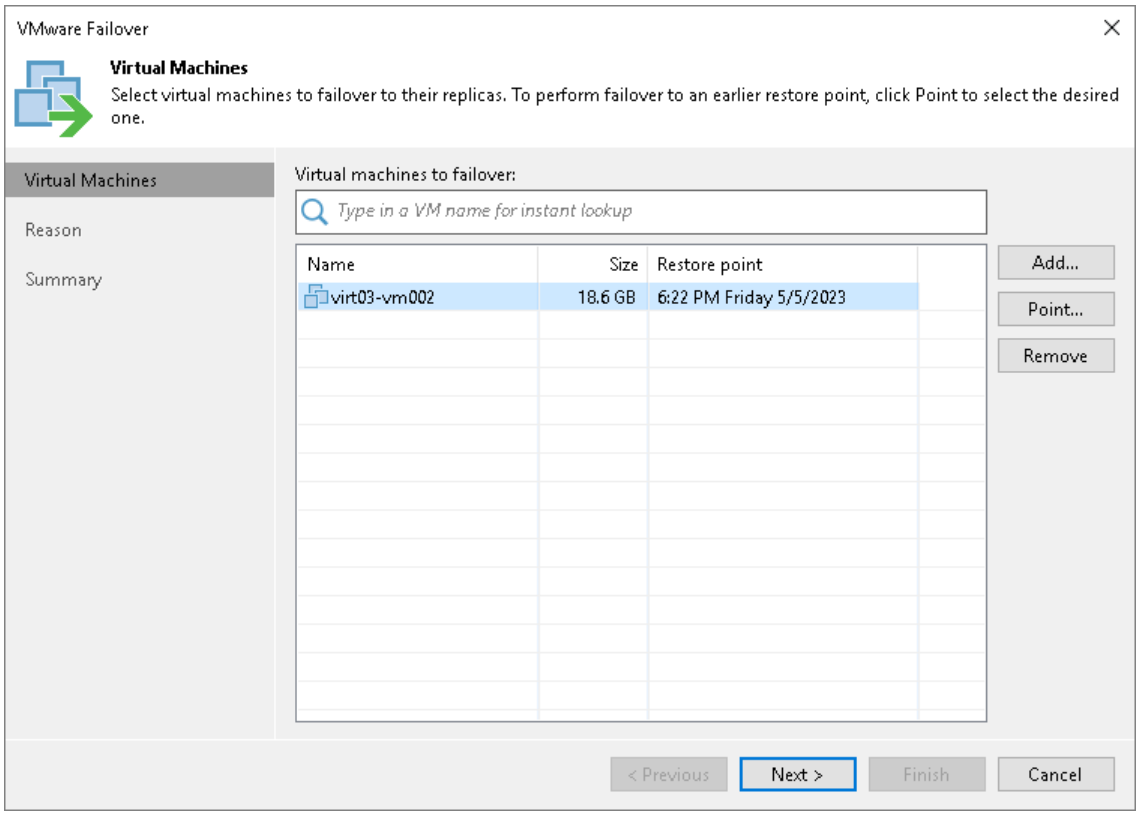

# Step 3. Select Restore Points

At the **Virtual Machines** step of the wizard select to which state of replicas you want to fail over:

1. In the **Virtual machines to failover** list, select the necessary VM and click **Point**.
2. In the **Restore Points** window, select whether you want to fail over to the latest available crash-consistent restore point, to the latest long-term application-consistent restore point or to a specific point in time.

   To restore to a short-term restore point, select a point in the green area. The darker the green, the more I/O load was produced on the source VM. To restore to a crash-consistent or application consistent long-term point, select a violet or turquoise vertical bar with a circle at the top.

   If you fail over to a specific point in time, use the right and left arrows on the keyboard to select the required restore point. To quickly find a long-term restore point, click a link that shows a date. In the opened window, you will see a calendar where you can select the necessary day. In the **Timestamp** section, you will see long-term restore points created during the selected day.

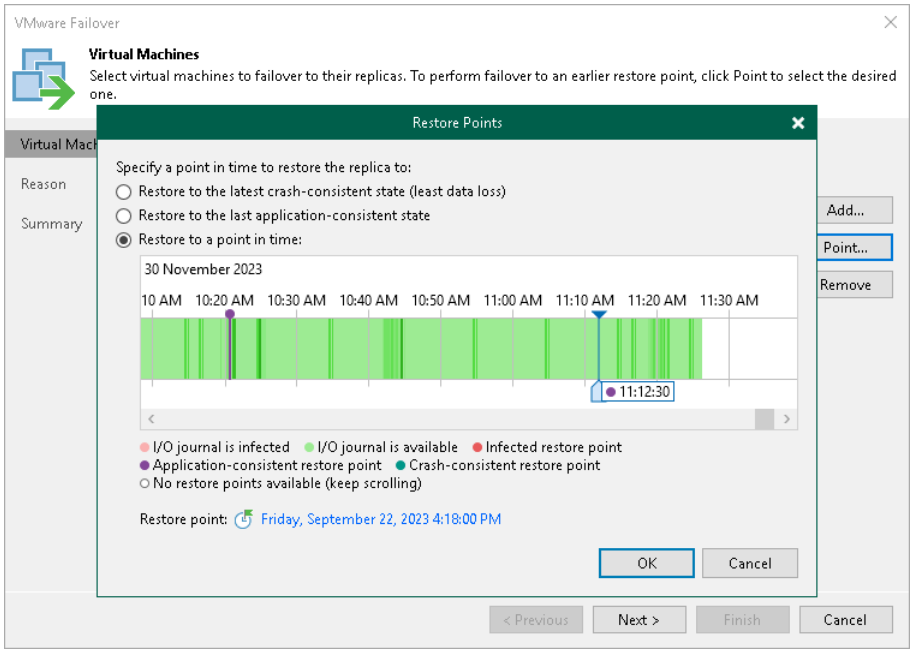

## Step 4. Specify Failover Reason

At the **Reason** step of the wizard, enter a reason for failing over to the replicas. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

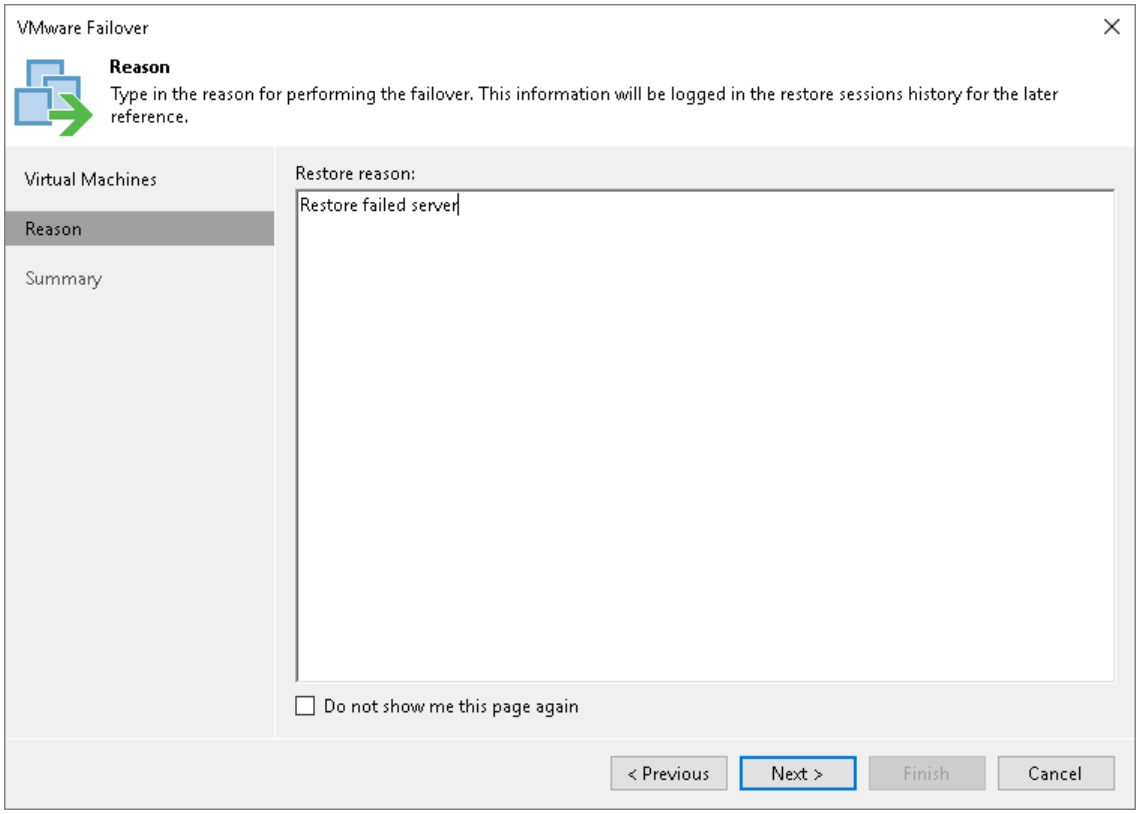

## Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the failover task and click **Finish** to exit the wizard.

During the failover process, Veeam Backup & Replication will request the target host to to power on the replicas.

# What You Do Next

Failover is an intermediate step that needs to be finalized. You can finalize failover in the following ways:

- Perform permanent failover.
- Undo failover.
- Perform failback.

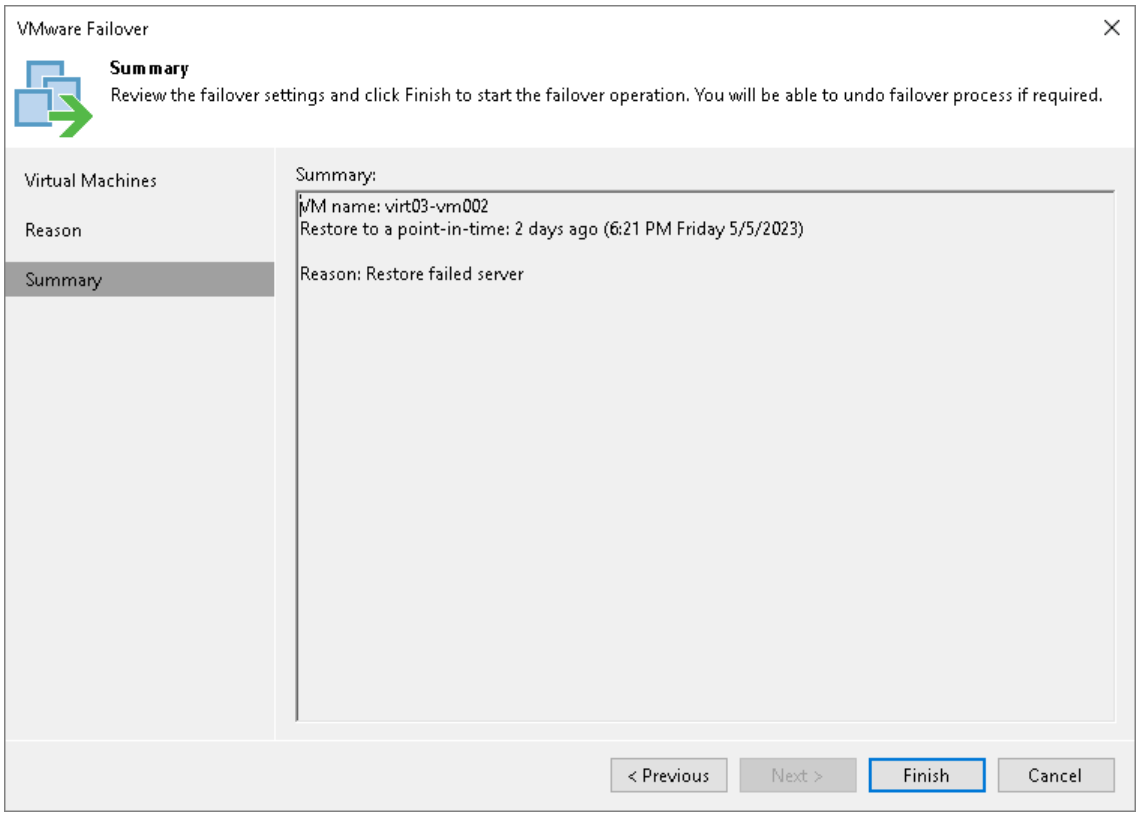

# Failover Plans

A failover plan helps you perform failover for dependent VMs one by one. In the failover plan, you define the order in which VMs must be processed and an interval of time for which Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list. For more information on the failover plan, see Failover Plans.

You can add a VM to an existing failover plan or to a new plan as described in section Creating Failover Plans. For more information on how to manage failover plans, see Running Failover Plans and Undoing Failover by Failover Plans.

# Permanent Failover

Permanent failover is one of the ways to finalize failover. When you perform permanent failover, you permanently switch processes from the source VM to its replica. As a result of permanent failover, the VM replica stops acting as a replica and starts acting as the production VM.

The permanent failover operation is performed in the following way:

1. Veeam Backup & Replication powers off the replica.
2. Veeam Backup & Replication removes short-term and long-term restore points of the replica from the replication chain and deletes associated files from the datastore. Changes that were written to the protective virtual disks (*<disk_name>-interim.vmdk*) are committed to the replica to bring the replica to the most recent state.
3. Veeam Backup & Replication removes the replica from the list of replicas in the Veeam Backup & Replication console.
4. To protect the replica from corruption after permanent failover is complete, Veeam Backup & Replication reconfigures the current CDP policy by adding the source VM to the list of exclusions. Note that other policies and jobs are not modified automatically. When the CDP policy starts, the source VM is skipped from processing. As a result, no data is written to the working VM replica.

## Performing Permanent Failover

For more information on permanent failover, see Failover and Failback for CDP and Permanent Failover.

To perform permanent failover, do one of the following:

- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, select the necessary replica and click **Permanent Failover** on the ribbon.
- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, right-click the necessary replica and select **Permanent failover**.

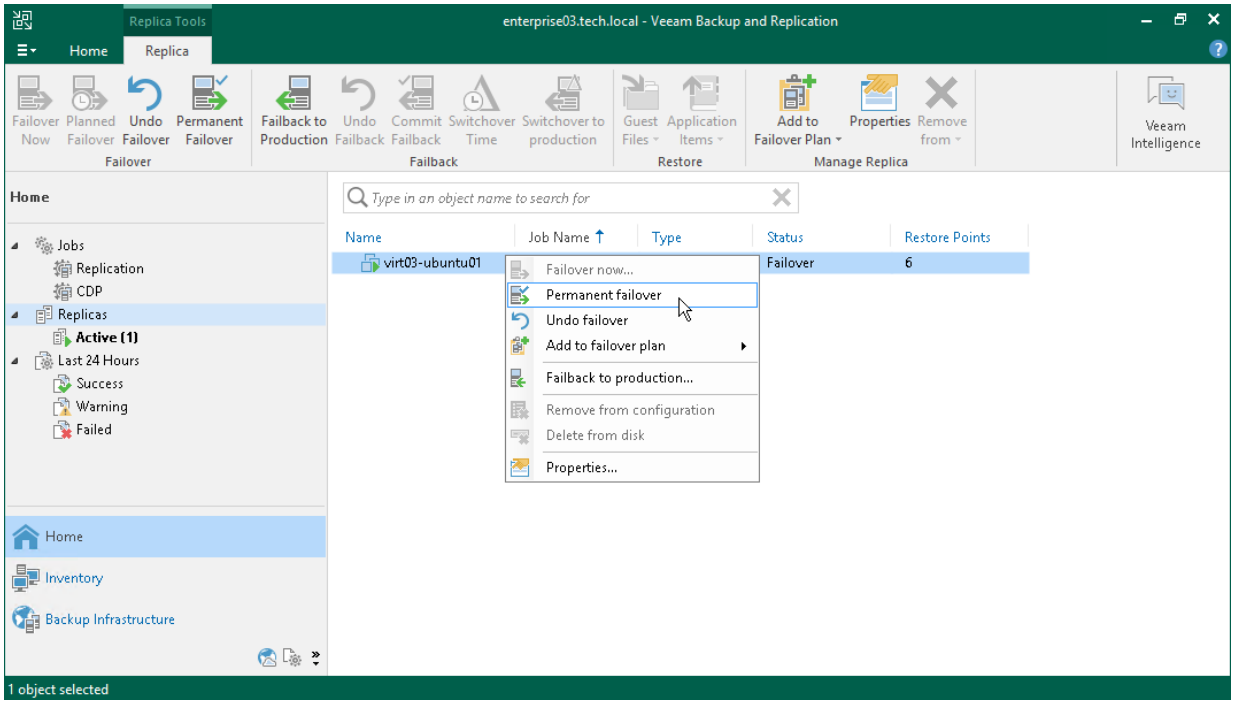

# Planned Failover

Planned failover is a process when you manually launch switching from a source VM to its replica with minimum interruption in operation. Planned failover is helpful when you know that your production VMs are about to go offline and you need to proactively switch the workload from the source VMs to their replicas. You can use the planned failover, for example, if you plan to perform datacenter migration, maintenance or software upgrade of the production VMs. You can also perform planned failover if you have noticed some signs of the approaching disaster.

As the procedure is designed to transfer the current workload to the replica, it does not suggest selecting a restore point to switch.

## How Planned Failover Works

The planned failover is performed in the following way:

1. Veeam Backup & Replication checks that the necessary backup infrastructure components are ready for failover.
2. If unreplicated changes exist, Veeam Backup & Replication copies them to the replica.
3. If unreplicated changes appear during the infrastructure check, Veeam Backup & Replication copies them to the replica.
4. The source VM is powered off.
5. If VMware Tools are installed on the VM, Veeam Backup & Replication tries to shut down the VM guest OS. If nothing happens after 15 minutes, Veeam Backup & Replication powers off the VM. If VMware Tools are not installed on the VM or the VM is suspended, Veeam Backup & Replication powers off the VM.
6. Veeam Backup & Replication copies the portion of last-minute changes to the replica. The replica becomes fully synchronized with the source VM.
7. The failover process triggers the creation of a long-term restore point.
8. The VM is failed over to its replica, to the created long-term restore point.
9. The VM replica is powered on.
10. All changes made to the replica disks while the replica is running in the *Failover* state are written to the protective virtual disks (*<disk_name>-interim.vmdk* files).

During the planned failover, Veeam Backup & Replication creates a helper restore point. This point stays in the replication chain and is deleted according to the configured retention settings. You can see this restore point in the list of restore points for the VM. You can use the restore point later to roll back to the necessary replica state.

## Finalizing Planned Failover

Failover is an intermediate step that needs to be finalized. You can perform the following operations:

- Undo failover.
- Perform permanent failover.
- Perform failback.

# Performing Planned Failover

For more information on planned failover, see Planned Failover and Failover and Failback for CDP.

To perform planned failover, use the **VMware Planned Failover** wizard.

## Before You Begin

Before you configure planned failover, check the following prerequisites:

- You can perform failover for VMs that have been successfully replicated at least once.
- Replicas must be in the *Ready* state.
- The CDP policy that created the replicas must be enabled.

# Step 1. Launch VMware Planned Failover Wizard

To launch the **VMware Planned Failover** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware vSphere** > **Restore from replica** > **Entire replica** > **Planned failover to a replica**.
- Open the **Home** view, expand the **Replicas** node. In the working area, select one or more VMs and click **Planned Failover** on the ribbon. You can also right-click one of the selected VMs and click **Planned Failover**.
- Open the **Inventory** view. In the working area, select one or more VMs and right-click one of the selected VMs and click **Restore** > **Planned Failover**.

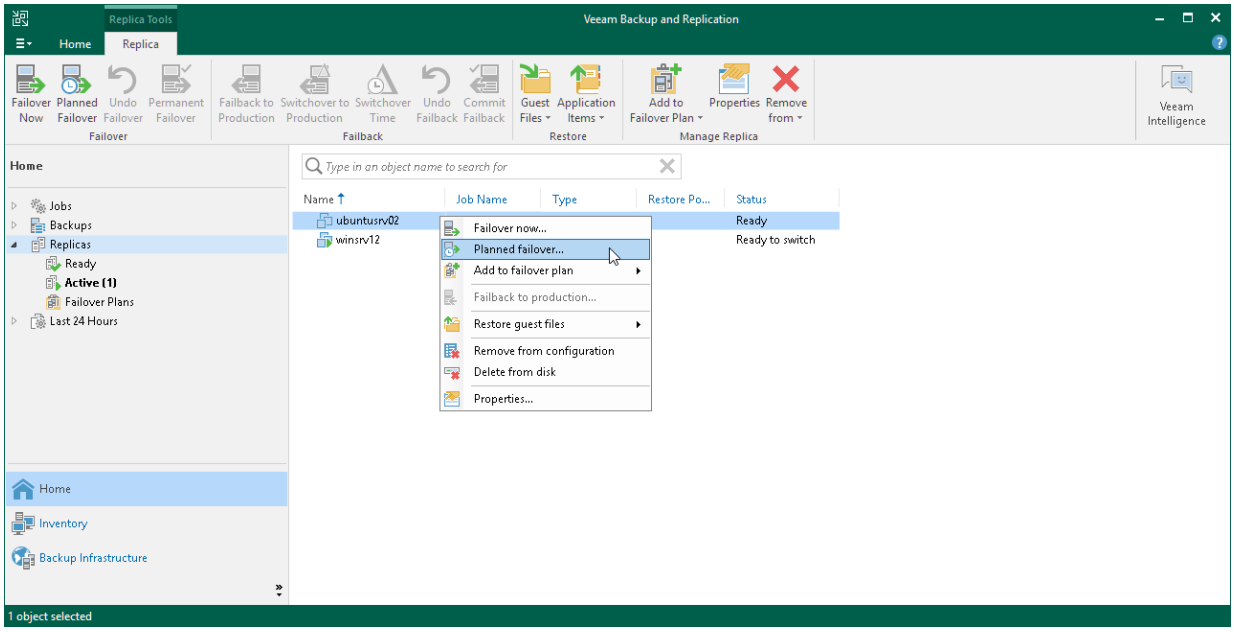

# Step 2. Select VMs

At the **Virtual Machines** step of the wizard, you can modify a list of VMs from which you fail over. To add VMs or VM containers, click **Add** > **From infrastructure** if you want to add VMs from the virtual infrastructure, or **Add** > **From replicas** if you want to add VMs from existing replicas. Then select the necessary VMs or VM containers. If you select VM containers, Veeam Backup & Replication will expand them to a plain VM list.

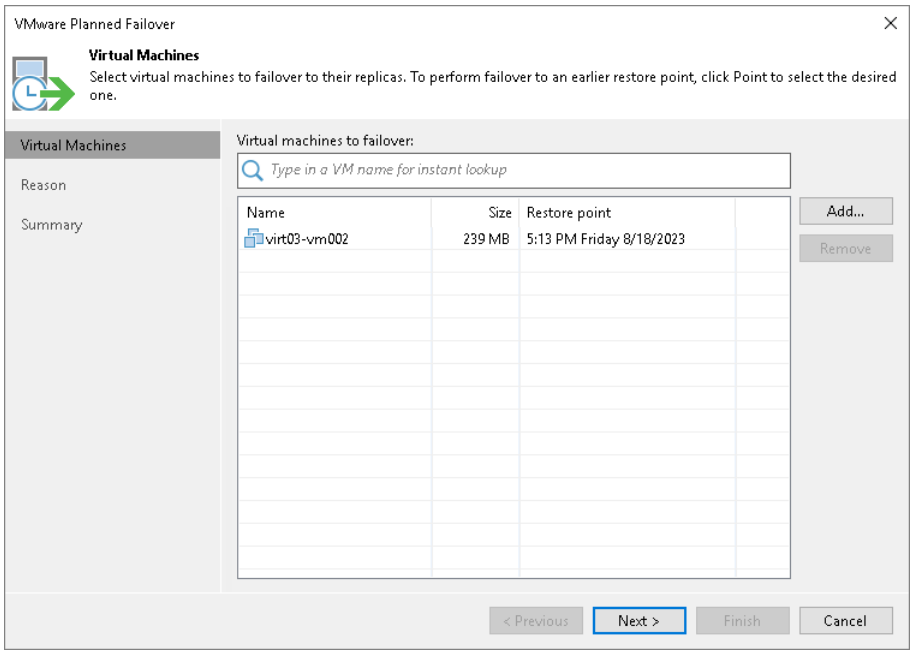

# Step 3. Specify Failover Reason

At the **Reason** step of the wizard, enter a reason for failing over to the replicas. The information you provide will be saved in the session history and you can reference it later.

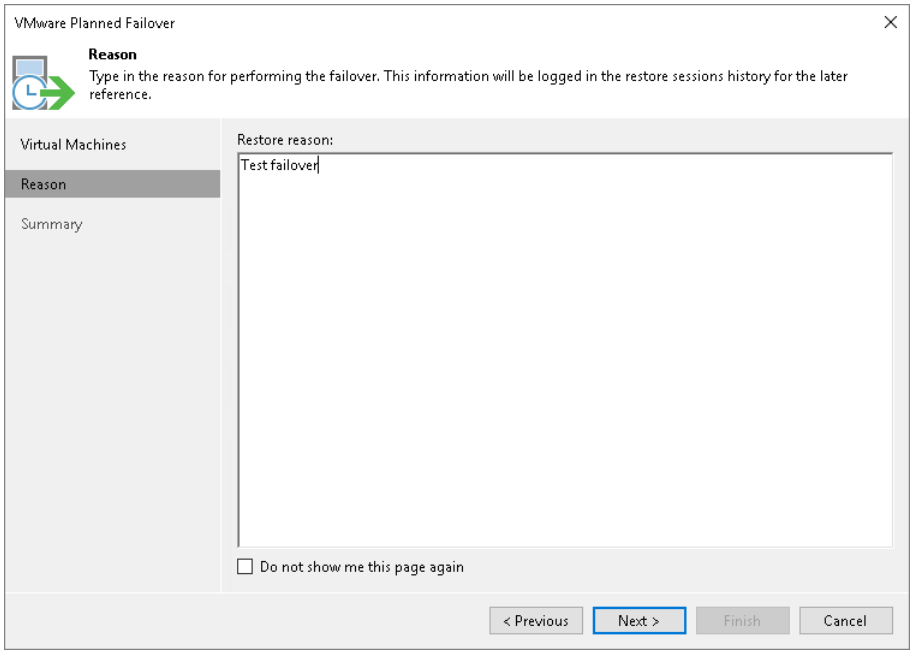

## Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the failover task and click **Finish** to exit the wizard.

During the failover process, Veeam Backup & Replication will request the target host to to power on the replicas.

# What You Do Next

Failover is an intermediate step that needs to be finalized. You can finalize failover in the following ways:

- Perform permanent failover.
- Undo failover.
- Perform failback.

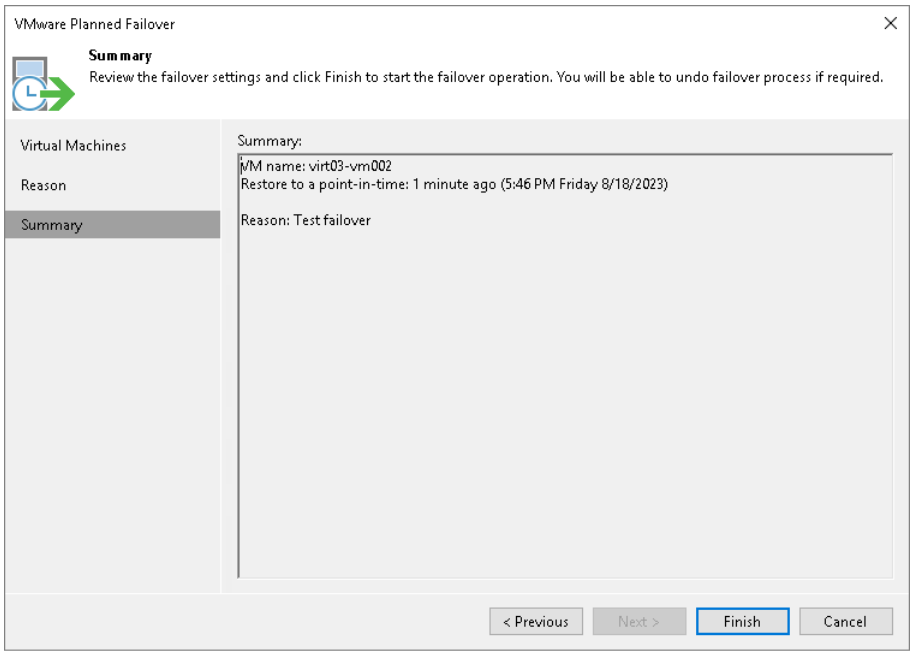

# Failover Undo

Failover undo is one of the ways to finalize failover. When you undo failover, you switch back from a replica to the source VM. Veeam Backup & Replication discards all changes made to the replica while it was in the *Failover* state.

The failover undo operation is performed in the following way:

1. Veeam Backup & Replication reverts the replica to its pre-failover state. To do this, Veeam Backup & Replication powers off the replica and gets it back to the latest restore point in the replication chain. Changes that were written to the protective virtual disks (*<disk_name>-interim.vmdk*) while the replica was in the *Failover* state are discarded.
2. The state of the replica gets back to *Ready*, and Veeam Backup & Replication resumes replication activities for the source VM on the source host.

## Undoing Failover

For more information on failover undo, see Failover and Failback for CDP and Failover Undo.

To undo failover:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, select the necessary replica and click **Undo Failover** on the ribbon. Alternatively, right-click the necessary replica and select **Undo Failover**.
4. In the displayed window, click **Yes** to confirm the operation.

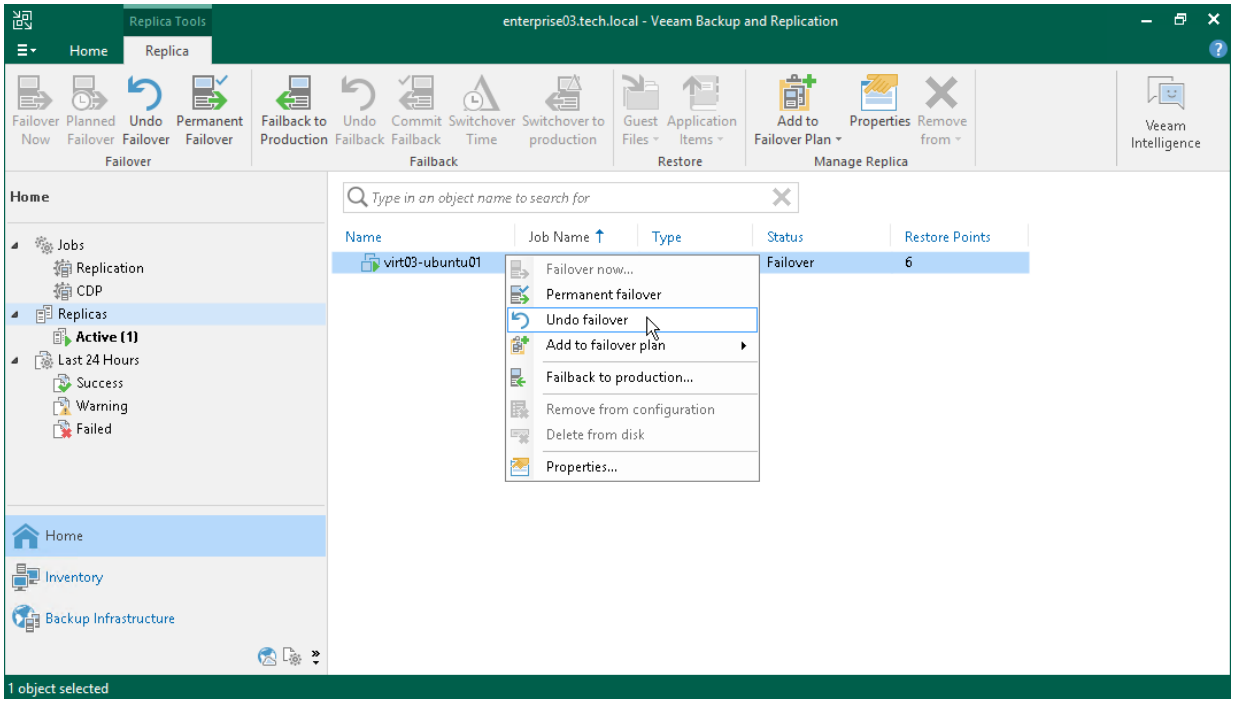

# Failback

Failback is one of the ways to finalize failover. When you perform failback, you switch back to the production VM from a replica, shift I/O processes from the disaster recovery site to the production site.

Veeam Backup & Replication provides you the following options to perform failback:

- You can fail back to the source VM in the original location.
- You can fail back to a VM already recovered to a new location. This VM must be recovered before you perform failback. For example, you can recover the VM from a backup.
- You can fail back to a VM recovered from a replica to a new location, or to any location but with different settings. The VM will be recovered from the replica during the failback process.

The first two options help you decrease recovery time and the use of the network traffic because Veeam Backup & Replication needs to transfer only differences between the source/recovered VM and replica. For the third option, Veeam Backup & Replication needs to transfer the whole VM data, including its configuration and virtual disk content. Use the third option if there is no way to use the source VM or restore it from a backup.

Veeam Backup & Replication performs failback in two phases:

- **First phase**: Veeam Backup & Replication synchronizes the state of the production VM (the source VM, an already recovered VM or a VM that will be recovered from the replica) with the current state of its replica. This phase may take a lot of time especially if VM is large. While Veeam Backup & Replication performs the first phase of failback, replicas are still up and running, users can access these VMs and perform daily routine tasks as normal.
- **Second phase**: Veeam Backup & Replication switches all processes from the replica to the production VM, turns off the replica and also sends to the production VM changes made to the replica since the end of the first phase.

  The time when the second phase starts depends on how you want to switch from the replica to the production VM. You can switch to the production VM automatically, at the scheduled time or manually. If you select to switch automatically, the second phase will start right after the first phase finishes. If you select to switch at the scheduled time or manually, the second phase will start at the time you want.

The process of failing back to the source VM or an already recovered VM differs from the process of failing back to a VM recovered from a replica:

- How failback to the source VM and already recovered VM works.
- How failback to a VM recovered from a replica works.

## How Failback to Source VM or Already Recovered VM Works

When you fail back to the source VM or an already recovered VM, Veeam Backup & Replication performs the following operations during the first phase:

1. If the source VM is running, Veeam Backup & Replication powers it off.

2. Veeam Backup & Replication calculates the difference between disks of the production VM and disks of the replica in the *Failover* state. Difference calculation helps Veeam Backup & Replication understand what data needs to be transferred to the production VM to synchronize its state with the state of the replica.

   [For VMware vSphere prior version 7.0] If you fail back to the original VM in the original location and you have enabled the **Quick rollback** option, difference calculation can be performed much faster than without this option enabled. For more information on quick rollback, see Quick Rollback.

3. Veeam Backup & Replication transfers the data that was detected at the previous step to the production VM. The transferred data is written to the production VM.

4. Veeam Backup & Replication changes the state of the replica from *Failover* to *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the following operations:

1. The guest OS of the replica is shut down or the replica is powered off.

   If VMware Tools are installed on the replica, Veeam Backup & Replication tries to shut down the replica guest OS. If nothing happens after 15 minutes, Veeam Backup & Replication powers off the replica. If VMware Tools are not installed on the VM or the VM is suspended, Veeam Backup & Replication powers off the VM. The replica remains powered off until you commit failback or undo failback.

2. Veeam Backup & Replication calculates the difference between disks of the production VM and disks of the replica. Difference calculation helps Veeam Backup & Replication understand what data was changed while the replica was in the *Ready to switch* state.

3. Sends data changed on the replica while it was in the *Ready to switch* state to the production VM.

4. The state of the replica is changed from *Ready to switch* to *Failback*.

5. [If you fail back to a recovered VM] Veeam Backup & Replication updates the ID of the source VM in the Veeam Backup & Replication configuration database. The ID of the source VM is replaced with the ID of the recovered VM.

6. If you have selected to power on the production VM after failback, Veeam Backup & Replication powers on the production VM on the host.

# How Failback to VM Recovered from Replica Works

When you fail back to a VM recovered from a replica, Veeam Backup & Replication performs the following operations during the first phase:

1. Veeam Backup & Replication requests vCenter Server to create on the target host an empty VM with the same configuration as the replica. vCenter Server registers the created production VM.

2. Veeam Backup & Replication transfers data of the replica to the production VM to update the production VM state to the replica state.

3. Veeam Backup & Replication changes the state of the replica from *Failover* to the *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the same operations as described in section How Failback to Source VM or Already Recovered VM Works except for the step 2. When you fail back to a VM recovered from a replica, Veeam Backup & Replication does not calculate the difference between disks.

# Quick Rollback

Quick rollback helps you significantly reduce the failback time. You can use quick rollback if you fail back from a replica to the source VM in the source location.

During failback, Veeam Backup & Replication calculates differences between disks of the source VM and disks of the replica. With the quick rollback option enabled, Veeam Backup & Replication compares only those disk sectors that have changed during the replica was in the *Failover* state instead of comparing entire disks. To get information about the changed disk sectors, Veeam Backup & Replication uses VMware vSphere Changed Block Tracking (CBT).

As a result of enabling quick rollback, difference calculation becomes much faster. After the differences are calculated, Veeam Backup & Replication performs failback in a regular way: transport changed blocks to the source VM, powers off the replica and synchronizes the source VM with the replica once again.

# Requirements for Quick Rollback

To perform quick rollback, make sure that the following requirements are met:

- You fail back to the source VM in the original location.
- Do not use quick rollback if the problem occurred at the VM hardware level, storage level or due to a power loss.

  Use quick rollback if you fail back to the source VM that had a problem at the guest OS level — for example, there was an application error or a user accidentally deleted a file on the source VM guest OS.
- CBT must be enabled for the source VM.

# Limitations for Quick Rollback

During the first replication job session after failback with quick rollback, CBT on the original VM is reset. Due to that Veeam Backup & Replication will read data of the entire VM.

# Performing Failback

For more information on failback, see Failover and Failback for CDP and Failback.

To perform failback, use the **Failback** wizard.

## Before You Begin

Before you perform failback, check the following prerequisites:

- VMs for which you plan to perform failback must be successfully replicated at least once.
- A replica from which you plan to fail back and a target VM to which you fail back must not have snapshots at the moment failback runs.
- A replicas from which you want to fail back must be in the *Failover* state.

# Step 1. Launch Failback Wizard

To launch the **Failback** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from replica** > **Entire replica** > **Failback to production**.
- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, right-click the necessary replica and select **Failback to production**. Alternatively, click **Failback to Production** on the ribbon.

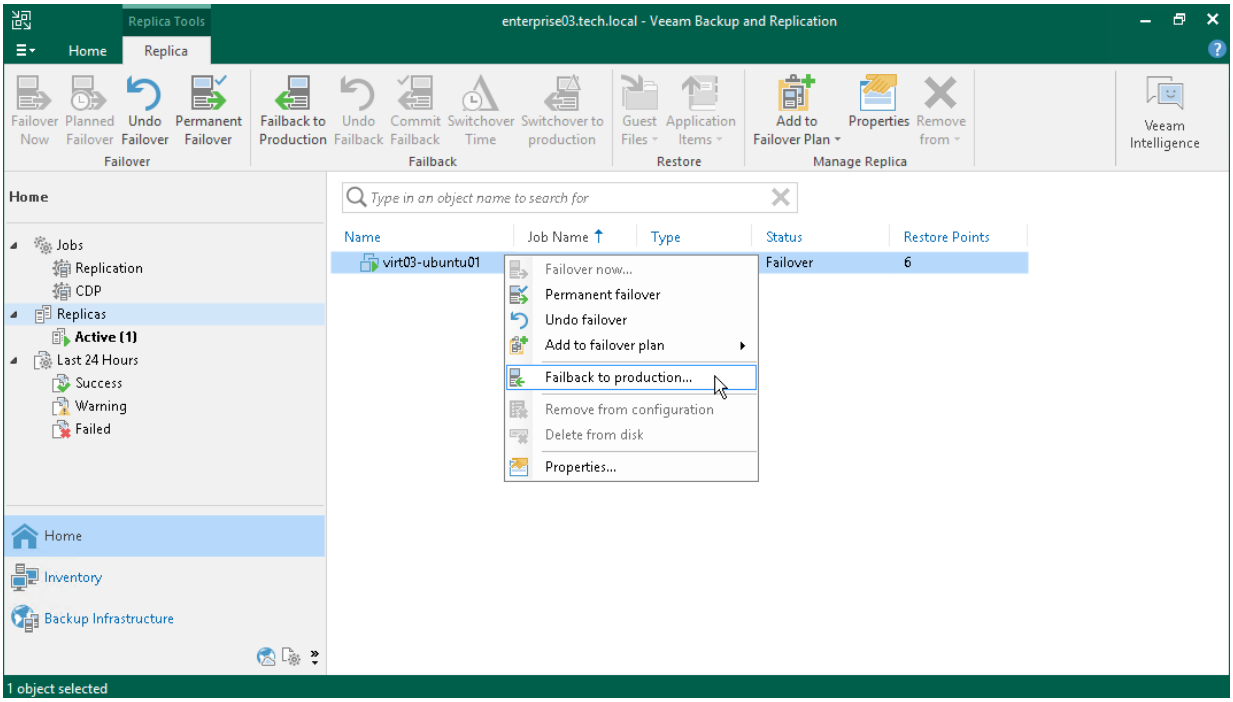

## Step 2. Select Replicas

At the **Replica** step of the wizard, select replicas from which you want to fail back.

To update the list of replicas that are ready for failback (replicas in the *Failover* state), click **Populate**.

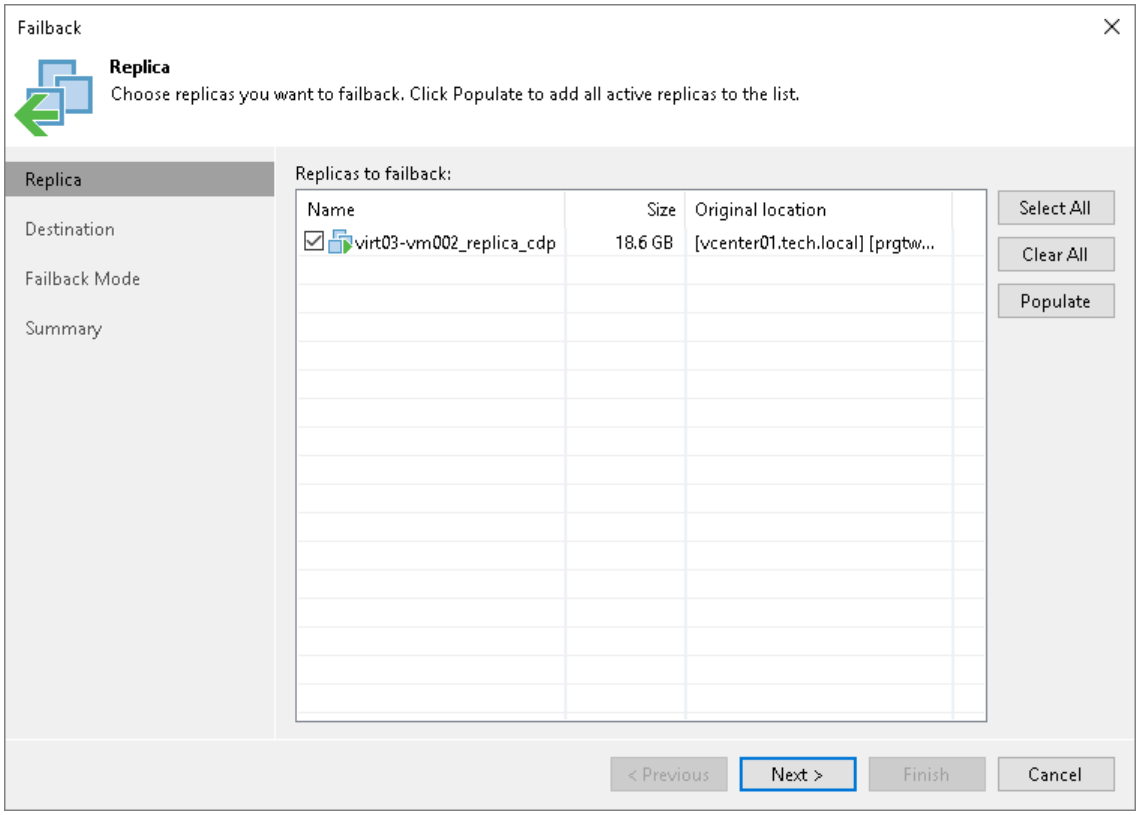

## Step 3. Select Destination

At the **Destination** step of the wizard, select a failback destination and backup proxies for data transfer during failback:

1. Select a destination for failback. Veeam Backup & Replication supports the following options:
   - **Failback to the original VM** – select this option if you want to fail back to the source VMs that reside on the source hosts. Veeam Backup & Replication will synchronize the state of the source VMs with the current state of their replicas to apply any changes that occurred to the replicas while running in the DR site.

     If you select this option, you will proceed to the Summary step of the wizard.
   - **Failback to the original VM restored in a different location** – select this option if the source VMs have already been recovered to a new location, and you want to switch to the recovered VMs from their replicas. Veeam Backup & Replication will synchronize the state of the recovered VMs with the current state of the replicas to apply any changes that occurred to the replicas while running in the DR site.

     If you select this option, you will proceed to the Target VM step of the wizard.
   - **Failback to the specified location** – select this option if you want to recover VMs from replicas. You can recover VMs to a new location, or to any location but with different settings (such as network settings, virtual disk type, configuration file path and so on). Select this option if there is no way to fail back to the source VM or an already recovered VM.

     If you select this option, the wizard will include additional steps.

   If you select one of the first two options, Veeam Backup & Replication will send to the source/recovered VMs only differences between the existing virtual disks. Veeam Backup & Replication will not send replica configuration changes such as different IP address or network settings (if replica Re-IP and network mapping were applied), new hardware or virtual disks added while the replicas were in the *Failover* state.

   If you select **Failback to the specified location**, Veeam Backup & Replication will send to the specified location whole VM data, including VM configurations and virtual disk content.
2. To select which backup proxies will be used for data transfer, click **Pick backup proxies for data transfer**. By default, Veeam Backup & Replication selects proxies automatically.

   If VMs and their replicas reside in different sites, select at least one backup proxy in the production site and one proxy in the disaster recovery site. If VMs and replicas reside in the same site, you can use the same backup proxy as the source and target one.

   We recommend you to select at least two backup proxies in each site to ensure that failback will be performed in case one proxy fails or looses the network connection.

3. [For failback to the original VMs] If you want to fasten failback, and the source VMs had problems at the guest OS level, select the **Quick rollback** check box. For more information on quick rollback, its requirements and limitations, see Quick Rollback.

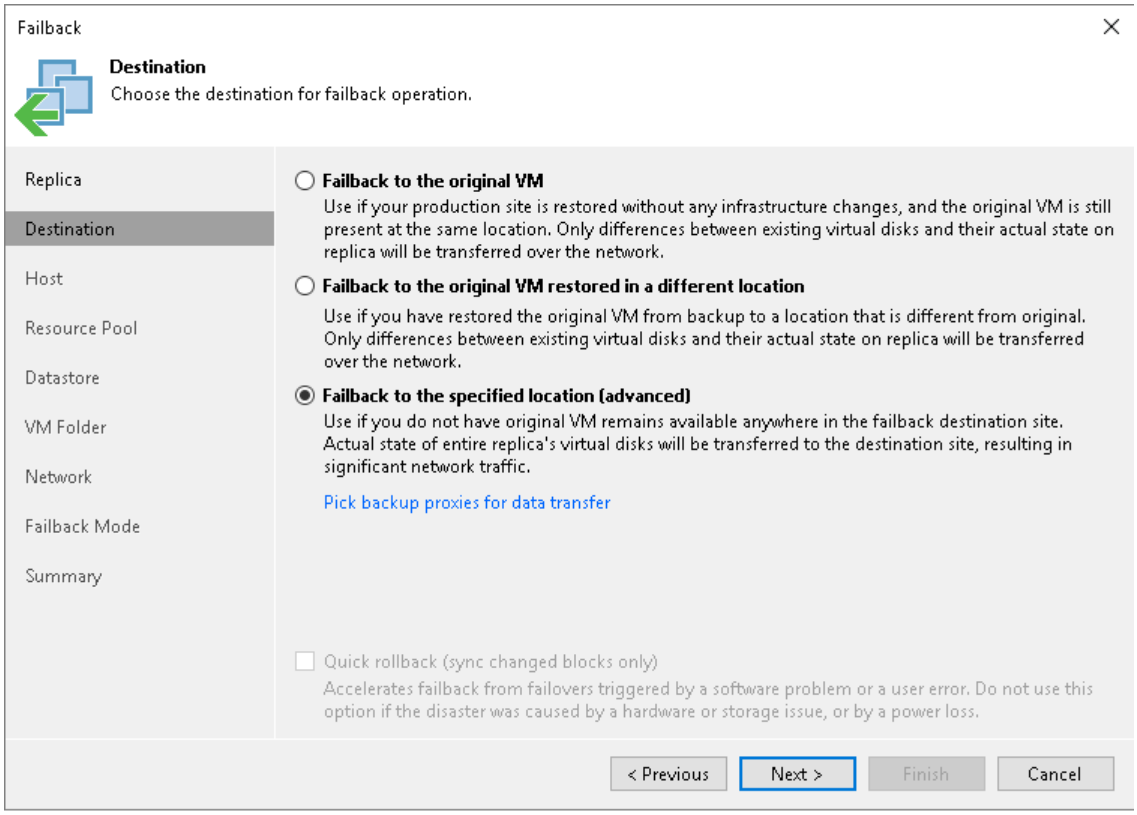

# Step 4. Select Hosts or Clusters

The **Host** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Host** step of the wizard, specify names for the recovered VMs and destination where the recovered VMs will be registered. You can select hosts or clusters as the destination. To specify these options, select one or multiple VMs and use **Host** or **Name** button.

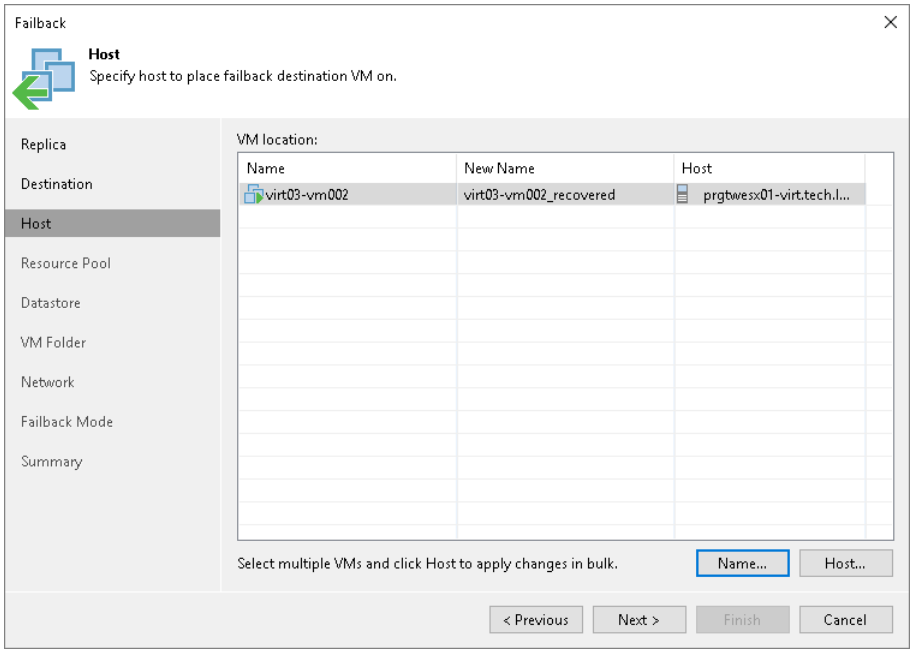

# Step 5. Select Resource Pools

The **Resource Pool** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Resource Pool** step of the wizard, select resource pools to which the recovered VMs will be added. To do this, select VMs that you want to add to the same resource pool, click **Pool** and select the necessary resource pool in the **Select Resource Pool** window.

As an alternative, you can select a vApp to which the restored VM will be included. To find the necessary vApp, at the left bottom corner of the **Select Resource Pool** window, click the resource pool icon ( ) and select *VirtualApp*.

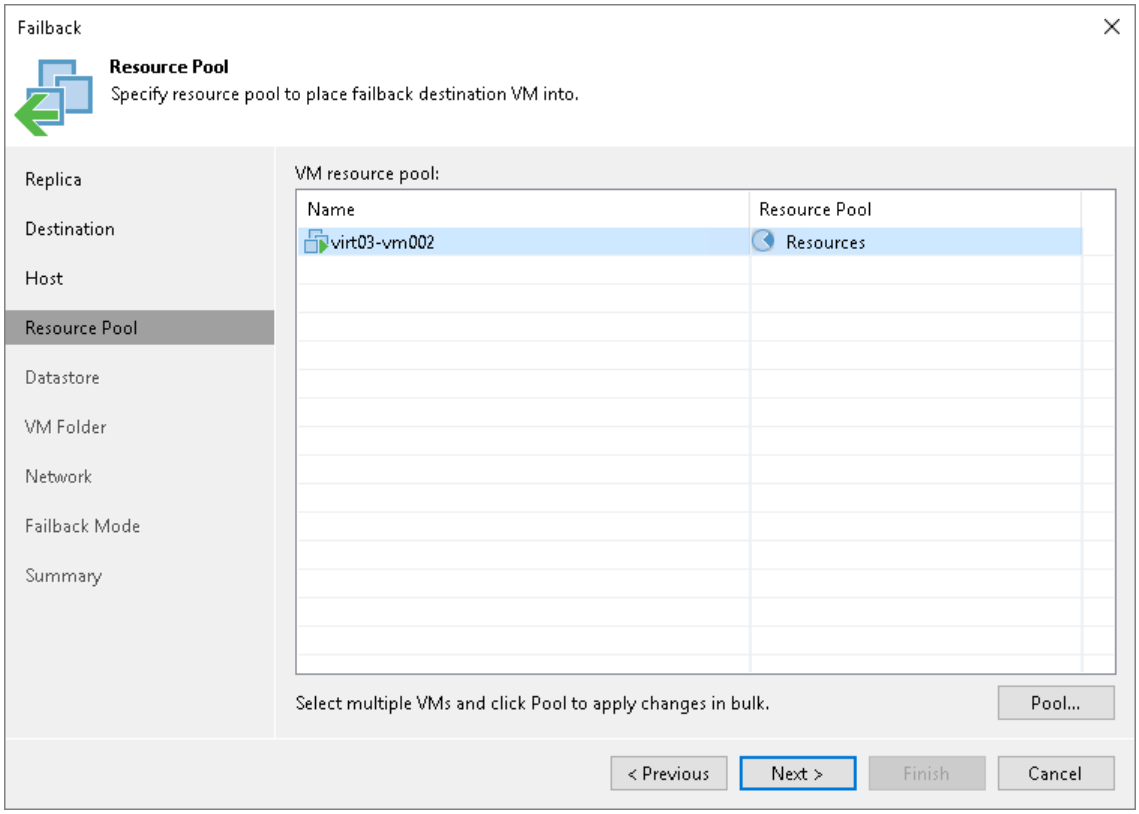

# Step 6. Select Datastores

The **Datastore** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Datastore** step of the wizard, specify datastores where you want to store configuration files and disk files of VMs that will be recovered. Also, you can change disk types.

1. To change a datastore where VM files will be stored, select the necessary VMs and click **Datastore**. From the list of available datastores, select the necessary datastore.

   If configuration and disk files of VMs must be placed to different datastores, select files of the necessary type, click **Datastore** and select the necessary datastore.

2. To change disk type settings, select the necessary disk files and click **Disk Type**. In the **Disk Type Settings** window, select the necessary disk type.

   By default, Veeam Backup & Replication preserves disk types of the source VMs.

   > **NOTE**
   >
   > You can change disk types only for VMs with Virtual Hardware version 7 or later.

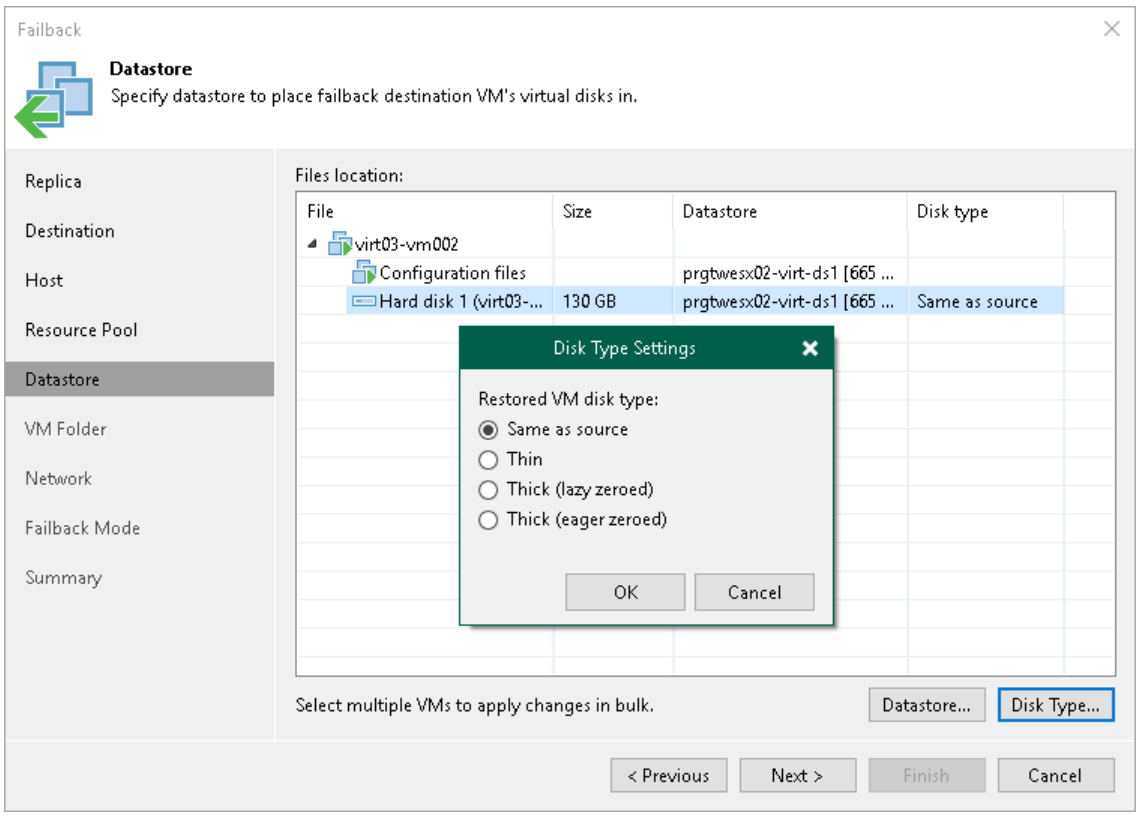

# Step 7. Select Folders

The **VM Folder** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **VM Folder** step of the wizard, specify folders in the target datastores where all files of the recovered VMs will be stored.

If you want the recovered VMs to have the same tags as the source VMs, select the **Restore VM Tags** check box.

> **NOTE**
>
> Consider the following:
>
> - You can select destination folders only if you recover VMs to destinations other than standalone hosts.
> - You can recover VM tags only if you recover VMs to their original locations, and the source VM tags are still available on the source vCenter Server.

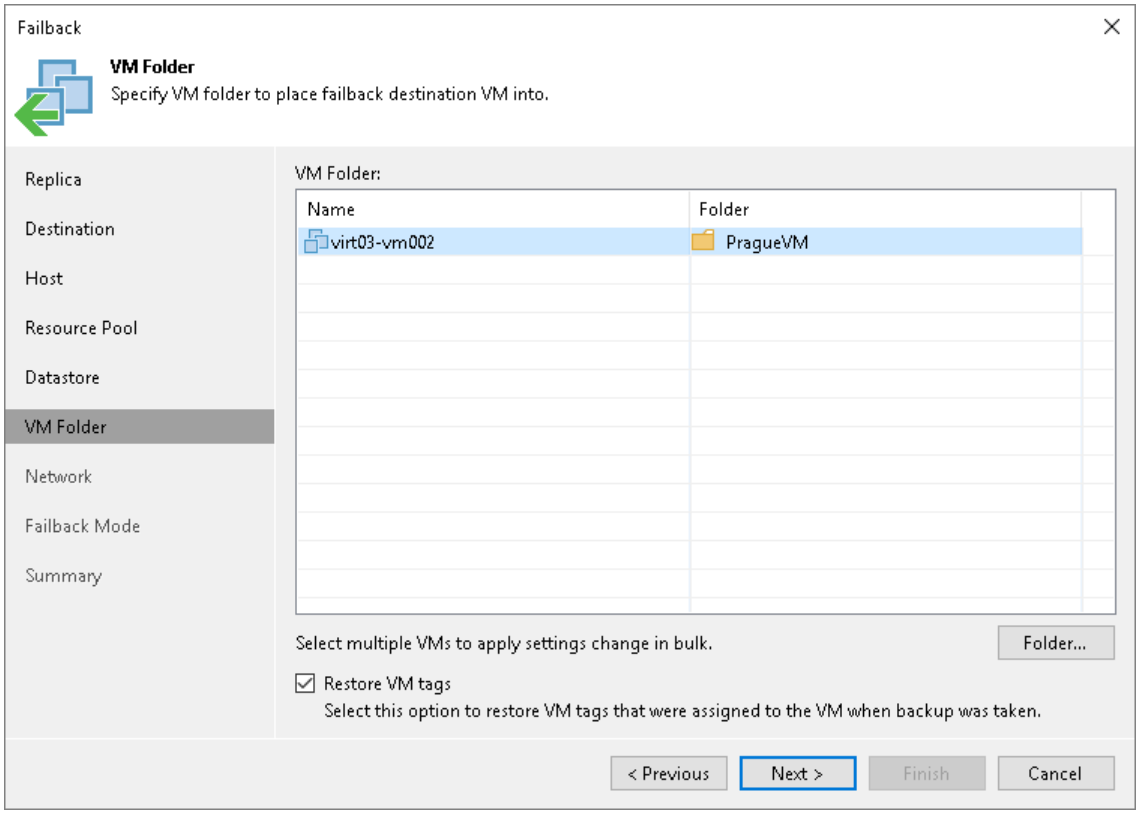

# Step 8. Configure Network Mapping

The **Network** step is available if you have selected the **Failback to the specified location** option at the Destination step. This step applies if you fail back to VMs recovered to new locations, and if networks in those locations differ from networks in the disaster recovery (DR) site.

At the **Network** step of the wizard create a network mapping table. This table maps networks in the DR site to networks in the site where the recovered VMs will reside. Veeam Backup & Replication will use the network mapping table to update configuration files of VMs on the fly, during the failback process.

To change networks to which the restored VMs will be connected:

1. In the **Network connections** list, select the necessary VMs and click **Network**.

   If VMs are connected to multiple networks, select networks which you want to map.

2. In the list of available networks, select a network to which the recovered VMs will be connected.

If you do not want to connect the recovered VMs to any virtual network, select the necessary VMs and click **Disconnect**.

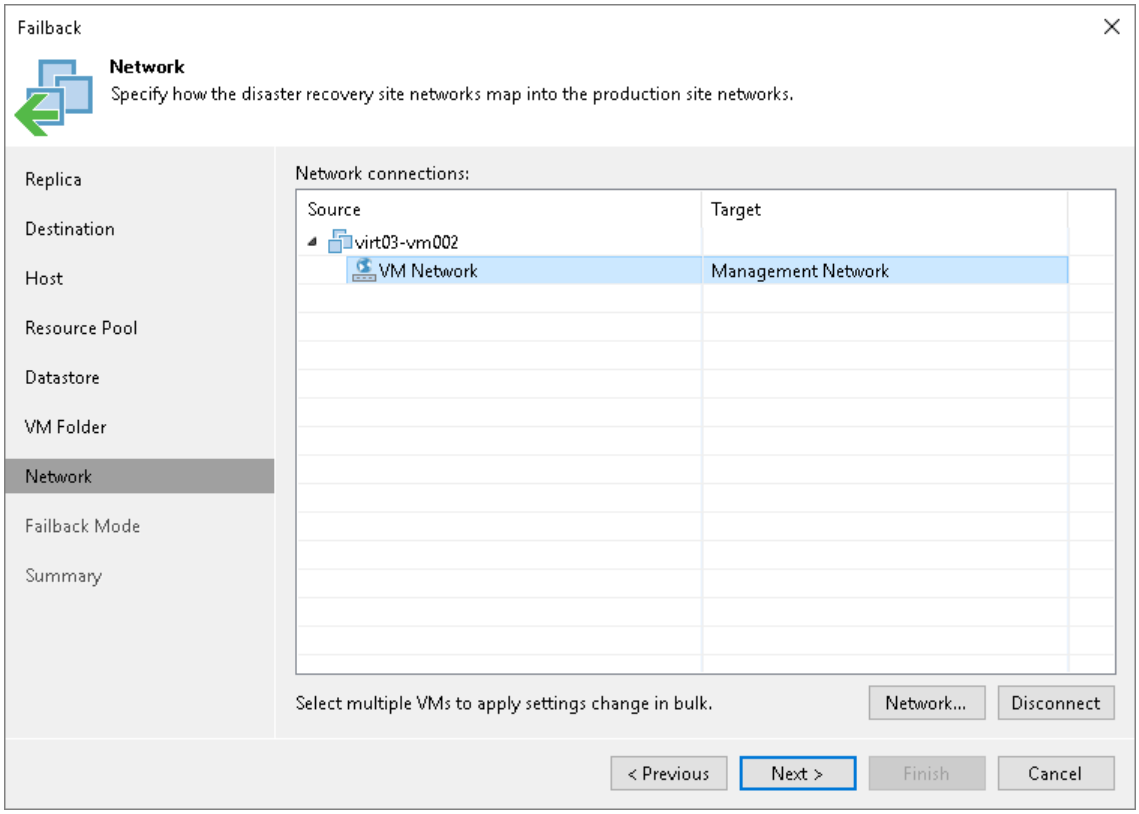

# Step 9. Map Replicas to Restored VMs

The **Target VM** step is available if you have selected the **Failback to the original VM restored in a different location** option at the Destination step.

At the **Target VM** step of the wizard, specify to which VMs you want to fail back from replicas. These VMs must be already restored from backups in the required location.

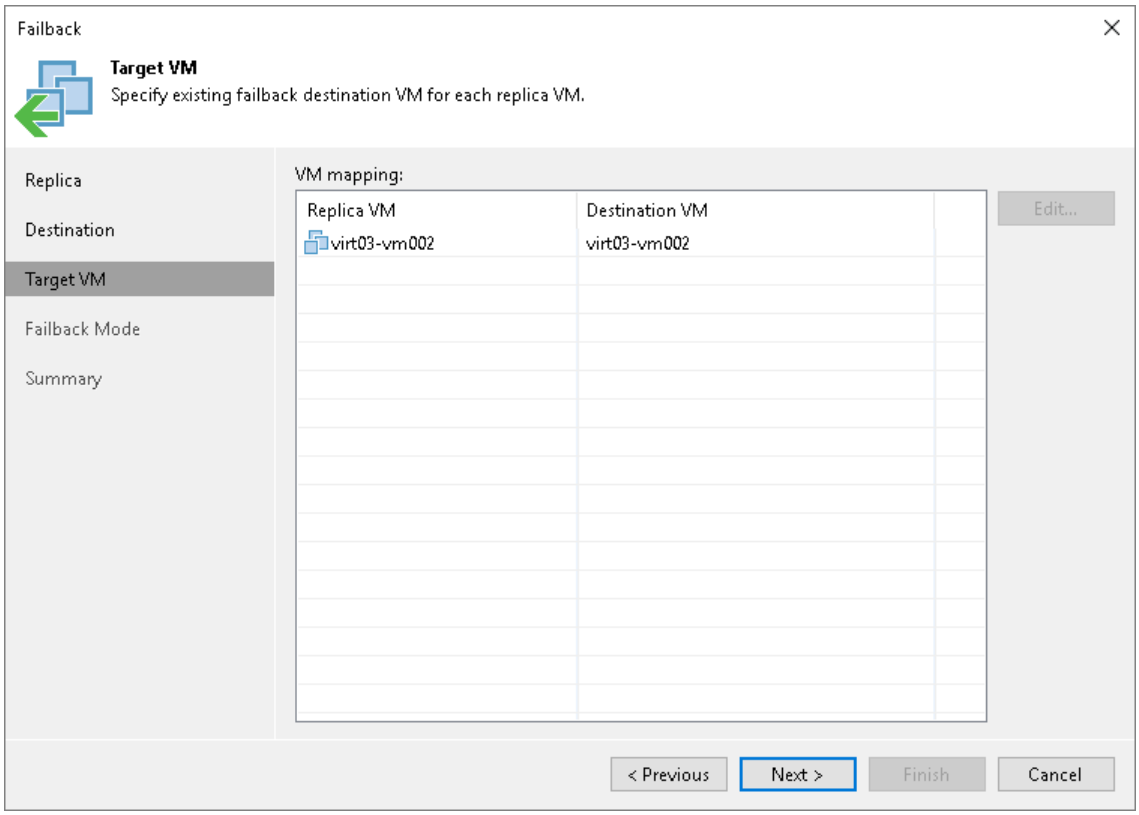

# Step 10. Schedule Switch to Production VMs

At the **Failback Mode** step of the wizard, specify when switch from replicas to production VMs must be performed:

- Select **Auto** if you want Veeam Backup & Replication to perform the switch automatically right after the state of the production VMs is synchronized with the state of their replicas.
- Select **Scheduled** if you want Veeam Backup & Replication to perform the switch at a specific time.
- Select **Manual** if you want to perform the switch manually.

If you select the **Scheduled** or **Manual** option, you can further reset/set the scheduled time or switch to the production VM manually. For more information, see Changing Switching Time and Switching to Production VMs Manually.

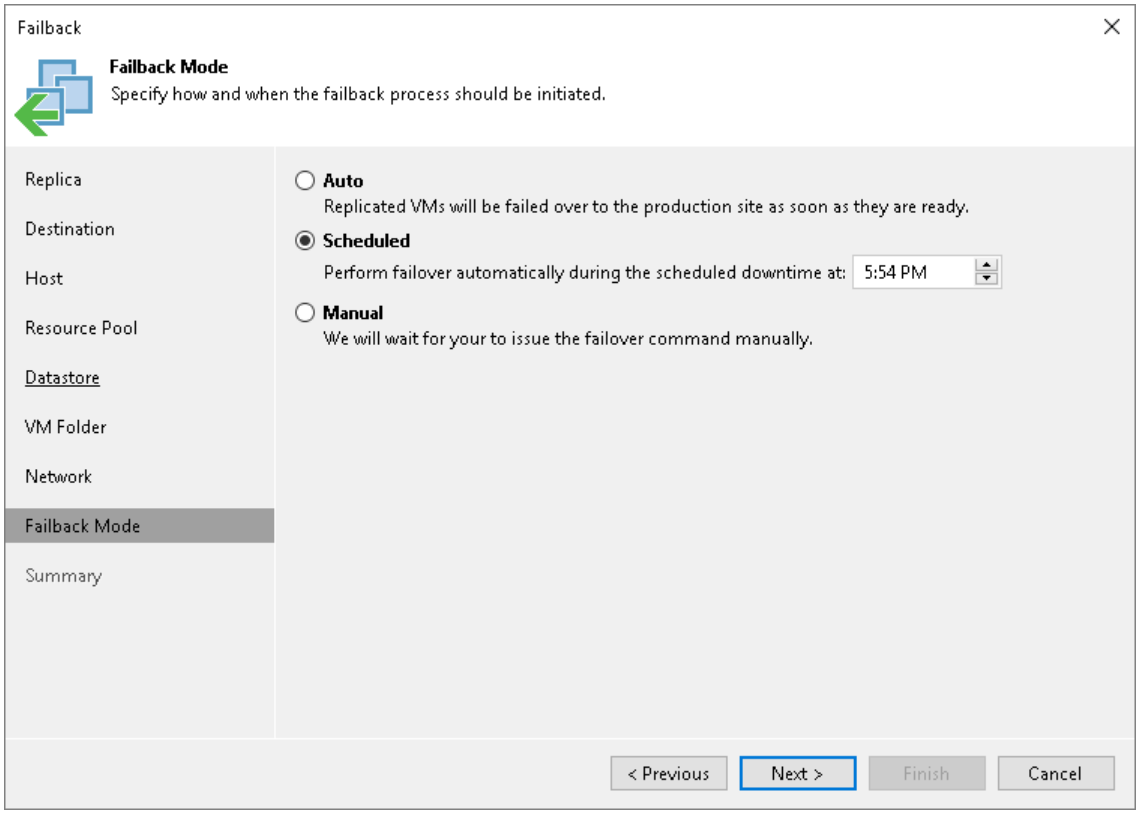

### Step 11. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured failback settings and click **Finish**.

If you want to power on the production VMs right after the switch to production operation is performed, select the **Power on target VM after restoring** check box.

## What You Do Next

Failback is an intermediate step that needs to be finalized. You can finalize failback in the following ways:

- Commit failback.
- Undo failback.

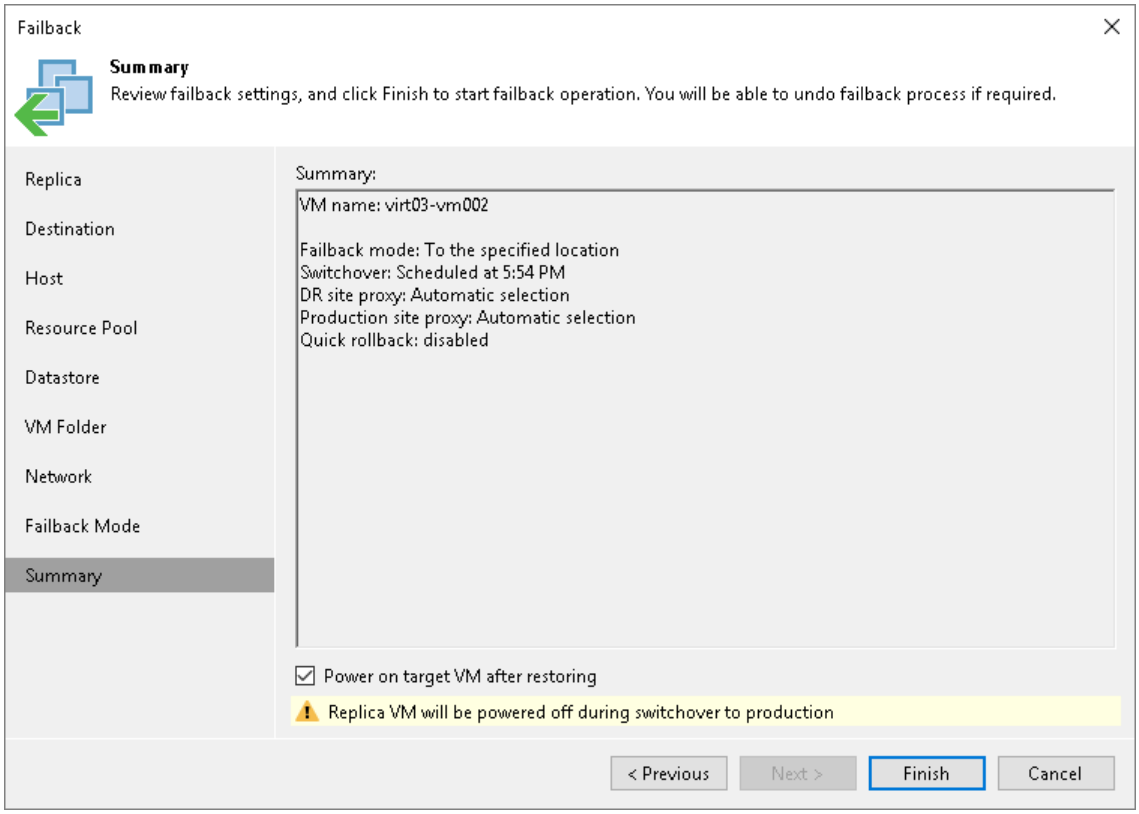


## Switching to Production VMs Manually

The following instructions apply if you have selected to switch from replicas to production VMs manually or at the scheduled time at the **Failback Mode** step of the **Failback** wizard.

To switch to a production VM from its replica, do the following:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. Right-click a replica in the *Ready to switch* state and select **Switchover to production**.

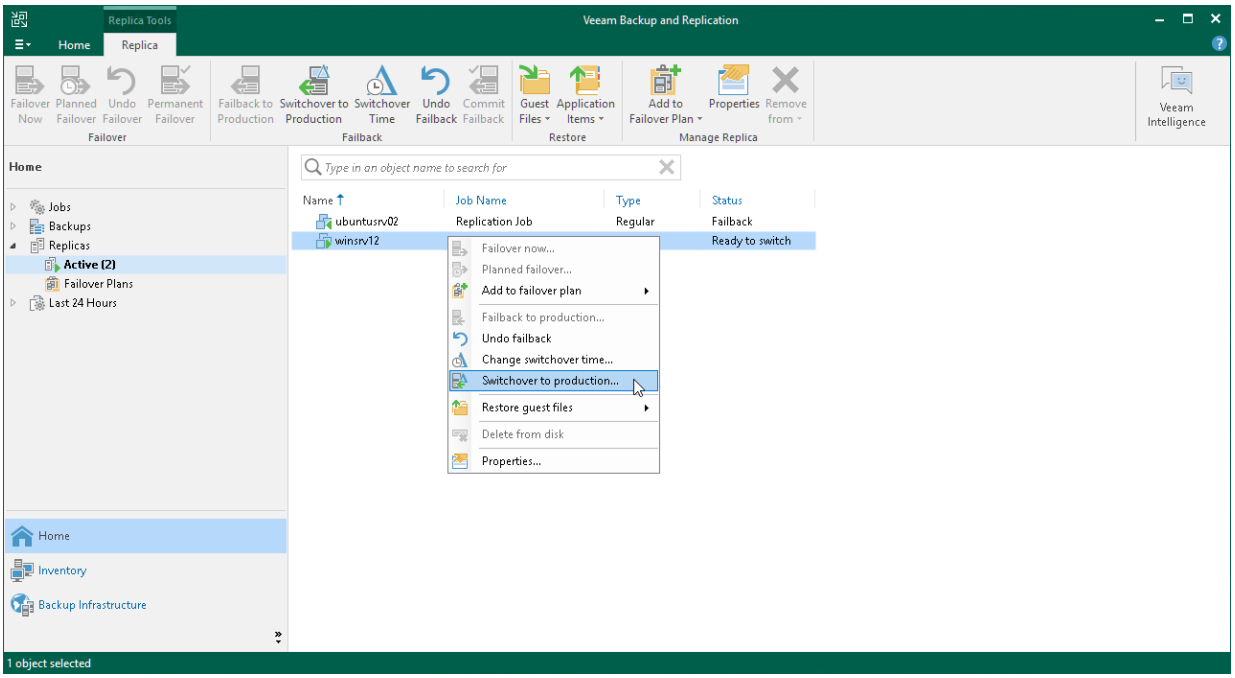


# What You Do Next

After you switch to the production VM, you must finalize failback. You can finalize failback in the following ways:

- Commit failback
- Undo failback

# Changing Switching Time

The following instructions apply if you have selected to switch from replicas to production VMs manually or at the scheduled time at the **Failback Mode** step of the **Failback** wizard.

To change the time when Veeam Backup & Replication will switch from replicas to production VMs:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. Right-click a replica in the *Ready to switch* state and select **Change switching time.**

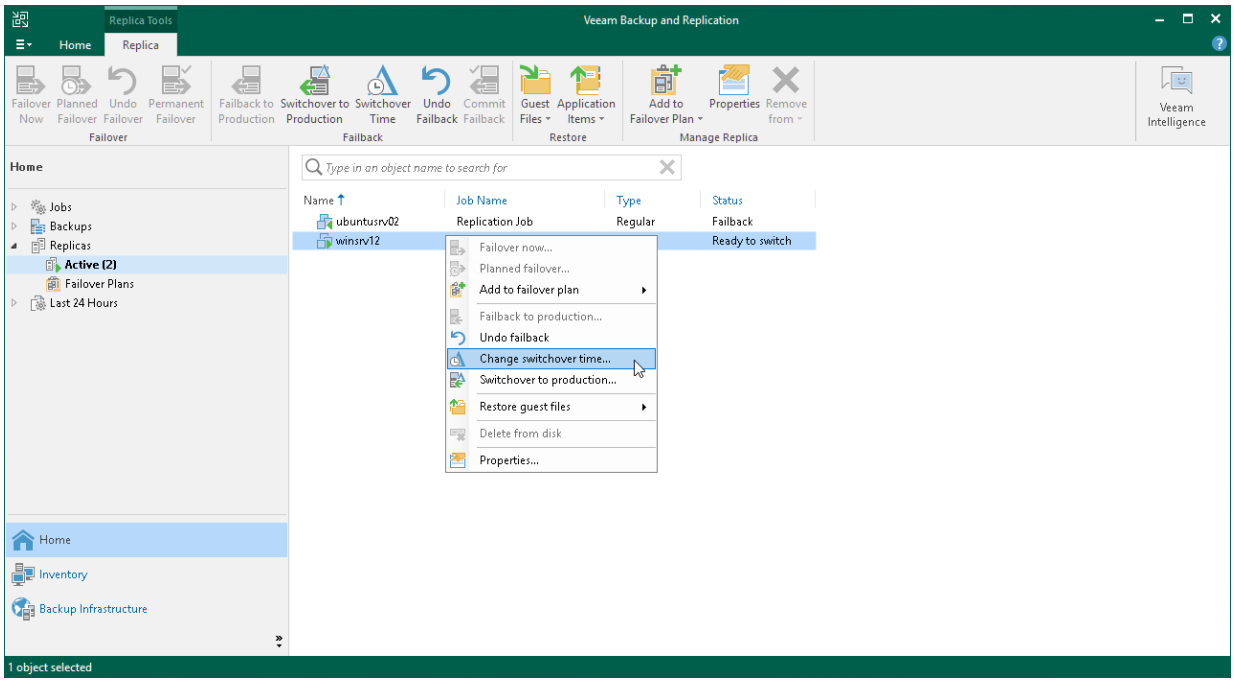

# Failback Commit

Failback commit is one of the ways to finalize failback. When you commit failback, you confirm that the VM to which you failed back (the production VM) and also changes sent to it during failback work as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production VM.

> **NOTE**
>
> If you have selected to switch to the production VM manually, you must first perform the switchover.

The failback commit operation is performed in the following way:

1. Veeam Backup & Replication changes the state of the replica from *Failback* to *Ready*.
2. Further operations depend on whether you have failed back to the source VM or recovered VM:
   - If you have failed back to a VM recovered from a backup or replica, Veeam Backup & Replication reconfigures all existing jobs where the source VM is present and adds the source VM to the list of exclusions. The recovered VM takes the role of the source VM and is included into all jobs instead of the excluded VM. When the CDP policy starts, Veeam Backup & Replication processes the recovered VM instead of the former source VM.
   - If you have failed back to the source VM, the CDP policy is not reconfigured. When the CDP policy starts, Veeam Backup & Replication still processes the source VM.

## Committing Failback

For more information on failback commit, see Failover and Failback for CDP and Failback Commit.

To commit failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary replica and click **Commit Failback** on the ribbon. As an alternative, you can right-click the replica and select **Commit failback**.

> **IMPORTANT**
>
> If you have failed back to a VM with IDE disks, you must manually power off this VM before committing the failback.

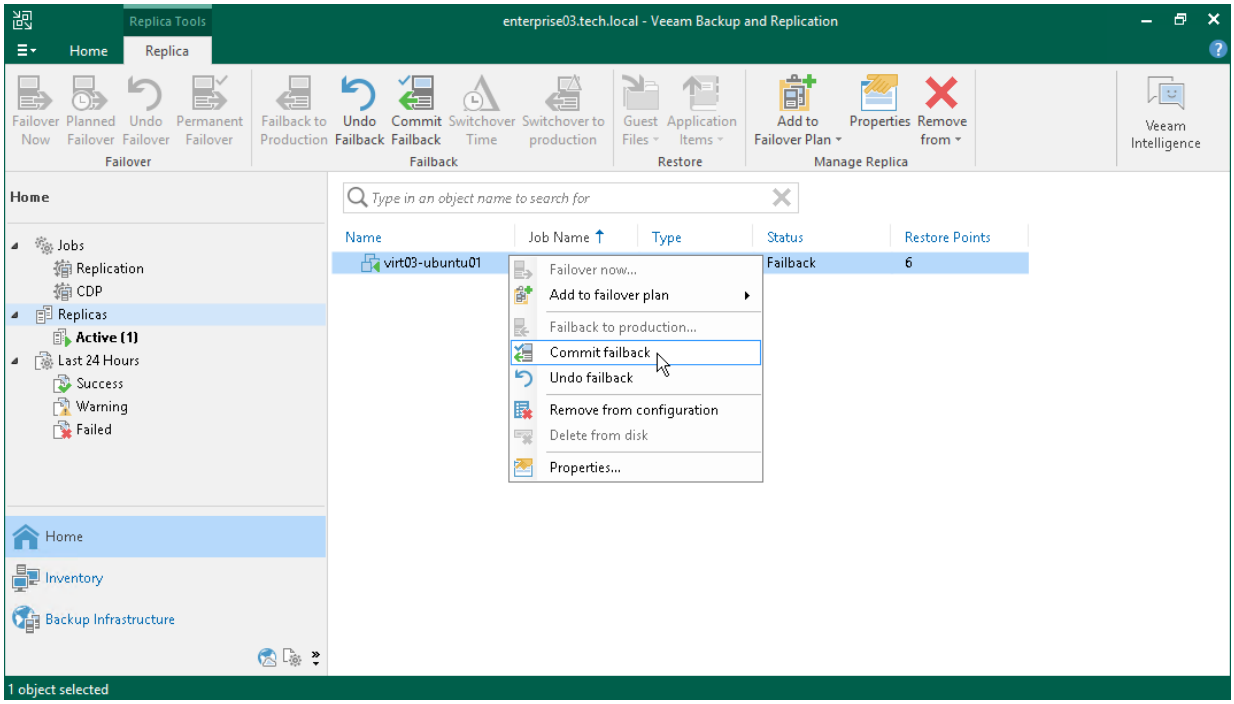
Replica Tools
enterprise03.tech.local - Veeam Backup and Replication
Home
Replica
Failover Now
Planned Failover
Undo Failover
Permanent Failover
Failover
Failback to Production
Undo Failback
Commit Failback
Switchover Time
Switchover to production
Failback
Guest Files
Application Items
Restore
Add to Failover Plan
Properties
Remove from
Manage Replica
Veeam Intelligence
Home
Jobs
Replication
CDP
Replicas
Active (1)
Last 24 Hours
Success
Warning
Failed
Home
Inventory
Backup Infrastructure
Type in an object name to search for
Name
Job Name
Type
Status
Restore Points
virt03-ubuntu01
Failback
6
Failover now...
Add to failover plan
Failback to production...
Commit failback
Undo failback
Remove from configuration
Delete from disk
Properties...
1 object selected

# Failback Undo

Failback undo is one of the ways to finalize failback. When you undo failback, you confirm that the VM to which you failed back (the production VM) and changes sent to it during failback work in a wrong way and you want to get back to the replica.

The failback undo operation is performed in the following way:

1. Veeam Backup & Replication powers off the production VM.
2. Veeam Backup & Replication reverts the replica to its pre-failback state.
3. Veeam Backup & Replication powers on the replica and changes the replica state from *Failback* to *Failover*.

## Undoing Failback

For more information on failback undo, see Failover and Failback for CDP and Failback Undo.

To undo failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary replica and click **Undo Failback** on the ribbon. Alternatively, you can right-click the necessary replica and select **Undo Failback**.

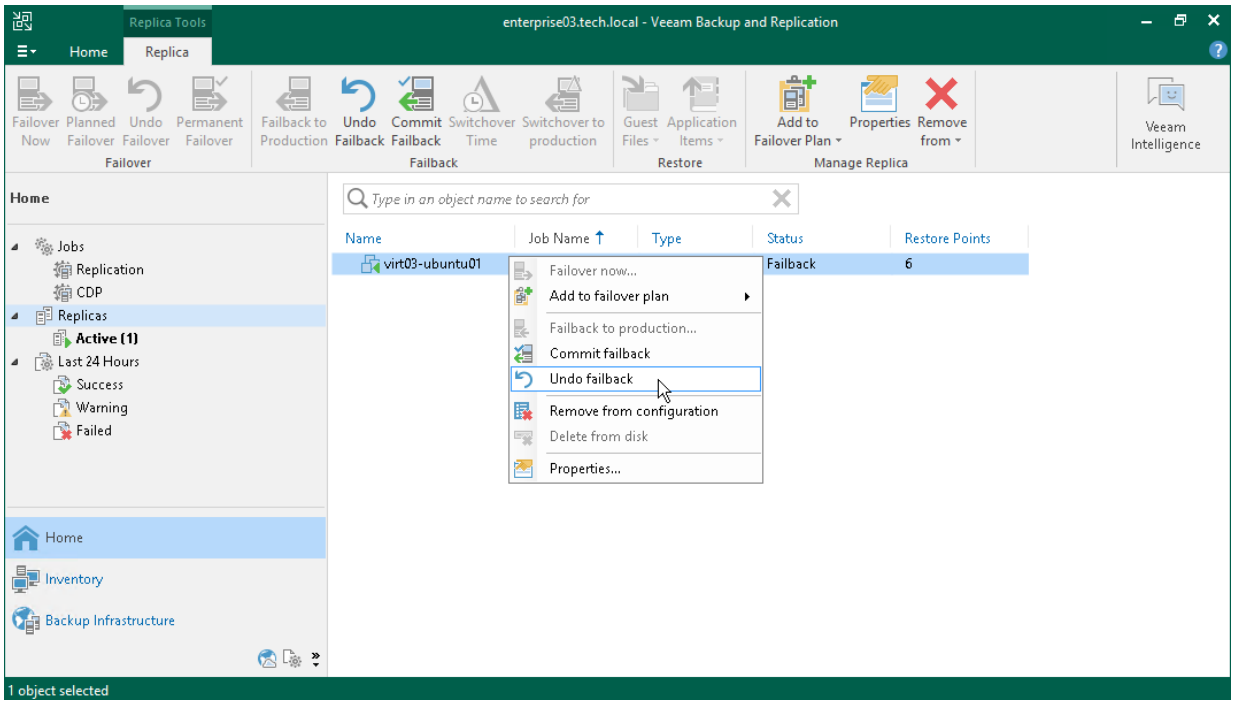

# Unstructured Data Backup

With Veeam Backup & Replication you can easily back up and restore unstructured data: content of SMB (CIFS) and NFS file shares, Windows- and Linux-based file servers, NAS filers and object storage systems (S3 and Azure Blob).

# Backup Infrastructure for Unstructured Data Backup

To protect your unstructured data, you can use your existing Veeam Backup & Replication infrastructure. To do so, configure the following components:

- Unstructured Data Source:
  - File shares
  - Object storage
- General-purpose backup proxies
- Cache repository

- Storage repositories

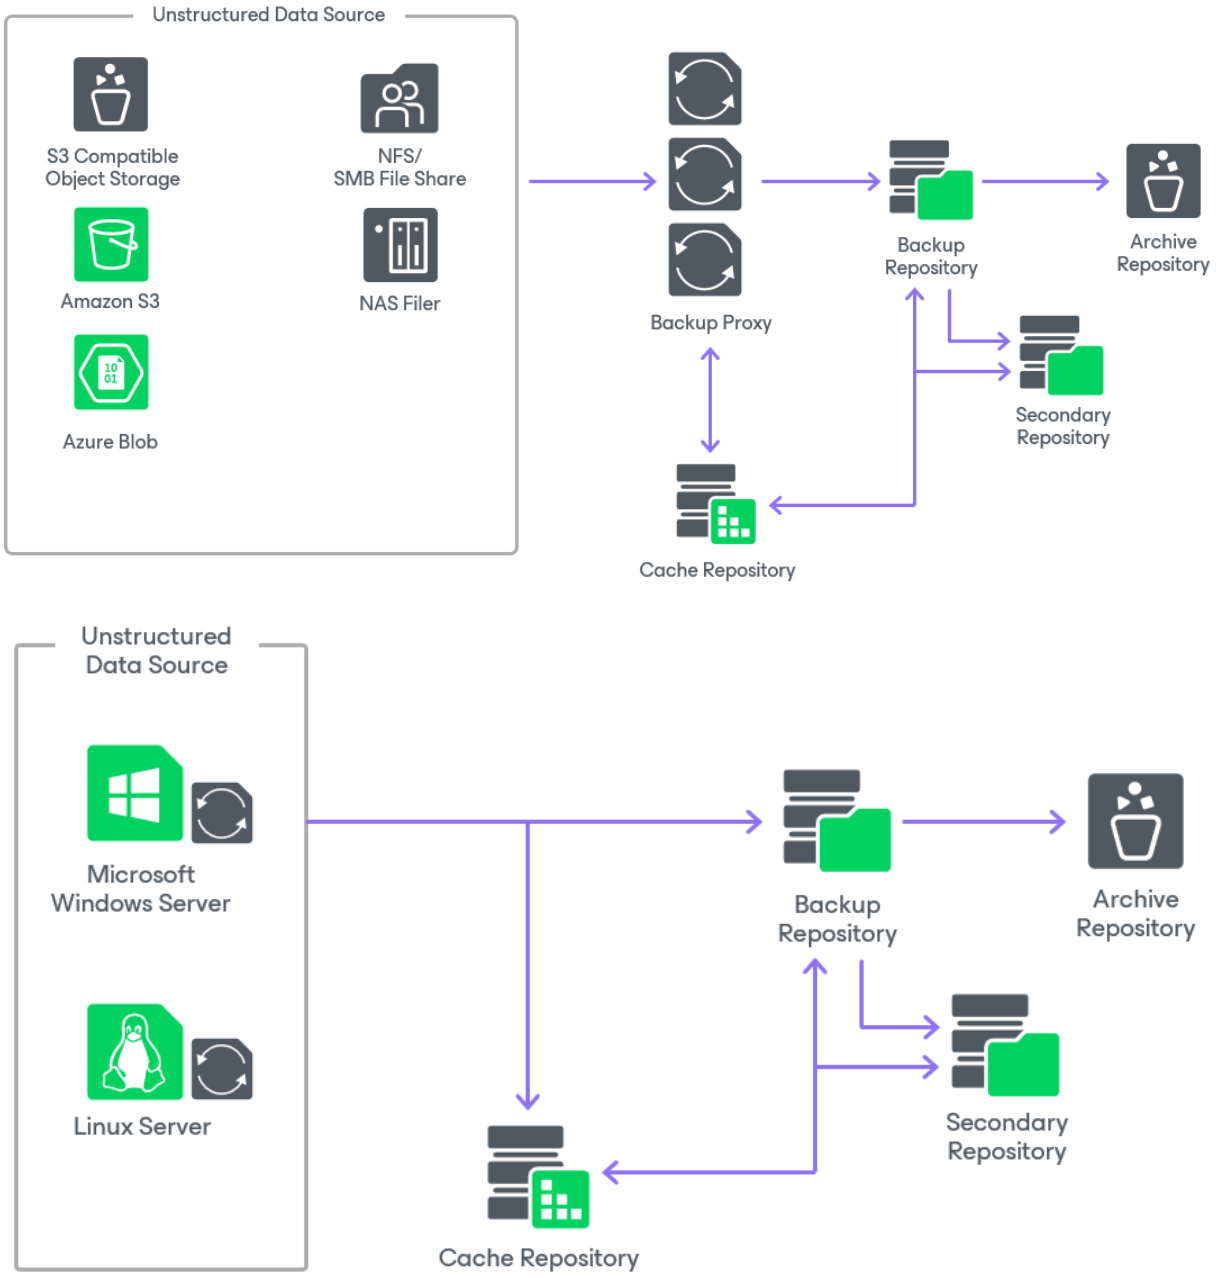


For system requirements for unstructured data backup components, see the System Requirements section.

> **NOTE**
>
> Backup of the content of file shares and object storage repositories is not supported by Veeam Cloud Connect.

To learn how backup components interact during the unstructured data backup, see the How Unstructured Data Backup Works section.

# Unstructured Data Source

Unstructured data supported as a source for backup by Veeam Backup & Replication includes:

- File shares

- Object storage

### File Shares

A file share is a storage device or data source available to multiple hosts through a computer network.

For supported file shares, requirements and limitations, see Supported Platforms and Applications.

File backup jobs in Veeam Backup & Replication can read data from the following sources:

- SMB (CIFS) path
- NFS path
- Path to the storage snapshot folder
- VSS snapshot

> **NOTE**
>
> Consider the following limitations:
>
> - Reading from VSS snapshots on SMB shares is available only under certain conditions, listed in this Veeam KB article.
> - Reading from VSS snapshots on DFS shares is not supported.
> - Names of NTFS Alternate Data Streams (ADS) may be unencrypted if both the file being backed up and its ADS are empty.
> - Different encoding systems for file paths within one file share are not supported. If the file or directory names in your environment have special symbols and use encoding other than en-us.UTF-8, you need to manually specify the encoding of each file share. For more information, see this Veeam KB article.

To learn how to add file shares to the inventory of the virtual infrastructure, see the Adding Unstructured Data Source section.

### Object Storage

An object storage is storage based on either a cloud solution or an S3 compatible on-premises storage solution.

For supported object storage, requirements and limitations, see Supported Platforms and Applications.

To learn how to add object storage as a source for backup to the inventory of the virtual infrastructure, see the Adding Object Storage section.

## General-Purpose Backup Proxies

A general-purpose backup proxy is an architecture component that sits between the unstructured data source and other components of the backup infrastructure. In case of Microsoft Windows and Linux servers, the role of the general-purpose backup proxy is assigned to the backup server itself instead of a dedicated server. The backup proxy operates as a data mover that transfers data between the data source and the backup repository. The backup proxy processes jobs and delivers backup and restore traffic.

For more information on general-purpose backup proxies, their requirements, limitations and deployment, see the General-Purpose Backup Proxies section.

After you configure the backup proxy, choose it to process the backup traffic from unstructured data sources, as described in the Adding NFS File Share, Adding SMB File Share, Adding S3 Compatible Object Storage, Adding Amazon S3 Object Storage, and Adding Microsoft Azure Blob Storage sections.

# Cache Repository

A cache repository is a storage location where Veeam Backup & Replication keeps temporary metadata and uses it to reduce the load on the data source during the backup procedure. The cache repository keeps track of all objects that have changed between each backup session. This allows performing incremental backups from the unstructured data source fast and efficiently. If you store your unstructured data backups on an object storage repository, the cache repository also stores active metadata. For more information, see the Data Structure in Backup, Archive and Secondary Repositories and Unstructured Data Backups in Object Storage Repositories sections.

You can assign the role of a cache repository to a backup repository added to the Veeam Backup & Replication infrastructure. To assign this role, select the backup repository as a cache repository when adding an unstructured data source.

> **NOTE**
>
> You can not assign the role of a cache repository to deduplicating storage appliances.

To minimize the network load during backup, locate the cache repository closer to the backup proxy in the computer network: at the best, they should be located on one machine.

# Storage Repositories

A **backup repository** is a main storage location where Veeam Backup & Replication keeps all versions of backed up files for the configured period and metadata files. Backups stored in the backup repository can be used to quickly restore the entire file share to the state as of a specific restore point.

[Optional] If you want to retain specific files for a longer period of time, you can use cheaper devices for archive purposes. To enable file archiving, configure Veeam Backup & Replication to move backup files and metadata files from the backup repository to an **archive repository**. By default, usage of the archive repository is disabled and, after the retention period for the backup repository is over, backup files are deleted.

[Optional] If you want to store a copy of the unstructured data backup in a different repository, you can configure a **secondary repository** where Veeam Backup & Replication will copy all backups created in the backup repository. The secondary repository can have its own retention policy and encryption settings. By default, no secondary repository is configured.

The following table describes which roles can be assigned to different storage types.

| Storage Type | Backup Repository | Archive Repository | Secondary Repository |
|---|---|---|---|
| Microsoft Windows server | ✓ | ✓ | ✓ |
| Linux server | ✓ | ✓ | ✓ |
| Hardened linux server | ✓ | ✓ | ✓ |
| SMB (CIFS) share [1] | ✓ | ✓ | ✓ |
| NFS share | ✓ | ✓ | ✓ |

| Storage Type | Backup Repository | Archive Repository | Secondary Repository |
|---|---|---|---|
| Dell Data Domain with Data Domain Boost (DDBoost) license | ✓ | ✓ | ✓ |
| ExaGrid | ✓ | ✓ | ✓ |
| HPE StoreOnce with Catalyst license [2] | ✓ | ✓ | ✓ |
| Quantum DXi | ✓ | ✓ | ✓ |
| Fujitsu ETERNUS CS800 | ✓ | ✓ | ✓ |
| Infinidat InfiniGuard | ✓ | ✓ | ✓ |
| Scale-out backup repository (SOBR) [3, 4] | ✓ | × | ✓ |
| Object storage repository [5] | ✓ | ✓ | ✓ |
| Repository with rotated drives | ✓ | × | ✓ |
| Veeam Cloud Connect repository | × | × | × |

[1] If you use a Dell PowerScale (formerly Isilon) storage system in the **CIFS Share Access** mode, make sure that you have assigned your service account to the built-in **BackupAdmin** role within PowerScale. Otherwise, the access to the share will be denied.

[2] If you plan to use HPE StoreOnce storage appliances, consider the following recommendations for optimal performance:

- A StoreOnce system can have multiple Catalyst stores, and large backup loads (exceeding 1PB) should be spread across more than one Catalyst store on the same StoreOnce system.
- Do not include Catalyst stores in a SOBR intended for unstructured data backups. This will reduce the global deduplication of the StoreOnce system.

[3] An object storage repository added as a capacity tier in a scale-out backup repository cannot be used for storing unstructured data backups. To archive unstructured data backup files to an object storage repository, assign the object storage repository as an archive repository when you create a file backup job.

[4] SOBR consisting of object storage repositories cannot be used as a target backup repository for file backup jobs.

[5] Consider the following limitations:

- Amazon S3 Glacier and Azure Blob Storage Archive Tier are not supported for unstructured data backup.
- Amazon S3 Snowball Edge and Azure Databox are not supported as archive repositories for unstructured data backup, but you can use them as backup repositories, secondary repositories, or targets for copying file share backups.
- You cannot use S3 compatible repositories with multiple buckets as a target repository for Unstructured Data Backup. This includes backup, archive, or secondary repositories.

You can create two object storage repositories pointing to the same cloud folder/bucket and use these repositories for storing both unstructured data backups and Capacity Tier backups at the same time: one object storage repository will be used to store unstructured data backups, the other one – to store virtual and physical machine backups as a capacity tier in a single SOBR. However, these object storage repositories (mapped to the same cloud folder) must not be used across multiple Veeam Backup & Replication servers for the same purposes as it leads to unpredictable system behavior and inevitable data loss.

# Deployment of Backup, Archive and Secondary Repositories

To use a storage, which is already added to the Veeam Backup & Replication Backup Infrastructure, to store unstructured data backups, define it as a target storage when creating a file backup job or an object storage backup job:

- Configure the backup repository at the **Backup Repository** step of the wizard when creating file backup jobs and object storage backup jobs.
- Configure the backup repository at the **Archive Repository** step of the wizard when creating file backup jobs and object storage backup jobs.
- Configure the backup repository at the **Secondary Target** step of the wizard when creating file backup jobs and object storage backup jobs.

If the required storage is not added as a backup repository in your Veeam Backup & Replication Backup Infrastructure, add it as described in the Backup Repositories section.

# Adding Unstructured Data Source

You must add sources of unstructured data, which you plan to protect with the unstructured data backup, to the inventory of the virtual infrastructure.

You can add the following types to the inventory of your virtual infrastructure:

- Windows-managed or Linux-managed file server
- File shares
  - SMB file share
  - NFS file share
- Enterprise NAS system as a NAS filer
- Object storage

## Adding File Server

Before you add a Windows- or Linux-managed server as a file server to the inventory of the virtual infrastructure, consider the following:

- This server must meet requirements listed in the Platform Support section.
- You must have this server added in **Backup Infrastructure**.

  For more information on how to add servers, see the Adding Microsoft Windows Servers and Adding Linux Servers section.
- If you plan to use a dedicated cache repository, make sure it is added in **Backup Infrastructure**.
- Data from managed servers is transferred directly to the repository without a proxy server.

## Step 1. Launch New File Server Wizard

To launch the **New File Server** wizard:

1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **Unstructured Data** node and select **Add unstructured data source**.
   - Select the **Unstructured Data** node and click **Add Data Source** on the ribbon.
   - Select the **Unstructured Data** node and click **Add Data Source** in the working area.
3. In the **Add Unstructured Data Source** window, click **File server**.

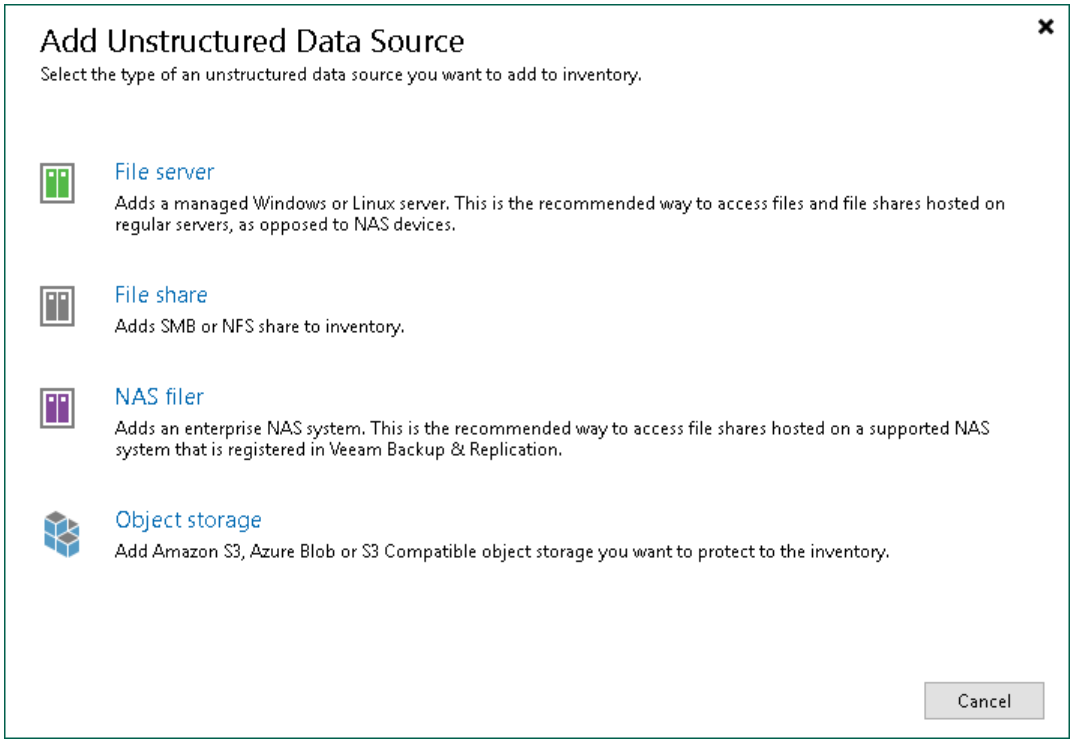

## Step 2. Add Managed Server

At the **File Server** step of the wizard, choose the server, which you want to use as a file share server. Select it from the **Managed Server** drop-down list.

> **NOTE**
>
> If you plan not only to back up the Linux-managed file server, but also to restore files to it, use an account with root access when adding the server to the backup infrastructure.

If the drop-down list does not display the required server, you must add it to the backup infrastructure. To add the server, do the following:

1. Click **Add New**.
2. Select **Microsoft Windows Server** or **Linux Server**.
3. Add a new Windows or Linux server to the backup infrastructure as described in the Adding Microsoft Windows Servers and Adding Linux Servers sections.
4. Select the newly added server from the **Managed Server** drop-down list.

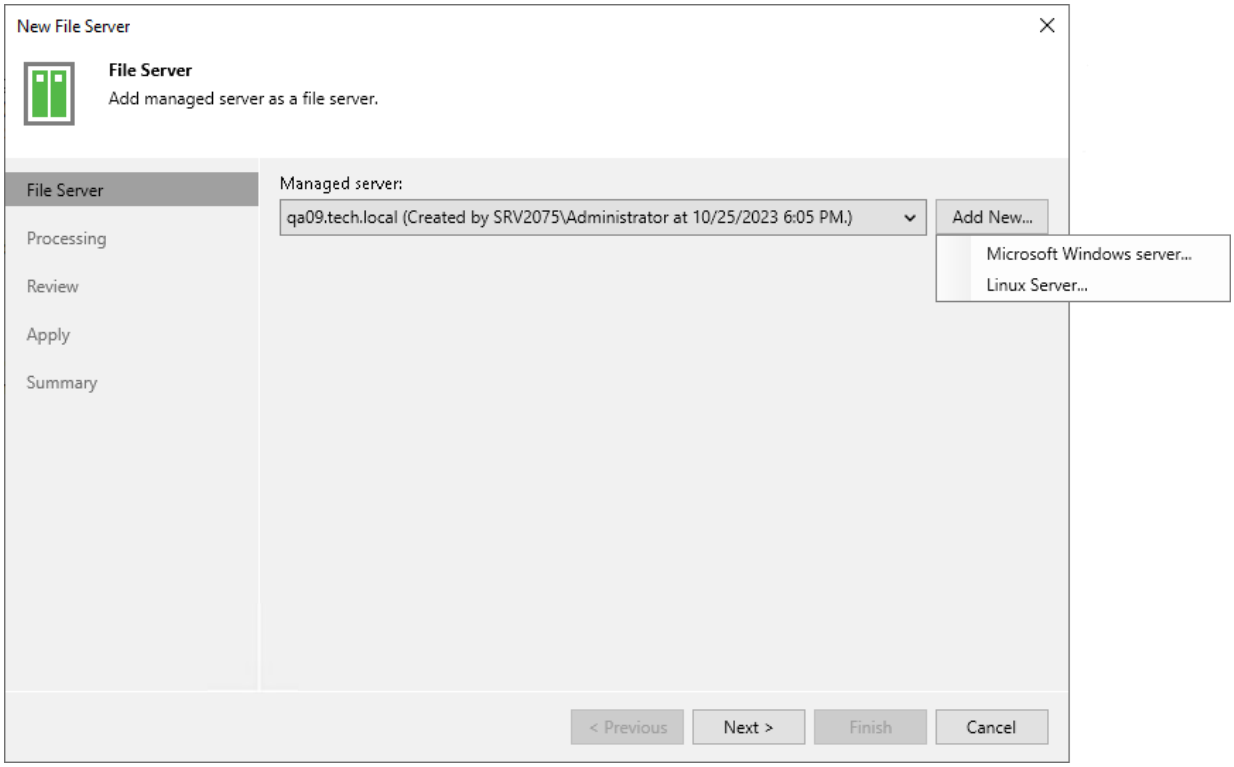

## Step 3. Specify File Server Processing Settings

At the **Processing** step of the wizard, define file server processing settings:

1. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located as close to the source file share as possible.

   If you change the cache repository for an existing file server whose backups are stored in the object storage, Veeam Backup & Replication will prompt you to either attach migrated metadata, copy metadata from the previous cache repository, or download metadata manually from the archive repository. For more information on storing unstructured data backups in the object storage and changing the cache repository, see the Unstructured Data Backups in Object Storage Repositories section.

2. Use the **Backup I/O control** slider to define how fast the backup proxies can read data from the source file server. This setting is based on the number of the parallel threads that can be used by proxies configured for processing the file server.

| I/O Control | Number of Proxies | Threads per Task |
|---|---|---|
| **Lower Impact** | 1 | 1 |
| **Below Normal** | 1 | 4 |
| **Normal** | 2 | 8 |
| **Above Normal** | 4 | 16 |
| **Faster Backup** | Unlimited | 16 |

   If resources of your file server are limited, it is recommended that you select the **Lower impact** option. If your file server is powerful enough, select the **Faster backup** option.

3. Click **Next** to save the configured settings.

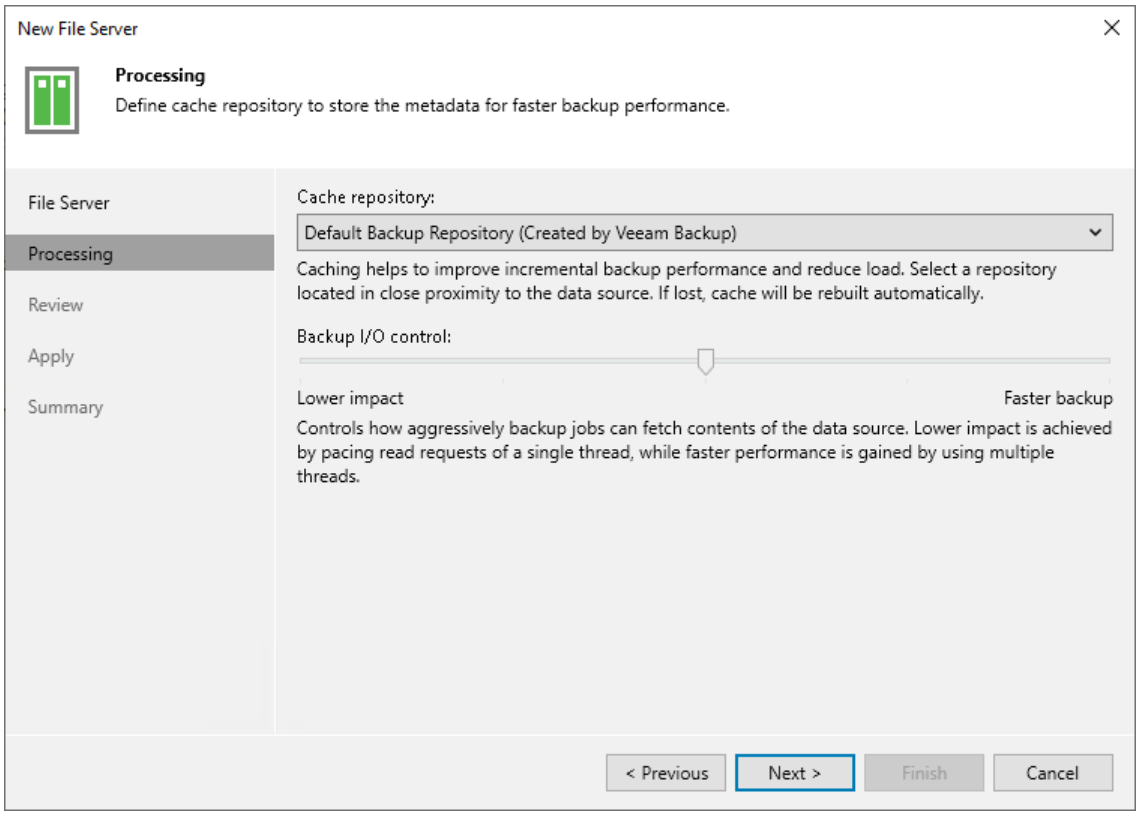

## Step 4. Review Components to Install

At the **Review** step of the wizard, review what Veeam Backup & Replication components are already installed on the server and click **Apply** to start installation of missing components.

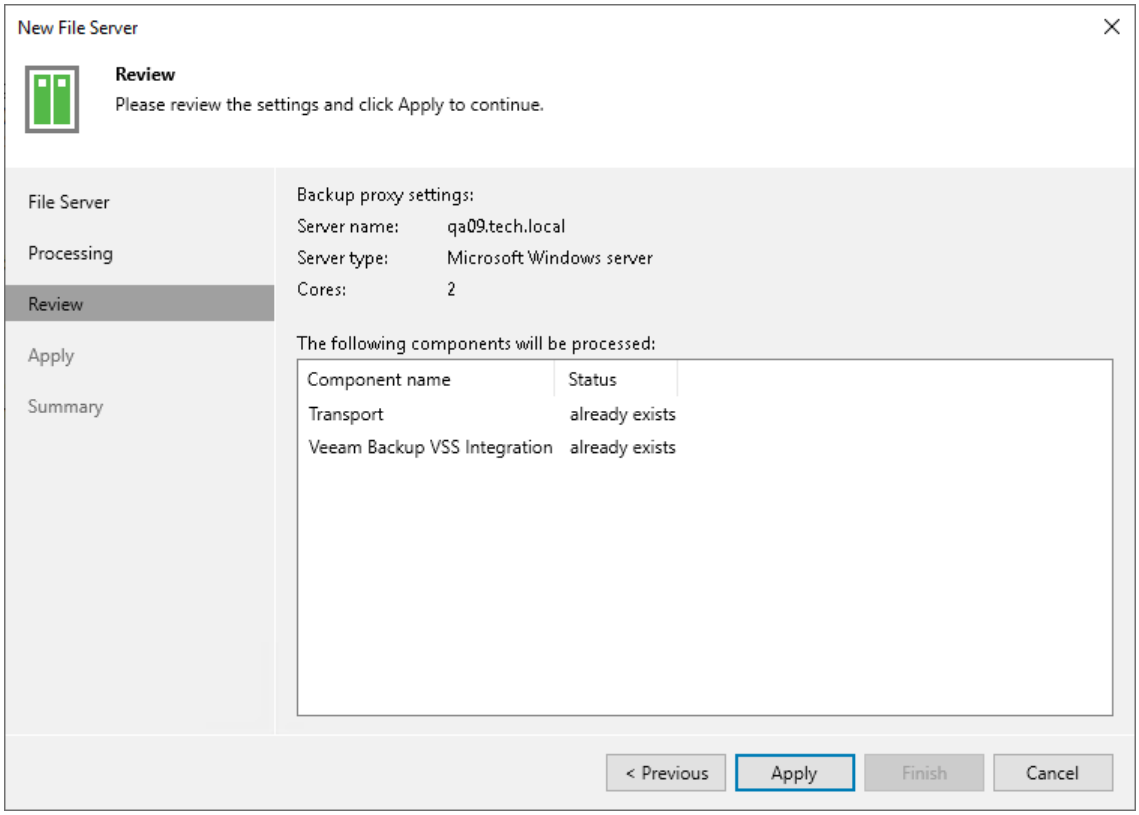

## Step 5. Apply File Share Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components. Click **Next** to complete the procedure of the file share role assignment to the managed file server.

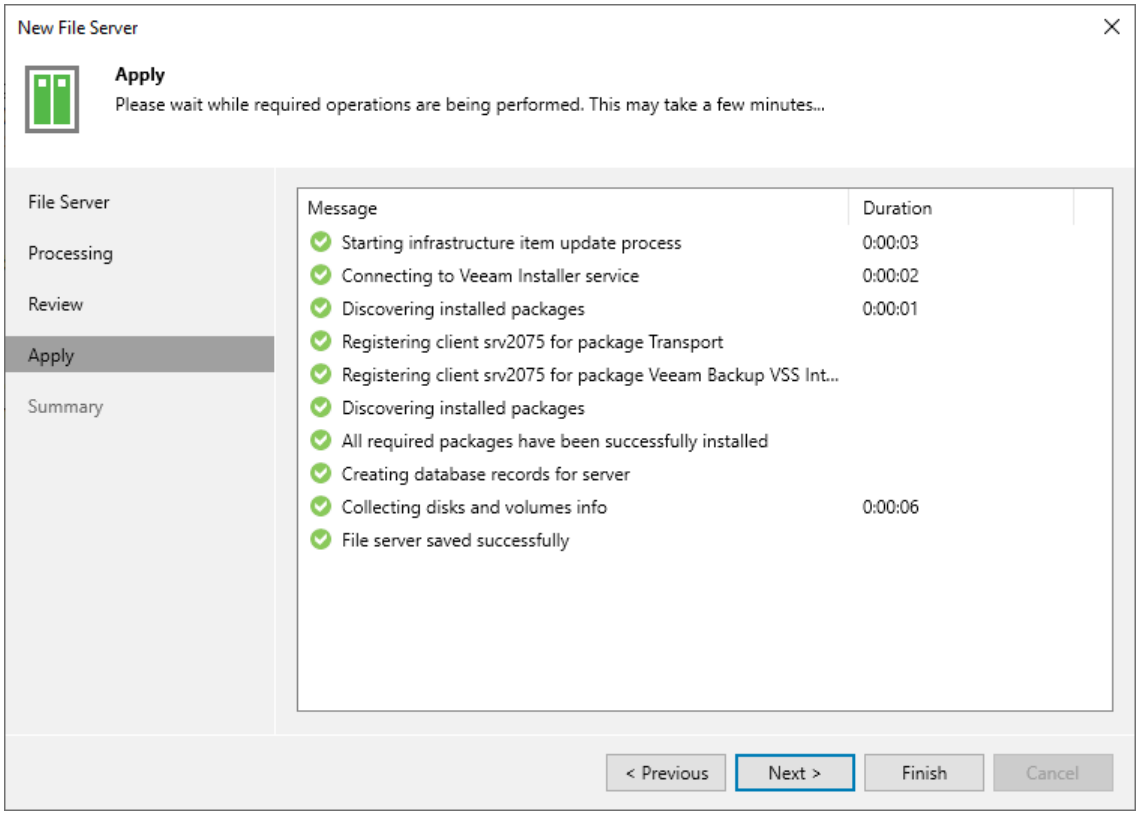

## Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added file server and click **Finish to exit the wizard**.

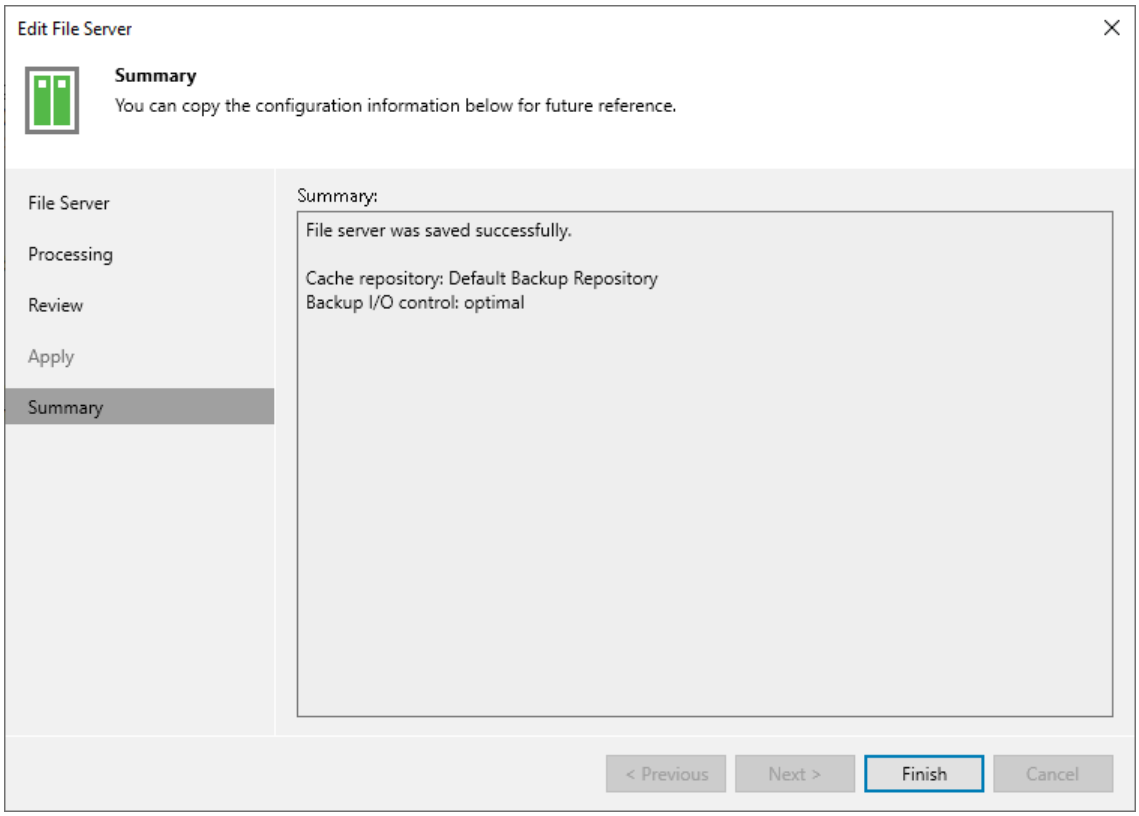

## Step 7. Specify Settings for Connected Volumes

> **NOTE**
>
> This feature is available for Windows-based servers only.

Before you specify settings for a managed server added as a file server, you must rescan its volumes. During volume rescan, Veeam Backup & Replication retrieves information about disks and volumes that are currently connected to the server and writes this information to the configuration database.

Veeam Backup & Replication automatically performs volume rescan every 4 hours. You can also start volume rescan manually:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed servers**.
3. In the working area, select the server and click **Rescan** on the ribbon. Alternatively, you can right-click the host and select **Rescan**.

After you add a managed server as a file server to the inventory of the virtual infrastructure, you can configure the following settings for it:

- Specify volume-specific settings.
- Enable or disable failover to direct backup if a snapshot is not available.

# Volume-Specific Settings

You can define volume-specific settings for the file server: select what VSS provider must be used for snapshot creation and specify the maximum number of concurrent snapshots that must exist for the volume.

To specify volume-specific settings:

1. Open the **Inventory** view.
2. In the inventory pane, select **Unstructured Data - File Servers**.
3. In the working area, select the server and click **Manage Volumes** on the ribbon. Alternatively, you can right-click the managed server and select **Manage volumes**.
4. Select the volume in the list and click **Edit**.
5. Specify VSS provider settings for the volume:
   - To take a VSS snapshot of a specific volume, Veeam Backup & Replication uses one of VSS providers available for this volume. To explicitly define what VSS provider must be used for the volume, select the VSS provider from the **VSS provider** list.
   - You can simultaneously store 4 snapshots of one volume. To change this number, specify the **Max snapshots** value. It is not recommended that you increase the number of snapshots for a slow storage. Many snapshots existing at the same time may cause VM processing failures.

6. Click **OK** to save the changes.

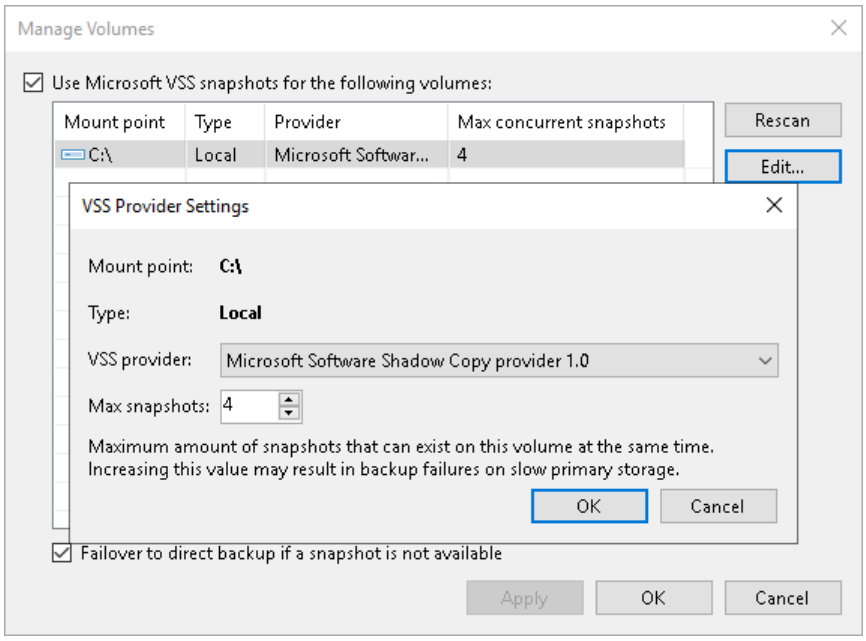


# Failover to Direct Backup Settings

If snapshots for volumes are not available, Veeam Backup & Replication may automatically fail over to direct backup.

By default, the failover option is enabled. To disable failover to direct backup:

1. Open the **Inventory** view.
2. In the inventory pane, select **Unstructured Data - File Servers**.
3. In the working area, select the server and click **Manage Volumes** on the ribbon. Alternatively, you can right-click the managed server and select **Manage volumes**.
4. In the **Manage Volumes** window, clear the **Failover to direct backup if a snapshot is not available** check box.

> **IMPORTANT**
>
> We do not recommend to clear the **Failover to direct backup if a snapshot is not available** check box, as it may result in the file backup job failure if a snapshot is not available for the volume.

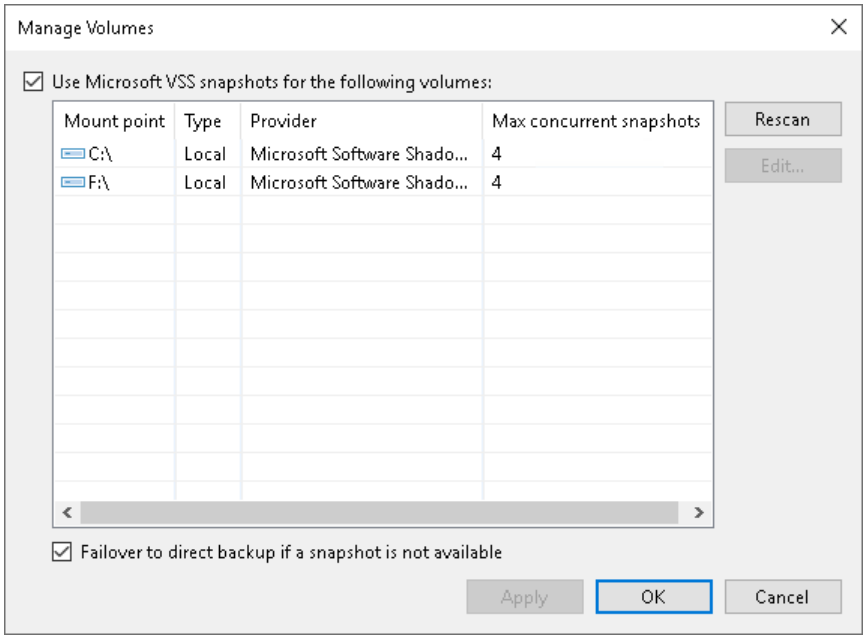


# Adding File Share

You can add NFS and SMB file shares as a source of unstructured data available for protection by Veeam Backup & Replication.

## Adding SMB File Share

Before you add an SMB file share to the inventory of the virtual infrastructure, consider the following:

- The file share meets requirements listed in the Platform Support section.
- If you plan to use a dedicated proxy server or cache repository, make sure these components are added in **Backup Infrastructure**.

## Step 1. Launch New File Share Wizard

To launch the **New File Share** wizard:

1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **File Shares** node and select **Add File Share**.
   - Select the **File Shares** node and click **Add File Share** on the ribbon.
   - Select the **File Shares** node and click **Add File Share** in the working area.
3. In the **Add File Share** window, click **SMB share**.

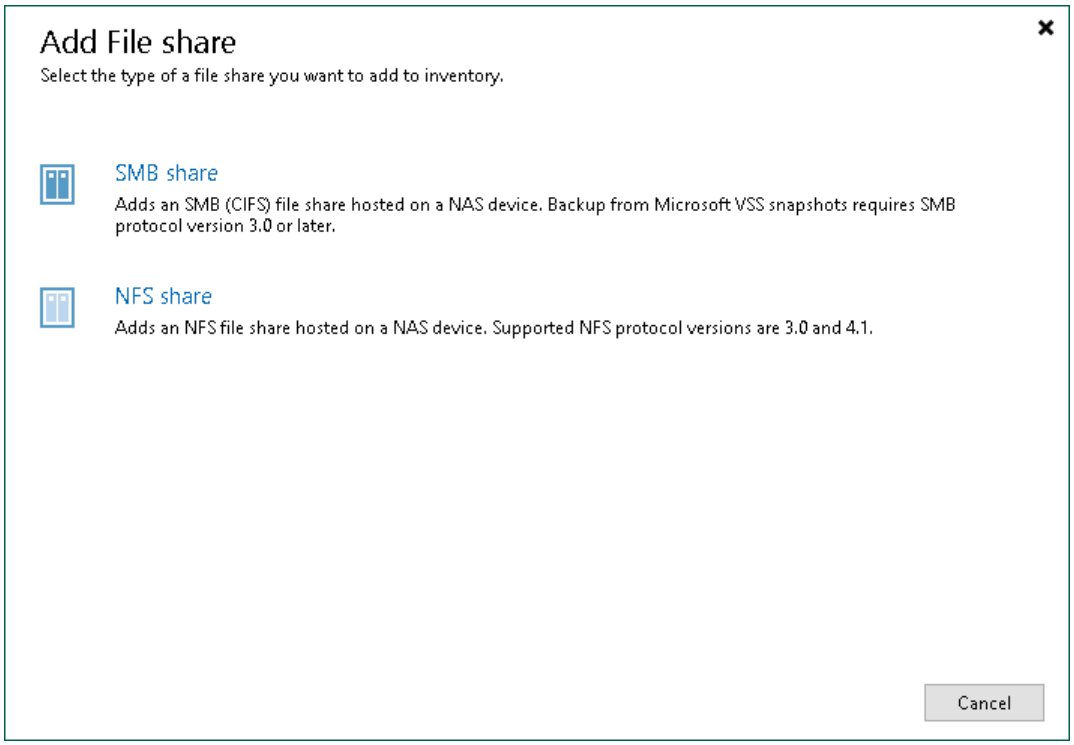

## Step 2. Specify Path to SMB File Share and Access Credentials

At the **SMB File Share** step of the wizard, specify access settings for the SMB file share:

1. In the **SMB server of file share** field, specify the path to an SMB file share in the *||server|folder* format. You can specify the IPv4 or IPv6 address of the server. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in the IPv6 Support section.

   You can also click **Browse** and select the shared folder from the list of available network shares.

   You can add the root server folder in the *||server* format to protect all SMB file shares residing on this server. After that, create a single file backup job to protect the added server, as described in the Creating File Backup Jobs section. Then all SMB file shares added on this server will be automatically processed with the file backup job and protected. If you previously had several separate non-root shared folders residing on the same server and want to switch to using a single root shared folder to cover the same shares, you do not have to run full backups to update data of protected shares. Instead, you can convert existing backups and update existing file backup jobs to protect single root shared folders comprising all other non-root shared folders residing on the same server. To learn more about the conversion, see the Converting Backups from Non-Root to Root Shared Folders section. Perform the conversion with extreme caution.

2. If you must specify user credentials to access the shared folder, select the **This share requires access credentials** check box. From the **Credentials** drop-down list, select a credentials record for a user account that has **Full Control** permissions on the shared folder.

   To access the SMB share, you must use an account that meets either of the following requirements:

   - If you only plan to back up the share, you can use an account with read-only permissions.
   - If you plan not only to back up the share, but also to restore files to it, use an account with read/write permissions.

   > **NOTE**
   >
   > Accessing the SMB file share with credentials in the User Principal Name format (*user@domain.xxx*) is not supported.

   If you have not set up credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see the Managing Credentials section.

   > **NOTE**
   >
   > If the **This share requires access credentials** check box is not selected, Veeam Backup & Replication uses the computer account of the proxy server to access the file share.

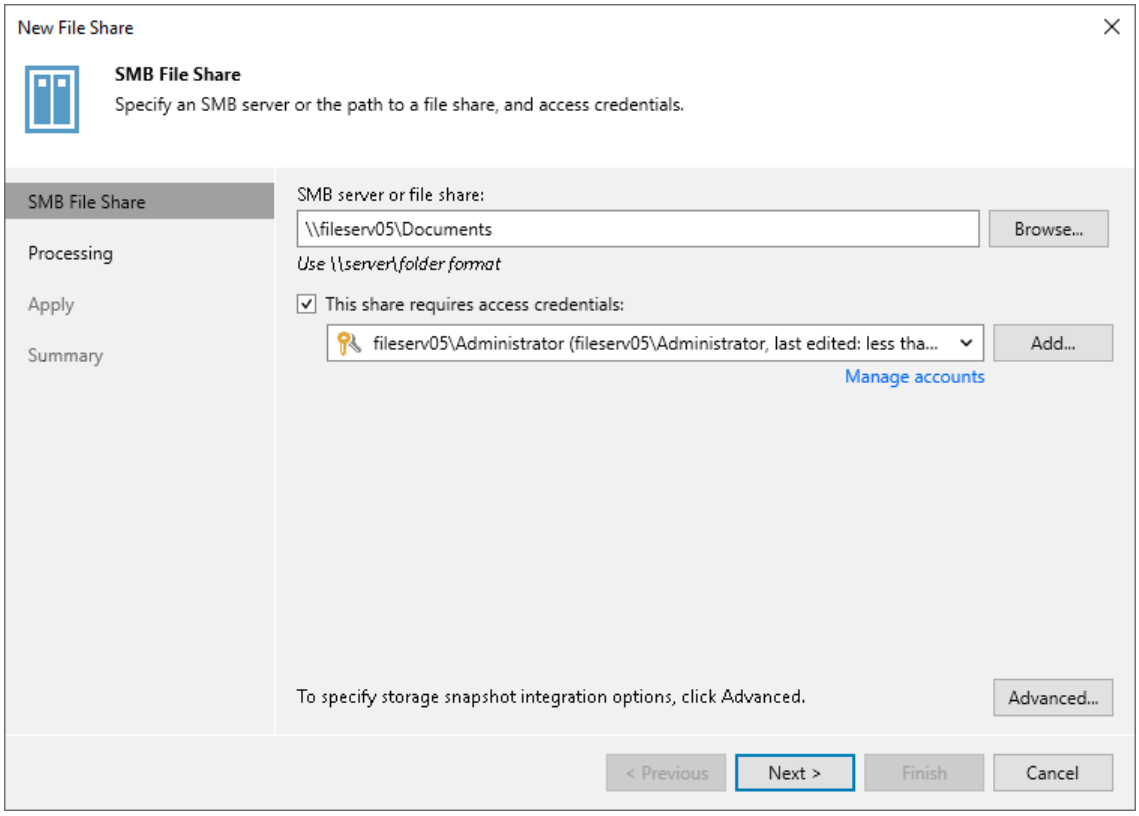
New File Share
SMB File Share
Specify an SMB server or the path to a file share, and access credentials.
SMB File Share
Processing
Apply
Summary
SMB server or file share:
\\fileserv05\Documents
Browse...
Use \\server\folder format
This share requires access credentials:
fileserv05\Administrator (fileserv05\Administrator, last edited: less tha...
Add...
Manage accounts
To specify storage snapshot integration options, click Advanced.
Advanced...
< Previous
Next >
Finish
Cancel

## Step 3. Specify Advanced SMB File Share Settings

You can instruct Veeam Backup & Replication to back up data from Microsoft VSS snapshots or native storage snapshots. During backup jobs, Veeam Backup & Replication will read data of shared files and folders from snapshots, which speeds up backup operations and improves RPOs.

To define if Veeam Backup & Replication will use snapshots for backups:

1. At the **SMB File Share** step of the wizard, click **Advanced**.
2. In the **Advanced** window, select one of the following options:
   - To ignore the snapshot functionality, select **Backup directly from the file share**.

     Veeam Backup & Replication will ignore locked files and folders. When creating a backup job, you can configure notifications to list files and folders that are skipped during the backup procedure. For more information see the Notification Settings section.
   - To back up files from Microsoft VSS snapshots, select **Backup from a Microsoft VSS snapshot**.

     If you select this option, make sure that the file share and the backup proxy used for the file backup job support SMB protocol version 3.0 or later.
   - To back up files from the native storage snapshot, select **Backup from a storage snapshot at the following path** and specify the path in the *||server|snapshotfolder|snapshotname* format to the snapshot stored on the SMB file share. You can specify the IPv4 or IPv6 address of the server. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in the IPv6 Support section.

     If you select this option, you can additionally use custom scripts written by you, for example, to create a snapshot before the backup and remove it after the backup. You can define these scripts when creating a new file backup job, as described in the Script Settings section.

> **NOTE**
>
> Consider that Veeam Backup & Replication does not take snapshots itself, but it can use a snapshot taken by the storage system. File backup jobs do not trigger the storage snapshot creation and deletion automatically. You can specify the folder where the storage snapshot is stored. In this case file backup jobs can access this folder and read data from the storage snapshot.

3. Select **Failover to direct backup if snapshot is not available** if you want Veeam Backup & Replication to read data for backup directly from the file share when the snapshot is unavailable. If you do not select the option and the snapshot is unavailable, the file backup job will stop with a failure.

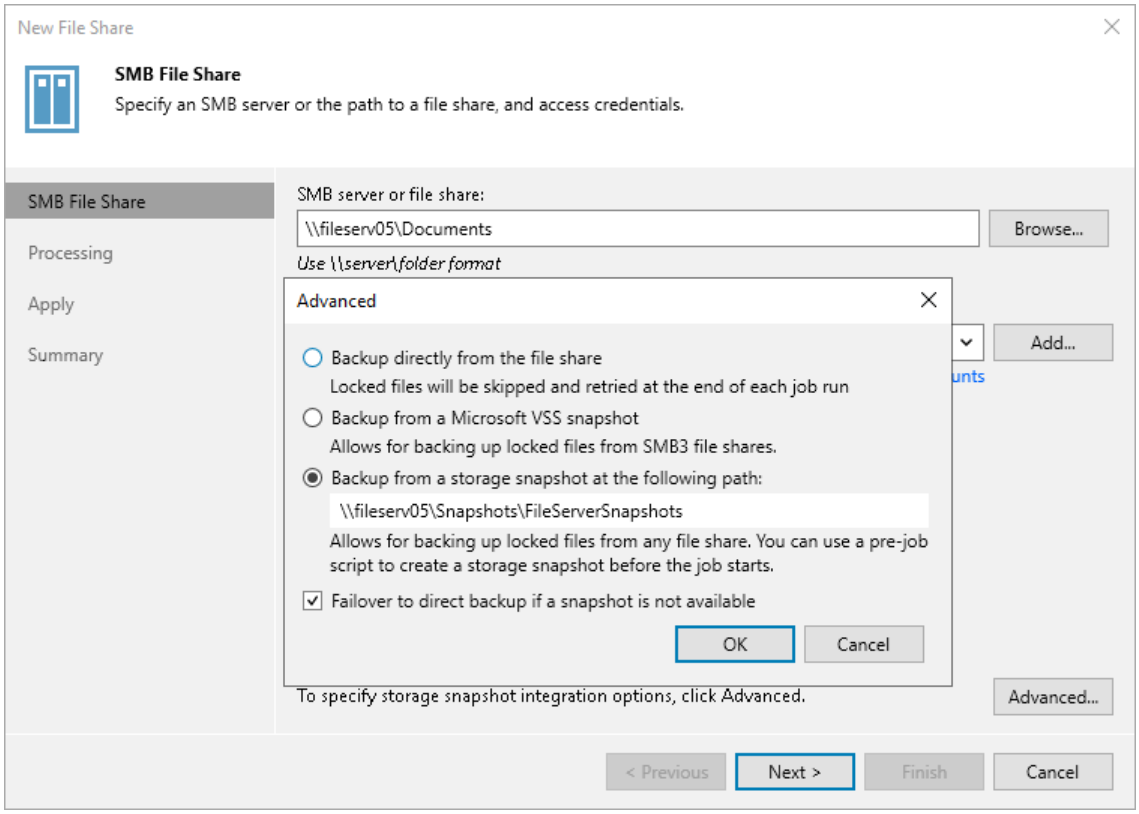

## Step 4. Specify File Share Processing Settings

At the **Processing** step of the wizard, do the following:

1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
2. In the **Backup Proxy** window, select proxy servers:
   - If you select **All proxies**, Veeam Backup & Replication will use all available backup proxies for file backup. The number of proxies in use defines the number of data threads that transfer data from the file share to the backup repository. The more data transfer threads Veeam Backup & Replication uses, the higher is the data transfer speed.

     If the file share is used as a source for a file to tape backup job, the tape server utilized for this job is added as yet another backup proxy when creating a file to tape backup job. This backup proxy has the highest priority over all others and is used by default if it has access rights to the file share. For details on file to tape backup jobs, see the File Backup to Tape section in the Veeam Backup & Replication User Guide.
   - If you select **Use the selected backup proxies only**, you can explicitly specify backup proxies that Veeam Backup & Replication must use for file backup.

     It is recommended that you select at least two backup proxies to ensure that the backup jobs start even if one of the proxies fails or loses its connectivity to the source file share. The more proxies you select, the more data transfer threads Veeam Backup & Replication will use for backup jobs, thus improving performance.

     Even if the file share is used as a source for file to tape backup jobs, Veeam Backup & Replication will use only proxies selected in the list to process the backup data traffic.

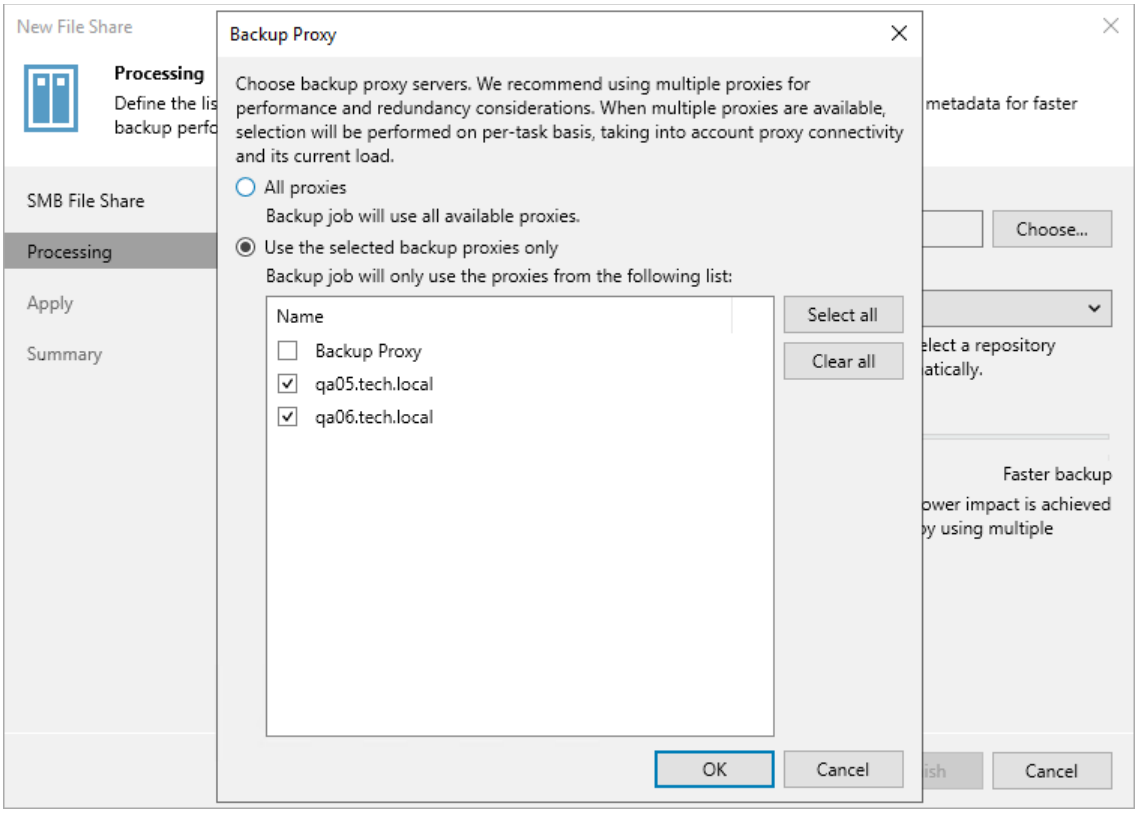

3. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located in the close proximity to the source file share and backup proxies.

   If you change the cache repository for an existing file share whose backups are stored in object storage, Veeam Backup & Replication will prompt you to either attach migrated metadata, copy metadata from the previous cache repository, or download metadata manually from the archive repository. For more information on storing NAS backups in the object storage and changing the cache repository, see the Unstructured Data Backups in Object Storage Repositories section.

4. Use the **Backup I/O control** slider to define how fast backup proxies can read data from the source file share. This setting is based on the number of parallel threads that can be used by proxies configured for processing the file share.

| I/O Control | Number of Proxies | Threads per Task |
|---|---|---|
| **Lower Impact** | 1 | 1 |
| **Below Normal** | 1 | 4 |
| **Normal** | 2 | 8 |
| **Above Normal** | 4 | 16 |
| **Faster Backup** | Unlimited | 16 |

   If resources of your NAS device are limited, it is recommended that you select the **Lower impact** option. If your NAS device is powerful enough, select the **Faster backup** option.

5. Click **Apply** to save the configured settings.

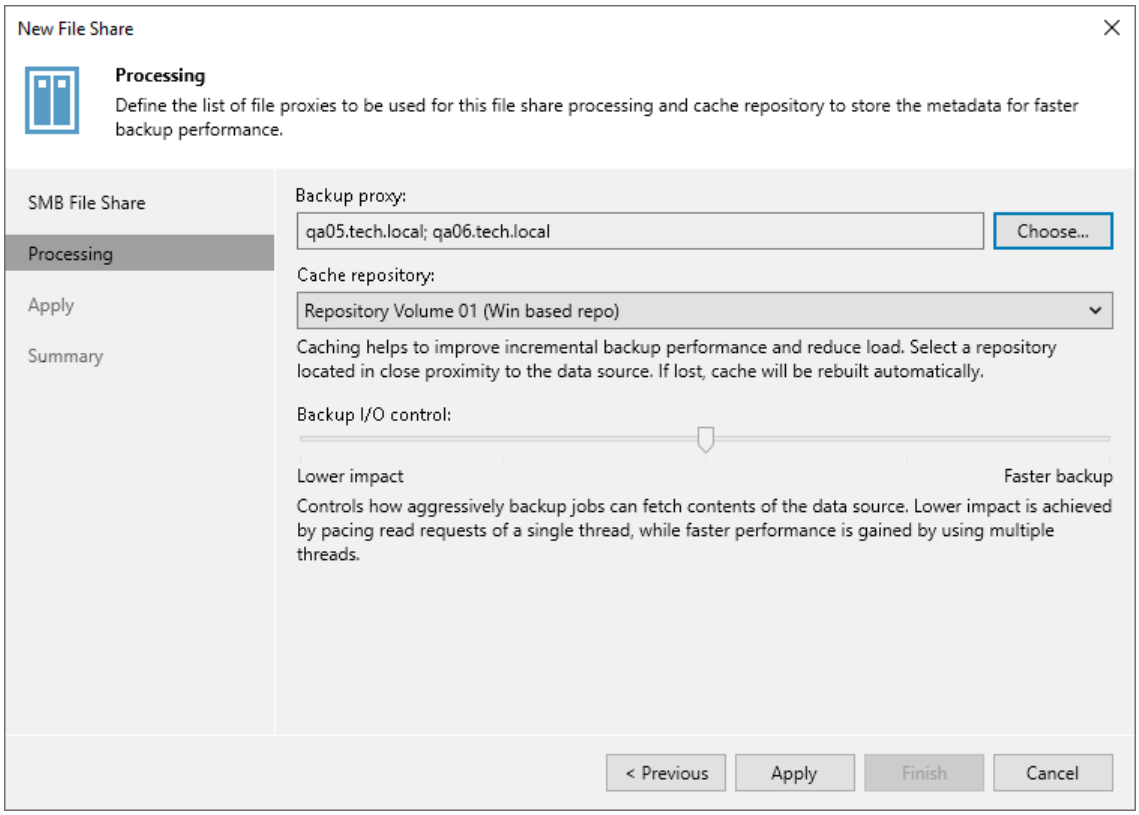

## Step 5. Apply File Share Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components and adds the SMB file share to the backup infrastructure. Click **Next** to proceed.

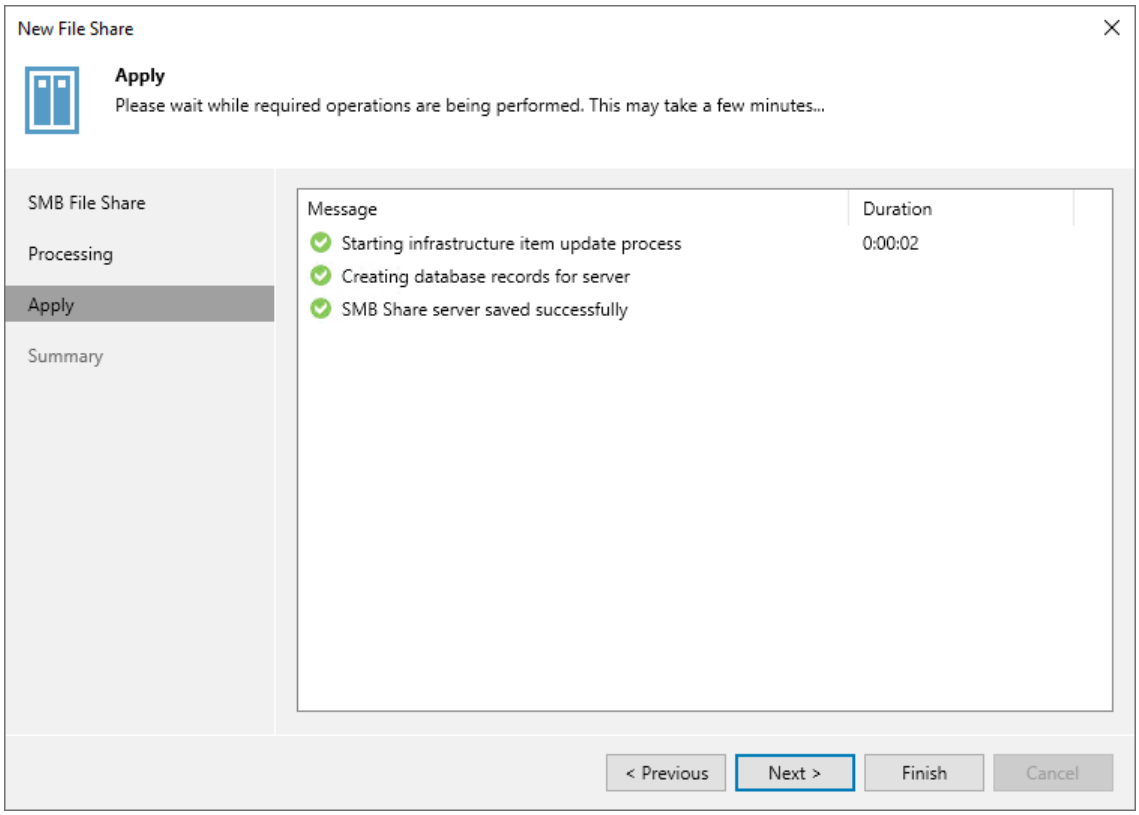

### Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added SMB share and click **Finish** to exit the wizard.

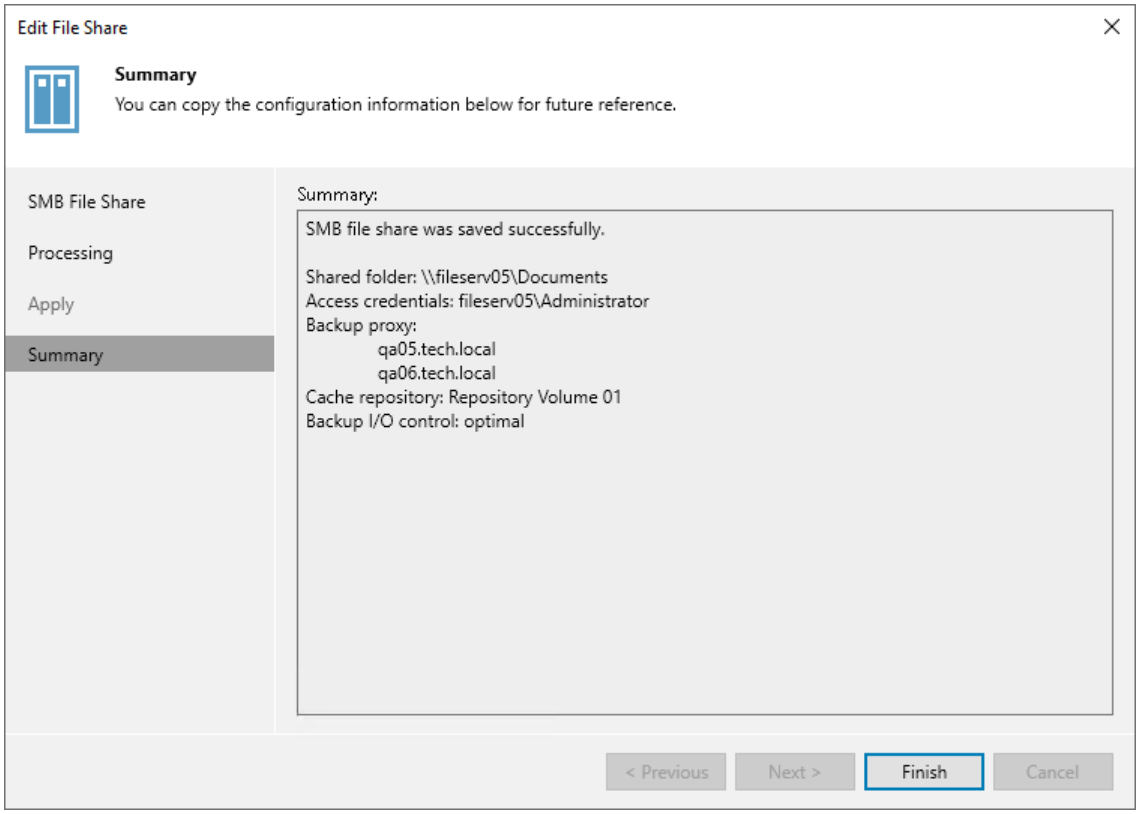


## Adding NFS File Share

Before you add an NFS file share to the inventory of the virtual infrastructure, consider the following:

- The file share meets requirements listed in Platform Support.
- If you plan to use a dedicated backup proxy server or cache repository, make sure these components are added in **Backup Infrastructure**.

## Step 1. Launch New File Share Wizard

To launch the **New File Share** wizard:

1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **File Shares** node and select **Add File Share**.
   - Select the **File Shares** node and click **Add File Share** on the ribbon.
   - Select the **File Shares** node and click **Add File Share** in the working area.
3. In the **Add File Share** window, click **NFS share**.

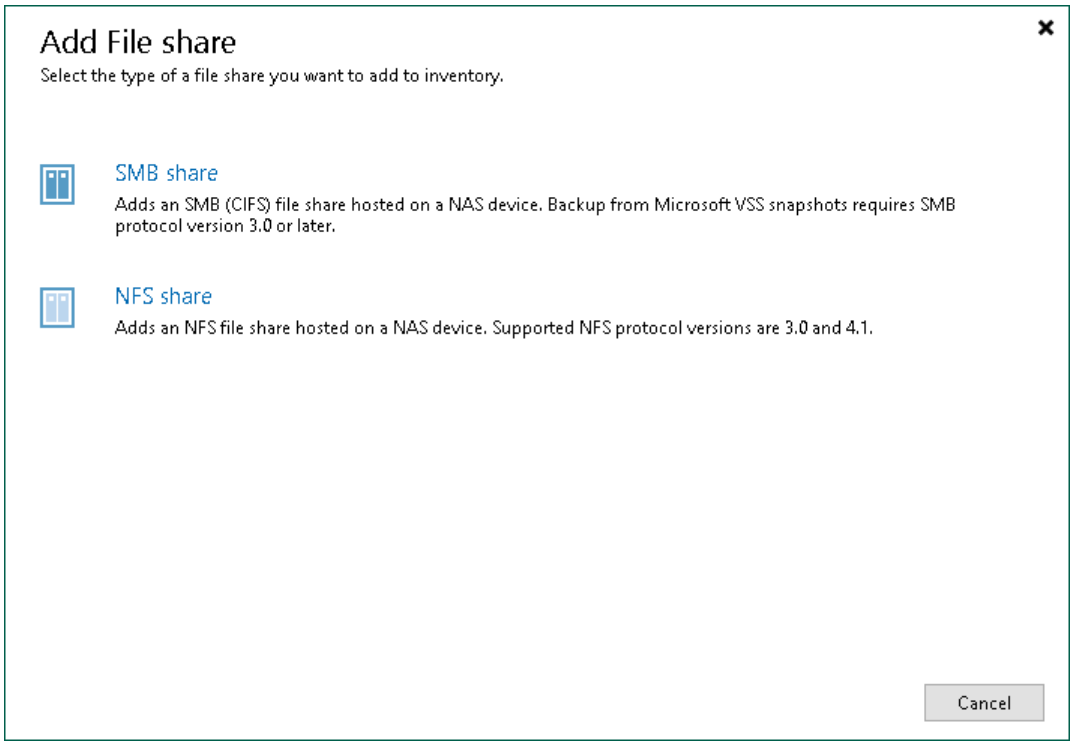

## Step 2. Specify Path to NFS File Share

At the **NFS File Share** step of the wizard, specify the path to an NFS file share in the *server:/folder* format.

You can add the root server folder in the *server:/format* to protect all NFS file shares residing on this server. You can specify the IPv4 or IPv6 address of the server. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in the IPv6 Support section.

After that, create a single file backup job to protect the added server, as described in the Creating File Backup Jobs section. Then all NFS file shares added on this server will be automatically processed with the file backup job and protected. If you previously had several separate non-root shared folders residing on the same server and want to switch to using a single root shared folder to cover the same shares, you do not have to run full backups to update data of protected shares. Instead, you can convert existing backups and update existing file backup jobs to protect single root shared folders comprising all other non-root shared folders residing on the same server. To learn more about the conversion, see the Converting Backups from Non-Root to Root Shared Folders section. Perform the conversion with extreme caution.

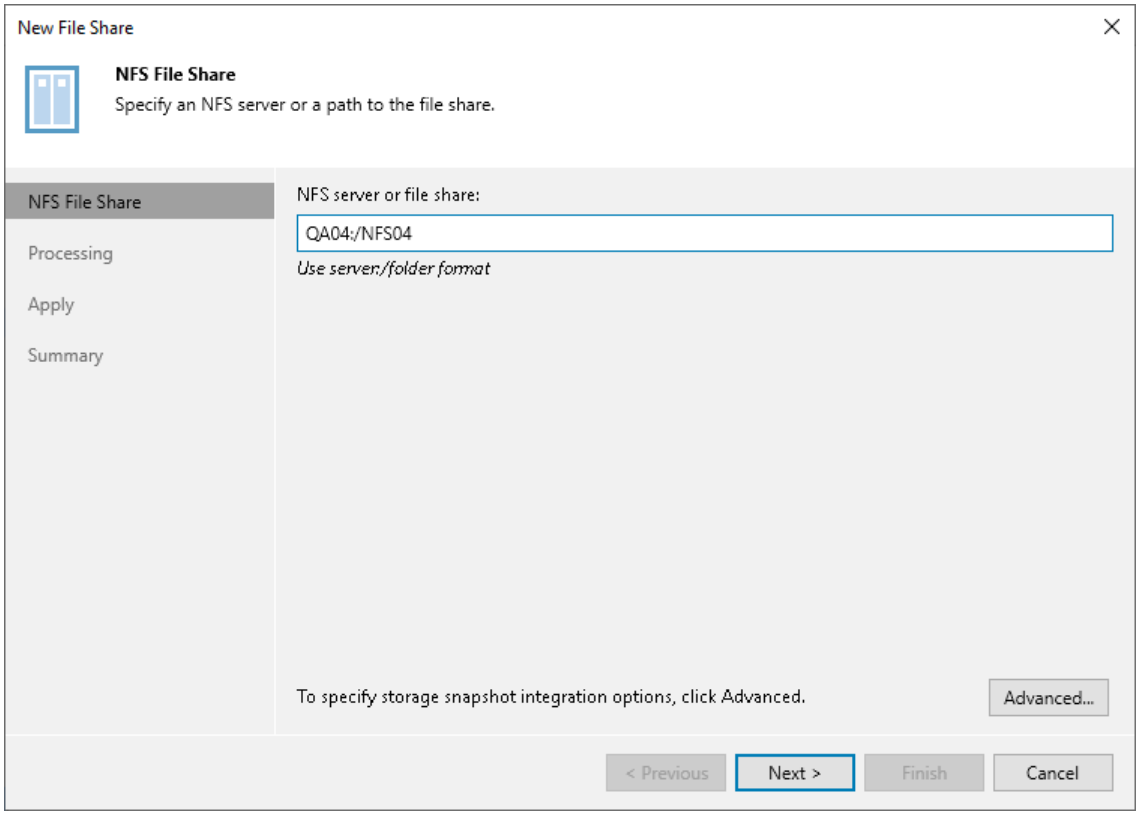

## Step 3. Specify Advanced NFS File Share Settings

You can instruct Veeam Backup & Replication to back up data from native storage snapshots. During backup jobs, Veeam Backup & Replication will read data of shared files and folders from snapshots, which speeds up backup operations and improves RPOs.

To define if Veeam Backup & Replication will use snapshots for backups:

1. At the **NFS File Share** step of the wizard, click **Advanced**.
2. In the **Advanced** window, select one of the following options:
   - To ignore the snapshot functionality, select **Backup directly from the file share**.
     Veeam Backup & Replication will ignore locked files and folders. When creating a backup job, you can configure notifications to list files and folders that are skipped during the backup procedure. For more information see the Notification Settings section.
   - To back up files from the native storage snapshot, select **Backup from a storage snapshot at the following path** and specify the path in the *server:/snapshotfolder/snapshotname* format to the snapshot stored on the NFS file share. You can specify the IPv4 or IPv6 address of the server. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section the IPv6 Support section.

     If you select this option, you can additionally use custom scripts written by you, for example, to create a snapshot before the backup and remove it after the backup. You can define these scripts when creating a new file backup job, as described in the Script Settings section.

> **NOTE**
>
> Consider the following:
>
> - Veeam Backup & Replication does not take snapshots itself, but it can use a snapshot taken by the storage system. File backup jobs do not trigger the storage snapshot creation and deletion automatically. You can specify the folder where the storage snapshot is stored. In this case file backup jobs can access this folder and read data from the storage snapshot.
> - The file backup to tape functionality does not support backup from storage snapshots for NFS file shares.

3. Select **Failover to direct backup if a snapshot is not available** if you want Veeam Backup & Replication to read data for backup directly from the file share when the snapshot is unavailable. If you do not select the option and the snapshot is unavailable, the file backup job will stop with a failure.

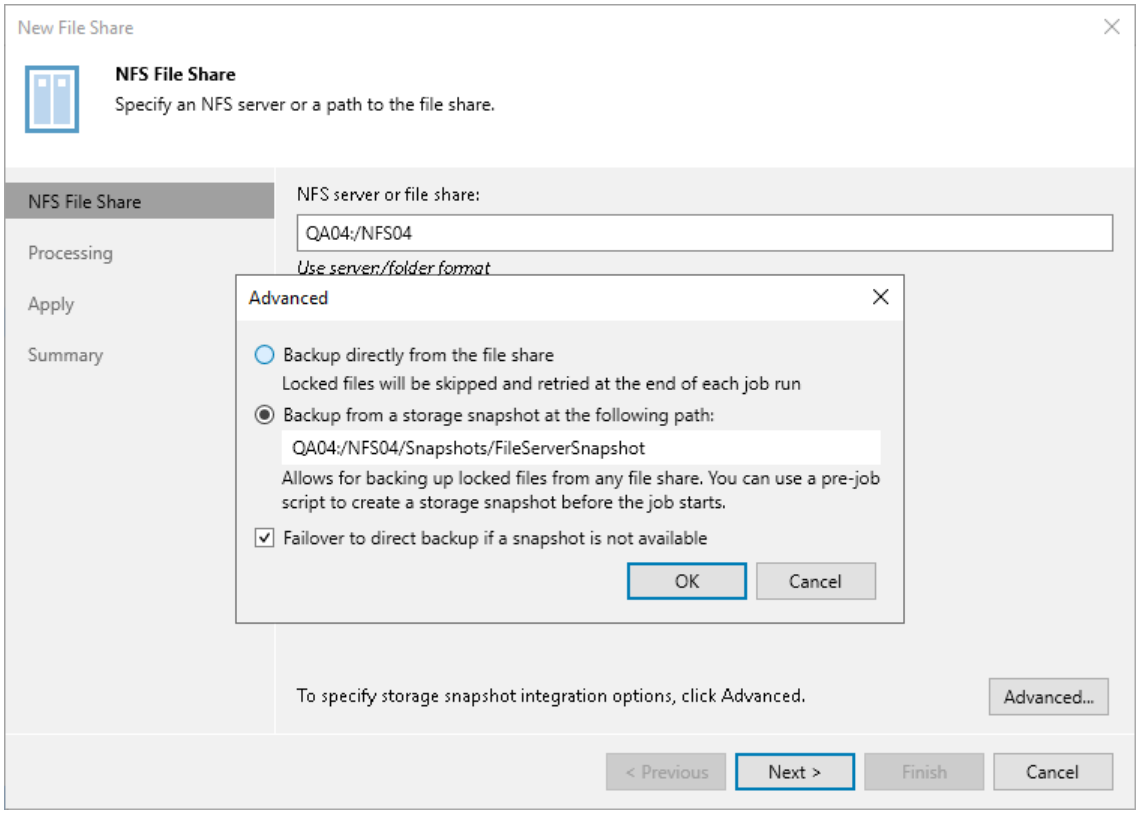

## Step 4. Specify File Share Processing Settings

At the **Processing** step of the wizard, do the following:

1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
2. In the **Backup Proxy** window, select proxy servers:
   - If you select **All proxies**, Veeam Backup & Replication will use all available backup proxies for file backup. The number of proxies in use defines the number of data threads that transfer data from the file share to the backup repository. The more data transfer threads Veeam Backup & Replication uses, the higher is the data transfer speed.

     If the file share is used as a source for a file to tape backup job, the tape server utilized for this job is added as yet another backup proxy when creating a file to tape backup job. This backup proxy has the highest priority over all others and is used by default if it has access rights to the file share. For details on file to tape backup jobs, see the File Backup to Tape section in the Veeam Backup & Replication User Guide.
   - If you select **Use the selected backup proxies only**, you can explicitly specify backup proxies that Veeam Backup & Replication must use for file backup.

     It is recommended that you select at least two backup proxies to ensure that the backup jobs start even if one of the proxies fails or loses its connectivity to the source file share. The more proxies you select, the more data transfer threads Veeam Backup & Replication will use for backup jobs, thus improving performance.

     Even if the file share is used as a source for file to tape backup jobs, Veeam Backup & Replication will use only proxies selected in the list to process the backup data traffic.

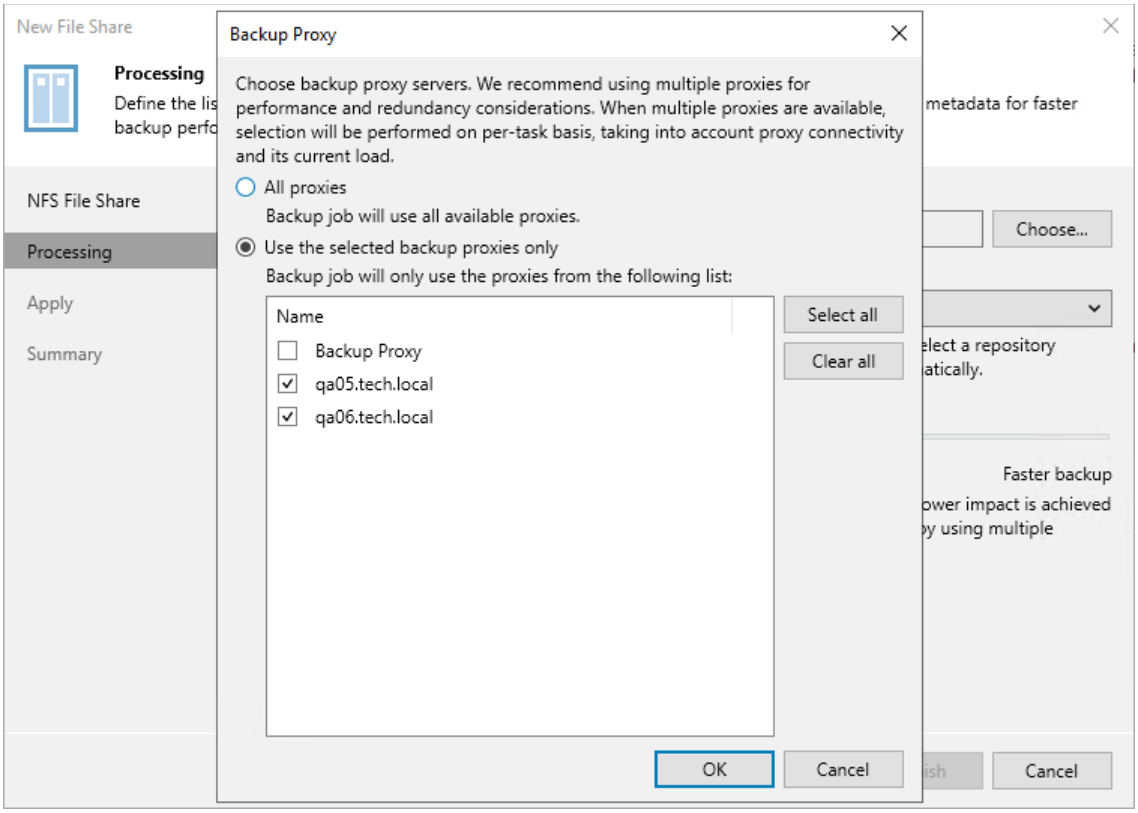

3. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located as close to the source file share as possible.

   If you change the cache repository for an existing file share whose backups are stored in object storage, Veeam Backup & Replication will prompt you to either attach migrated metadata, copy metadata from the previous cache repository, or download metadata manually from the archive repository. For more information on storing NAS backups in the object storage and changing the cache repository, see the Unstructured Data Backups in Object Storage Repositories section.

4. Use the **Backup I/O control** slider to define how fast backup proxies can read data from the source file share. This setting is based on the number of parallel threads that can be used by proxies configured for processing the file share.

| I/O Control | Number of Proxies | Threads per Task |
|---|---|---|
| **Lower Impact** | 1 | 1 |
| **Below Normal** | 1 | 4 |
| **Normal** | 2 | 8 |
| **Above Normal** | 4 | 16 |
| **Faster Backup** | Unlimited | 16 |

   If resources of your NAS device are limited, it is recommended that you select the **Lower impact** option. If your NAS device is powerful enough, select the **Faster backup** option.

5. Click **Apply** to save the configured settings.

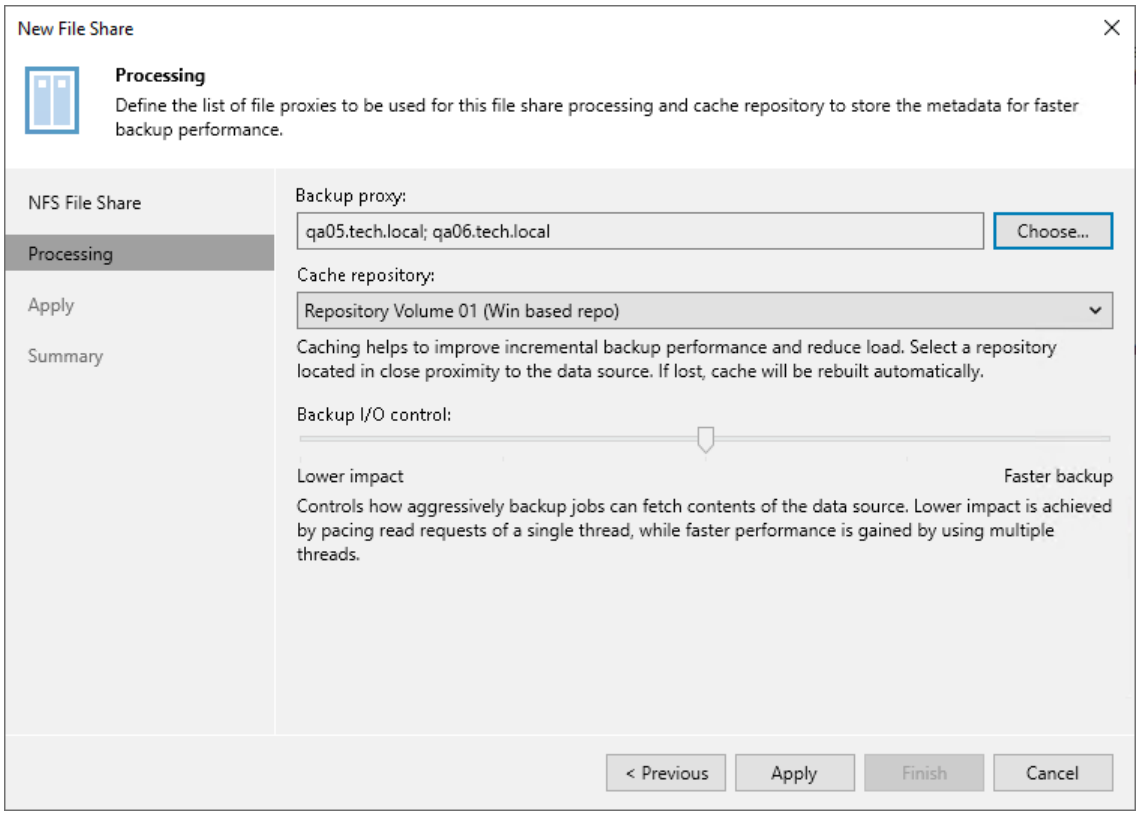

## Step 5. Apply File Share Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components and adds the NFS file share to the backup infrastructure. Click **Next** to proceed.

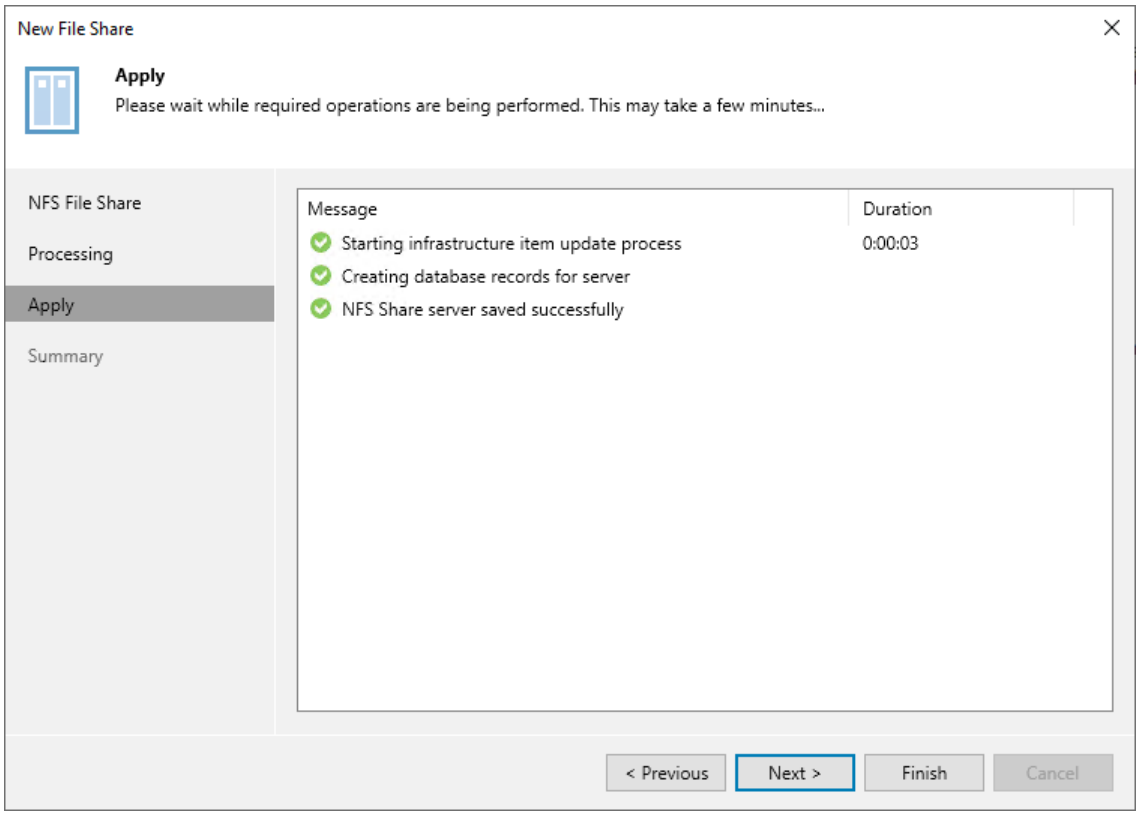

### Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added NFS share and click **Finish** to exit the wizard.

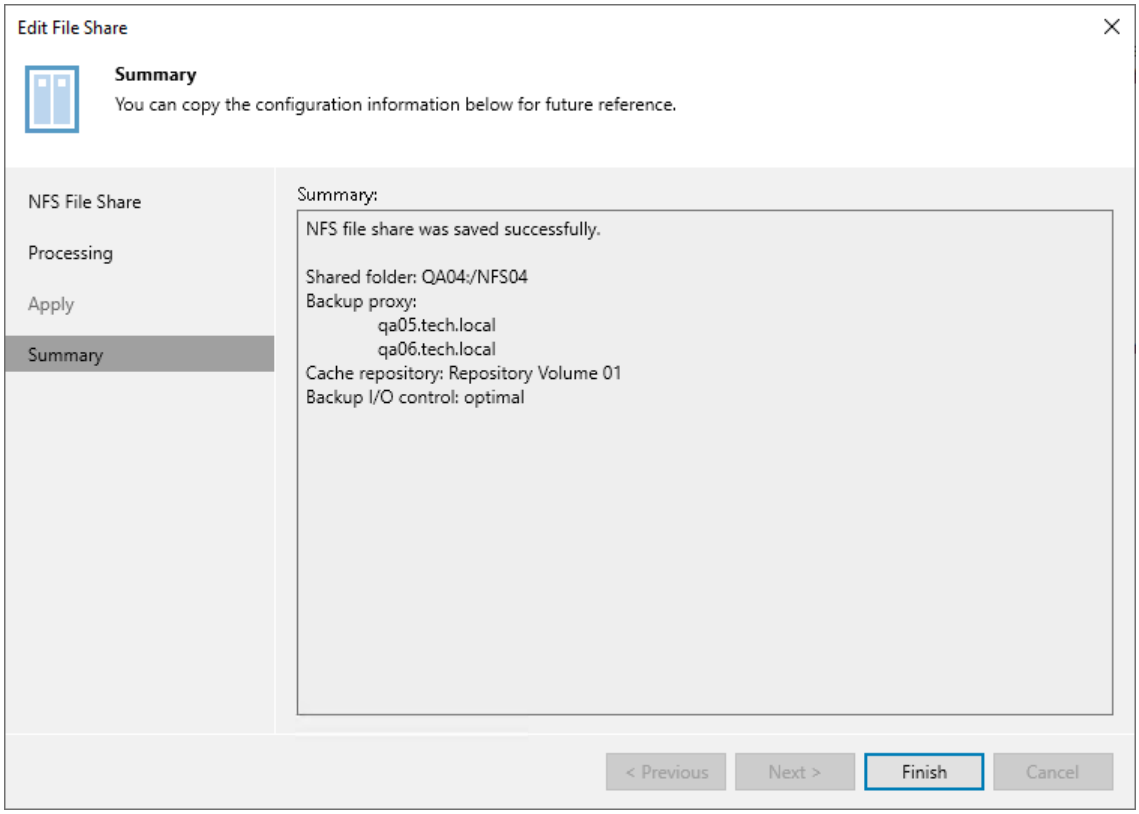


# Adding Enterprise Storage System as NAS Filer

Before you add an enterprise storage system as a NAS filer to the inventory of the virtual infrastructure, consider the following:

- The NAS device meets requirements listed in section Supported Platforms and Applications.
- If you plan to use a dedicated backup proxy or cache repository, make sure these components are added in **Backup Infrastructure**.
- When a storage virtual machine (SVM) has both the NFS and SMB protocols enabled, but the SVM does not export any NFS shares, file backup jobs configured to protect NFS file shares on that SVM fail. To fix this, either disable the NFS protocol on the SVM, or disable NFS protocol for the storage in Veeam Backup & Replication. You can do the latter by clearing the **NFS** check box at the **NAS Filer** step of the wizards described in the Adding NetApp Data ONTAP, Adding Lenovo ThinkSystem DM/DG Series, Adding Dell PowerScale (formerly Isilon) or Adding Nutanix Files Storage sections.
- Native changed files tracking (SnapDiff) is not supported for NetApp ONTAP versions from 9.10.1 to 9.10.1P10.
- Native changed files tracking (SnapDiff) is not supported for NetApp ONTAP FlexGroup volumes.

## Step 1. Launch New NAS Filer Wizard

To launch the **New NAS Filer** wizard:

1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **Unstructured Data** node and select **Add unstructured data source**.
   - Select the **Unstructured Data** node and click **Add Data Source** on the ribbon.
   - Select the **Unstructured Data** node and click **Add Data Source** in the working area.
3. In the **Add Unstructured Data Source** window, click **NAS filer**.

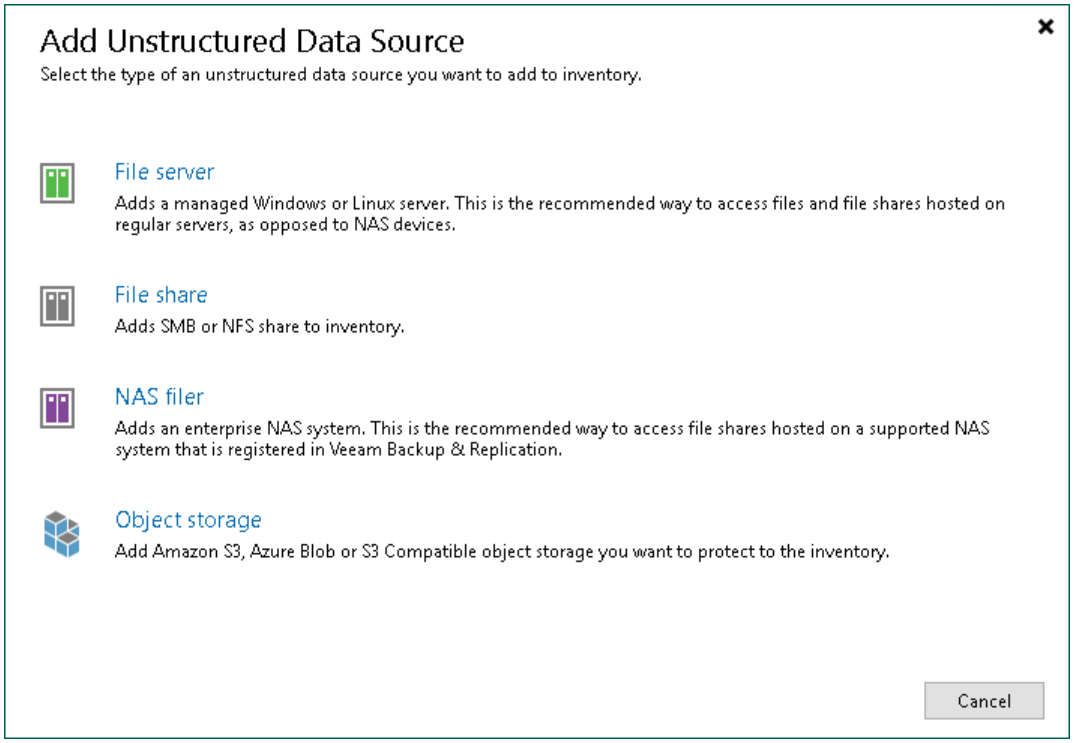

## Step 2. Select NAS Device

At the **NAS Filer** step of the wizard, choose the NAS device, which you want to use as a NAS filer where protected file shares reside.

1. Select the required NAS device from the **Select NAS filer** drop-down list.

   If the drop-down list does not display the required device, you must add it to the storage infrastructure. To add the NAS device, click **Add New** and follow the instructions described in the Adding NetApp Data ONTAP, Adding Lenovo ThinkSystem DM/DG Series, Adding Dell PowerScale, or Adding Nutanix Files Storage sections, depending on the type of the storage system you use.

2. If you must specify user credentials to access the NAS filer, select the **Use the following account to access the NAS filer** check box. From the **Credentials** drop-down list, select a credentials record for a user account that has **Full Control** permissions on the file system.

   [For Nutanix Files Storage] When adding a Nutanix Files Storage as a NAS filer, use credentials of an AD user account with administrator privileges added as a file server admin or backup admin. Otherwise, Veeam Backup & Replication cannot access SMB file shares.

   If you have not set up credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Credentials Manager.

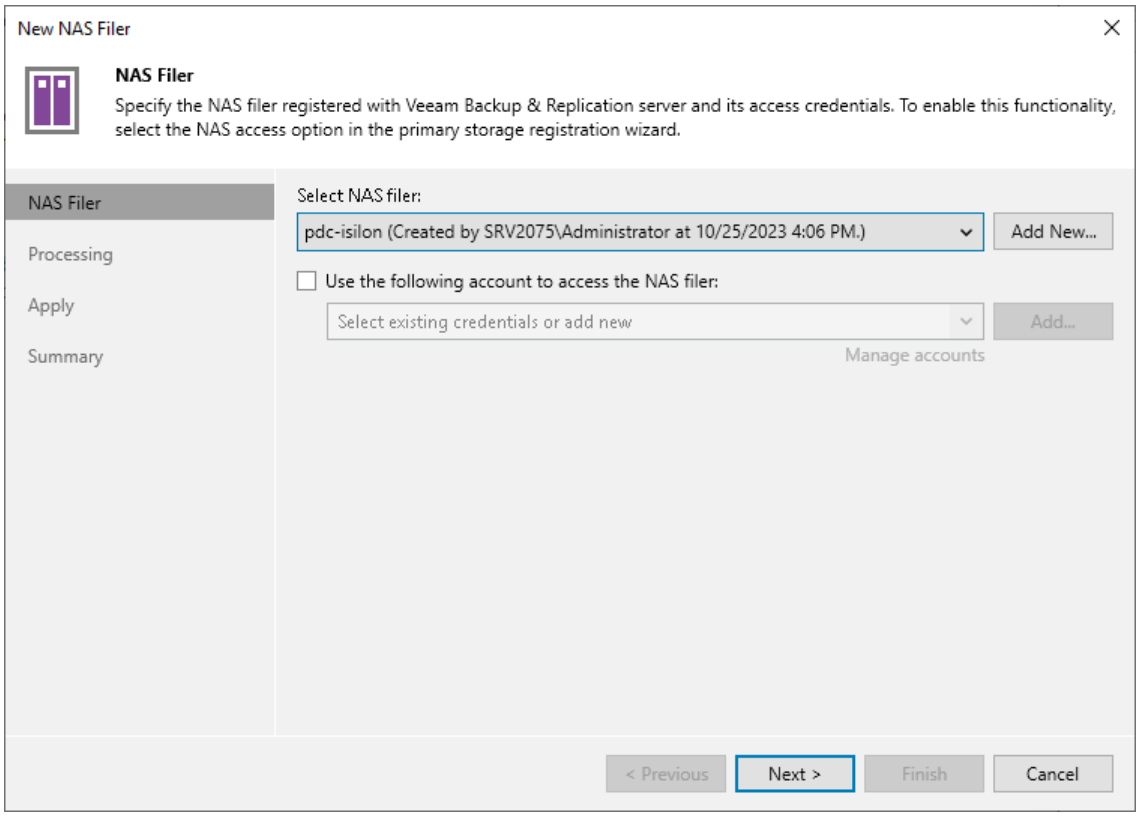


# Integration with Dell PowerScale (formerly Isilon)

To provide proper integration of Veeam Backup & Replication with Dell PowerScale (formerly Isilon) to implement the NAS backup functionality, consider the following:

- The Dell PowerScale system must be licensed for SnapshotIQ. Otherwise NAS integration will not work.

- The PowerScale service account is used to add the Dell PowerScale system to Veeam Backup & Replication, as described in the Adding Dell PowerScale section in the Storage System Snapshot Integration Guide. This account is used only for registering the storage system in the backup infrastructure and further storage rescans. You cannot use it for NAS file share backup and restore.
- Starting from Dell PowerScale 9.5, Veeam Backup & Replication supports SmartConnect.
- We recommend that you create a custom **VeeamStorageIntegration** role to group all necessary privileges. This role must be created in the System access zone. For more information on access zones, see the Dell documentation.

  Add the following privileges to the **VeeamStorageIntegration** role:

| Privilege ID | Read Only | Comments |
|---|---|---|
| ISI_PRIV_LOGIN_PAPI | True | Used to log in to the Platform API and the web administration interface. |
| ISI_PRIV_AUTH | True | Used to configure external authentication providers, including root-level accounts. |
| ISI_PRIV_DEVICES | True | Used to create new roles and assign privileges. |
| ISI_PRIV_JOB_ENGINE R/W | False | Used for array Change File Tracking report scheduling for changelist API call. |
| ISI_PRIV_NETWORK | True | Used to configure network interfaces. |
| ISI_PRIV_NFS R/W | False | Used to update export rules. |
| ISI_PRIV_SMB | True | Used to configure the SMB server. |
| ISI_PRIV_SNAPSHOT R/W | False | Used to schedule, take, and view snapshots. |

  For more information on privileges, see the PowerScale documentation.

  You can use the following command to check privileges added to the **VeeamStorageIntegration** role:

```
isi auth roles view VeeamStorageIntegration
```

  Make sure that you make the service account to be used for NAS integration a member of the **VeeamStorageIntegration** role.
- To enable backup and restore of SMB file shares, additionally do the following:
  - Grant minimum read access to SMB file shares for the account that will be used for NAS integration. That will allow backing up SMB file shares.

- Grant "run as root" access to SMB file shares for the account that will be used for NAS integration. That will allow restoring SMB file shares.

Consider that the PowerScale BackupAdmin role is not enough for backup and restore. It is mainly used for SMB access and enables backup and restore of files from `/ifs`. It does not propagate down to subfolders and therefore cannot be used to backup anything, but the system access zone.

- To correctly use SmartConnect for NAS backup, consider the following:

  a. Enable SmartConnect on the storage system and configure DNS as described in the Dell documentation.

  b. Add the storage system with the SmartConnect service IP (SSIP) to the backup server.

  c. Add general-purpose backup proxies to the backup infrastructure. We recommend adding as many general-purpose backup proxies as many Dell Powerscale storage nodes you plan to use for backup.

  d. If SmartConnect updates DNS infrequently, general-purpose backup proxies can end up working with the same storage node. This can happen because SmartConnect only provides the IP address of that node. To avoid this situation, do the following:

     i. Open a support ticket and ask for the registry value. This registry value delays the start of the backup job for some seconds for each general-purpose backup proxy. This allows SmartConnect to distribute different IP addresses for the proxies.

     ii. [If you still have issues after setting key value] Distribute the load manually. For this, edit the hosts file of each general-purpose backup proxy and map the cluster FQDN name to a storage node IP address. For each proxy, use a different storage node IP.

Dell PowerScale SmartConnect helps distribute the load among storage nodes: SmartConnect provides clients different IP addresses from different nodes to work with the storage cluster in parallel. If you do not use SmartConnect, Veeam Backup & Replication will use single IP address during backup, which can cause significant delays.

## Step 3. Specify File Share Processing Settings

At the **Processing** step of the wizard, define NAS filer processing settings:

1. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located as close to the NAS filer as possible.

   > **NOTE**
   >
   > Consider that you cannot use a Linux-based server as a cache repository to process content of the protected shares on enterprise storage systems if you enable the **Use native changed files tracking** option.

   If you change the cache repository for an NAS filer whose backups are stored in the object storage, Veeam Backup & Replication will prompt you to either attach migrated metadata, copy metadata from the previous cache repository, or download metadata manually from the archive repository. For more information on storing file backups in the object storage and changing the cache repository, see the Unstructured Data Backups in Object Storage Repositories section.

2. Use the **Backup I/O control** slider to define how fast the backup proxy can read data from the source NAS filer. This setting is based on the number of parallel threads that can be used by the proxy configured for processing the file shares on the NAS filer. If resources of your storage system are limited, it is recommended that you select the **Lower impact** option. If your storage system is powerful enough, select the **Faster backup** option.

3. Select the **Use native changed files tracking** check box if you want to use the file change tracking technology provided by the storage system manufacturer.

4. Click **Next** to save the configure settings.

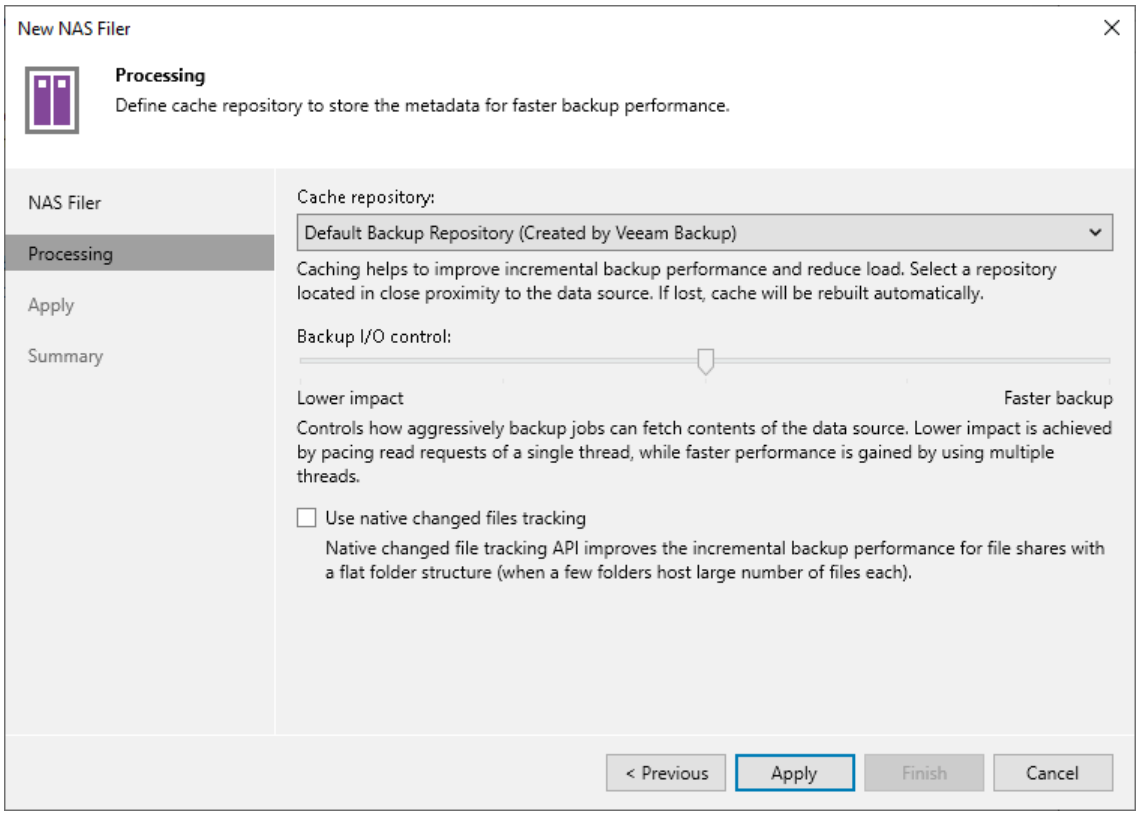

## Step 4. Apply File Share Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components. Click **Next** to complete the procedure of adding the storage system as a NAS filer.

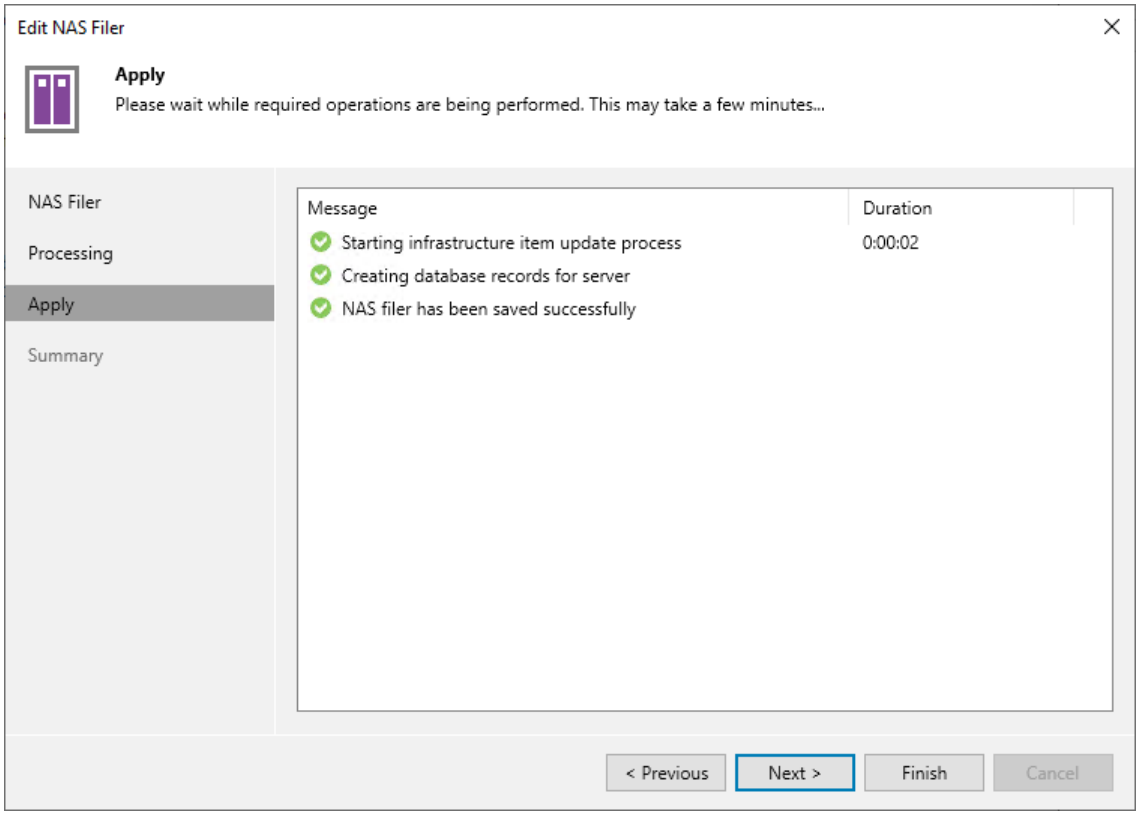

## Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the storage system added as a NAS filer and click **Finish** to exit the wizard.

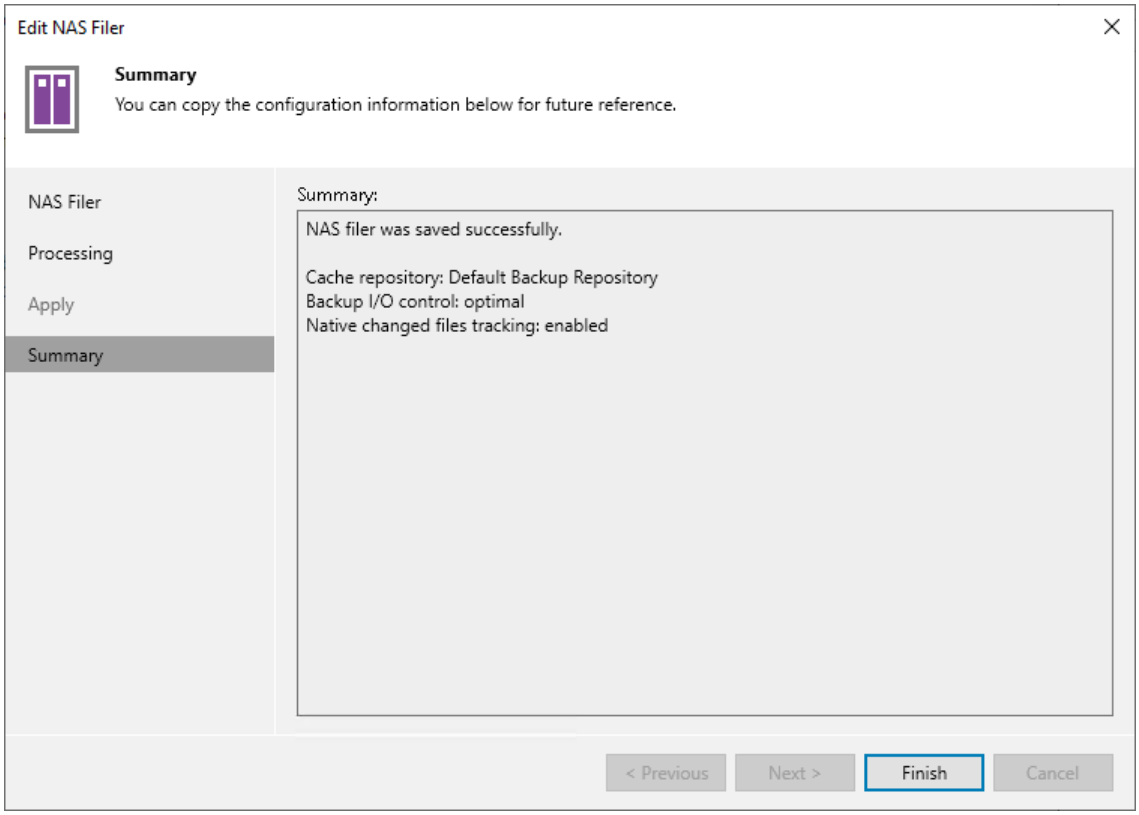


# Adding Object Storage

You can add various object storage systems as sources of unstructured data available for protection by Veeam Backup & Replication.

## Adding S3 Compatible Object Storage

Before you add an S3 compatible object storage to the inventory of the virtual infrastructure, consider the following:

- The object storage meets requirements and limitations listed in the Platform Support section.
- If you plan to use dedicated proxy servers, make sure these components are added in the Backup Infrastructure.

## Step 1. Launch New Object Storage Wizard

To launch the **New Object Storage** wizard:

1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **Unstructured Data** node and select **Add unstructured data source**.
   - Select the **Unstructured Data** node and click **Add Data Source** on the ribbon.
   - Select the **Unstructured Data** node and click **Add Data Source** in the working area.
   - If at least one object storage is added as the unstructured data source, open the **Object Storage** node, right-click empty-space in the working area and click **Add object storage**. Alternatively, you can right-click the **Object Storage** node, and click **Add object storage**.
3. In the **Add Unstructured Data Source** window, select **Object storage** > **S3 Compatible**.

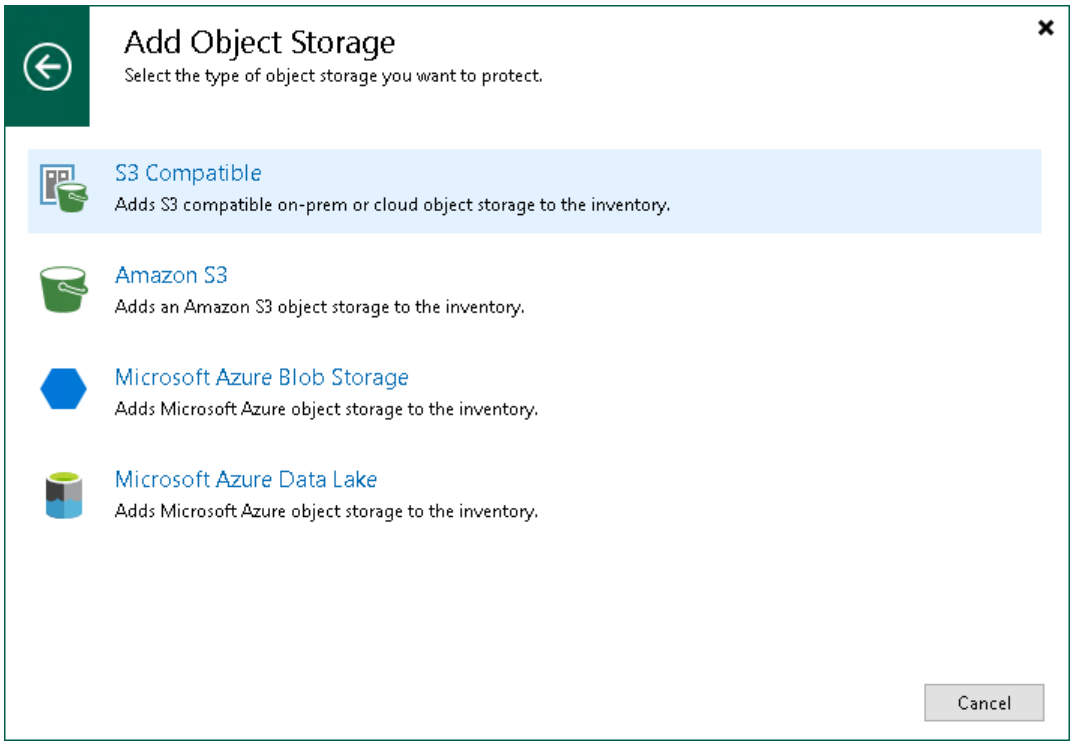

## Step 2. Specify Account Settings

At the **Account** step of the wizard, specify a friendly name and connection settings of your object storage:

1. In the **Friendly name** field, specify a name you want to assign to your object storage. This name will display in the list of you object storage repositories in the inventory of the virtual infrastructure.
2. In the **Service point** field, specify a service point address of your object storage.
3. From the **Credentials** drop-down list, select user credentials to access your S3 compatible object storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.
4. In the **Region** field, specify a region.

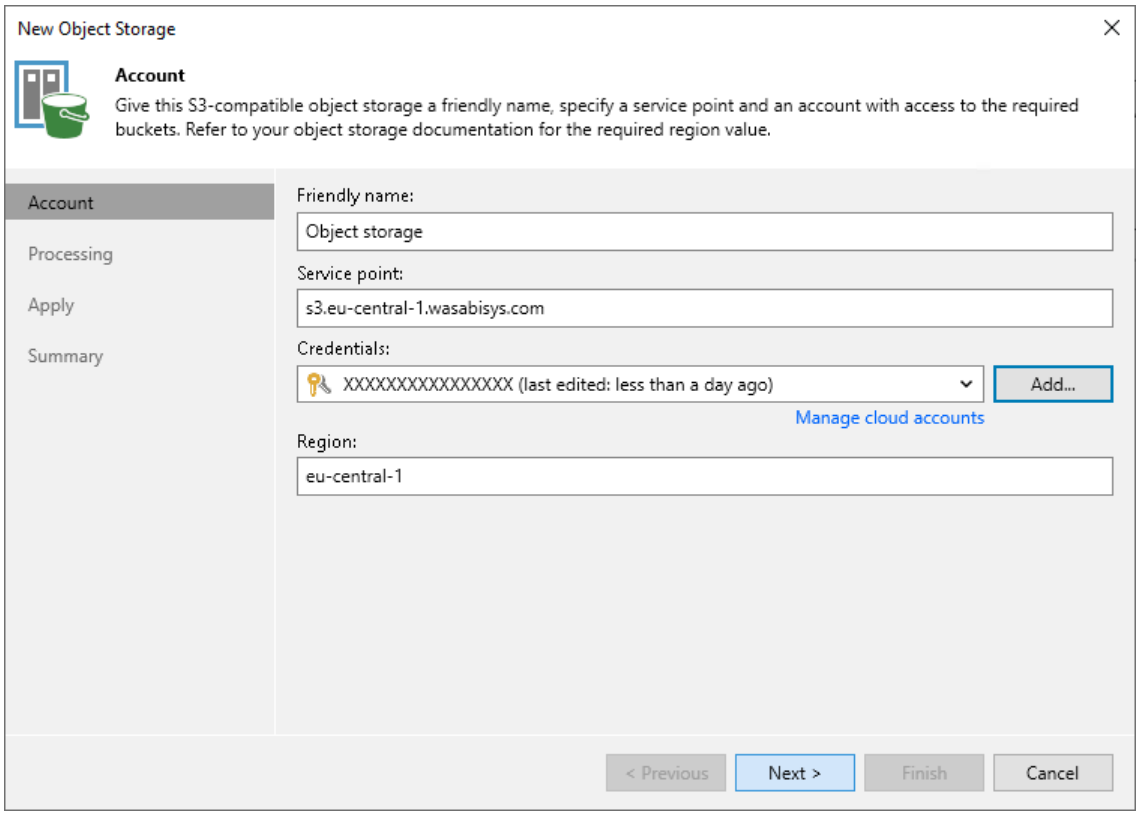

## Step 3. Specify Object Storage Processing Settings

At the **Processing** step of the wizard, do the following:

1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
2. In the **Backup Proxy** window, select proxy servers:
   - If you select **All proxies**, Veeam Backup & Replication will use all available backup proxies for the object storage backup. The number of proxies in use defines the number of data threads that transfer data from the object storage to the backup repository. The more data transfer threads Veeam Backup & Replication uses, the higher is the data transfer speed.

     If the object storage is used as a source for an object to tape backup job, the tape server utilized for this job is added as yet another backup proxy when creating an object to tape backup job. This backup proxy has the highest priority over all others and is used by default if it has access rights to the object storage. For details on object to tape backup jobs, see the Object Storage Backup to Tape section in the Veeam Backup & Replication User Guide.
   - If you select **Use the selected backup proxies only**, you can explicitly specify backup proxies that Veeam Backup & Replication must use for the object storage backup.

     It is recommended that you select at least two backup proxies to ensure that the backup jobs start even if one of the proxies fails or loses its connectivity to the source object storage. The more proxies you select, the more data transfer threads Veeam Backup & Replication will use for backup jobs, thus improving performance.

     Even if the object storage is used as a source for object to tape backup jobs, Veeam Backup & Replication will use only proxies selected in the list to process the backup data traffic.

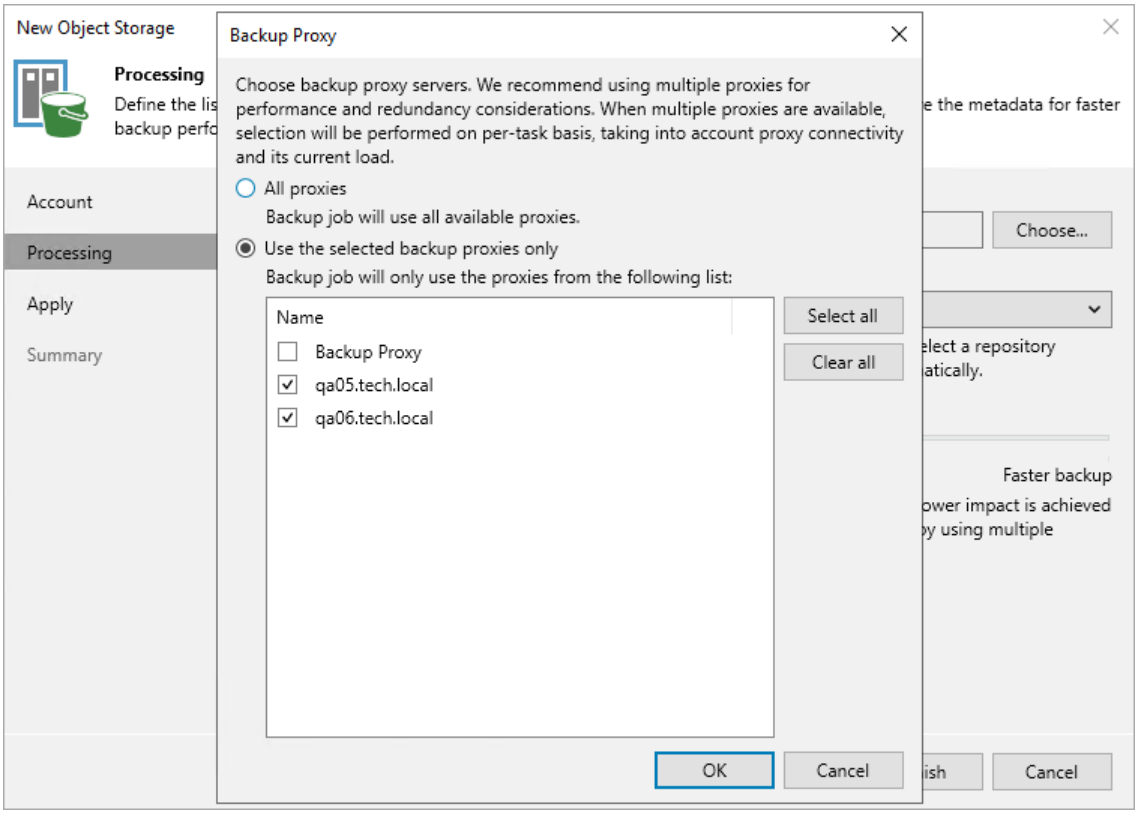

3. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located in the close proximity to the source object storage and backup proxies.

   If you change the cache repository for an existing object storage whose backups are stored in another object storage, Veeam Backup & Replication will prompt you to either attach migrated metadata, copy metadata from the previous cache repository, or download metadata manually from the archive repository. For more information, see the Unstructured Data Backups in Object Storage Repositories section.

4. Use the **Backup I/O control** slider to define how fast backup proxies can read data from the source object storage. This setting is based on the number of parallel threads that can be used by proxies configured for processing the object storage.

| I/O Control | Number of Proxies | Threads per Task |
|---|---|---|
| **Lower Impact** | 1 | 1 |
| **Below Normal** | 1 | 4 |
| **Normal** | 2 | 8 |
| **Above Normal** | 4 | 16 |
| **Faster Backup** | Unlimited | 16 |

If resources of your object storage source are limited, it is recommended that you select the **Lower impact** option. If your object storage source is powerful enough, select the **Faster backup** option.

5. Click **Apply** to save the configured settings.

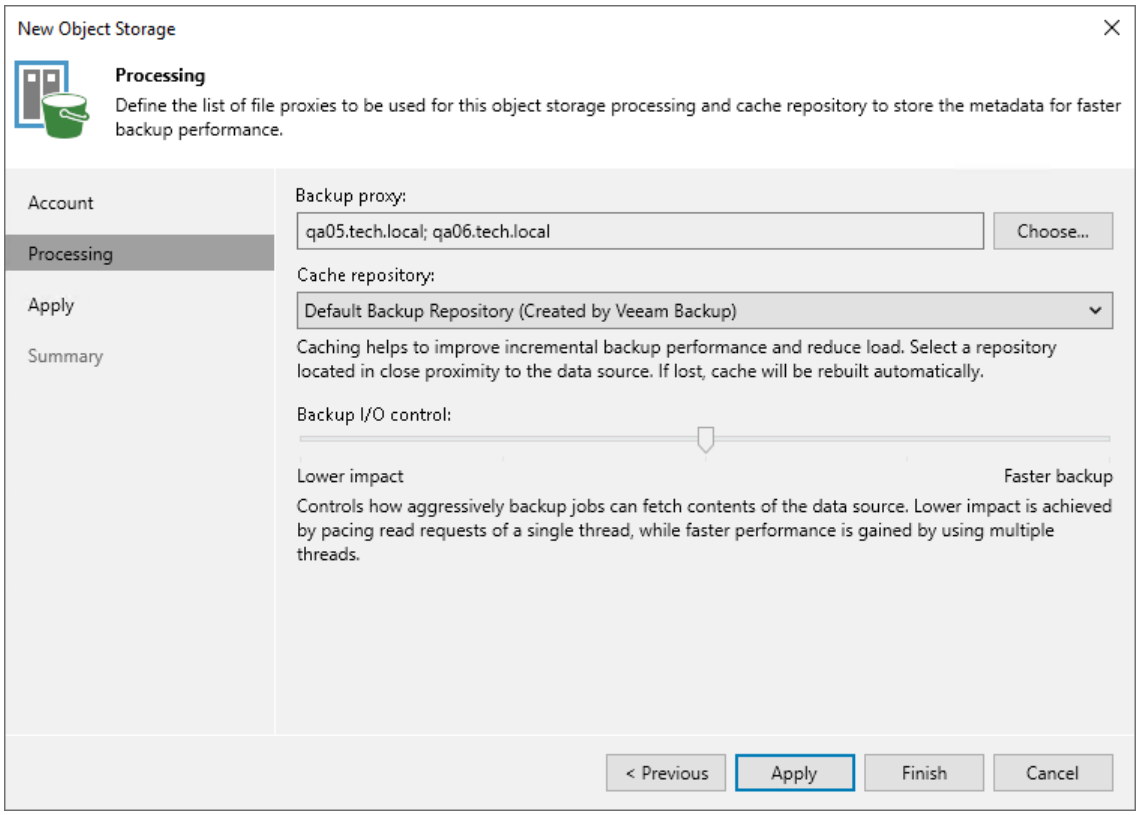

## Step 4. Apply Object Storage Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components and adds the object storage to the inventory of the virtual infrastructure. Click **Next** to proceed.

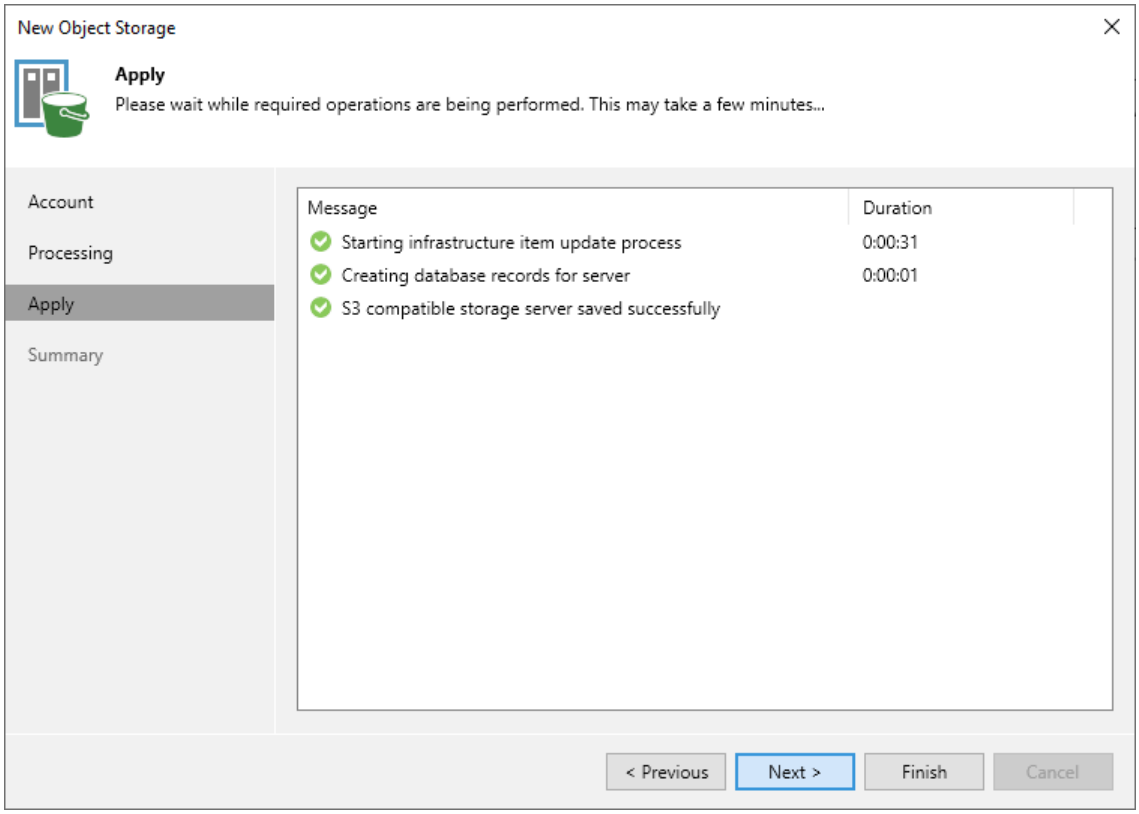

### Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added object storage and click **Finish** to exit the wizard.

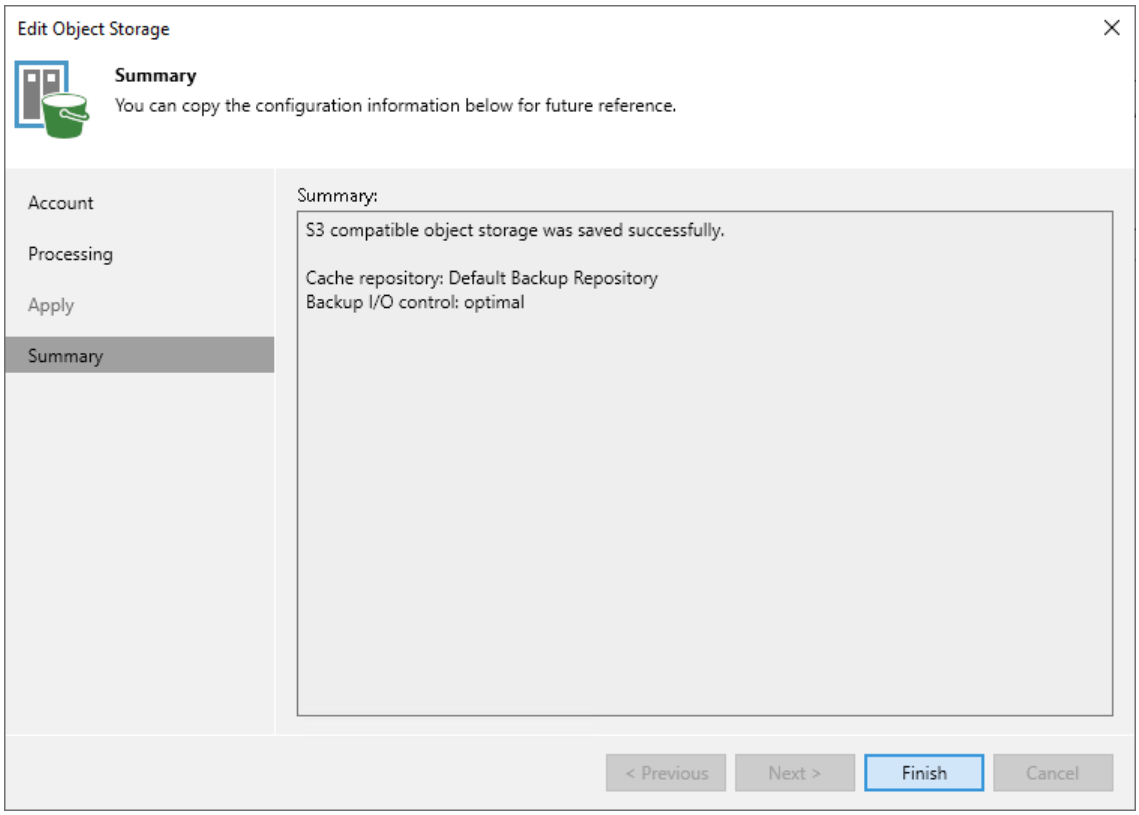


## Adding Amazon S3 Object Storage

Before you add an Amazon S3 object storage to the inventory of the virtual infrastructure, consider the following:

- The object storage meets requirements and limitations listed in the Platform Support section.
- If you plan to use dedicated proxy servers, make sure these components are added in the Backup Infrastructure.

## Step 1. Launch New Object Storage Wizard

To launch the **New Object Storage** wizard:

1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **Unstructured Data** node and select **Add unstructured data source**.
   - Select the **Unstructured Data** node and click **Add Data Source** on the ribbon.
   - Select the **Unstructured Data** node and click **Add Data Source** in the working area.
   - If at least one object storage is added as the unstructured data source, open the **Object Storage** node, right-click empty-space in the working area and click **Add object storage**. Alternatively, you can right-click the **Object Storage** node, and click **Add object storage**.
3. In the **Add Unstructured Data Source** window, select **Object storage** > **Amazon S3**.

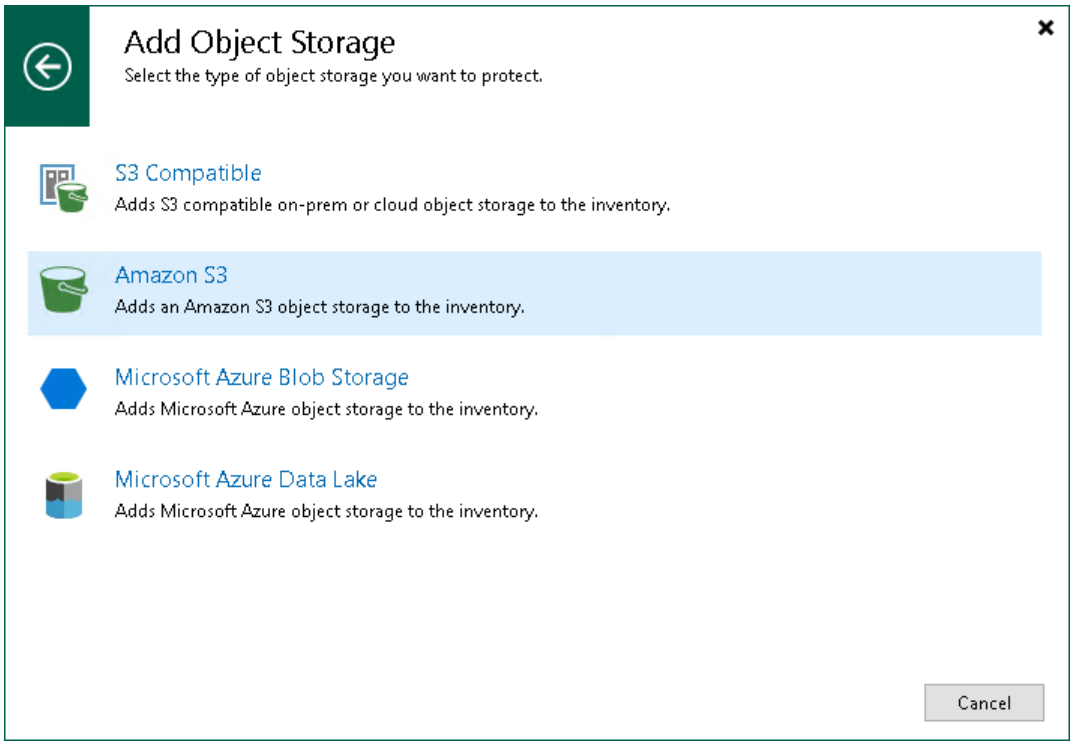

## Step 2. Specify Account Settings

At the **Account** step of the wizard, specify a friendly name and connection settings of your object storage:

1. In the **Friendly name** field, specify a name you want to assign to your object storage. This name will display in the list of you object storage repositories in the inventory of the virtual infrastructure.
2. From the **Credentials** drop-down list, select user credentials to access your Amazon S3 object storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.
3. From the **AWS region** drop-down list, select the AWS region where the Amazon S3 bucket is located.
4. From the **Data center region** drop-down list, select a region.

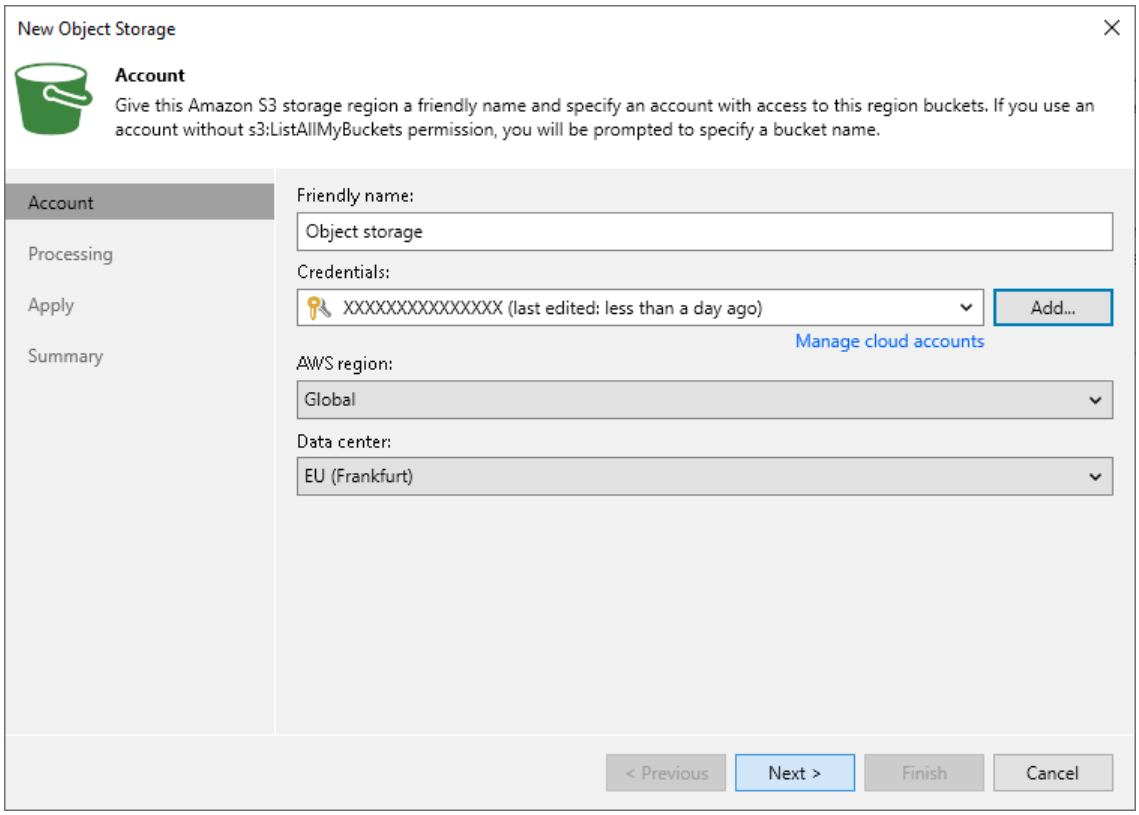

## Step 3. Specify Object Storage Processing Settings

At the **Processing** step of the wizard, do the following:

1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
2. In the **Backup Proxy** window, select proxy servers:
   - If you select **All proxies**, Veeam Backup & Replication will use all available backup proxies for the object storage backup. The number of proxies in use defines the number of data threads that transfer data from the object storage to the backup repository. The more data transfer threads Veeam Backup & Replication uses, the higher is the data transfer speed.

     If the object storage is used as a source for an object to tape backup job, the tape server utilized for this job is added as yet another backup proxy when creating an object to tape backup job. This backup proxy has the highest priority over all others and is used by default if it has access rights to the object storage. For details on object to tape backup jobs, see the Object Storage Backup to Tape section in the Veeam Backup & Replication User Guide.
   - If you select **Use the selected backup proxies only**, you can explicitly specify backup proxies that Veeam Backup & Replication must use for the object storage backup.

     It is recommended that you select at least two backup proxies to ensure that the backup jobs start even if one of the proxies fails or loses its connectivity to the source object storage. The more proxies you select, the more data transfer threads Veeam Backup & Replication will use for backup jobs, thus improving performance.

     Even if the object storage is used as a source for object to tape backup jobs, Veeam Backup & Replication will use only proxies selected in the list to process the backup data traffic.

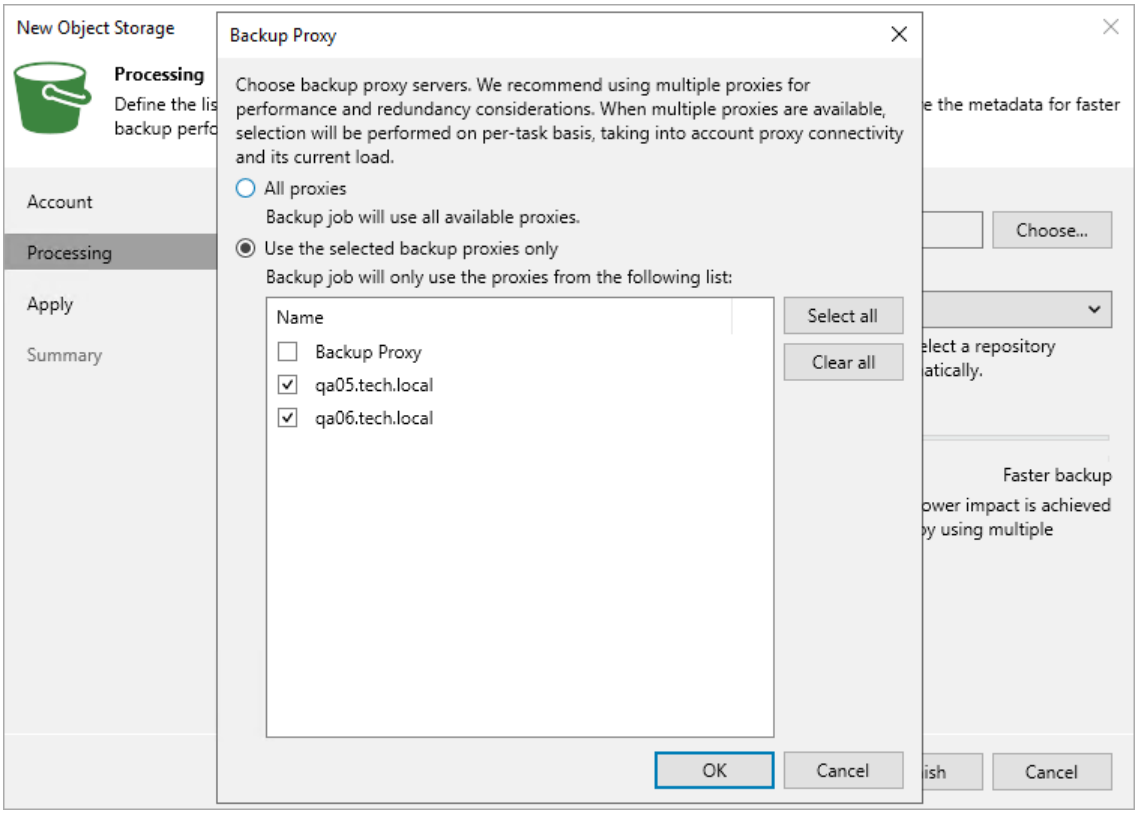

3. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located in the close proximity to the source object storage and backup proxies.

   If you change the cache repository for an existing object storage whose backups are stored in another object storage, Veeam Backup & Replication will prompt you to either attach migrated metadata, copy metadata from the previous cache repository, or download metadata manually from the archive repository. For more information, see the Unstructured Data Backups in Object Storage Repositories section.

4. Use the **Backup I/O control** slider to define how fast backup proxies can read data from the source object storage. This setting is based on the number of parallel threads that can be used by proxies configured for processing the object storage.

| I/O Control | Number of Proxies | Threads per Task |
|---|---|---|
| **Lower Impact** | 1 | 1 |
| **Below Normal** | 1 | 4 |
| **Normal** | 2 | 8 |
| **Above Normal** | 4 | 16 |
| **Faster Backup** | Unlimited | 16 |

   If resources of your object storage source are limited, it is recommended that you select the **Lower impact** option. If your object storage source is powerful enough, select the **Faster backup** option.

5. Click **Apply** to save the configured settings.

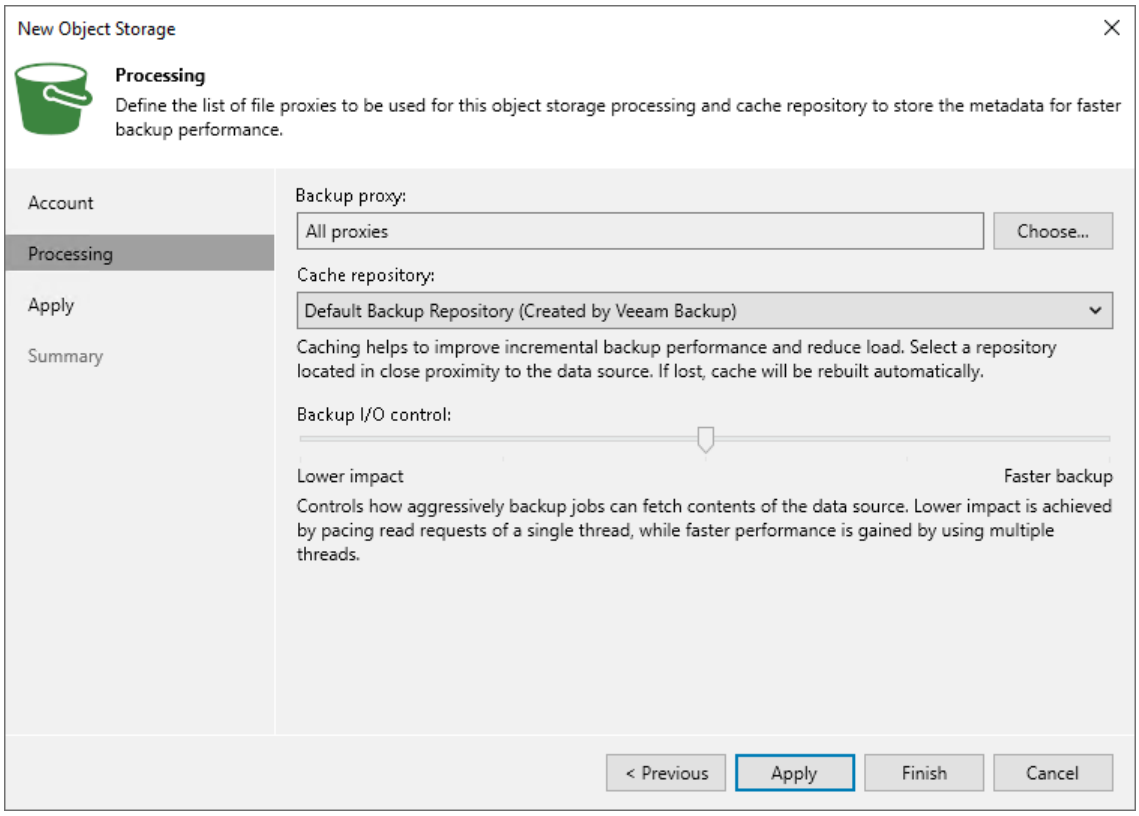

## Step 4. Apply Object Storage Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components and adds the object storage to the inventory of the virtual infrastructure. Click **Next** to proceed.

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components and adds the object storage to the backup infrastructure. Click **Next** to proceed.

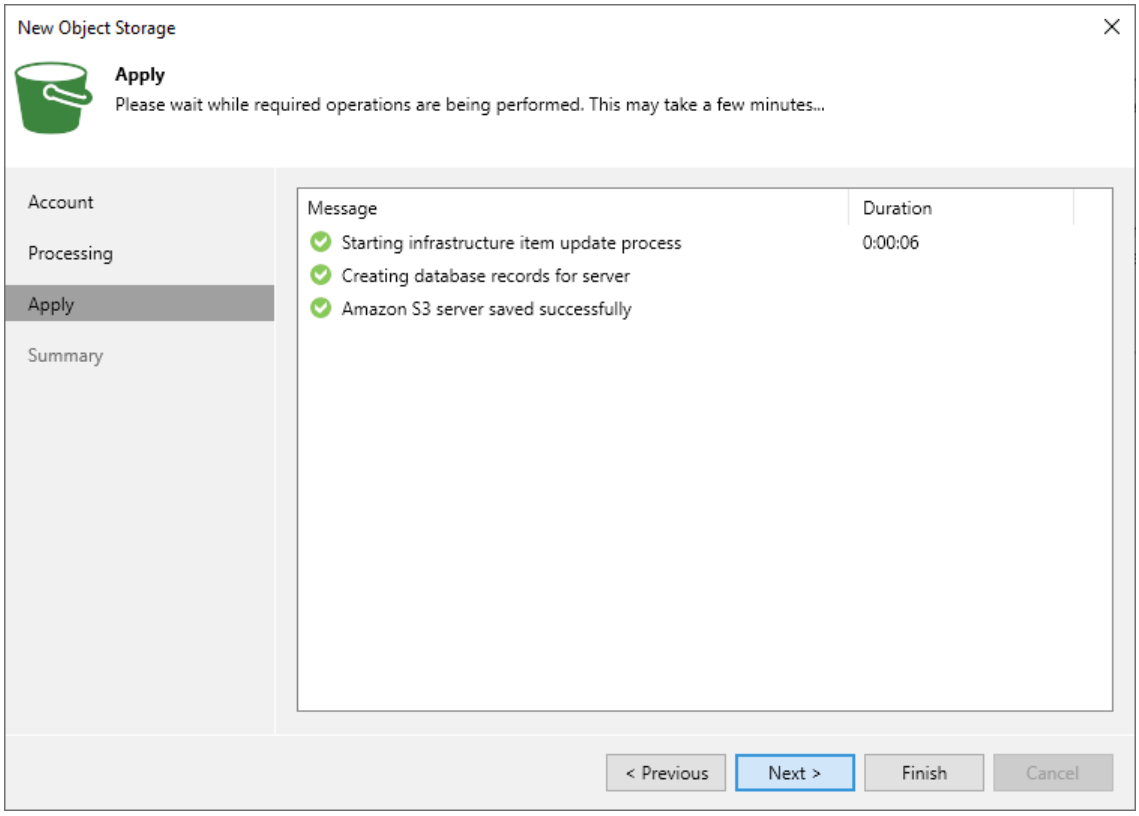

### Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added object storage and click **Finish** to exit the wizard.

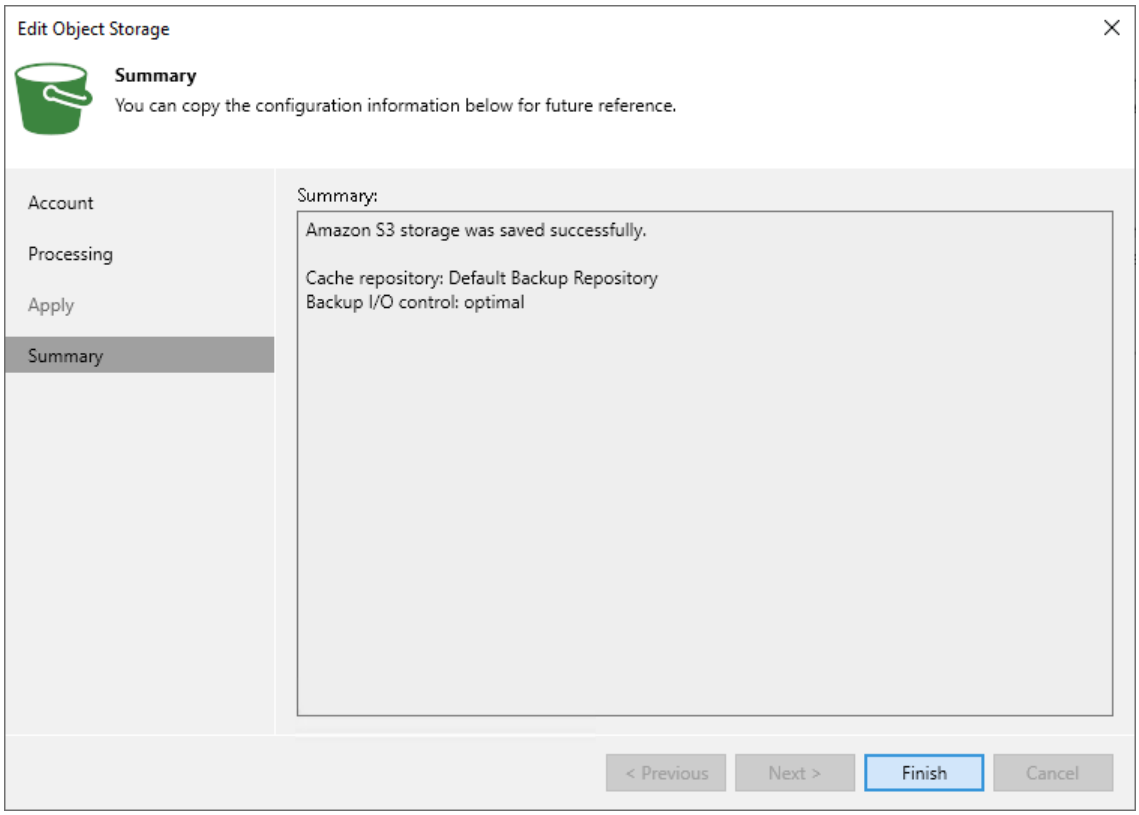


## Adding Microsoft Azure Object Storage

You can add the following types of Microsoft Azure object storage repositories:

- Microsoft Azure Blob Storage.
- Microsoft Azure Data Lake.

Before you add a Microsoft Azure object storage to the inventory of the virtual infrastructure, consider the following:

- The object storage meets requirements and limitations listed in the Platform Support section.
- If you plan to use dedicated proxy servers, make sure these components are added in the Backup Infrastructure.

## Step 1. Launch New Object Storage Wizard

To launch the **New Object Storage** wizard:

1. Open the **Inventory** view.
2. Do one of the following:
   - In the inventory pane, right-click the **Unstructured Data** node and select **Add unstructured data source**.
   - Select the **Unstructured Data** node and click **Add Data Source** on the ribbon.
   - Select the **Unstructured Data** node and click **Add Data Source** in the working area.
   - If at least one object storage is added as the unstructured data source, open the **Object Storage** node, right-click empty space in the working area and click **Add object storage**. Alternatively, you can right-click the **Object Storage** node, and click **Add object storage**.
3. In the **Add Unstructured Data Source** window, select **Object storage** and specify the type of the Microsoft Azure object storage repository:
   - **Microsoft Azure Blob Storage**
   - **Microsoft Azure Data Lake**

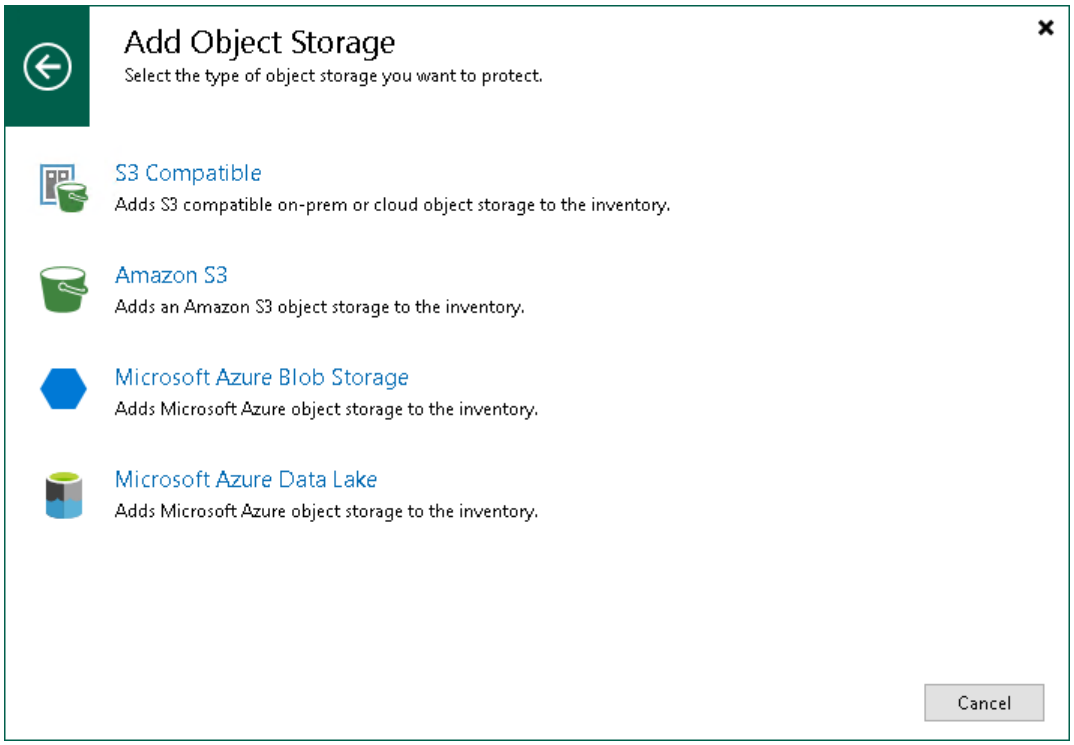

## Step 2. Specify Account Settings

At the **Account** step of the wizard, specify a friendly name and connection settings of your object storage:

1. In the **Friendly name** field, specify a name you want to assign to your object storage. This name will display in the list of you object storage repositories in the inventory of the virtual infrastructure.
2. From the **Credentials** drop-down list, select user credentials to access your Microsoft Azure Blob storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in section Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.
3. From the **Region** drop-down list, select an Azure region.

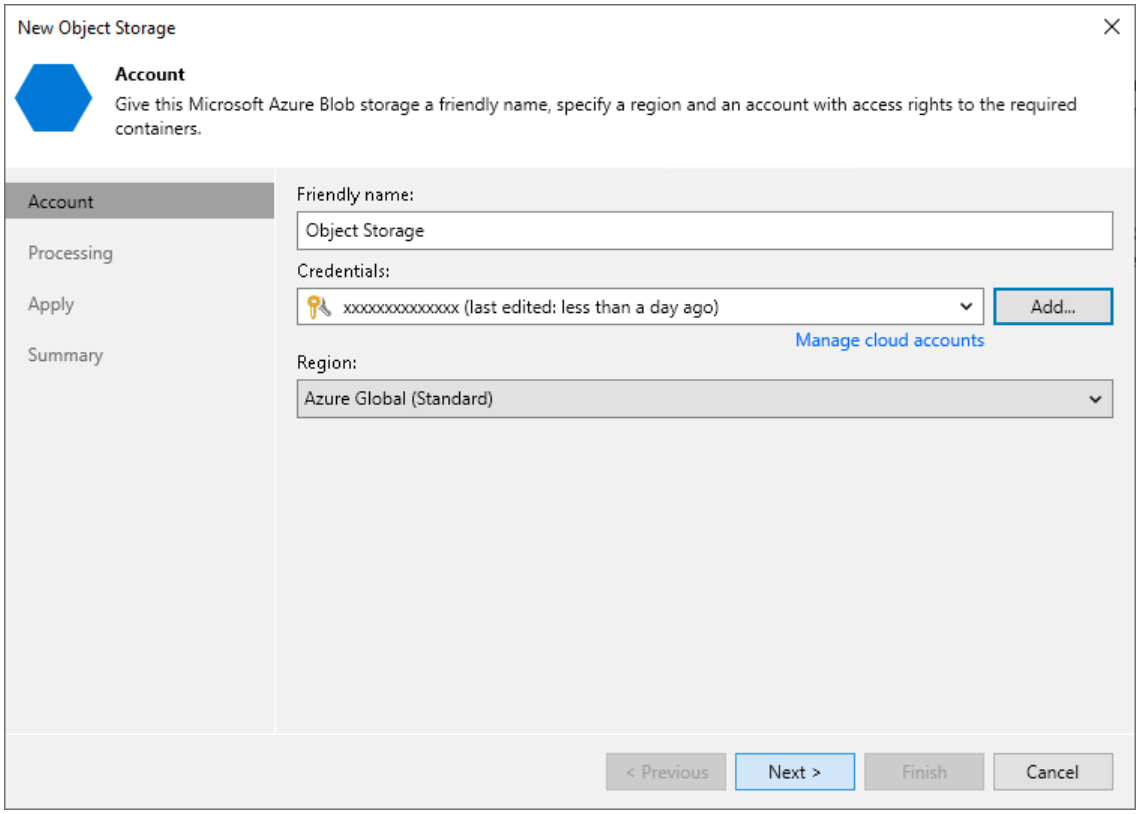

## Step 3. Specify Object Storage Processing Settings

At the **Processing** step of the wizard, do the following:

1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
2. In the **Backup Proxy** window, select proxy servers:
   - If you select **All proxies**, Veeam Backup & Replication will use all available backup proxies for the object storage backup. The number of proxies in use defines the number of data threads that transfer data from the object storage to the backup repository. The more data transfer threads Veeam Backup & Replication uses, the higher is the data transfer speed.

     If the object storage is used as a source for an object to tape backup job, the tape server utilized for this job is added as yet another backup proxy when creating an object to tape backup job. This backup proxy has the highest priority over all others and is used by default if it has access rights to the object storage. For details on object to tape backup jobs, see the Object Storage Backup to Tape section in the Veeam Backup & Replication User Guide.
   - If you select **Use the selected backup proxies only**, you can explicitly specify backup proxies that Veeam Backup & Replication must use for the object storage backup.

     It is recommended that you select at least two backup proxies to ensure that the backup jobs start even if one of the proxies fails or loses its connectivity to the source object storage. The more proxies you select, the more data transfer threads Veeam Backup & Replication will use for backup jobs, thus improving performance.

     Even if the object storage is used as a source for object to tape backup jobs, Veeam Backup & Replication will use only proxies selected in the list to process the backup data traffic.

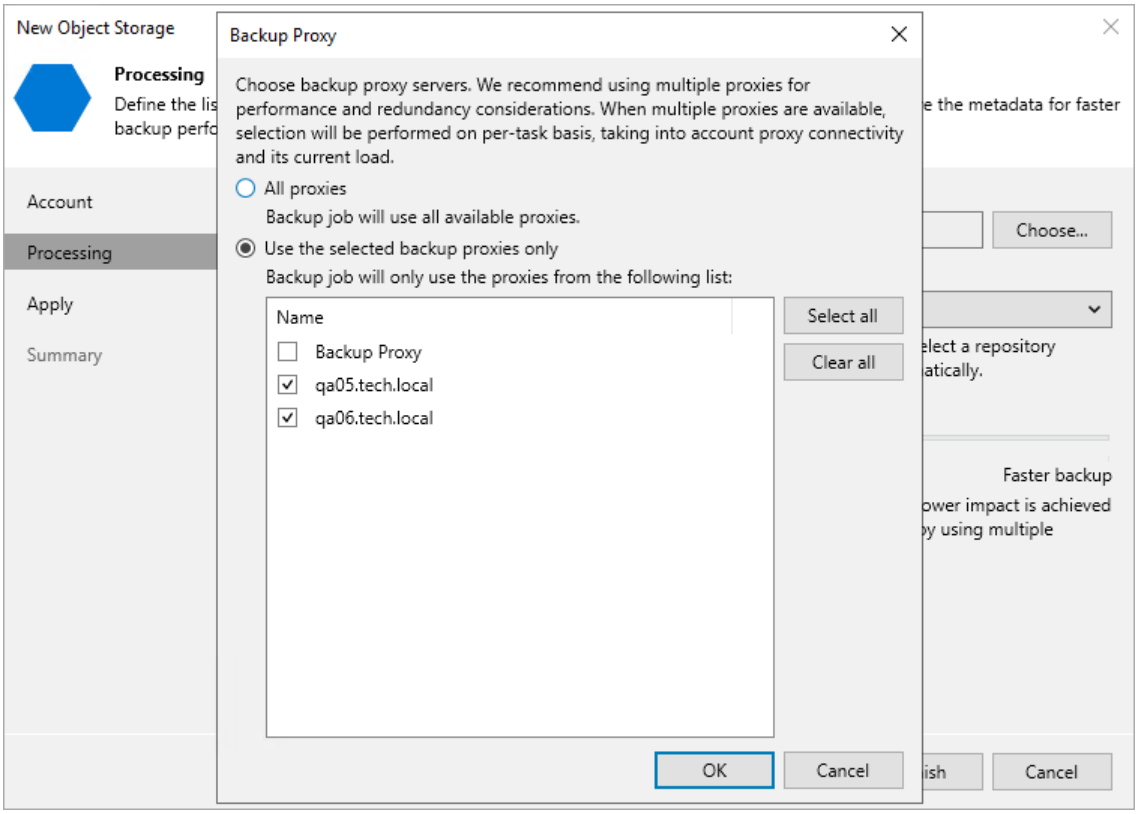

3. From the **Cache repository** drop-down list, select a cache repository where temporary cache files must be stored. This repository must be located in the close proximity to the source object storage and backup proxies.

   If you change the cache repository for an existing object storage whose backups are stored in another object storage, Veeam Backup & Replication will prompt you to either attach migrated metadata, copy metadata from the previous cache repository, or download metadata manually from the archive repository. For more information, see the Unstructured Data Backups in Object Storage Repositories section.

4. Use the **Backup I/O control** slider to define how fast backup proxies can read data from the source object storage. This setting is based on the number of parallel threads that can be used by proxies configured for processing the object storage.

| I/O Control | Number of Proxies | Threads per Task |
|---|---|---|
| **Lower Impact** | 1 | 1 |
| **Below Normal** | 1 | 4 |
| **Normal** | 2 | 8 |
| **Above Normal** | 4 | 16 |
| **Faster Backup** | Unlimited | 16 |

If resources of your object storage source are limited, it is recommended that you select the **Lower impact** option. If your object storage source is powerful enough, select the **Faster backup** option.

5. Click **Apply** to save the configured settings.

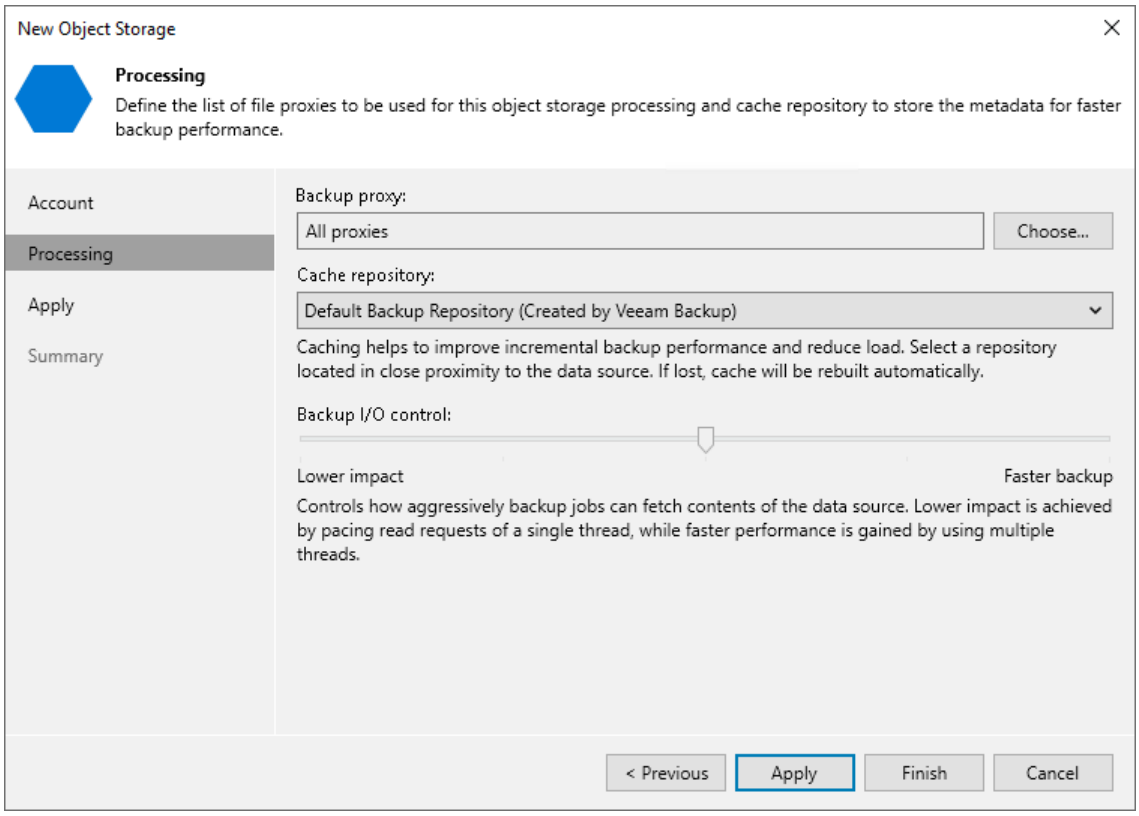

## Step 4. Apply Object Storage Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all required components and adds the object storage to the inventory of the virtual infrastructure. Click **Next** to proceed.

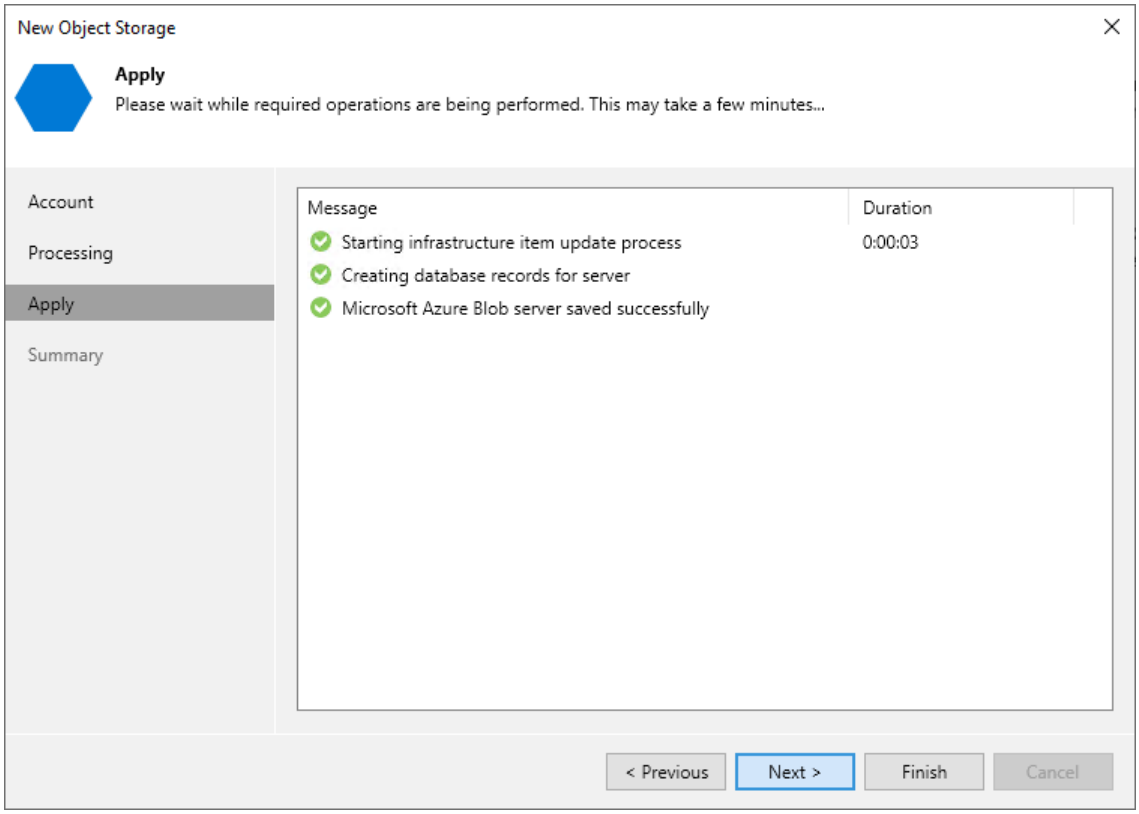

## Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added object storage and click **Finish** to exit the wizard.

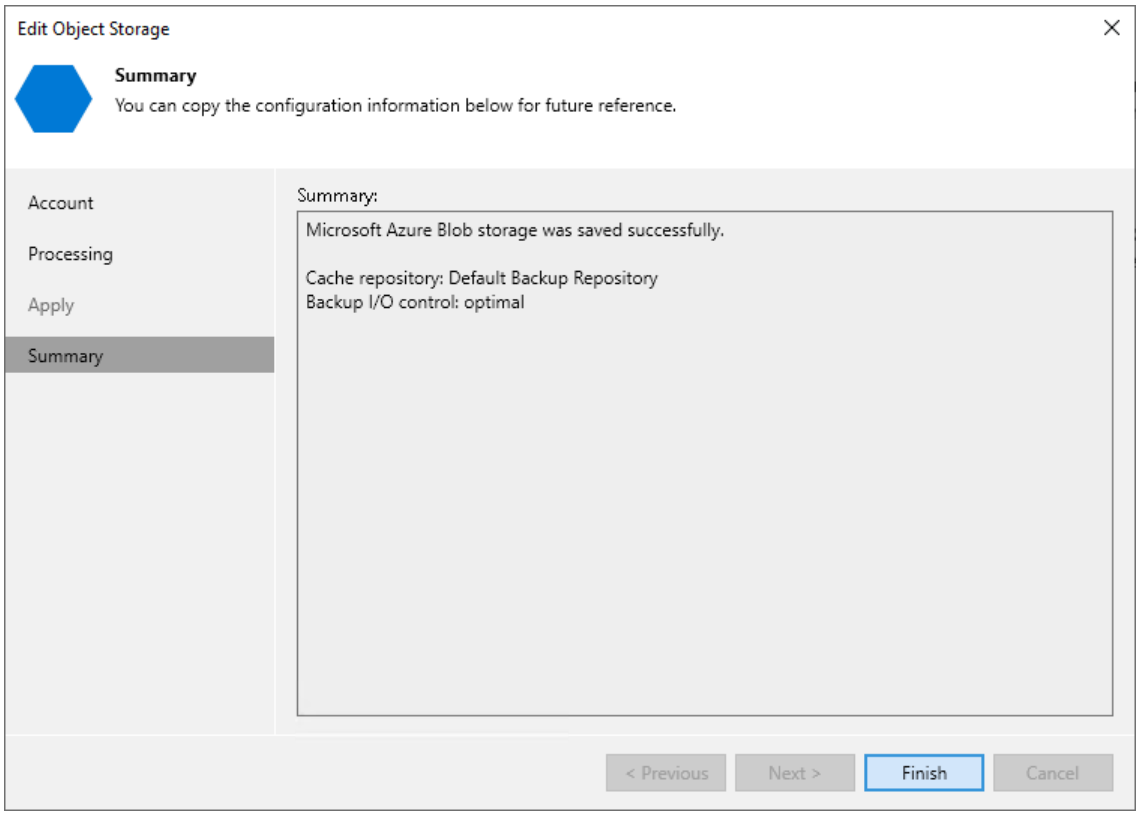

# How Unstructured Data Backup Works

Veeam Backup & Replication performs backup of unstructured data to the backup storage in the following way:

1. When a new backup job session starts, Veeam Backup & Replication assigns a backup proxy to process the unstructured data.
2. The backup proxy enumerates files and folders on the file share or object storage repository and creates a cyclic redundancy check (CRC) tree.
3. The backup proxy transfers the CRC tree to the cache repository.
4. The cache repository saves the CRC tree.

   When the cache repository receives a new CRC tree structure from the proxy, it compares it with the CRC tree created during the previous run of the backup session. If any source files or folders have changed since the previous backup session run, the cache repository instructs the backup proxy to start reading changed data from the source file share or object storage repository.
5. The backup proxy reads new data from the source file share or object storage repository.
6. The backup proxy creates data packages and transfers them to the target backup repository.

   Data packages comprise backup data files (each 64 MB in size) and metadata files that contain names and versions of backup files and allocation of data in backup files.
7. Veeam Backup & Replication checks file versions in the backup repository against retention settings and moves backup data from the backup repository to the archive repository if necessary.

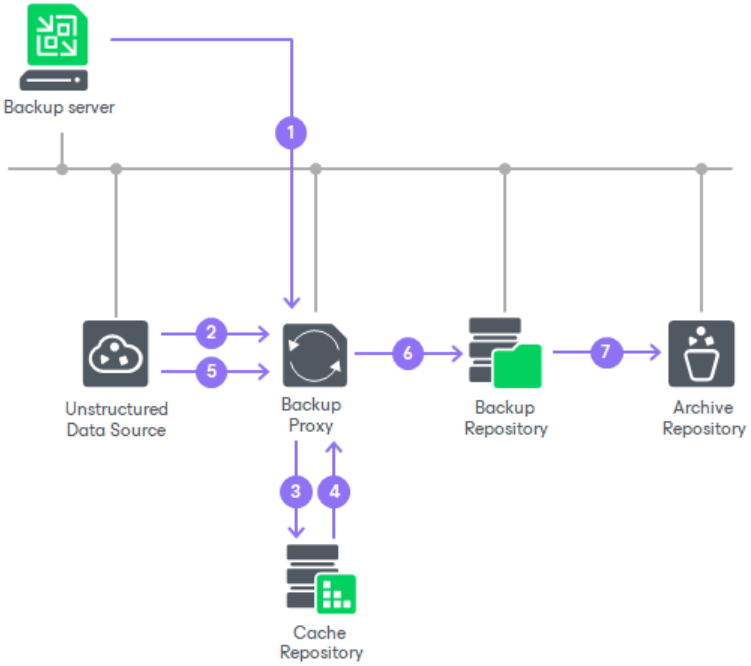

# Data Structure in Backup, Archive and Secondary Repositories

Backup, archive and secondary repositories store file and object storage backups as objects in VBLOB files (up to 64 MB each) of unstructured data. They also store metadata about the original objects on unstructured data sources and their structure.

When you run the file backup job or object storage backup for the first time, Veeam Backup & Replication creates a full backup for objects on the data source. During subsequent backup job sessions, Veeam Backup & Replication copies only objects that have changed since the last backup job session.

Although the file backup job or object storage backup job first creates a full backup and afterwards incremental backups, Veeam Backup & Replication does not create a separate file for each backup job run as it does during VM backup. Instead of this, it consistently creates multiple VBLOB files accompanied by metadata files that track all the changes on the unstructured data source.

In course of time, as the retention is applied to data blocks in VBLOB files, some blocks are marked as outdated and due for deletion. When the number of outdated blocks in a VBLOB file reaches a certain, so called transformation threshold, Veeam Backup & Replication reorganizes the data in the VBLOB file. The transformation threshold value is different for each object storage type and for each role that the object storage is used in (backup repository or archive repository). By default the threshold is set to the optimum percentage of changed blocks. You can learn the default threshold value for the object storage used as a backup repository from Veeam Backup & Replication logs, from Veeam Customer Support or by running the Get-VBRNASObjectStorageTransformThreshold PS cmdlet. If required, you can update the transformation threshold value for the object storage used as a backup repository by using the Set-VBRNASObjectStorageTransformThreshold PS cmdlet. Do that with maximum caution.

Veeam Backup & Replication uses the following structure for storing the file backup data and object storage backup data in the backup repository:

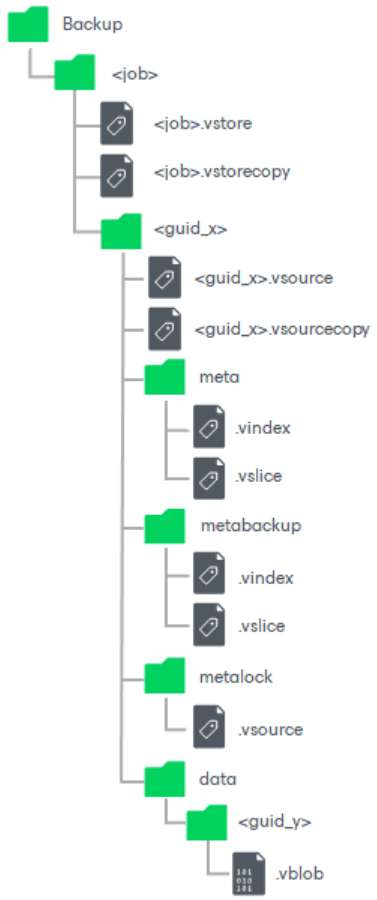


| Folder/File | | | | | | Description |
|---|---|---|---|---|---|---|
| **Backup** | | | | | | Folder in the repository that is dedicated to store backups. |
| | **<job>** | | | | | Folder that contains all data backed up by a specific file backup job or object storage backup job. Data for each job will be placed to its own directory. |
| | | **<job>.vstore** | | | | XML metadata file that describes the entire backup file for a specific file backup job or object storage backup job. |

| Folder/File | | | | | | Description |
|---|---|---|---|---|---|---|
| | | <job>.vstorecopy | | | | Copy of the XML metadata file that describes the entire backup file for a specific file backup job or object storage backup job. |
| | | <guid_x> | | | | Folder that contains all data for a single source. Data for each source will be placed to its own directory. |
| | | | <guid_x>.vsource | | | XML metadata file that describes the single source object. |
| | | | <guid_x>.vsourcecopy | | | Copy of the XML metadata file that describes the single source object. |
| | | | meta | | | Folder with binary metadata files that describe the content of the backup.<br><br>If backup files are stored on an object storage repository, this folder is missing from the repository itself and is instead stored on the cache repository specified for the selected backup repository.<br><br>If backup files are stored on a scale-out backup repository, this folder on each performance extent contains metadata for data stored on the same extent of the scale-out backup repository. For more information, see Unstructured Data Backups in Scale-Out Repositories.<br><br>If the repository where you store backups is immutable, this folder and files in it do not have the immutability lock and may be updated at every backup job run. |

| Folder/File | | | | | | Description |
|---|---|---|---|---|---|---|
| | | | | **.vindex** | | Binary metadata that describes backup files (names and versions). |
| | | | | **.vslice** | | Binary metadata that describes allocation of data in VBLOB backup files. |
| | | | **metabackup** | | | Folder with a replica of binary metadata files.<br>If backup files are stored on a scale-out backup repository, this folder on each performance extent contains replica of metadata stored on other extents of the scale-out backup repository. For more information, see Unstructured Data Backups in Scale-Out Repositories. |
| | | | **metalock** | | | Folder with XML data files that describe immutable metadata replica.<br>This folder exists only for backups in immutable repositories. |
| | | | | **.vsource** | | XML metadata file that describes a single immutable metadata replica. |
| | | | **data** | | | Folder with binary data. |
| | | | | **<guid_y>** | | 1 GB basket that stores VBLOB backup files. |

| Folder/File | | | | | | Description |
|---|---|---|---|---|---|---|
| | | | | | .vblob | By default, a classic 64 MB file that stores data from the file share backup or object storage backup.<br><br>If necessary, you can convert unstructured data backup files into the single file basket format recommended for storing data on HPE StoreOnce storage appliances. For more information, see description of the `Convert-VBRNASBackupStorageFormat` cmdlet in the Veeam PowerShell Reference. |

# Unstructured Data Backup Retention Scenarios

There can be a number of backup retention scenarios depending on the configuration of backup and archive repositories. In this section, you can find example cases that illustrate file and object storage backup retention with different settings.

### Case 1

Only 1 file version is created. The file does not change.

File version 1 always remains in the backup repository and is not moved to the archive repository even if this behavior is enabled and configured in the retention policy settings.

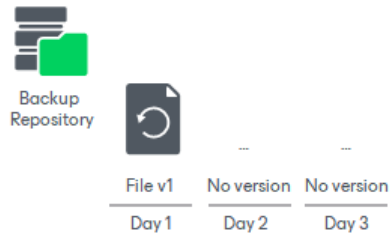


### Case 2

Retention for the backup repository is set to 5 days. No archive repository is configured. The file changes once a day. The backup is performed once a day.

On day 6, file version 6 is added to the backup repository, file version 1 is deleted by retention.

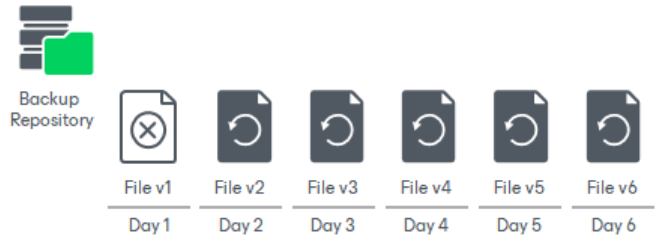


### Case 3

Retention for the backup repository is set to 3 days. The file changes every hour. The backup is performed 2 times a day.

On day 4, versions 7 and 8 are added to the backup repository, file versions 1 and 2 added to the backup repository on day 1 are deleted by retention.

### Case 4

Retention for the backup repository is set to 3 days. The file changes once a day.

On day 3, the source file is deleted from the source share, the backup repository considers file version created on this day as deleted.

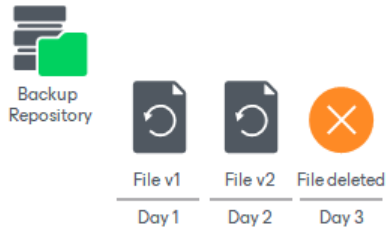


On day 4, the backup repository still detects the file as deleted, file version 1 is deleted from the backup repository by retention.

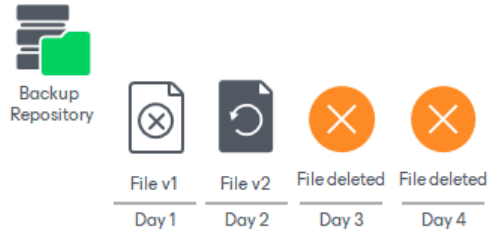


On day 5, the backup repository still detects the file as deleted, file version 2 is deleted from the backup repository by retention.

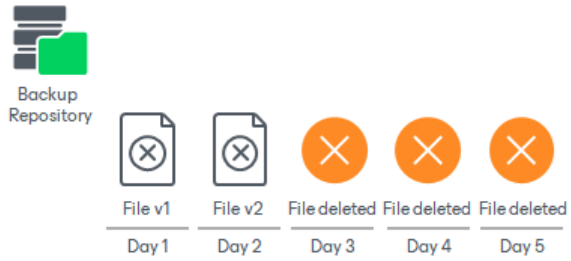


Thus, no file versions are stored in the backup repository for this file any longer.

### Case 5

Retention for the backup repository is set to 5 days. The archive repository is enabled with default settings. The file changes every day. The backup is performed once a day.

On day 6, file version 6 is added to the backup repository, file version 1 is moved to the archive repository by retention.

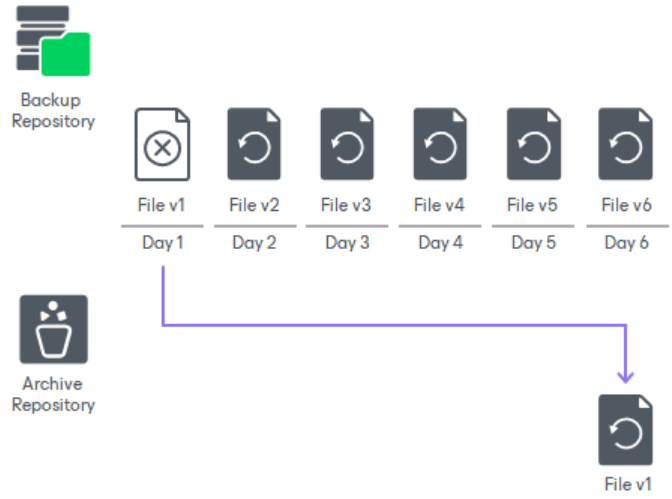


### Case 6

Retention for the backup repository is set to 3 days. The archive repository is enabled with DOCX files to be excluded from archiving. The files change once a day. The backup is performed once a day.

On day 4, file versions created on day 1 are removed from the backup repository. File version 1 for DOCX file is deleted, file version 1 for XLSX file (non-DOCX) is moved to the archive repository.

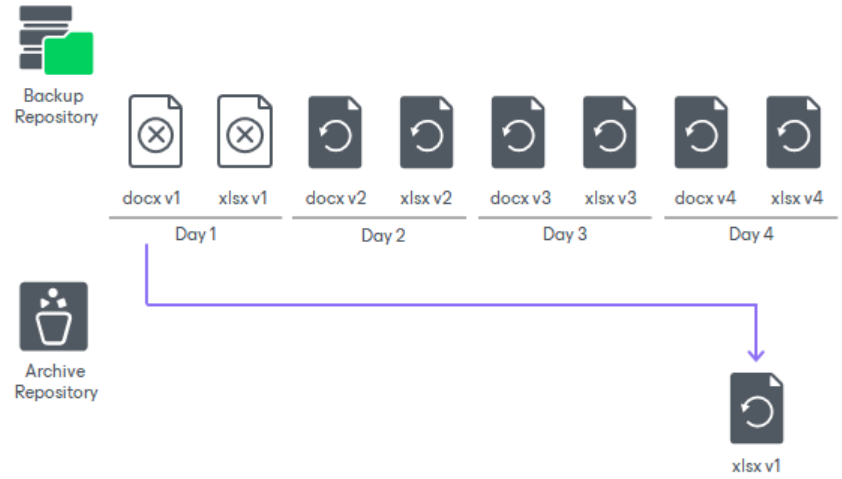


### Case 7

Retention for the backup repository is set to 4 days. The archive repository is enabled and configured to keep 3 versions of active files and 2 versions of deleted files.

On day 8, file version 8 is added to the backup repository, file version 4 is moved from the backup repository to the archive repository to keep file versions for 4 days, file version 1 is deleted from the archive repository to keep 3 file versions of the active file (versions 2, 3, 4).

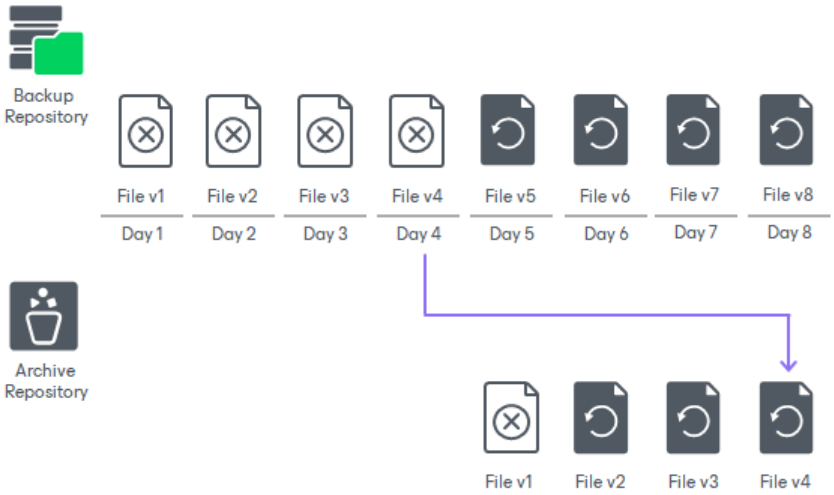


On day 9, the file is removed from the source, file version 9 (denoting the missing file) is added to the backup repository, file version 5 is moved from the backup repository to the archive repository, file versions 2 and 3 are deleted from the archive repository to keep 2 file versions of the deleted file (versions 4 and 5).

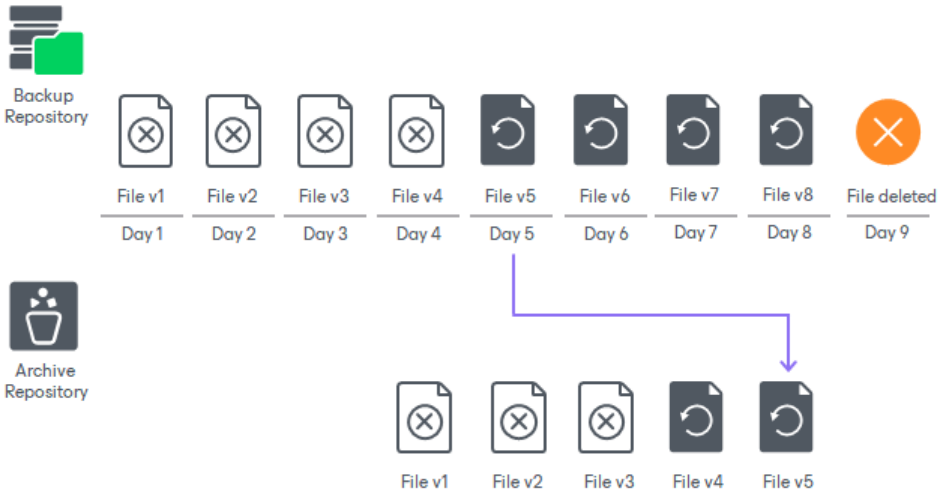


On day 10 and 11, file versions 6 and 7 are successively moved from the backup repository to the archive repository. File versions 4 and 5 are deleted from the archive repository.

On day 12, file version 8 (the last file version) is moved from the backup repository to the archive repository, file version 6 is deleted from the archive repository. After that, versions 7 and 8 are stored in the archive repository further on.

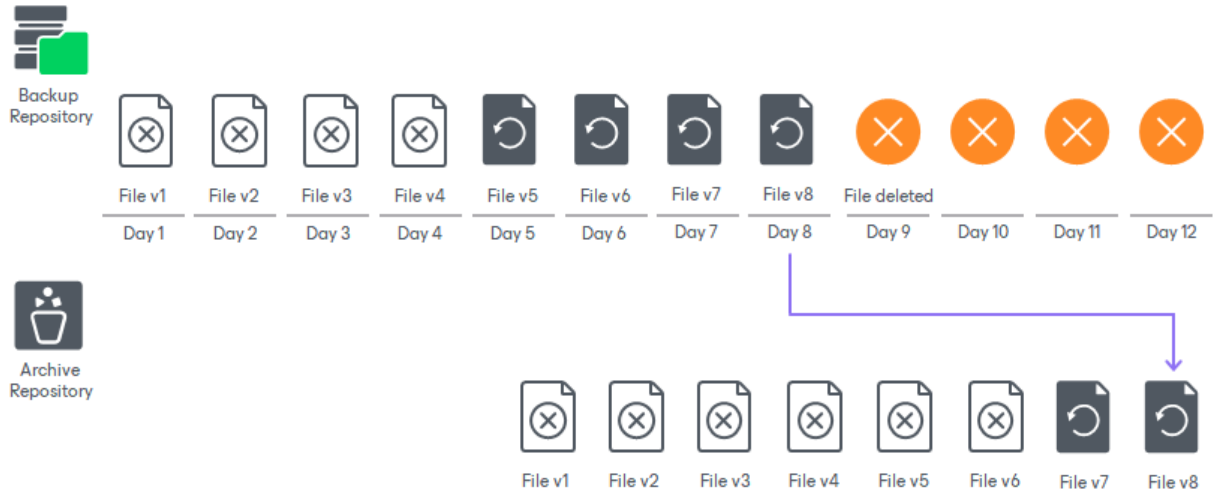


### Case 8

Retention for the backup repository which is a hardened repository is set to 30 days, the immutability for it is set to 14 days. The file changes once a day.

On day 3, the source file is deleted from the source share, the backup repository considers the file version created on this day as deleted.

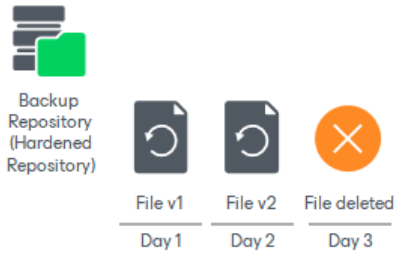


According to retention settings, Veeam Backup & Replication will keep all file versions for 30 days. After that, it will start marking them for deletion. On day 31, file version 1 is marked for deletion from the backup repository by retention, but as the repository is immutable it will still keep this file version.

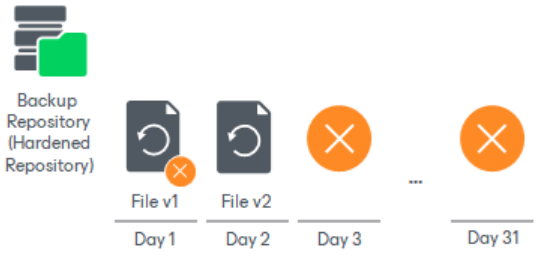

On day 32, file version 2 is marked for deletion from the backup repository by retention, but as the repository is immutable it will keep this file version as well.

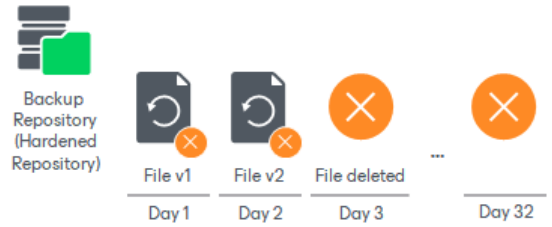


On day 45, when the immutability lock is released after 14 days (14 days of immutability configured for the repository), file version 1 is deleted from the backup repository. File version 2 is marked for deletion, but is still immutable.

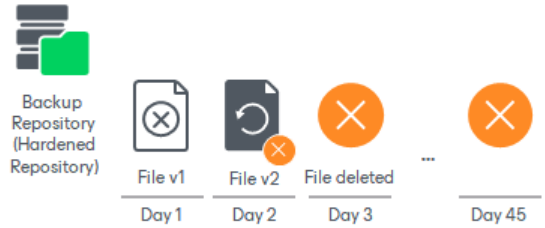


Finally, on day 46, the immutability lock is released for file version 2 and it is deleted from the backup repository.

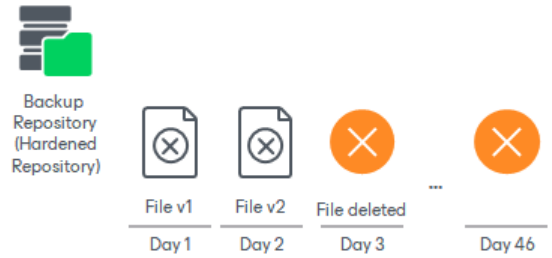


Thus, no file versions are stored in the backup repository for this file any longer.

### Case 9

Retention for the backup repository which is an immutable object storage is set to 30 days, the immutability for it is set to 14 days, 10 days of immutability are added automatically as a Block Generation period described in the File Backups in Immutable Repositories section. The file changes once a day.

On day 3, the source file is deleted from the source share, the backup repository considers file version created on this day as deleted.

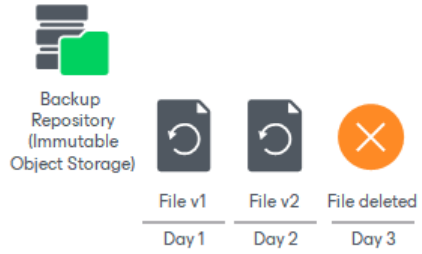

According to retention settings, Veeam Backup & Replication will keep all file versions for 30 days. After that, it will start marking them for deletion. On day 31, file version 1 is marked for deletion from the backup repository by retention, but as the repository is immutable it will still keep this file version.

On day 32, file version 2 is marked for deletion from the backup repository by retention, but as the repository is immutable it will keep this file version as well.

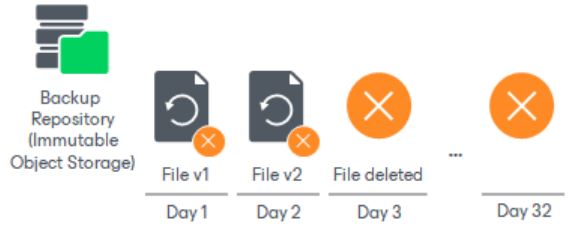


On day 55, when the immutability lock is released after 24 days (14 days of the immutability period configured for the repository plus 10 days of the Block Generation period), file version 1 is deleted from the backup repository. File version 2 is marked for deletion, but is still immutable.

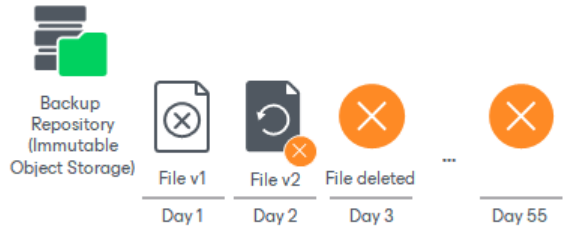


Finally, on day 56, the immutability lock is released for file version 2 and it is deleted from the backup repository.

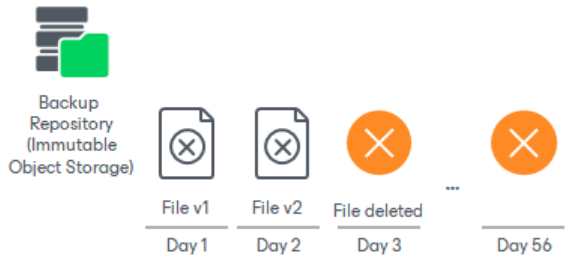


Thus, no file versions are stored in the backup repository for this file any longer.

# Unstructured Data Backups in Object Storage Repositories

## How Unstructured Data Backup to Object Storage Repository Works

You can select an object storage repository added in your backup infrastructure as a backup repository for storing unstructured data backups.

For unstructured data backups stored on non-object storage repositories, Veeam Backup & Replication stores active metadata, metadata copy, and the data itself all next to each other on the repository, as described in the Data Structure in Backup, Archive and Secondary Repositories section.

The main peculiarity of storing unstructured data backups in an object storage repository is keeping active metadata in the cache repository. This metadata is actively used during backup, restore, merge, transform, or health check operations. All the changes of the active metadata is then replicated to the metadata replica that is stored on the repository next to the data.

The metadata replica is self-sufficient: if anything happens to the source unstructured data and the cache repository, you still will be able to restore data from the backup stored in the object storage repository.

### Types of Connection to Object Storage Repository

Depending on your infrastructure and the activity type, the source general-purpose backup proxy can connect to the object repository through one of the connection types:

- **Direct**. In this mode, the general-purpose backup proxy writes directly to the object storage. This connection type is used, for example, for backup and restore sessions when your target object storage repository is located in your local infrastructure.

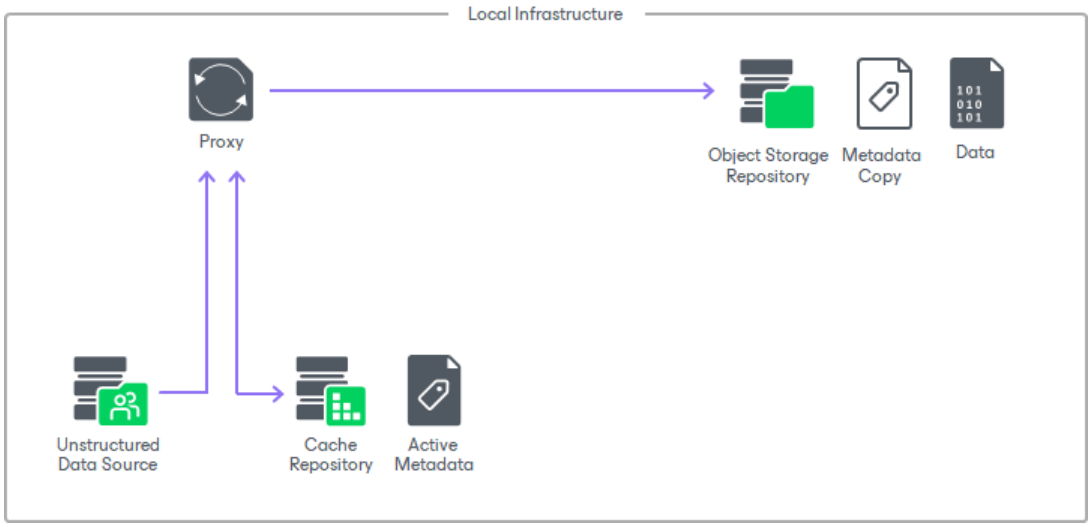

- **Through gateway servers**. In this mode, the backup server automatically selects the most suitable gateway server from the preconfigured list. This connection type is used, for example, for backup copy sessions when the source backup is stored on an SMB or NFS share, or if the target object storage repository is a cloud repository.

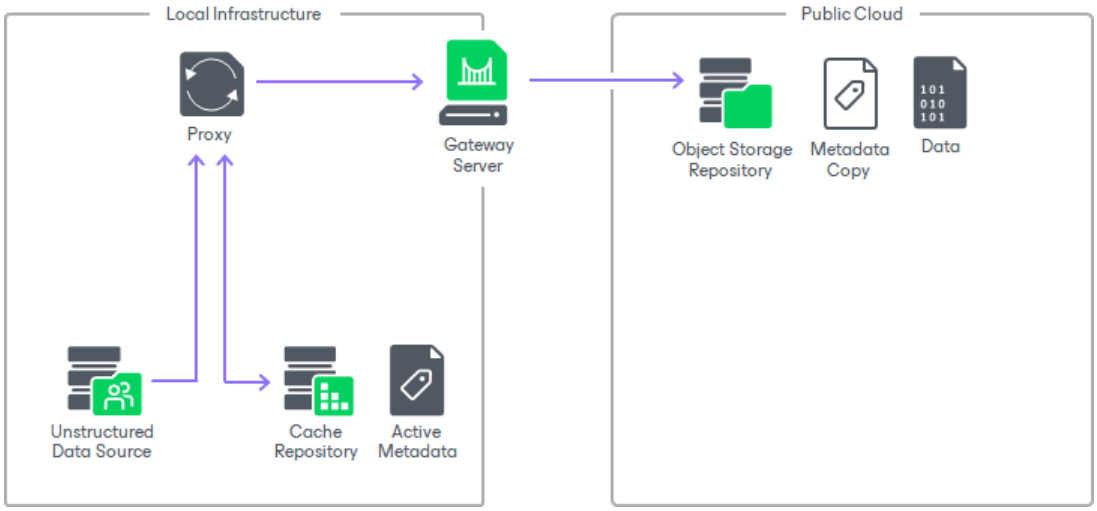


The connection type is configured at the **Account** step of the New Object Storage Repository wizard.

## Helper Appliance in Unstructured Data Backup

Unstructured data is forever forward incremental and it does not create periodic full backups. To avoid constantly downloading data to and uploading data from the cloud object storage for performing transform and health check operations, you can configure a helper appliance. To learn how to configure the helper appliance, see the **Mount Server** step in Adding S3 Compatible Object Storage, Adding Amazon S3 Storage, Adding Google Cloud Object Storage, Adding IBM Cloud Object Storage, Adding Azure Blob Storage, Adding Wasabi Cloud Object Storage.

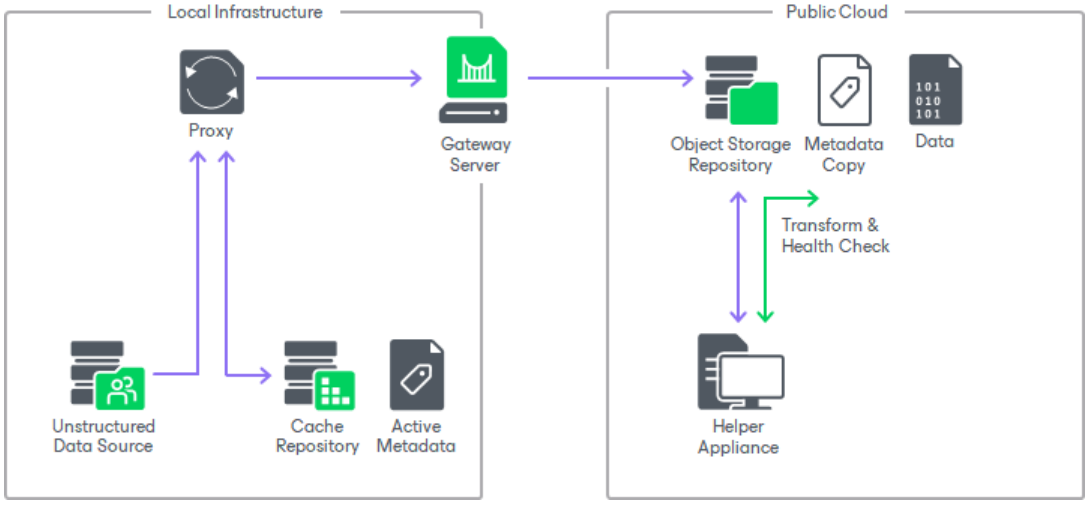

# Changing Cache Repository

If your unstructured data source is protected with an unstructured data backup job that backs up data to an object storage and you decide to change the cache repository for this data source, Veeam Backup & Replication will prompt you to:

- **Attach manually migrated data**. You can manually copy metadata from the old cache repository to a new one. After that, when changing the cache repository for the unstructured data source, click **Attach** in the displayed window.
- **Copy metadata from the previous cache repository**. If you click **Copy** when changing the cache repository for the unstructured data source, Veeam Backup & Replication will automatically copy all metadata from the old cache repository to a new one.
- **Download metadata from the archive repository**. If the old cache repository is not available, you can click **Download** when changing the cache repository for the unstructured data source to automatically download all metadata to a new cache repository.

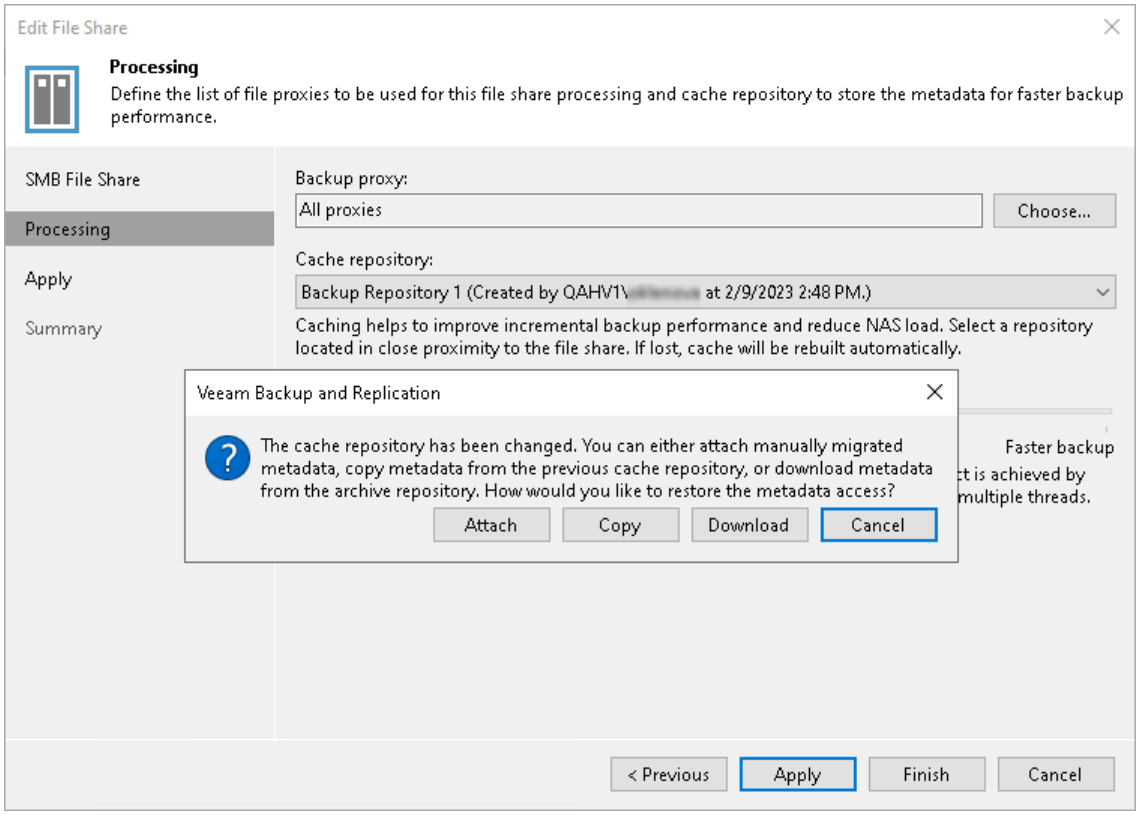

# Unstructured Data Backups in Immutable Repositories

All data stored in immutable repositories cannot be altered or removed until the immutability period expires and the immutability lock for the data is released. The only exception is active metadata for backups stored in immutable hardened repositories and immutable deduplicating storage appliances (HPE StoreOnce and Dell Data Domain).

For more information on the structure of unstructured backup data in backup repositories, see the Data Structure in Backup, Archive and Secondary Repositories section.

## How Immutability for Unstructured Backup Works

Immutability settings for a repository storing unstructured data backups apply to the entire backup of the unstructured data source. This includes all backup files protecting this source. The immutability countdown for the backup files begins when Veeam Backup & Replication marks some old file versions and their data blocks for deletion, as per the backup job retention settings. Files and data blocks are deleted from the backup repository when their immutability lock is released.

> **TIP**
>
> You can roll back an unstructured data backup stored on an immutable backup repository to a point in time state by using the Sync-VBRUnstructuredBackupState PowerShell cmdlet.

### Unstructured Data Backups in Hardened Repositories and Immutable Deduplicating Storage Appliances

The immutability lock period for unstructured data backups in hardened repositories and deduplicating storage appliances (HPE StoreOnce and Dell Data Domain) is set by the backup repository immutability settings.

For an example of how this mechanism works, see **Case 8** in the Unstructured Data Backup Retention Scenarios section. For more information on hardened repositories and on how to configure them, see the Hardened Repository section. For more information on how to configure immutability for HPE StoreOnce, see the HPE StoreOnce section, for Dell Data Domain — the Dell Data Domain section.

### File Backups in Immutable Object Storage Repositories

For general information on immutability for object storage repositories in Veeam Backup & Replication and configuration details, see the Immutability for Object Storage Repositories section.

In addition to the set immutability period for each object storage repository, Veeam Backup & Replication automatically adds up to 10 days to the immutability expiration date to reduce I/O operations and associated costs. This period is known as Block Generation. It is applied automatically and does not require configuration. For example, if the immutability period is 14 days, Veeam Backup & Replication automatically adds 10 days to specific objects to reduce I/O operations with the data blocks over time, totaling 24 days of immutability.

For an example of how this mechanism works, see **Case 9** in the Unstructured Data Backup Retention Scenarios section. For more information, see the Block Generation section for object storage repositories.

IMPORTANT

We strongly recommend the following configuration for storing unstructured data backups in immutable object storage repositories:

- Set the immutability period for the immutable object storage repository to a maximum of 14 days.
- Configure file backup jobs or object storage backup jobs that use the immutable object storage repository to run no more than once per day.

Otherwise, the unstructured data backup may consume excessive storage space.

These settings are crucial for cloud object storage repositories where storage costs can be significant.

# Metadata of Unstructured Data Backups in Immutable Storage Repositories

When creating unstructured data backups, Veeam Backup & Replication generates two sets of metadata: active metadata and metadata replica. The metadata replica is always stored in the backup repository alongside the data. The active metadata storage location varies depending on the backup repository type:

- In **hardened repositories** or **deduplicating storage appliance (HPE StoreOnce or Dell Data Domain)**, the active metadata is stored in the immutable repository with the locked metadata replica and locked backup data.
- In **object storage repositories**, the active metadata is stored in the cache repository.

Data in immutable repositories is locked temporarily and cannot be modified or removed until the immutability period ends. The metadata replica that is stored next to the data is also immutable and unchangeable. However, the active metadata does not have the immutability lock, changes actively during every backup session and keeps the up-to-date state of the unstructured data backup. Thus, it can be used by Veeam Backup & Replication to track changes in the unstructured data backup.

Once in 30 days (if the job runs once a day or once in several days) or as the metadata replica chain reaches 30 generations (if the job runs several times a day), Veeam Backup & Replication uses metadata replica files created during this period to generate a new locked metadata flat file.

After the immutability period for the metadata replica files, which were already transformed into the metadata flat file, ends, Veeam Backup & Replication removes them from the object storage repository.

The schemas in this section show examples of creating a locked metadata flat file for a job that runs once every day. After 30 generations of metadata replica files are created (that is after 30 days), Veeam Backup & Replication uses them to generate a new metadata flat file. The older metadata replica files that comprised it are marked for deletion and can be deleted based on the retention and immutability settings of the repository.

## Metadata in Hardened Repositories and Immutable Deduplicating Storage Appliances

If unstructured data backups are stored in an immutable hardened repository or deduplicating storage appliance (HPE StoreOnce or Dell Data Domain), the active metadata is stored on the immutable repository alongside with the locked metadata replica and locked backup data.

## Metadata in Immutable Object Storage Repositories

If unstructured data backups are stored in an immutable object storage such as Amazon S3, Microsoft Azure Storage, IBM Cloud Object Storage, Wasabi Cloud Object Storage, or S3-compatible object storage, the active metadata is stored on the cache repository.

# Unstructured Data Backups in Scale-Out Repositories

If you use a scale-out backup repository as a backup repository for storing unstructured data backups, by default Veeam Backup & Replication evenly distributes backup data among all the performance extents added to this repository. The backup data is accompanied by the metadata files. To provide an opportunity to restore data written to extents remaining available in case one or two of other extents are unavailable, Veeam Backup & Replication triplicates metadata when writing it to extents. Thus, every extent keeps metadata for data stored on this extent (in folder `meta`) and replica of metadata for data stored on adjacent extents (in folder `metabackup`). Even if two of three adjacent extents are lost, you can restore backup data stored on the third extent.

**NOTE**

Consider the following when storing unstructured data backups on a scale-out backup repository:

- The metadata redundancy approach protects against two extents out of three being completely lost.
- The metadata redundancy approach allows protecting metadata, not data.
- An unstructured data backup job consumes 1 task slot on every extent of the SOBR per each file share added to the job.
- The backup file placement policy set for the scale-out backup repository, as described in section Backup File Placement for Performance Tier, is ignored.

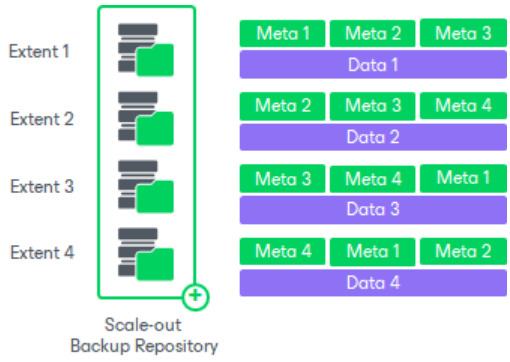

# Unstructured Data Backups in Deduplicating Storage Appliances

When you use a deduplicating storage appliance for storing unstructured data backups, it can store data and metadata. However, storing both data and metadata on the same deduplicating storage appliance is slow due to how deduplication works. To avoid this, we strongly recommend that you configure a scale-out backup repository with performance extents acting as metadata extents. This allows you to store metadata separately from the backup data.

The recommended backup infrastructure is a scale-out backup repository with a fast SSD drive for storing metadata, paired with a slower deduplicating storage appliance (Dell Data Domain, ExaGrid, HPE StoreOnce, Quantum DXi) for storing backup data.

Most often, when performing restore, merge, transform operations, Veeam Backup & Replication interacts with metadata rather than with the backup data. Processing metadata stored on an SSD is faster and more efficient than accessing large data arrays stored on a slow storage.

**IMPORTANT**

Consider the following when assigning roles to extents in a scale-out backup repository:

- An extent with the metadata role can be used for storing metadata for unstructured data backup jobs only. However, extents with the data role can be used by any job.
- Make sure that you assign both roles to extents in a scale-out repository. If an extent with one of the roles is missing, Veeam Backup & Replication cannot store backups on this scale-out repository.

**NOTE**

If the data role is assigned to an extent, Veeam Backup & Replication will also copy replica of metadata to this extent to provide the metadata redundancy. While the original metadata is available, Veeam Backup & Replication does not use the replica of metadata on data extents. If by some reason metadata stored on metadata extents is corrupted or lost, to restore it Veeam Backup & Replication will use the replica of metadata stored on data extents.

To assign the metadata or data role to extents in a scale-out backup repository, use the `Set-VBRRepositoryExtent` cmdlet, as described in the Veeam PowerShell Reference. If previously the role was not assigned to the extents for unstructured data backup or you changed the assigned role, during the next run of the file backup job or object storage backup job that writes backups to this scale-out repository, Veeam Backup & Replication will move metadata to the metadata extent, data — to the data extents.

For example, if you already have an existing backup and want to move its metadata to a specific extent that will operate as a metadata only extent, do the following:

1. Make sure that your license allows using a scale-out backup repository with object storage support. For more information, see Viewing License Information and Veeam Backup & Replication Feature Comparison.
2. Create a scale-out backup repository with the following extents: one extent that currently stores the backup (its data and metadata), another extent that will store metadata of the backup. Usually, it is a fast storage, for example, SSD-based. Let us assume that these extents are named “Backup Repository 1” and "File Backup Repository on SSD".
3. Run the `Set-VBRRepositoryExtent` cmdlet to assign the data role to the “Backup Repository 1” extent and the metadata role to the "File Backup Repository on SSD" extent.

```
Set-VBRRepositoryExtent –Extent “File Backup Repository on SSD” -Metadata
Set-VBRRepositoryExtent –Extent “Backup Repository 1” –Data
```

For more information, see the Veeam PowerShell Reference.

4. Run the file backup job and make sure that the metadata of the backup was moved to the metadata extent: the backup job session displays a line notifying of that.

To view the roles of the extents in a scale-out backup repository, do either of the following:

- Check the role of each extent (the **Role** column) in the list of extents under the certain scale-out repository in the **Backup Infrastructure** view.
- Run the `Get-VBRRepositoryExtent` cmdlet, as described in the Veeam PowerShell Reference.

# File Backup Integration with Storage Systems

There are two approaches in backing up file shares residing on enterprise NAS storage systems.

## Integration with Storage System as NFS or SMB File Share Server

You can add the storage system as a root folder of the server where NFS or SMB file shares reside. The procedure of configuring the file share protection in this case will consist of the following steps:

1. Add the storage system as an NFS file share to the inventory, as described in section Adding NFS File Share, or as an SMB fie share, as described in section Adding SMB File Share. As a file share path, specify the root server folder.

   When adding the storage system in this way, you cannot configure what containers, volumes or file shares will be available for further protection. Therefore, to configure them, you must carefully configure inclusion/exclusion settings when creating a file backup job.

   Consider that servers accessed by NFS (with file shares and folders within them) and servers accessed by SMB (with file shares and folders within them) are added to the inventory separately. For example, if the storage system IP address is 173.25.136.64, add an NFS share for this server by specifying its root folder as `173.25.136.64:/`, and add an SMB share for this server by specifying its root folder as `\\173.25.136.64`.

   > **NOTE**
   >
   > If you add a root server folder as a source for protection, hidden and admin file shares are skipped from processing by default. You can enable their processing with registry values. For more information, contact Veeam Customer Support.

2. Create a file backup job, as described in section Creating File Backup Jobs. As a source to protect, you can select the following entities:

   - whole server
   - file share residing on the server
   - separate folders within the share

   To protect all file shares residing on one server, you must add to the file backup job both NFS and SMB shares previously added to the inventory.

3. Configure what files and folders must be included in or excluded from processing by the file backup job. For more information on how to include/exclude files and folders from processing, see Select Files and Folders to Back Up.

## Integration with Storage System as NAS Filer

You can add the storage system as a NAS filer. This option is preferable if you want to leverage the backup from storage snapshots technology. For more information, see the NAS File Share Backup from Storage Snapshots section in the Storage System Snapshot Integration Guide.

> **NOTE**
>
> If you used to protect NFS and SMB file shares residing on the enterprise storage system and added as file shares in inventory, and now you want to protect them using benefits of NAS filer, you can convert backups created for existing SMB or NFS shares into the format of NAS filer shares. For more information, see Converting Backups from SMB or NFS Shares to NAS Filer Shares.

The procedure of configuring the file share protection in this case will consist of the following steps:

1. Depending on the type of the NAS system you use, add the storage system to the backup infrastructure, as described in the Adding NetApp Data ONTAP, Adding Lenovo ThinkSystem DM/DG Series, Adding Dell PowerScale, or Adding Nutanix Files Storage sections in the Storage System Snapshot Integration Guide.

   Depending on storage settings, the IP address for accessing the storage system can differ from one used for accessing it as a server where file shares reside. You can also use the DNS name of the server.

   When adding the storage system, make sure that you do not forget to perform the following steps:

   a. Enable the NAS filer role for the added storage system.

   b. Specify what protocols the storage should use as a NAS filer: NFS or SMB. Only file shares using the selected protocols will be displayed when you add the storage as a NAS filer and thus available for protection.

   c. Select storage volumes to analyze for the presence of newly added file shares. You can either configure Veeam Backup & Replication to analyze all storage volumes, or exclude some volumes from processing, or specify only certain volumes that will be processed. Only file shares on the selected storage volumes will be displayed when you add the storage as a NAS filer and thus available for protection.

   At this step, you must carefully consider what file shares on what volumes must be protected and through what protocols. Limiting the number of volumes reduces the storage load.

   > **NOTE**
   >
   > Hidden file shares on storage systems added as NAS filers are processed by default. You can use exclude masks to skip a hidden file share from processing or disable the processing of all hidden file shares with a registry value. For more information, contact Veeam Customer Support.

2. Add the configured storage system as a NAS filer to the inventory, as described in section Adding Enterprise Storage System as NAS Filer.

3. Create a file backup job, as described in section Creating File Backup Jobs. As a source to protect, you can select the following entities:

   - whole storage
   - container (for Dell PowerScale — access zone, for NetApp Data ONTAP — SVM, for Nutanix Files Storage — not applied)
   - volume
   - file share

   You cannot specify separate folders within file shares. Therefore, to configure files and folders to be protected, you must properly configure inclusion/exclusion settings.

4. Configure what files and folders must be included in or excluded from processing by the file backup job. For more information on how to include/exclude files and folders from processing, see Select Files and Folders to Back Up.

# Creating Backup Jobs for Protecting Unstructured Data

Depending on the source of the unstructured data, use one of the following backup job types to protect it:

- File backup jobs
- Object storage backup jobs

# Creating File Backup Jobs

To protect files and folders on the file share, configure a file backup job. The backup job defines how, where and when to back up data from the file share. One job can be used to protect one or more file shares. Jobs can be started manually or scheduled to run automatically at a specific time.

File backup jobs are used to protect the following sources of unstructured data:

- File servers (Windows and Linux)
- File shares (NFS and SMB (CIFS)
- Enterprise storage systems added as NAS filers

Before you create a file backup job, check prerequisites.

## Before You Begin

Before you create a file backup job, consider the following:

- Backup infrastructure components that will take part in the file backup process must be added to the backup infrastructure and properly configured. These include source file shares to back up, backup proxies, and all repositories, including cache, backup and archive repositories. For more information, see the Backup Infrastructure for Unstructured Data Backup section.
- The target backup repository must have enough free space to store created backup files. If you want to receive notifications on the repository running low on free space, configure global notification settings as described in the Specifying Other Notification Settings section.
- Make sure that repositories intended to store file share backups are not configured to store files in the WORM status. Otherwise, the backup jobs will fail when Veeam Backup & Replication cannot update the backup metadata files.
- If you plan to map a file backup job to a backup that already exists in the backup repository, you must perform the rescan operation for this backup repository. Otherwise, Veeam Backup & Replication will not be able to recognize backup files in the backup repository.

  For more information on how to rescan backup repositories, see the Rescanning Backup Repositories section.
- If you plan to use pre-job and post-job scripts, you must create scripts before you configure the file backup job.
- Antivirus software may significantly slow down file backup jobs. To improve performance, we recommend you exclude the `c:\Program Files (x86)\Veeam\Backup Transport\x64\VeeamAgent.exe` process from the antivirus scan on machines running the file backup proxy and backup repository roles. Keep in mind that it can weaken the security of these machines.

  > **NOTE**
  >
  > If the objects that you want to back up were marked by the antivirus software as infected, the file backup job will be finished with the *Warning* or *Failed* state. The state will depend on the backed-up object.
- Veeam Backup & Replication does not create the separate node in the inventory pane for the archived backups. If you use archive repositories and want to make sure that backups were moved to it, you can do it using backup properties. For more information, see Viewing Unstructured Data Backup Properties.

- If you plan to store backups in Veeam Data Cloud Vault, you must enable encryption.
- You cannot use the capacity tier of scale-out backup repositories as a target for file backup jobs.

# Step 1. Launch New File Backup Job Wizard

To launch the **New File Backup Job** wizard, do one of the following:

- On the **Home** tab, click **Backup Job** > **File Share**.
- Open the **Home** view. Right-click in the working area, and select **Backup** > **File share**.
- Open the **Home** view. In the inventory pane, right-click the **Jobs** node and select **Backup** > **File share**.
- You can quickly add the file share to an already existing job. Open the **Inventory** view. Under the **Unstructured Data** (**File Shares** – for version 12) node in the inventory pane, select **File Shares**. In the working area, right-click the file share you want to back up and select **Add to backup job** > *name of the job*.

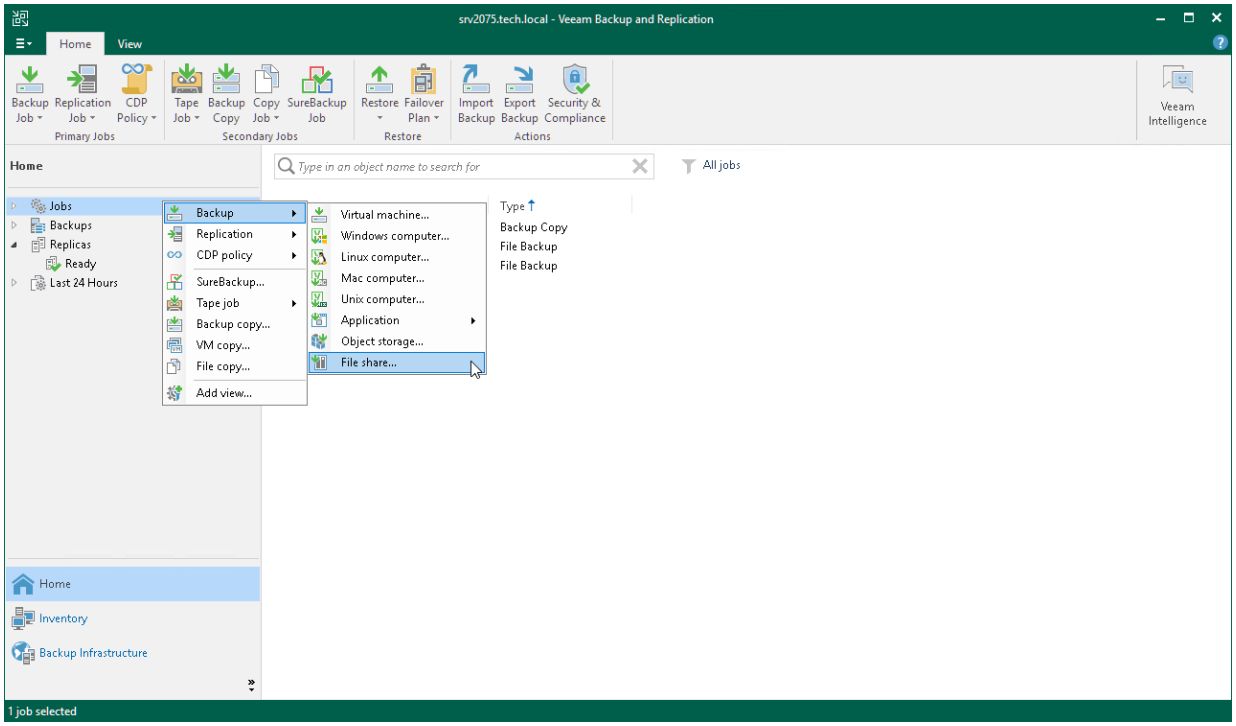

# Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a name and description for the file backup job.

1. In the **Name** field, enter a name for the file backup job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the job, date and time when the job was created.
3. Select the **High priority** check box if you want the resource scheduler of Veeam Backup & Replication to prioritize this job higher than other similar jobs and to allocate resources to it in the first place. For more information on job priorities, see the Job Priorities section.

> **TIP**
>
> In the list of jobs in the Veeam Backup & Replication console, jobs with the **High priority** option enabled are marked with a red flag ( ).

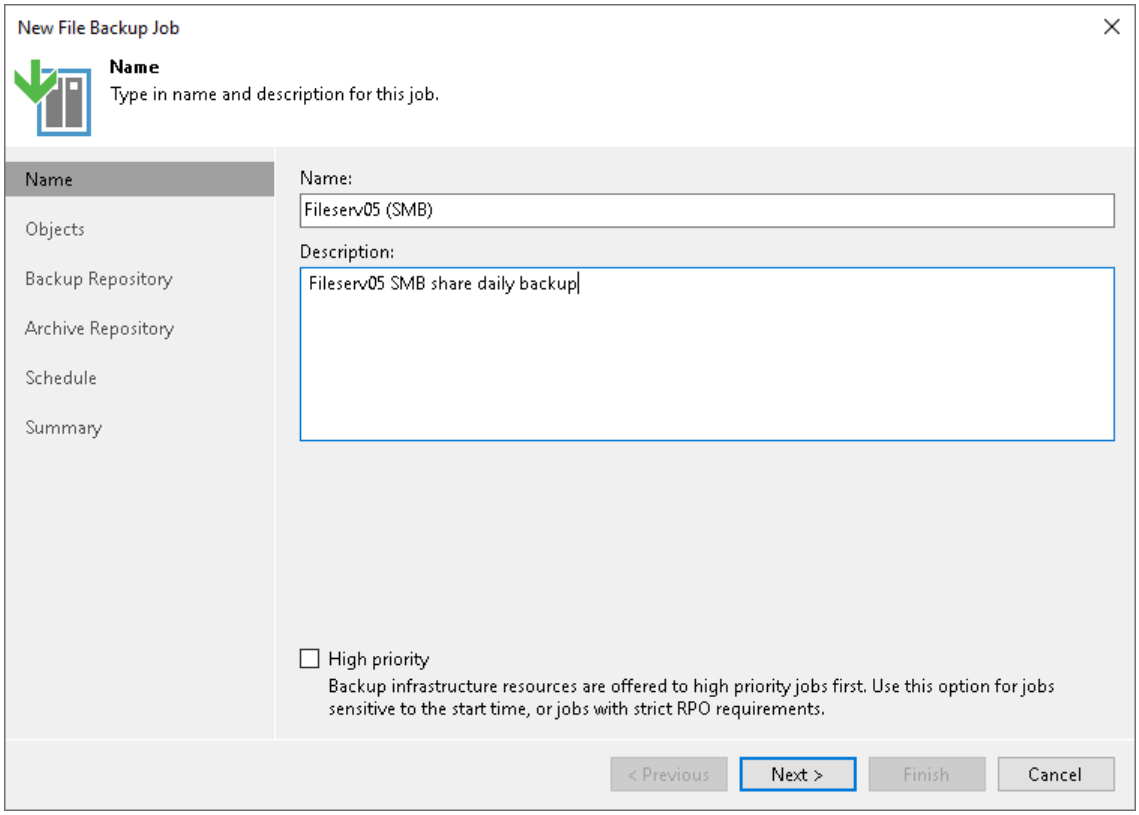

# Step 3. Select Files and Folders to Back Up

At the **Objects** step of the wizard, select files and folders that you want to back up.

1. Click **Add**.
2. From the **Server** list, select a file share on which the necessary files or folders reside.

   > **NOTE**
   >
   > If you plan to protect file shares residing on the enterprise storage system, you can choose between two different approaches. For more information, see the File Backup Integration with Storage Systems section.

3. In the **Objects** tree, select folders you want to back up.

   > **NOTE**
   >
   > If you select a NAS filer from the **Server** list, but the **Objects** tree is empty, make sure that the storage system rescan was performed and finished. For more information about the storage system rescan process, see the Storage Discovery Process section in the Storage System Snapshot Integration Guide.

   To select multiple folders, hold [Ctrl] and click necessary folders. Although different folders of the same share form separate records in the table, they will be processed by one job task.

   Consider the following:

   - Hard links in file shares are protected with content included.
   - Symbolic links in file shares are protected as links, without the content they refer to.

4. If you add a folder to the job, all the folder contents will be processed.

   If necessary, you can choose only specific files from the added folder.

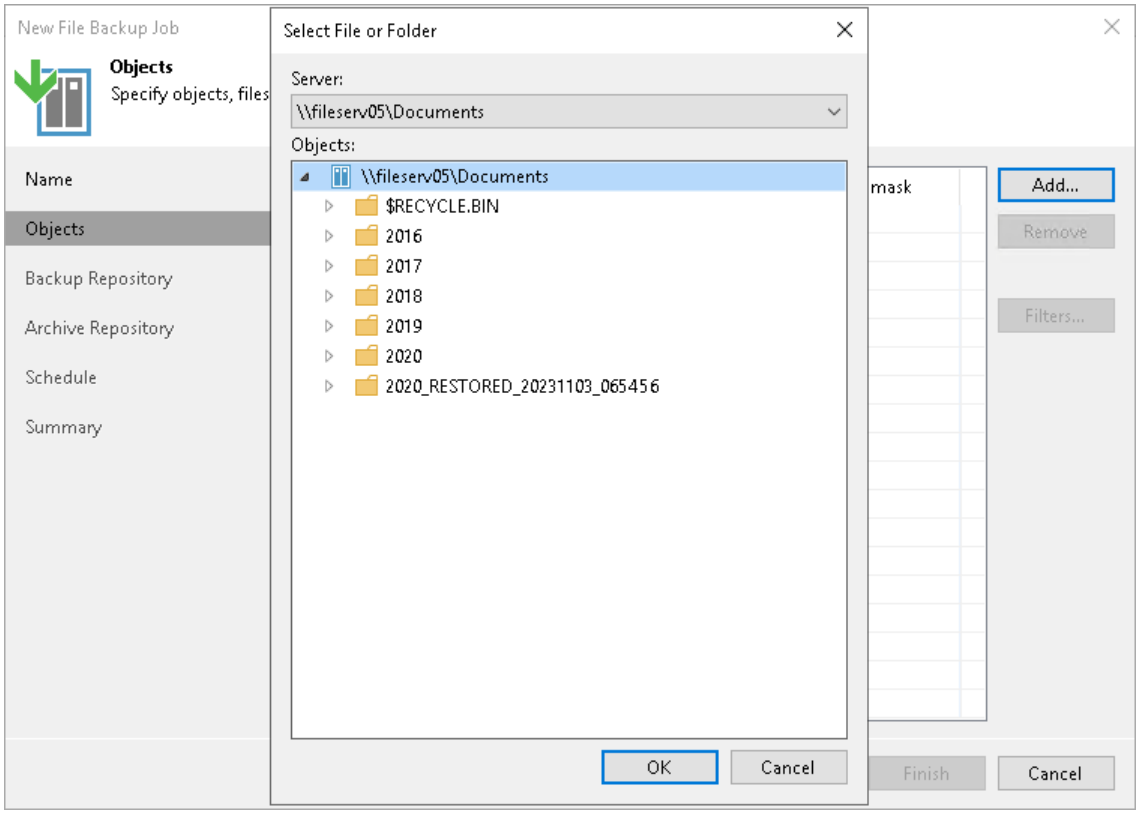


# Including Objects

To filter objects that you want to back up, you can specify the object name or extension mask to include into processing.

1. Select an object in the **Objects** list and click **Filters**.
2. In the **Include and exclude masks** window, use the **Include masks** section to include objects.
3. Click **Add**.
4. In the **Files by mask** window, specify the value associated with the object. For example:
   - To backup only PDF files, enter `*.PDF` in the **Extension** field and click **OK**.
   - To backup all files with name **sales_report.xlsx**, enter `sales_report.xlsx` in the **Extension** field and click **OK**.

5. Click **OK**.

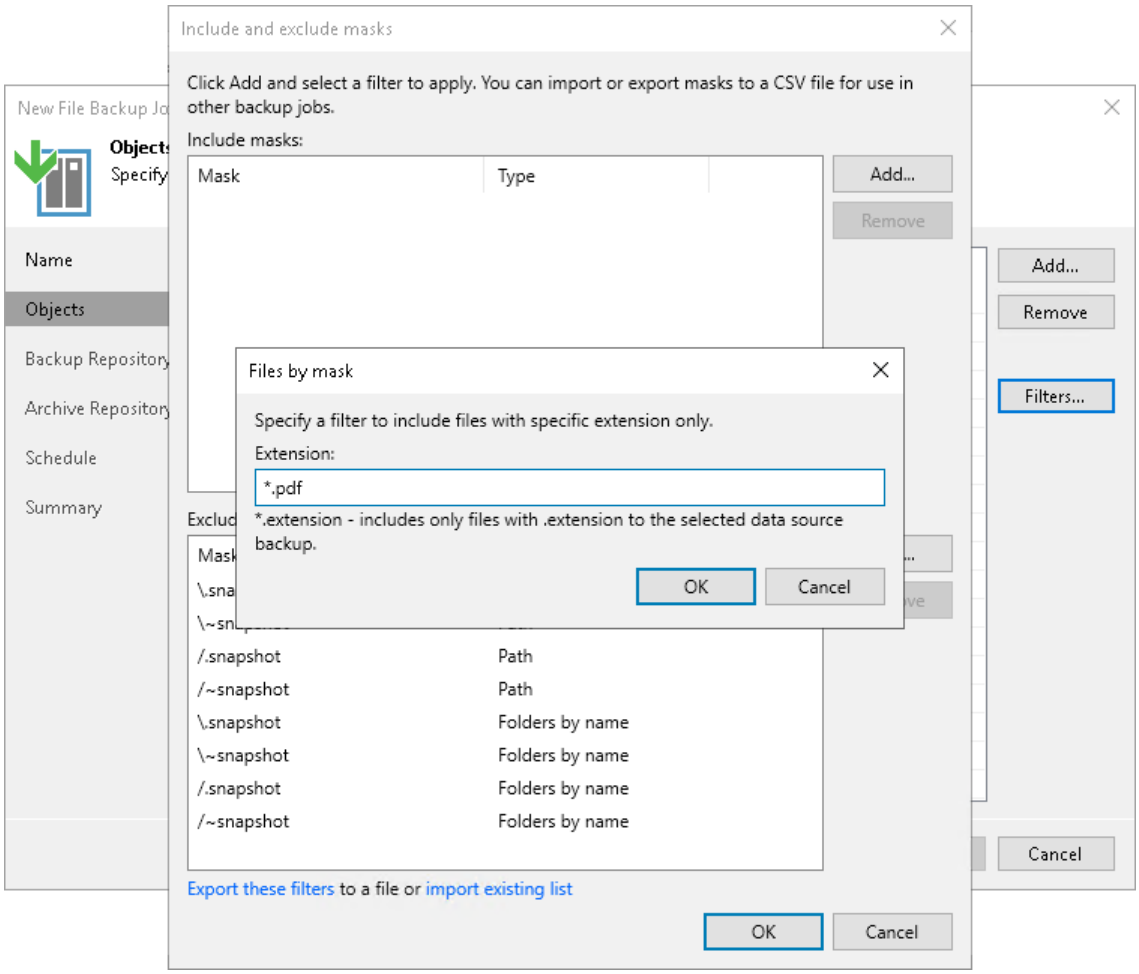


# Excluding Objects

To filter objects that you do not want to back up, you can specify the path or file mask to exclude from processing.

1. Select an object in the **Objects** list and click **Filters**.
2. In the **Include and exclude masks** window, use the **Exclude masks** field to exclude objects.
3. Click **Add**.
4. Depending on the way you want to exclude objects, select one of the following:
   - To exclude object by path, select **Object by path** and specify a filter to exclude objects based on their path.

     > **NOTE**
     >
     > Wildcards are not supported.

   - To exclude files by mask, click **Files by mask** and specify a filter to exclude files with specific extension.

- To exclude multiple objects by path, click **Multiple Objects by Path** and specify a filter to exclude folder with the matching name.

For example:

- To exclude PDF files from processing, select the **Files by mask** option, enter `*.PDF` in the **Extension** field and click **OK**.
- To exclude all files with name **pricelist.xlsx** from processing, select the **Files by mask** option, enter `pricelist.xlsx` in the **Extension** field and click **OK**.
- To exclude the folder from processing, select the **Multiple Objects by Path** option, enter `\folder` for SMB shares or `/folder` for NFS shares, and click **OK**.
- To exclude folder **2016** from processing, select the **Object by path** option, enter the full path to it in the **Path** field and click **OK**. For example, for an NFS file share this path looks like: `QA04:/NFS04/Documents/2016`, for an SMB file share — `\\fileserv05\Documents\2016`).

  Alternatively, you can specify a relative path to the folder or file to exclude. In case of the NAS filer, this is the only option to specify a path to exclude. For example, to exclude folders **call_records** (where call records are located) from all file shares residing on the NAS filer, in the **Path** field of the **Object by path** window, specify `/call_records` (this mask will exclude the **call_records** folder from processing for all NFS file shares on this NAS filer) and `\call_records` (this mask will exclude the **call_records** folder from processing for all SMB file shares on this NAS filer) and click **OK**.

  You can exclude a whole file share from processing. For example, you add the `\\Server\` SMB server to the file backup job, but you want to exclude the `\\Server\Sharing` file share from processing. To exclude this file share from processing, in the **Path** field of the **Object by path** window, enter the `\\Server\Sharing` path and click **OK**. The job will back up all file shares that Veeam Backup & Replication detects on this server, but will skip the excluded file share.

**NOTE**

Consider the following:

- Include and exclude masks are case insensitive.
- To exclude a junction from the NAS Filer backup, use the **Files by mask** option and specify the folder name. Note that all folders and files with the same name will be excluded as well.
- You cannot exclude a whole file share from processing if the storage system, where file shares reside, is added to Veeam Backup & Replication as a NAS filer.

  As a workaround, you can edit the storage in Storage Infrastructure and exclude the volumes on the **NAS Filer** step, as described in the Adding NetApp Data ONTAP, Adding Lenovo ThinkSystem DM/DG Series, Adding Dell PowerScale, or Adding Nutanix Files Storage sections, depending on the type of the storage system you use.
- You cannot use mask with * to specify folders to exclude from processing. For example, mask `QA04:/NFS04/Documents/201*` will not work.
- You cannot mix different exclusion options, for example, you cannot use a mask to exclude files with certain extensions from the specific folder. For example, `QA04:/NFS04/Documents/2016/*`.xlsx will not work.

**TIP**

You can follow this set of examples for specifying the path in the **Object by path** window:

- For SMB shares:
  - `\\server\share\folder\` includes/excludes the folder on the specific SMB server.
  - `\share` includes/excludes all shares with the matching name.
  - `\share\folder` includes/excludes only the specific folder in all shares with the matching name.
- For NFS shares:
  - `server:/share/folder/` includes/excludes the folder on the specific NFS server.
  - `/share/` includes/excludes all shares with the matching name.
  - `/share/folder/` includes/excludes only the specific folder in all shares with the matching name.

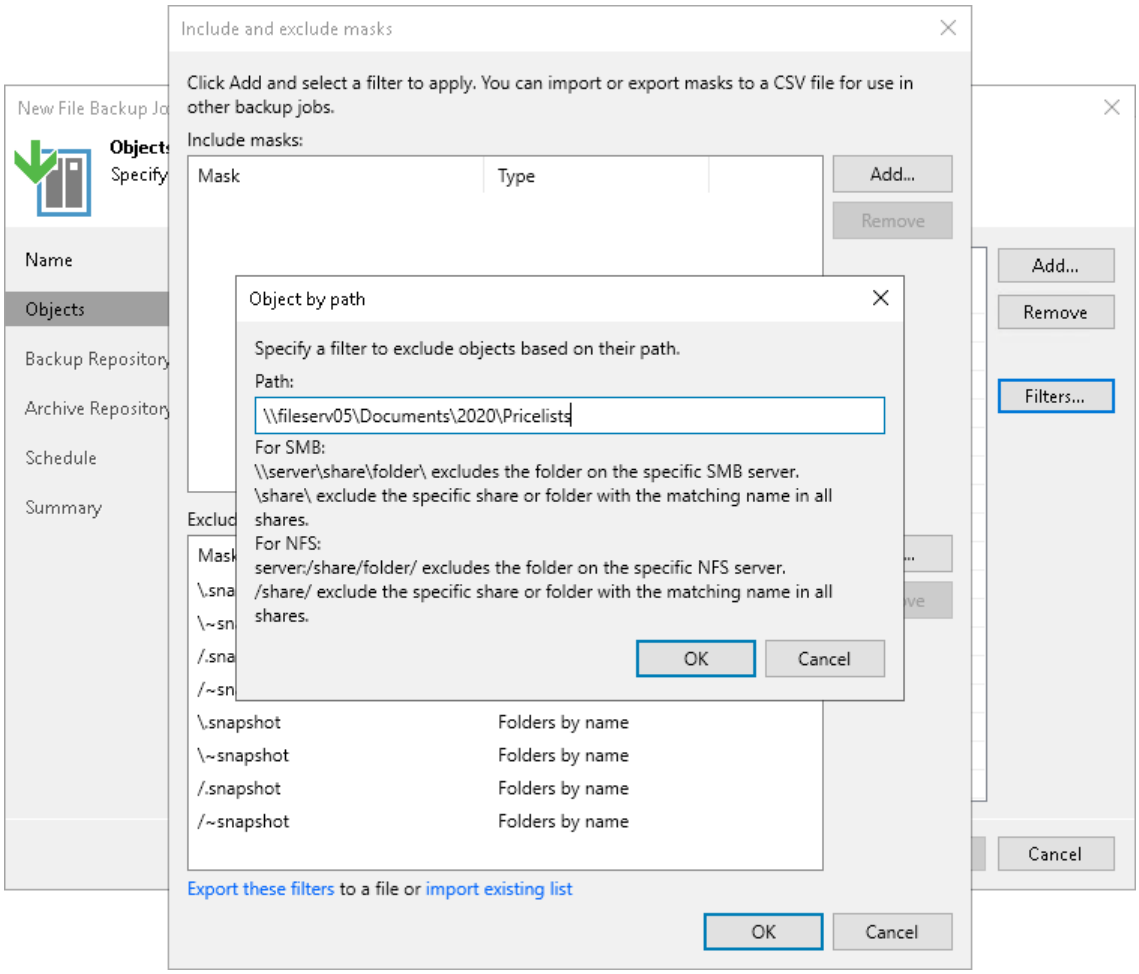


# Exporting and Importing Filters

If necessary, you can export and import your masks:

- To export a mask to a file, click the **Export these filters** link. In the **Export to file** window, specify a path to the necessary XML file. Click **OK**.
- To import existing masks from a file, click the **Import existing list** link. In the **Import masks from file** window, specify a path to the necessary XML file. Click **OK**.

# Step 4. Specify Backup Repository Settings

At the **Backup Repository** step of the wizard, define the primary backup repository where the file backup job must store backup files, and settings for moving files and folders to this repository. To learn what storage types you can assign the role of the backup repository to, see Storage Repositories in the **Backup Infrastructure for Unstructured Data Backup** section.

> **NOTE**
>
> Consider that if you use the option of limiting the number of file versions to keep configured in File Version Settings, Veeam Backup & Replication first applies those file-version retention settings and only after that applies time-based retention settings specified at this step.

1. From the **Backup repository** drop-down list, select a repository where backup files must be stored. When you select a backup repository, Veeam Backup & Replication automatically checks the amount of free space left. Make sure that you have enough free space to store backups.

   > **NOTE**
   >
   > If you plan to store backups in Veeam Data Cloud Vault, you must enable encryption.

2. You can map the job to a specific backup stored in the backup repository. Backup job mapping allows you to move backup files to a new backup repository and to point the job to existing backups on this new backup repository. You can also use backup job mapping if the configuration database got corrupted and you need to reconfigure backup job settings.

   To map the job to a backup, click the **Map backup** link. In the opened **Select Backup** window, select a backup in the backup repository. Backups can be easily identified by job names. To find the backup, you can also use the search field at the bottom of the window.

3. Use the **Keep all file versions for the last** field to specify how long copies of all recent file versions in the selected file share must be kept in the backup repository. You can restore the entire file share to any restore point within the period specified in this setting.

   If, for example, **Keep all file versions for the last** is set to 30 days, the backup repository will store all file versions that appeared at the file share during the last 30 days. At the scheduled time on the 31st day, the file backup job first backs up new file versions and saves them to the backup repository. Right after that, file versions older than 30 days (created on the 1st day) are either deleted from the backup repository or moved to the archive repository. File versions are moved to the archive repository, if at the Archive Repository step of the wizard you enable the **Archive file versions to the following archive repository** check box and configure the archive retention.

4. If you need to keep a copy of the backups in another repository, select the **Configure secondary destinations for this job** check box. That enables the Secondary Target step of the wizard.

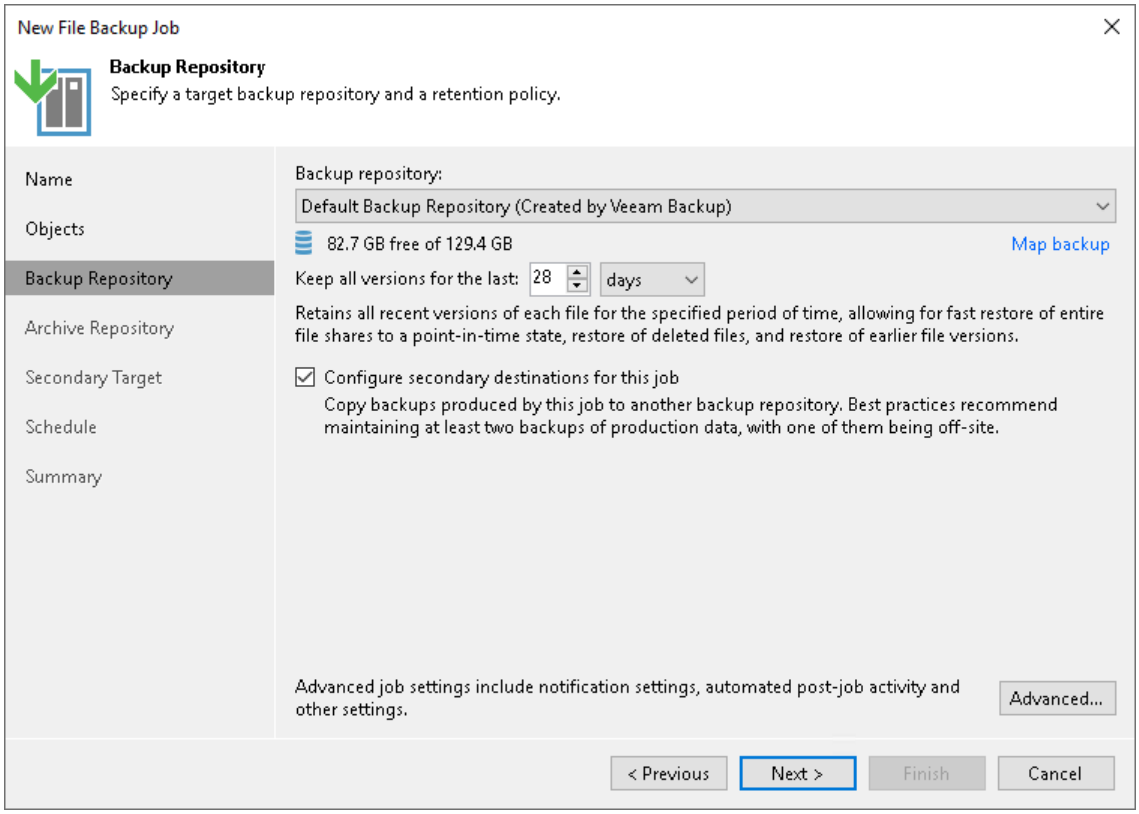

# Step 5. Specify Advanced Backup Settings

At the **Backup Repository** step of the wizard, specify advanced settings for the file backup job:

- File version settings
- ACL handling settings
- Storage settings
- Maintenance settings
- Script settings
- Notification settings

> **TIP**
>
> After you specify necessary settings for the backup job, you can save them as default settings. To do this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new backup job, Veeam Backup & Replication will automatically apply the default settings to the new job.

## File Version Settings

To configure how many file versions to keep for protected files, do the following:

1. At the **Backup Repository** step of the wizard, click **Advanced**.
2. On the **File Versions** tab, specify to what file versions the settings should apply:
   - Select **Keep all file versions** to keep all file versions for the time period specified in the main window at the **Storage** step.
   - Select **Limit the number of archived file versions only** to limit archived file versions to the numbers specified to the right of the **Active file versions limit** and **Deleted file versions limit** check boxes.
   - Select **Limit the number of both recent and archived file versions** to limit recent and archived file versions to the numbers specified to the right of the **Active file versions limit** and **Deleted file versions limit** check boxes.
3. After you choose what file versions to keep, specify how many file versions to keep:
   - Select **Active file versions limit** to keep the specified number of versions for files currently existing in the source file share. Specify how many file versions to store.
   - Select **Deleted file versions limit** to keep the specified number of versions for files deleted from the source file share. Specify how many file versions to store.

4. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

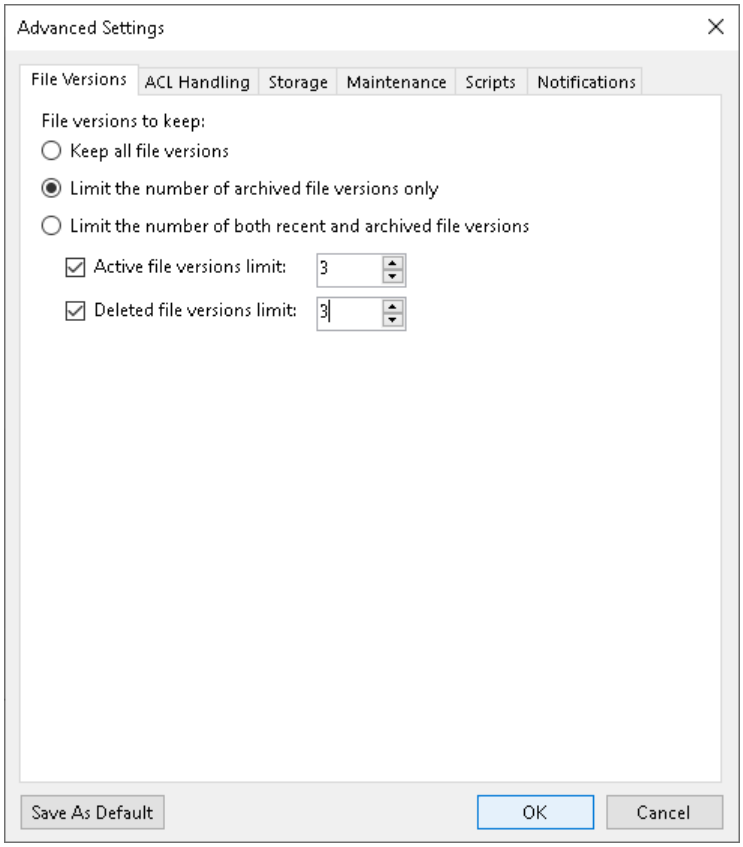


## ACL Handling Settings

To specify how the backup job will process permissions and attributes:

1. At the **Backup Repository** step of the wizard, click **Advanced**.
2. On the **ACL Handling** tab, specify how the backup job will process permissions and attributes:
   - Select **Folder-level only (recommended)** to back up permissions and attributes from folders only. The restored files will inherit permissions from the target folder.
   - Select **Files and folders (slower)** to back up permissions and attributes from both folders and individual files. This option can significantly reduce the backup performance.

   If you change the ACL handling settings for an existing file backup job, Veeam Backup & Replication will create a full backup file during the next job session.

3. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

> **NOTE**
>
> Consider that Veeam Backup & Replication does not collect ACL handling settings of the source file share root folder, so you cannot restore them. Before restoring an entire file share, you will have to specify required ACL handling settings for the root folder of the target file share.

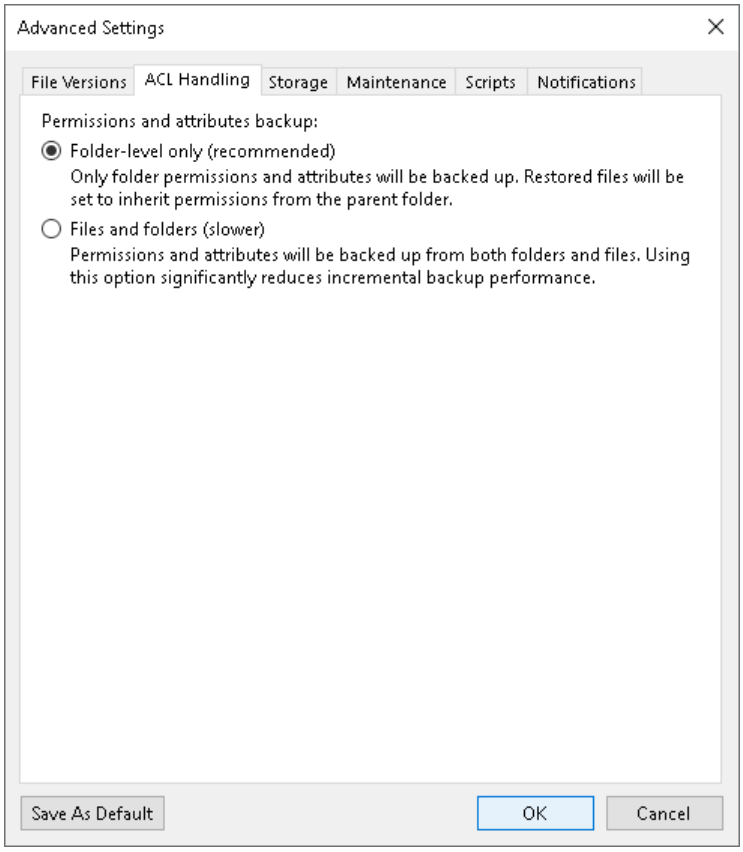


## Storage Settings

To specify advanced storage settings for the file backup job:

1. At the **Backup Repository** step of the wizard, click **Advanced**.
2. On the **Storage** tab, specify data reduction and encryption settings:
   - From the **Compression level** list, select a compression level for the backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*.
   - To encrypt backup files, select the **Enable backup file encryption** check box. In the **Password** field, select the encryption method you want to use — a password that will generate secret keys or an external KMS server that will generate KMS keys. For more information about data encryption, see How Data Encryption Works.

     If you have not created the password beforehand, click **Add** or use the *Manage passwords* link to specify a new password. For more information, see Password Manager.

     If you do not use password loss protection, Veeam Backup & Replication will display a warning about it. For more information, see Password Loss Protection.

> **NOTE**
>
> Consider the following:
>
> - If you enable encryption for an existing backup job, during the next job session Veeam Backup & Replication will back up all the files of the file share to a new backup file irrespective of whether they changed or not. The created backup files and subsequent backup files will be encrypted with the specified password.
> - If you enable encryption for an existing job, Veeam Backup & Replication does not encrypt the previous backup chain created by this job.
> - You can also use KMS keys for encryption. For more information, see Key Management System Keys.
> - If you plan to store backups in Veeam Data Cloud Vault, you must enable job encryption.

3. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

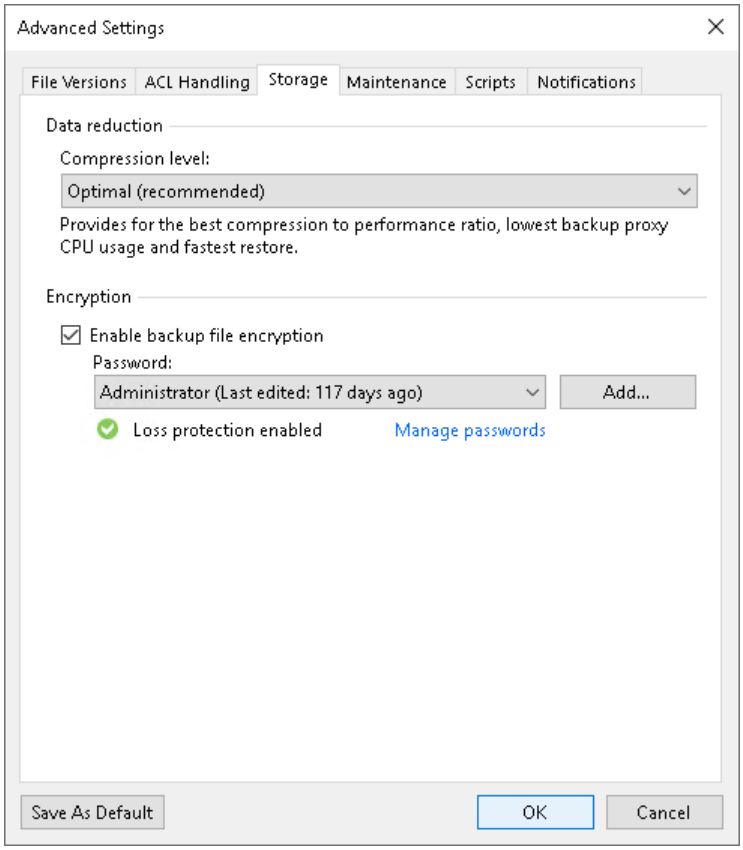


## Maintenance Settings

You can instruct Veeam Backup & Replication to periodically perform a health check for the backup. The health check helps make sure that the backup is consistent, and you will be able to restore data from it.

During the health check, Veeam Backup & Replication performs a CRC check for metadata and a hash check for data blocks in the file share backup files to verify their integrity. For more information, see the Performing Health Check and Repair for Unstructured Data Backups section.

To configure the health-check settings for the backup job:

1. At the **Backup Repository** step of the wizard, click **Advanced**.

2. On the **Maintenance** tab, select **Perform backup files health check** to enable the health check option. It allows ensuring that all data and metadata is backed up correctly.

3. Click **Configure** and specify the time schedule for the health check.

4. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

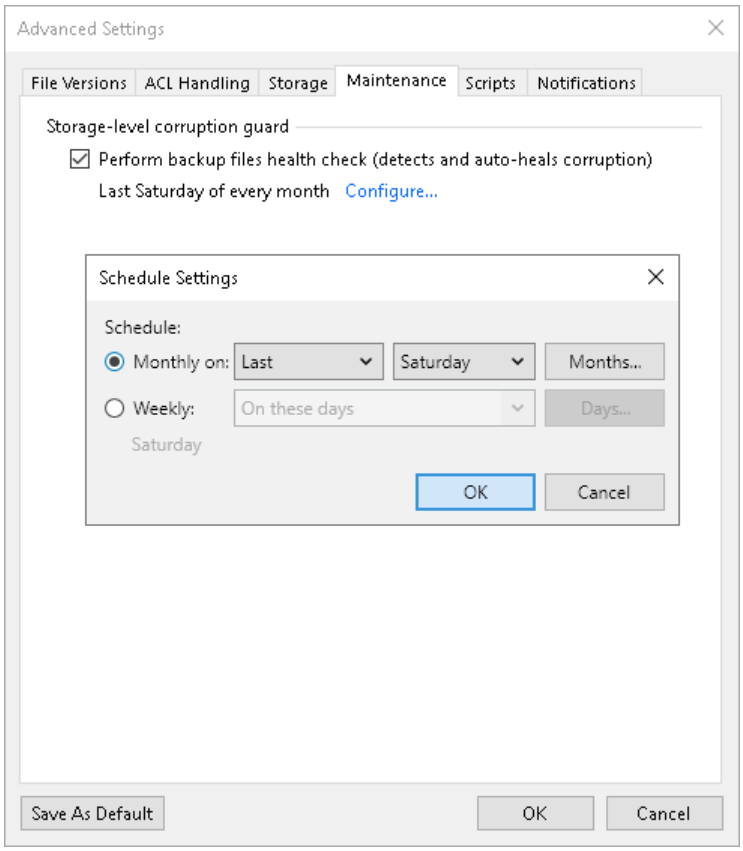


## Script Settings

You can configure custom scripts to run before or after the file backup job. For example, you can configure scripts to take a VSS snapshot before running the job and to delete it after completing the job.

To specify script settings for the backup job:

1. At the **Backup Repository** step of the wizard, click **Advanced**.

2. Click the **Scripts** tab.

3. If you want to execute custom scripts, select the **Run the following script before the job** and **Run the following script after the job** check boxes and click **Browse** to choose executable files from a local folder on the backup server. The scripts are executed on the backup server.

   You can select to execute pre- and post-backup actions after a number of backup sessions or on specific week days.

   - If you select the **Run scripts every <N> backup session** option, specify the number of the backup job sessions after which the scripts must be executed.

   - If you select the **Run scripts on the selected days only** option, click **Days** and specify week days on which the scripts must be executed.

4. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

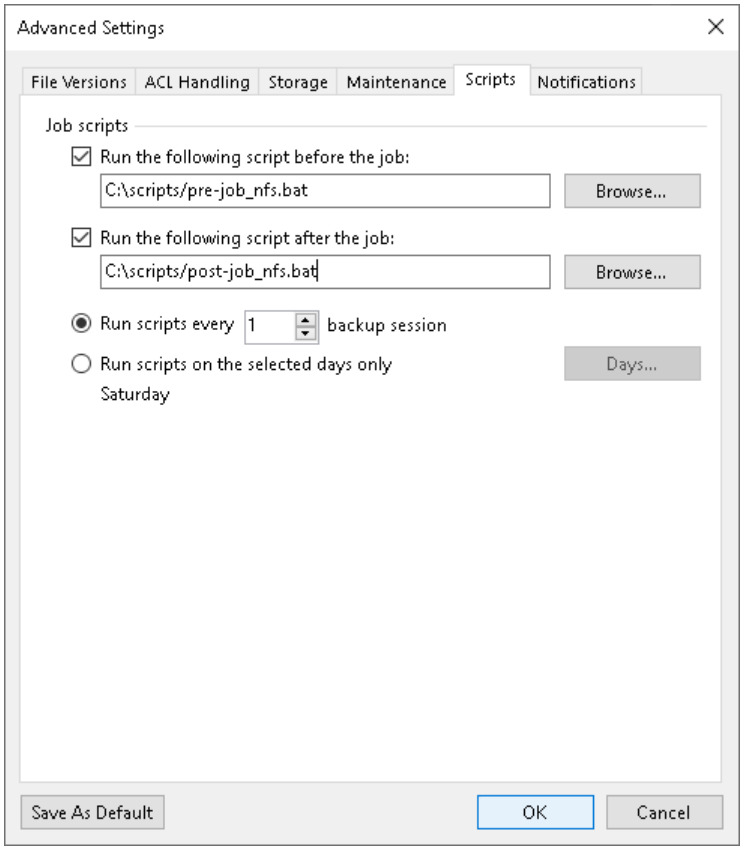


## Notification Settings

To specify notification settings for the backup job:

1. At the **Backup Repository** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully.

   SNMP traps will be sent if you specify global SNMP settings in Veeam Backup & Replication and configure software on recipient's machine to receive SNMP traps. For more information, see Specifying SNMP Settings.
4. Select the **Send email notifications to the following recipients** check box if you want to receive notifications about the job completion status by email. In the field under the check box, specify the recipient email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication. For more information, see Configuring Global Email Notification Settings.
5. You can choose to use global notification settings or specify custom notification settings:
   - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see Configuring Global Email Notification Settings.
   - To configure a custom notification for the job, select **Use custom notification settings specified below**. You can specify the following notification settings:

i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%Time%* (completion time), *%JobName%*, *%JobResult%*, *%ObjectCount%* (number of file shares in the job) and *%Issues%* (number of files shares in the job that have finished with the *Warning* or *Failed* status).

ii. Select the **Notify on success**, **Notify on warning**, and **Notify on error** check boxes to receive email notification if the job completes successfully, completes with a warning or fails.

iii. Select the **Suppress notifications until the last retry** check box to receive a notification about the final job status. If you do not enable this option, Veeam Backup & Replication will send one notification per every job retry.

6. Specify severity options if required:

   - Select the **Make file or folder processing issues result in a job warning** to receive a warning at the end of the job processing session if any issues with file or folder processing occur.
   - Select the **Make file or folder attribute's processing issues result in a job warning** to receive a warning at the end of the job processing session if any issues with processing of file or folder attributes occur.

7. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

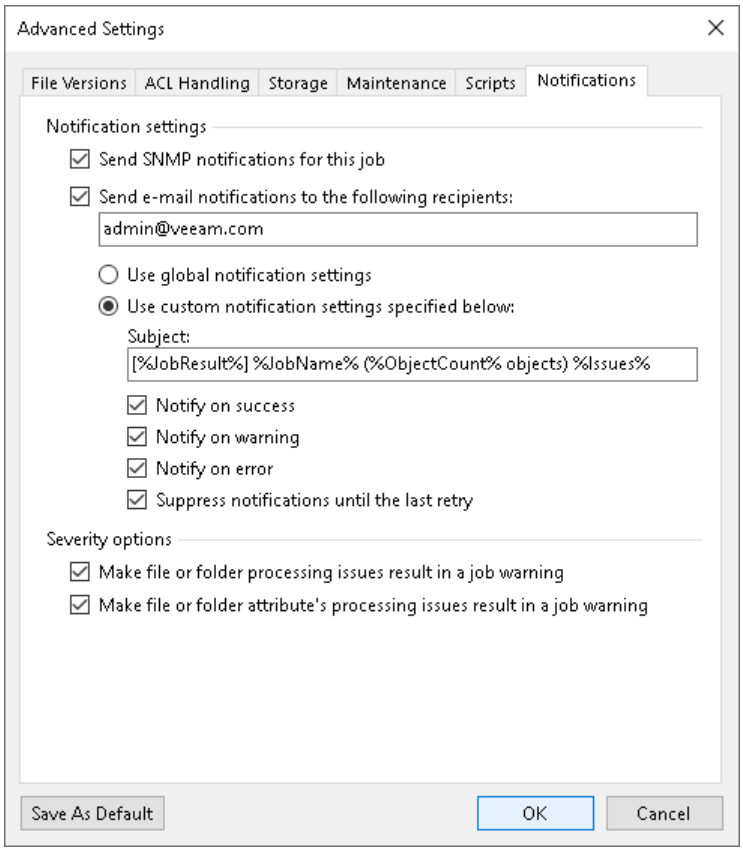

# Step 6. Specify Archive Repository Settings

At the **Archive Repository** step of the wizard, define the archive repository, where the file backup job must store backup files, and settings for moving or copying files and folders to this repository. To learn what storage types you can assign the role of the archive repository to, see Storage Repositories in the **Backup Infrastructure for Unstructured Data Backup** section.

1. If you need to keep versions of some files for a longer time after they are moved from the backup repository, you can configure archiving options to move files versions to the archive repository. You can also copy the recent file versions to the archive repository to store them according to the backup repository retention policy.

   To use the archive repository, select the **Archive file versions to the following archive repository** check box.

2. From the drop-down list under the **Archive file versions to the following archive repository** check box, select the storage to be used as a repository to store archived files and folders.

   By default, all files deleted from the backup repository will be moved to the archive repository. If you do not need all the files in the archive, you can choose what files to keep.

   > **NOTE**
   >
   > If you plan to store backups in Veeam Data Cloud Vault, you must enable encryption for file backup jobs.

3. If you need to keep the copy of the data stored in the backup repository also in the archive repository, select the **Archive recent file versions** check box.

   When you archive the recent file versions, Veeam Backup & Replication immediately copies all backed up files from backup repository to the archive repository and stores them according to the backup repository retention policy. For example, if backup repository is configured to store file versions for 20 days, archive repository will also store these file versions for 20 days.

   > **NOTE**
   >
   > The copy mode in unstructured data backup requires a license. Thus, this feature is not supported in the Veeam Backup & Replication Community (free) Edition. For details, see Veeam Editions Comparison.

4. To specify the number of months or years during which backup files must be retained, select the **Archive previous file versions for** check box and specify the period.

   The period setting denotes a time period starting from the creation of the backup files in the backup repository, not from the moment when the file versions are moved from the backup repository to the archive repository.

   When you archive the previous file versions, Veeam Backup & Replication moves backup and metadata files after their retention period is over from backup repository to the archive repository and stores them according to the archive repository retention policy. For example, if backup repository is configured to store backup files versions for 20 days, Veeam Backup & Replication moves backup files to the archive repository on the 21 day.

5. To specify what files must be archived or excluded from the archive, do the following:

   a. Click **Choose** to open the **File Archive Settings** window.

   b. Under **Files to archive**, specify what files must be archived:

- **All files** – select this option to archive all files moved from the backup repository to the archive repository.
- **All files except the following extensions** – select this option to exclude files with certain extensions from the selection to be archived. Specify extensions for files to exclude from the selection. Files with the specified extensions will not be copied from the backup repository to the archive repository.
- **Files with the following extensions only** – select this option to archive files with certain extensions only. Specify extensions for files to archive. Files with these extensions will be copied from the backup repository to the archive repository.

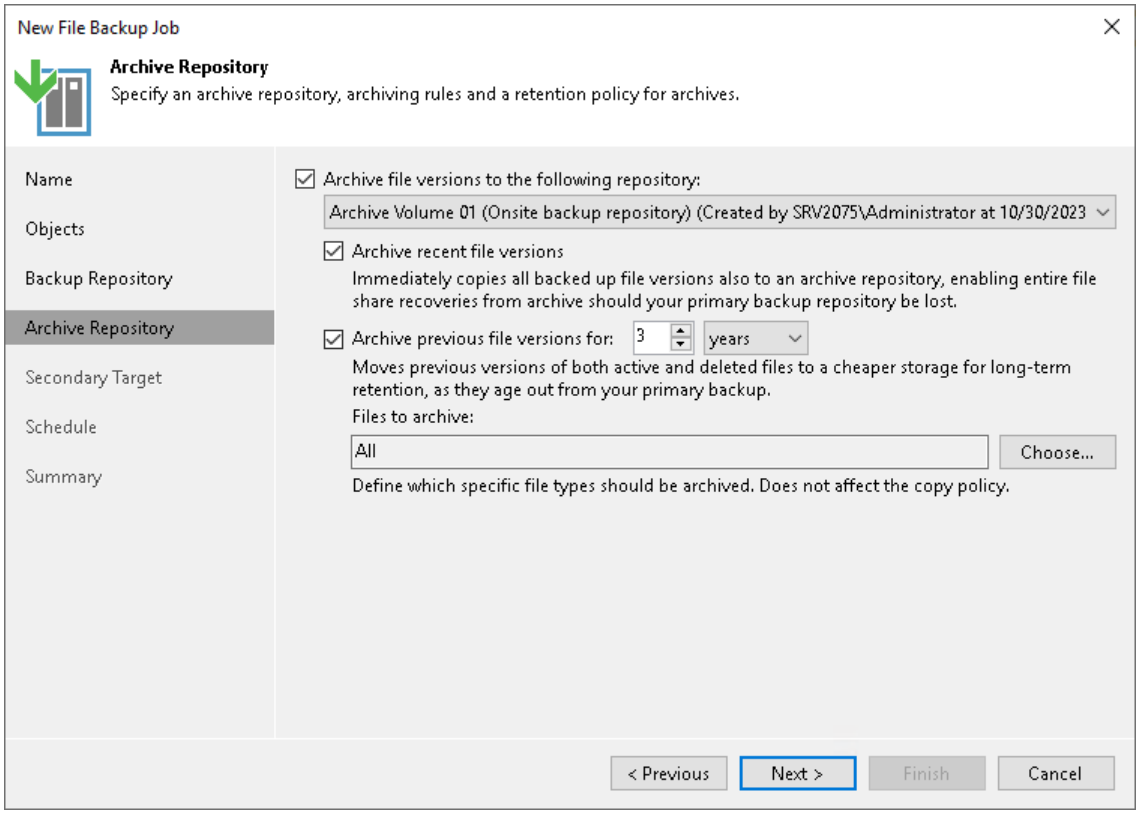

# Step 7. Specify Secondary Repository Settings

At the **Secondary Target** step of the wizard, you can specify a secondary repository that will be used to store additional copies of backup files from the backup storage for redundancy. To learn what storage types you can assign the role of the secondary repository to, see Storage Repositories in the **Backup Infrastructure for Unstructured Data Backup** section.

If you add a secondary repository, Veeam Backup & Replication will create a separate job for backup copy to it. The data copy process will start automatically after each primary job runs.

> **NOTE**
>
> This step is available, if you select the **Configure secondary destinations for this job** check box at the Backup Repository step of the wizard.

To add a secondary repository:

1. Click **Add**.
2. From the list of existing repositories, select a repository that will keep additional copy of the backup files. You can add several secondary repositories for copying files of the primary backup job. To quickly find the repository, use the search field at the bottom of the wizard.

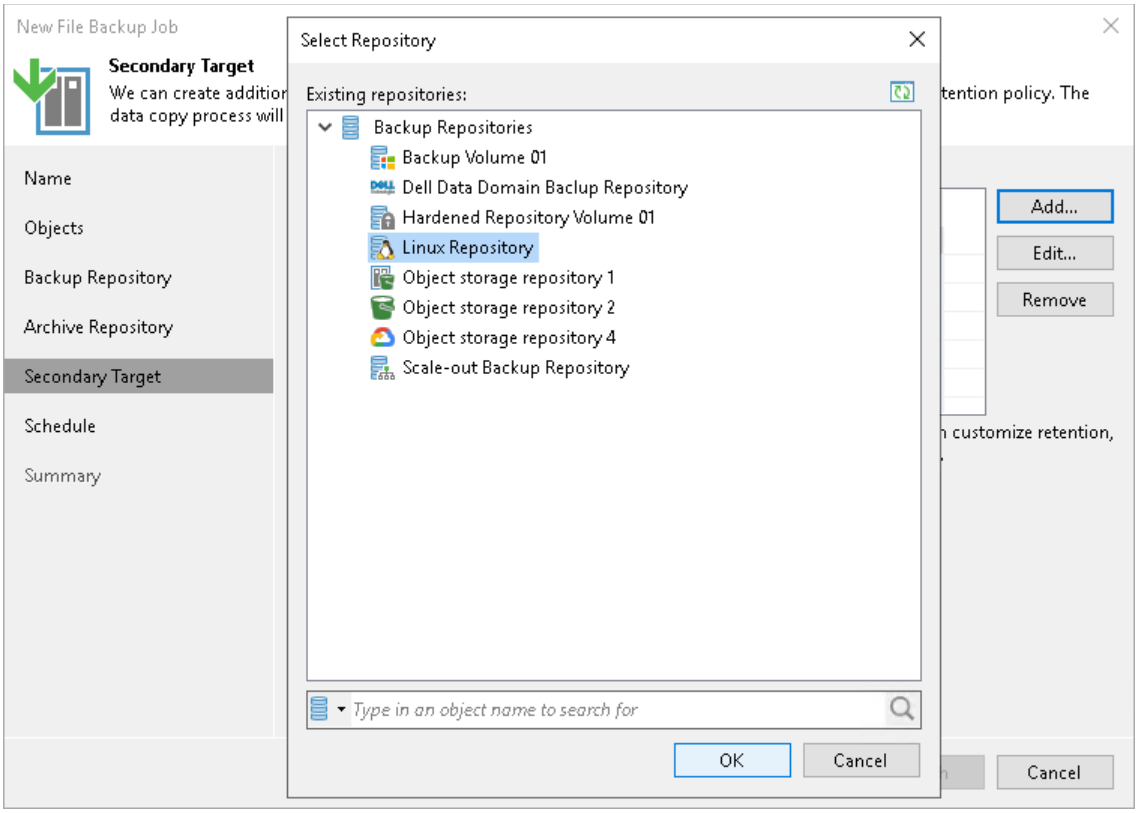


3. By default, retention and encryption settings for the secondary target repository are inherited from the primary job. To customize them, select the necessary repository in the **Secondary repositories** list and click **Edit**.
   - To enable custom retention settings:
     i. Select **Use custom retention policy**.
     ii. Specify how long all versions of each file will be kept in the secondary repository.

- To specify encryption settings different from those of the primary repository:

  i. Select **Use custom data encryption settings**.

  ii. In the **Password** field, select a password that you want to use for encryption. If you have not created the password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Password Manager.

  If the backup server is not connected to Veeam Backup Enterprise Manager and does not have the Veeam Universal License or a legacy socket-based Enterprise or Enterprise Plus license installed, you will not be able to restore data from encrypted backups in case you lose the password. Veeam Backup & Replication will display a warning about it. For more information, see Decrypting Data Without Password.

- Configure time intervals at which the data can be copied to the secondary repository.

  - If you select the **Any time (continuously)** option, Veeam Backup & Replication will copy backup files to the secondary repository as soon as the primary file backup job completes.
  - If you want to specify time periods when it is permitted to start copying backup files to the secondary repository, select the **During the following time periods only** option and configure allowed and prohibited hours. These periods do not work as the backup window, so they will not cause the file backup copy to fail.

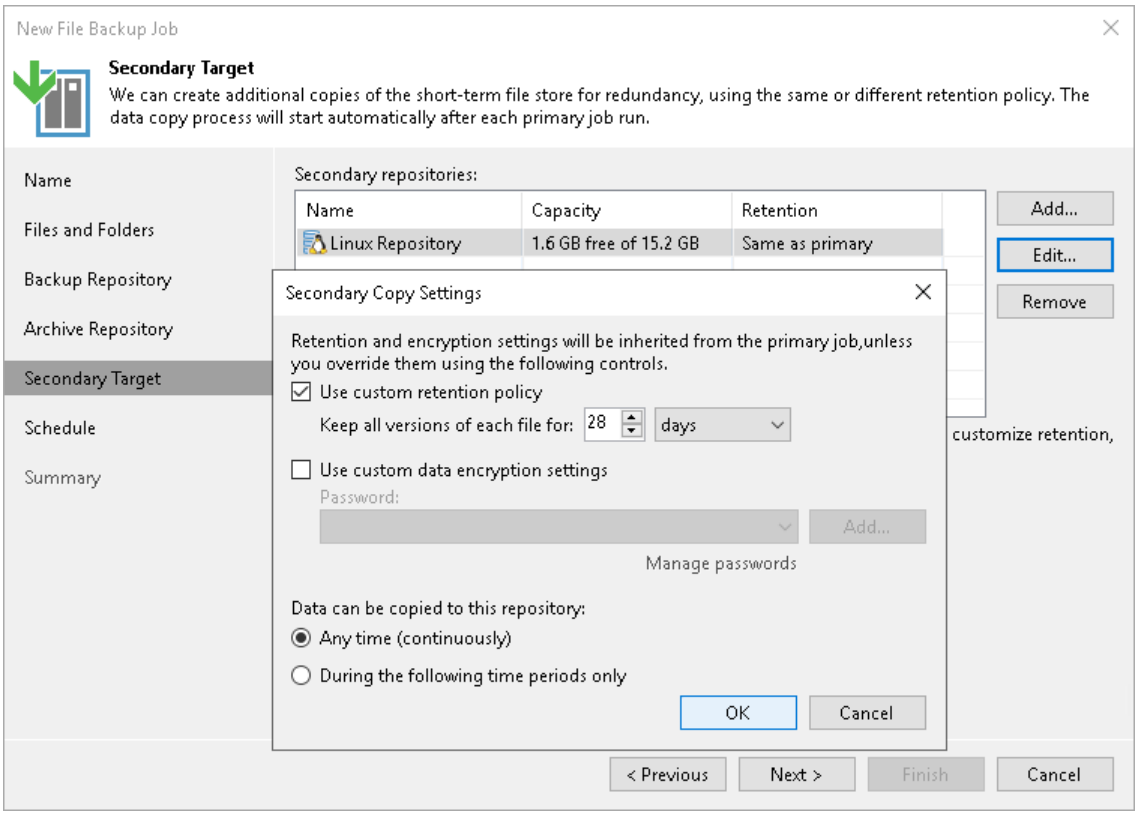

# Step 8. Define Job Schedule

At the **Schedule** step of the wizard, select to run the backup job manually or schedule the job to run on a regular basis.

To specify the job schedule:

1. Select the **Run the job automatically** check box. If this check box is not selected, you will have to start the job manually to create the file share backup.
2. Define scheduling settings for the job:
   - To run the job at specific time daily, on defined week days or with specific periodicity, select the **Daily at this time** option. Use the fields on the right to configure the necessary schedule.
   - To run the job once a month on specific days, select the **Monthly at this time** option. Use the fields on the right to configure the necessary schedule.

     > **NOTE**
     >
     > When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a specific time interval, select the **Periodically every** option. In the field on the right, select the necessary time unit: **Hours** or **Minutes**. Click **Schedule** and use the time table to define the permitted time window for the job. In the **Start time within an hour** field, specify the exact time when the job must start.

     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run the job by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.
   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right. A new backup job session will start as soon as the previous backup job session finishes.
   - To chain jobs, use the **After this job** field. In the common practice, jobs start one after another: when job *A* finishes, job *B* starts and so on. If you want to create a chain of jobs, you must define the time schedule for the first job in the chain. For the rest of the jobs in the chain, select the **After this job option** and choose the preceding job from the list.
3. In the **Automatic retry** section, define whether Veeam Backup & Replication must attempt to run the backup job again if the job fails for some reason. During a job retry, Veeam Backup & Replication processes failed file shares only. Enter the number of attempts to run the job and define time intervals between them. If you select continuous backup, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job runs.

4. In the **Backup window** section, define the time interval within which the backup job must complete. The backup window prevents the job from overlapping with production hours and ensures that the job does not provide unwanted overhead on the production environment. To set up a backup window for the job:

   a. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**.

   b. In the **Time Periods** window, define the allowed hours and prohibited hours for backup. If the job exceeds the allowed window, it will be automatically terminated.

   Veeam Backup & Replication saves the progress of the terminated file backup job and continues the backup process during the next run.

> **NOTE**
>
> The **After this job** function will automatically start a job if the first job in the chain is started automatically by schedule. If you start the first job manually, Veeam Backup & Replication will display a notification. You will be able to choose whether Veeam Backup & Replication must start the chained job as well.

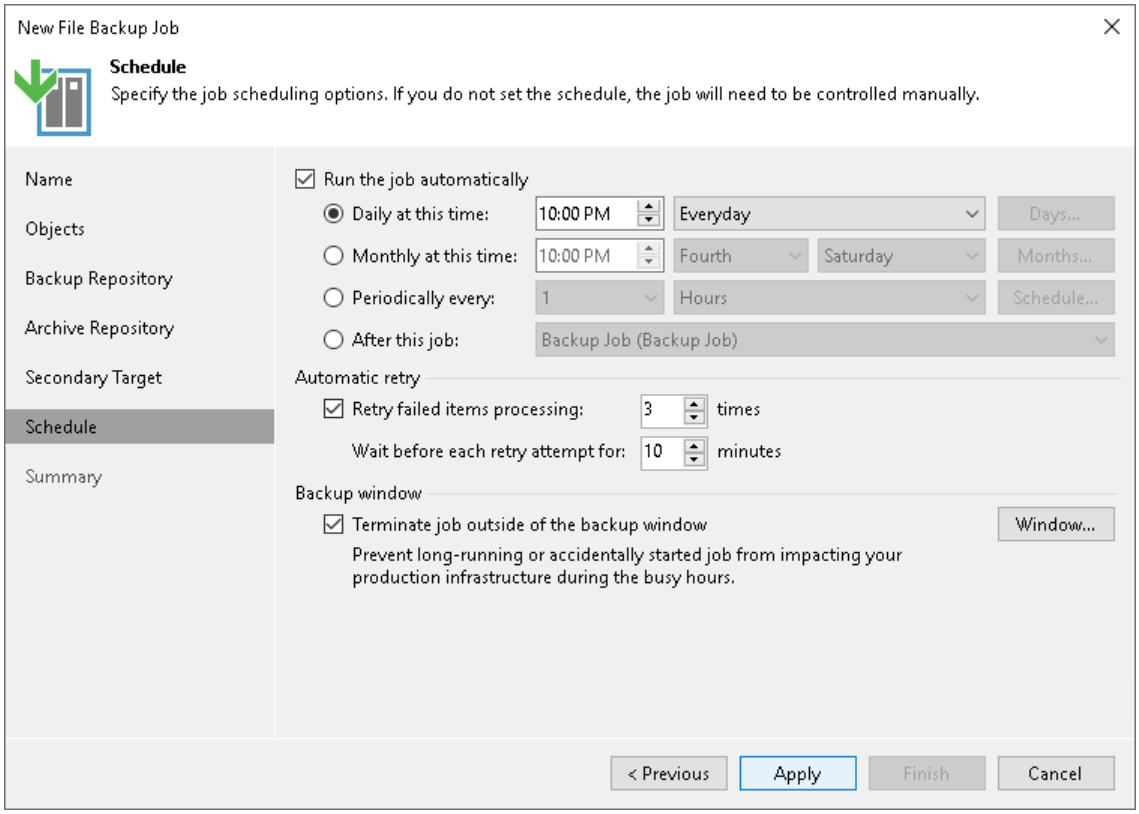

# Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of backup job configuration.

1. Review details of the backup job.
2. If you want to start the job right after you finish working with the wizard, select the **Run the job when I click Finish** check box.
3. Click **Finish** to close the wizard.

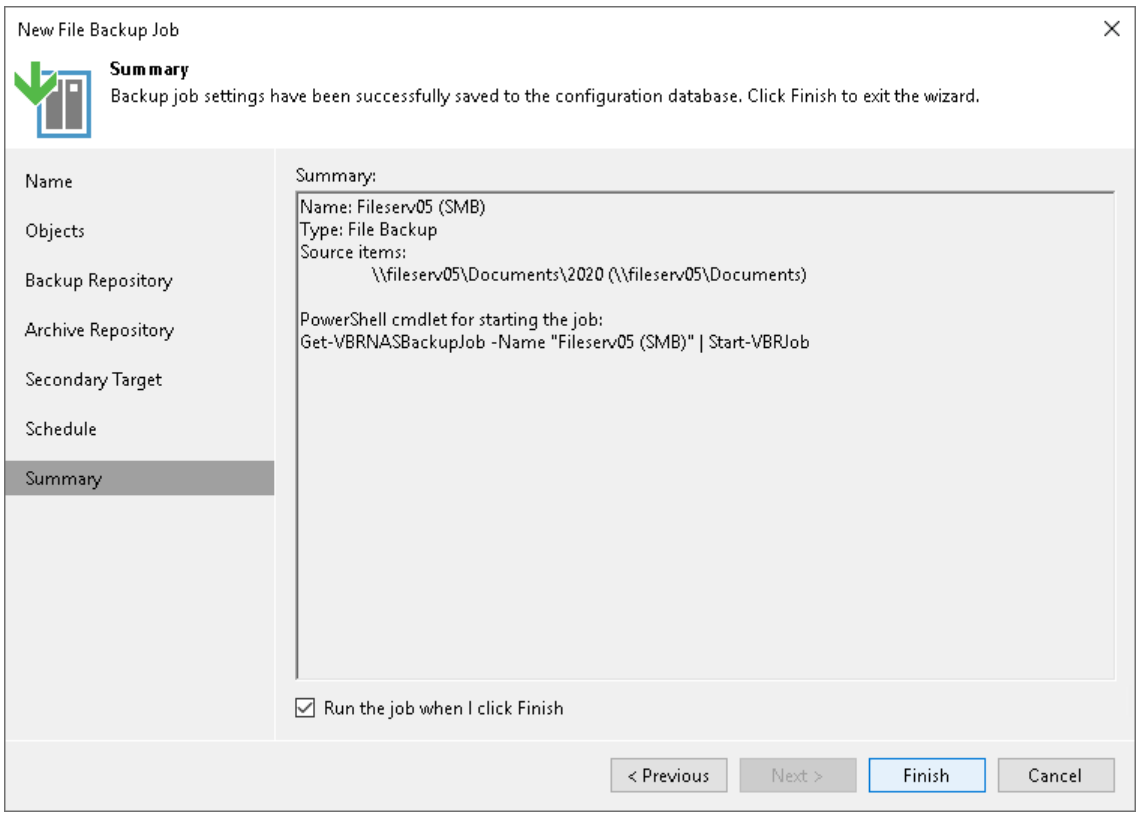

# Creating Object Storage Backup Jobs

To protect an object storage repository, configure an object storage backup job. The backup job defines how, where and when to back up data from the object storage. One job can be used to protect one or more object storage repositories. Jobs can be started manually or scheduled to run automatically at a specific time.

Object storage backup jobs are used to protect the following sources of unstructured data:

- S3 compatible object storage
- Amazon S3 object storage
- Microsoft Azure Blob storage including Azure Data Lake storage Gen2

Before you create an object storage backup job, check prerequisites.

## Before You Begin

This section contains information that you should consider before you create an object storage backup job.

## General Considerations

Before you create an object storage backup job, consider the following:

- Consider limitations listed in the Supported Platforms and Applications section.
- The account that you use to add Amazon S3 and S3 Compatible object storage as unstructured sources and back up data from them must be able to perform the actions listed in the Permissions section.
- Backup infrastructure components that will take part in the object storage backup process must be added to the backup infrastructure and properly configured. These include objects storage sources to back up, backup proxies, and all repositories, including cache, backup, archive, and secondary repositories. For more information, see the Backup Infrastructure for Unstructured Data Backup section.
- The target backup repository must have enough free space to store created backup files. If you want to receive notifications on the repository running low on free space, configure global notification settings as described in the Specifying Other Notification Settings section.
- Make sure that repositories intended to store object storage backups are not configured to store files in the Write Once Read Many (WORM) status. Otherwise, the backup jobs will fail when Veeam Backup & Replication cannot update the backup metadata files.
- If you plan to map an object storage backup job to a backup that already exists in the backup repository, you must perform the rescan operation for this backup repository. Otherwise, Veeam Backup & Replication will not be able to recognize backup files in the backup repository.

  For more information on how to rescan backup repositories, see the Rescanning Backup Repositories section.
- If you plan to use pre-job and post-job scripts, you must create scripts before you configure the object storage backup job.
- Antivirus software may significantly slow down object storage backup jobs. To improve performance, we recommend you exclude the `c:\Program Files (x86)\Veeam\Backup Transport\x64\VeeamAgent.exe` process from the antivirus scan on machines running the object storage backup proxy and backup repository roles. Keep in mind that it can weaken the security of these machines.

- If the object content and modification time are not changed, Veeam Backup & Replication will back up only attributes of this object.
- Veeam Backup & Replication does not back up version history of objects.
- Veeam Backup & Replication does not create the separate node in the inventory pane for the archived backups. If you plan to use archive repositories and want to make sure that backups were moved to it, you can do it using backup properties. For more information, see Viewing Unstructured Data Backup Properties.
- You cannot use the capacity tier of scale-out backup repositories as a target for object storage backup jobs.

# Limitations for Microsoft Azure Blob Storage

Before you create an object storage backup job, consider the following limitations for Microsoft Azure Blob storage:

- Veeam Backup & Replication backs up objects encrypted only with Microsoft-managed and customer-managed keys. If you create a blob using the Put Blob or Put Block List API method and specify a customer-provided key, Veeam Backup & Replication cannot back up such blobs. For more information about customer-provided keys, see this Microsoft article.
- Veeam Backup & Replication does not support backup of the objects with the *archive* access tier.
- Veeam Backup & Replication does not support the following backup options if you add the source Microsoft Azure Blob storage using the Microsoft Azure storage account with Microsoft Entra authorization:
  - Backup of ACL for Azure containers.
  - Backup of hierarchical namespace.

# Limitations for Amazon S3 Storage

Before you create an object storage backup job, consider the following limitations for Amazon S3 storage:

- Veeam Backup & Replication backs up objects encrypted only with Amazon S3 managed keys and AWS KMS keys. If you add an object to a bucket using the Put Object API method and specify a customer-provided key, Veeam Backup & Replication cannot back up such objects. For more information about customer-provided keys, see this Amazon article.
- Veeam Backup & Replication does not support backup of the objects with the *Amazon S3 Glacier Flexible Retrieval* and *Amazon S3 Glacier Deep Archive* storage classes.

# Step 1. Launch New Object Storage Backup Job Wizard

To launch the **New Object Storage Backup Job** wizard, do one of the following:

- On the **Home** tab, click **Backup Job** > **Object Storage**.
- Open the **Home** view. Right-click in the working area, and select **Backup** > **Object storage**.
- Open the **Home** view. In the inventory pane, right-click the **Jobs** node and select **Backup** > **Object storage**.

You can quickly add an object storage source to an already existing job. Open the **Inventory** view. Under the **Unstructured Data** node in the inventory pane, select **Object Storage**. In the working area, right-click the object storage source you want to back up and select **Add to backup job** > *name of the job*.

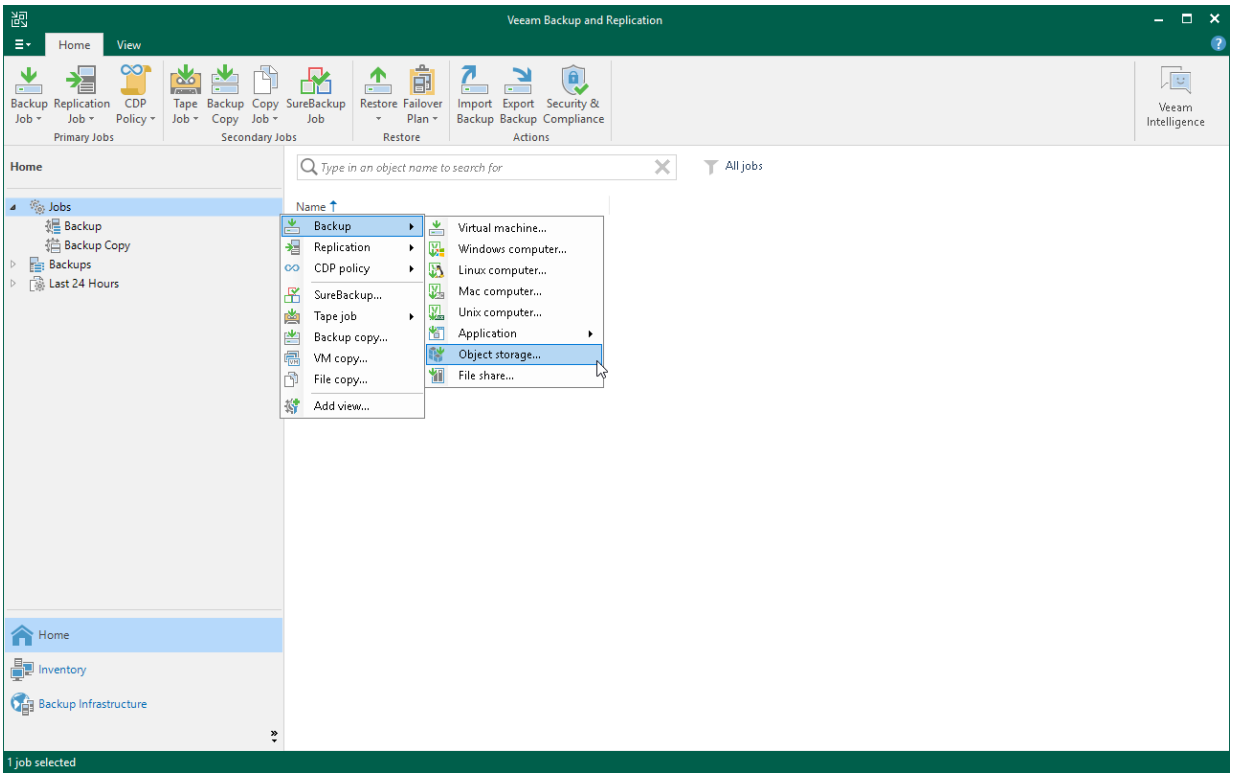

# Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a name and description for the object storage backup job.

1. In the **Name** field, enter a name for the object storage backup job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the job, date and time when the job was created.
3. Select the **High priority** check box if you want the resource scheduler of Veeam Backup & Replication to prioritize this job higher than other similar jobs and to allocate resources to it in the first place. For more information on job priorities, see Job Priorities.

> **TIP**
>
> In the list of jobs in the Veeam Backup & Replication console, jobs with the **High priority** option enabled are marked with a red flag ( ).

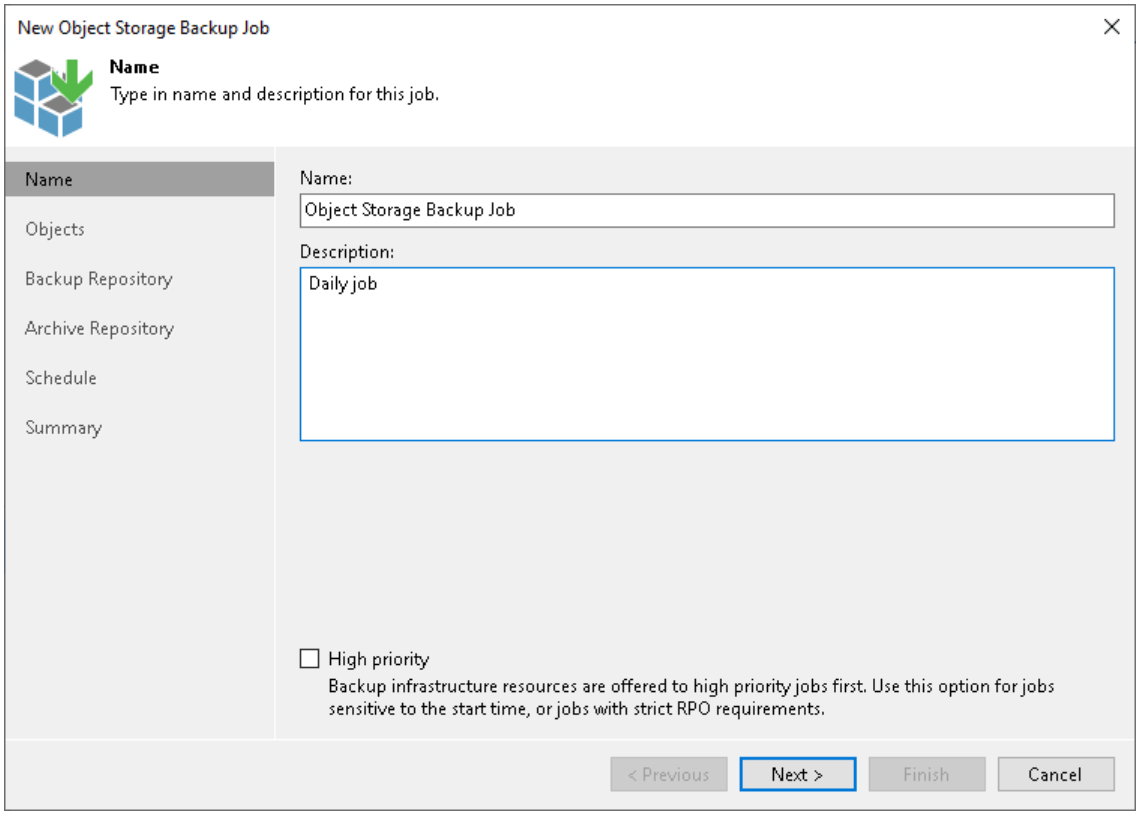

# Step 3. Select Objects to Back Up

At the **Objects** step of the wizard, select objects that you want to back up..

1. Click **Add**.
2. From the **Object Storage** drop-down list, select an object storage on which the necessary objects reside.

   This drop-down lists contains all object storage added to the inventory, as described in the Adding Object Storage section.
3. In the **Objects** tree, select objects you want to back up.

   To select multiple objects, hold [Ctrl] and click necessary folders.
4. If you specify a path to a bucket or container or prefix, all the contents will be processed.

   If necessary, you can choose only specific objects from the added prefix.

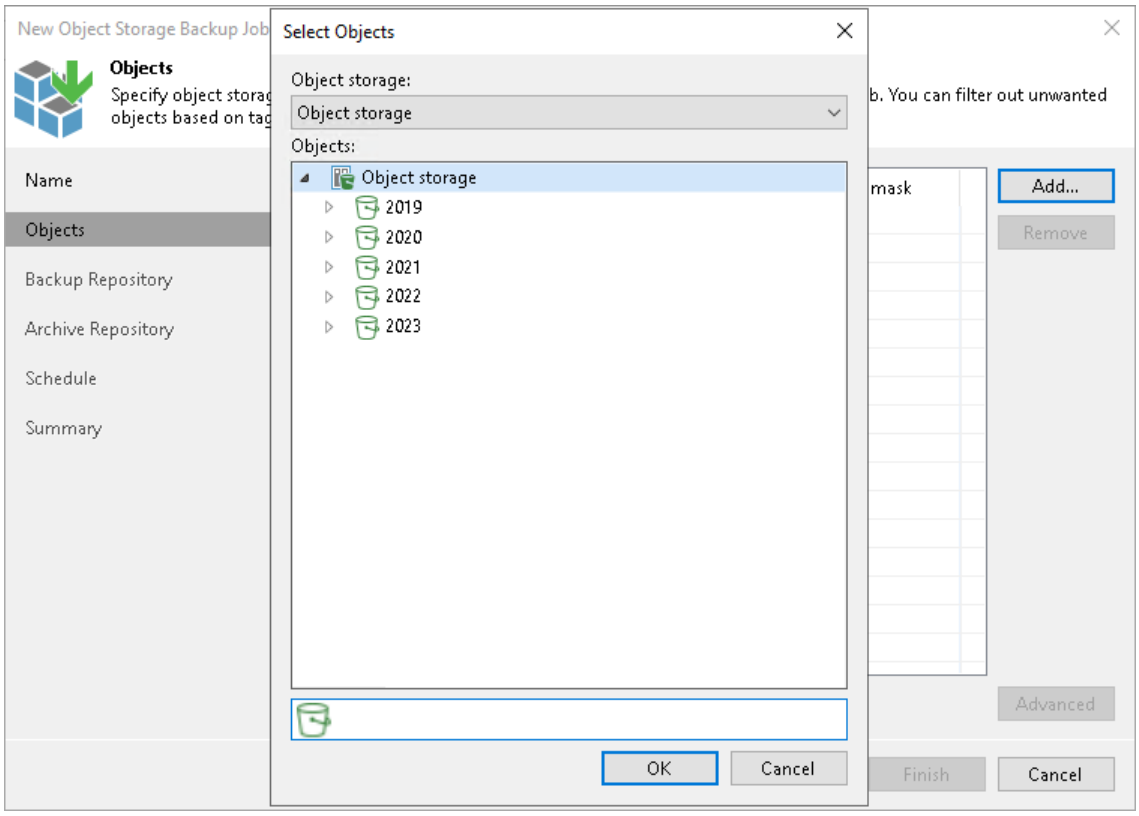


# Including Objects

To filter objects that you want to back up, you can specify the object name and extension masks.

1. Select an object in the **Objects** list and click **Filters**.
2. In the **Include and exclude masks** window, use **Include masks** field to include objects or prefixes.
3. Click **Add**.
4. In the **Include object by tag** window, specify the tag name and value associated with the object or prefix.
5. Click **OK**.

# Excluding Objects

To filter objects that you do not want to back up, you can specify the objects by a path or by a tag.

1. Select an object in the **Objects** list and click **Advanced**.
2. In the **Include and exclude masks** window, use the **Exclude masks** field to exclude objects.
3. Click **Add**.
4. Depending on the way you want to exclude objects, select one of the following:
   - To exclude objects by path, click **Objects by path...** and in the **Exclude objects by path** window, specify a path to the objects that you want to exclude.

     > **NOTE**
     >
     > Wildcards are not supported.

   - To exclude objects by tag, click **Objects by tag**... and in the **Exclude objects by tag** window specify the tag name and value associated with the object or prefix.

> **TIP**
>
> You can follow this examples for specifying the path in the **Exclude Object by path** window:
>
> - `bucketname` excludes the entire `bucketname` bucket.
> - `bucketname/myprefix/` excludes all objects in `bucketname` which names start with `myprefix`.

# Exporting and Importing Filters

If necessary, you can export and import your masks:

- To export a mask to a file, click the **Export these filters** link. In the **Export to file** window, specify a path to the necessary XML file. Click **OK**.
- To import existing masks from a file, click the **Import existing list** link. In the **Import masks from file** window, specify a path to the necessary XML file. Click **OK**.

# Step 4. Specify Backup Repository Settings

At the **Backup Repository** step of the wizard, define the primary backup repository, where the object storage backup job must store objects, and settings for moving objects to this repository. To learn what storage types you can assign the role of the backup repository to, see Storage Repositories in the **Backup Infrastructure for Unstructured Data Backup** section.

> **NOTE**
>
> Consider that if you use the option of limiting the number of object versions to keep configured in Object Version Settings, Veeam Backup & Replication first applies those object-version retention settings and only after that applies time-based retention settings specified at this step.

1. From the **Backup repository** drop-down list, select a repository where backup files must be stored. When you select a backup repository, Veeam Backup & Replication automatically checks the amount of free space left. Make sure that you have enough free space to store backups.
2. You can map the job to a specific backup stored in the backup repository. Backup job mapping allows you to move backup files to a new backup repository and to point the job to existing backups on this new backup repository. You can also use backup job mapping if the configuration database got corrupted and you need to reconfigure backup job settings.

   To map the job to a backup, click the **Map backup** link. In the opened **Select Backup** window, select a backup in the backup repository. Backups can be easily identified by job names. To find the backup, you can also use the search field at the bottom of the window.
3. Use the **Keep all versions for the last** field to specify how long copies of all recent object versions from the selected object storage bucket or container must be kept in the backup repository. You can restore the entire bucket or container to any restore point within the period specified in this setting.

   If, for example, **Keep all versions for the last** is set to 30 days, the backup repository will store all object versions that appeared in the object storage bucket or container during the last 30 days. At the scheduled time on the 31st day, the object storage backup job first backs up new object versions and saves them to the backup repository. Right after that, object versions older than 30 days (created on the 1st day) are either deleted from the backup repository or moved to the archive repository. Object versions are moved to the archive repository, if at the Archive Repository step of the wizard you enable the **Archive file versions to the following archive repository** check box and configure the archive retention.

4. If you need to keep a copy of backups in another repository, select the **Configure secondary destinations for this job** check box. That enables the Secondary Target step of the wizard.

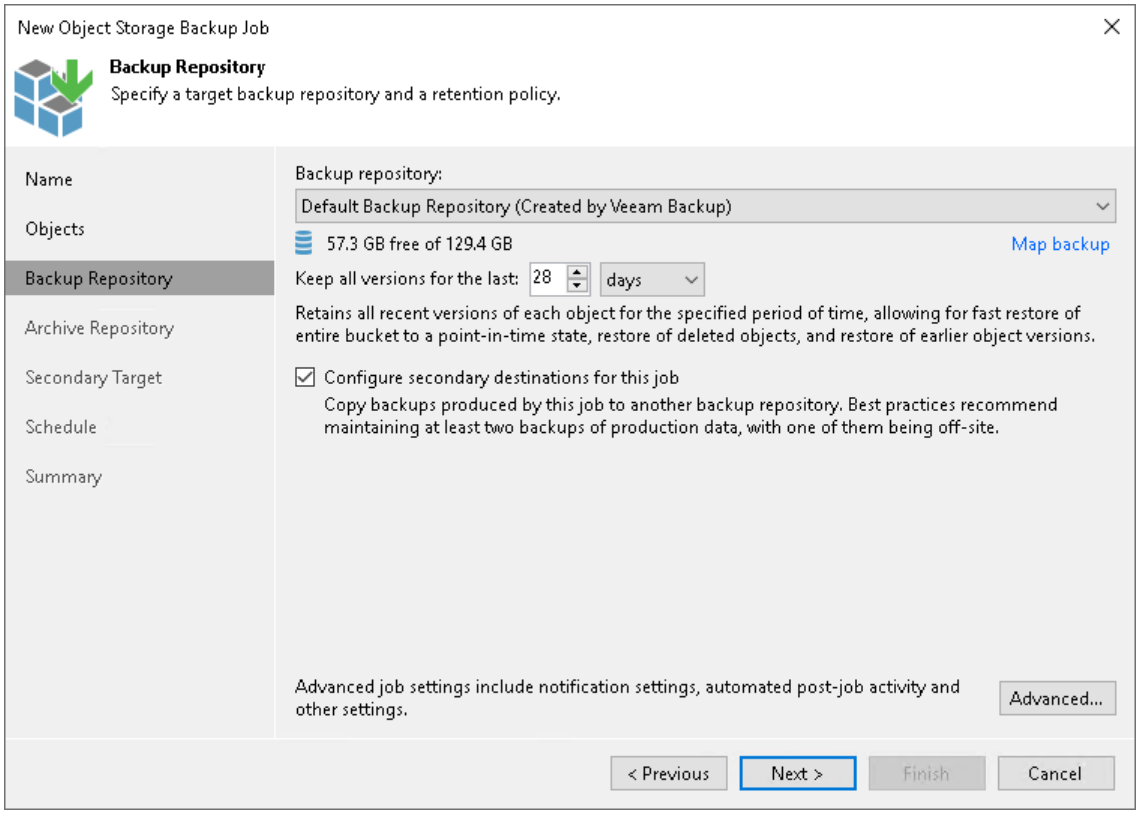

# Step 5. Specify Advanced Backup Settings

At the **Backup Repository** step of the wizard, specify advanced settings for the object storage backup job:

- Object version settings
- Storage settings
- Maintenance settings
- Script settings
- Notification settings

> **TIP**
>
> After you specify necessary settings for the backup job, you can save them as default settings. To do this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new backup job, Veeam Backup & Replication will automatically apply the default settings to the new job.

## Object Version Settings

To configure how many versions to keep for protected objects, do the following:

1. At the **Backup Repository** step of the wizard, click **Advanced**.
2. On the **Object versions** tab, specify to what object versions the settings should apply:
   - Select **Keep all versions** to keep all object versions for the time period specified at the Backup Repository step.
   - Select **Limit the number of archived versions only** to limit archived versions to the numbers specified to the right of the **Active versions limit** and **Deleted versions limit** check boxes.
   - Select **Limit the number of both recent and archived versions** to limit recent and archived versions to the numbers specified to the right of the **Active versions limit** and **Deleted versions limit** check boxes.
3. After you choose what object versions to keep, specify how many file versions to keep:
   - Select **Active versions limit** to keep the specified number of versions for objects currently existing in the object storage source. Specify how many object versions to store.
   - Select **Deleted versions limit** to keep the specified number of versions for objects deleted from the object storage source. Specify how many object versions to store.

4. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

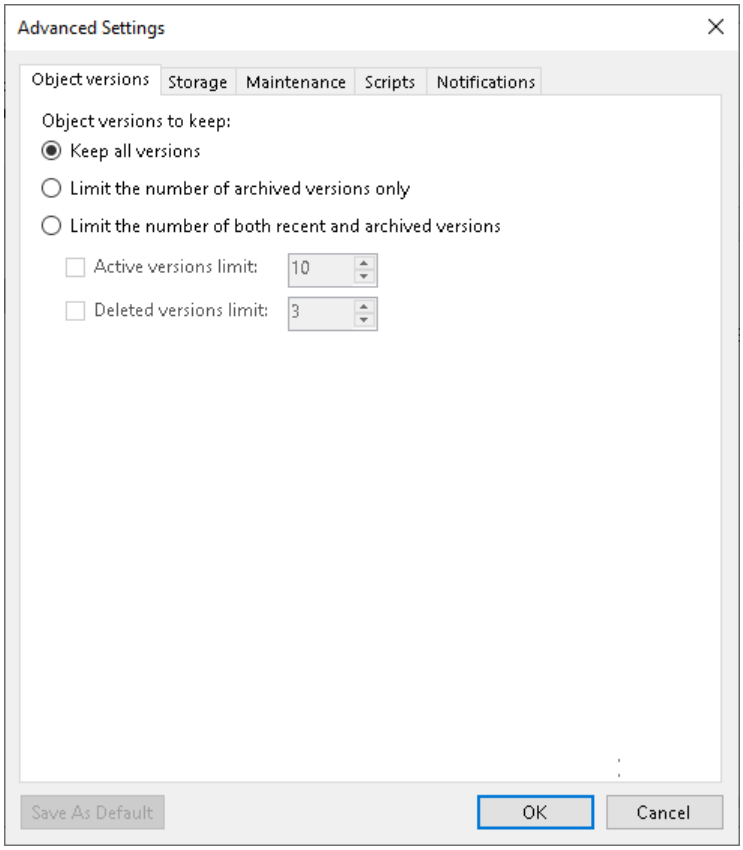


## Storage Settings

To specify advanced storage settings for the object storage backup job:

1. At the **Backup Repository** step of the wizard, click **Advanced**.
2. On the **Storage** tab, specify data reduction and encryption settings:
   - From the **Compression level** list, select a compression level for the backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*.
   - To encrypt backup files, select the **Enable backup file encryption** check box. In the **Password** field, select the encryption method you want to use — a password that will generate secret keys or an external KMS server that will generate KMS keys. For more information about data encryption, see How Data Encryption Works.

     If you have not created the password beforehand, click **Add** or use the *Manage passwords* link to specify a new password. For more information, see Password Manager.

     If you do not use password loss protection, Veeam Backup & Replication will display a warning about it. For more information, see Password Loss Protection.

> **NOTE**
>
> Consider the following:
>
> - If you enable encryption for an existing backup job, during the next job session Veeam Backup & Replication will back up all the files of the file share to a new backup file irrespective of whether they changed or not. The created backup files and subsequent backup files will be encrypted with the specified password.
> - If you enable encryption for an existing job, Veeam Backup & Replication does not encrypt the previous backup chain created by this job.
> - You can also use KMS keys for encryption. For more information, see Key Management System Keys.

3. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

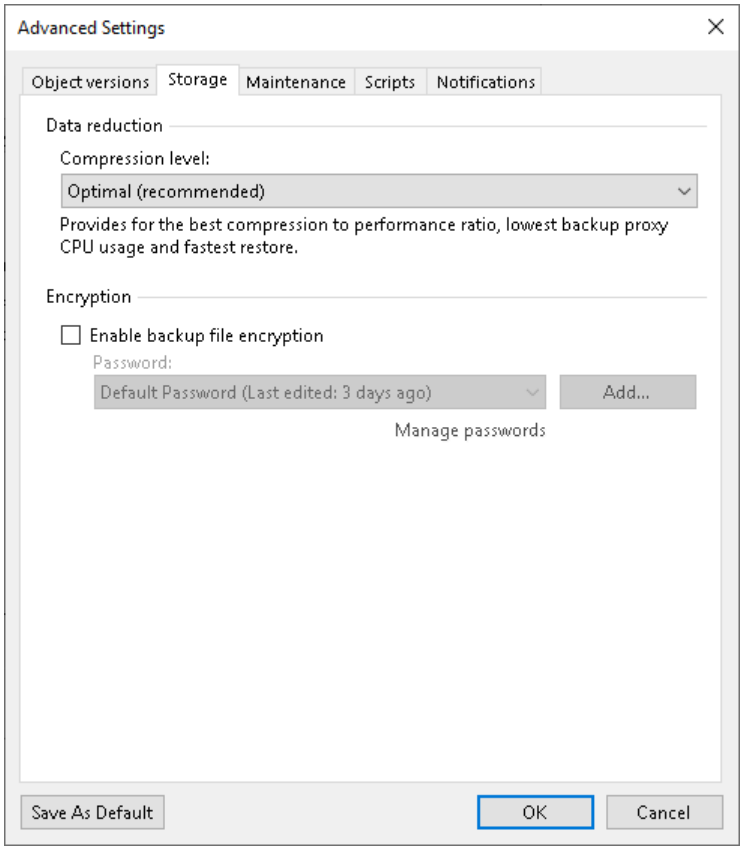


## Maintenance Settings

You can instruct Veeam Backup & Replication to periodically perform a health check for the backup. The health check ensures that the backup is consistent, and you will be able to restore data from it.

During the health check, Veeam Backup & Replication performs a CRC check for metadata and a hash check for data blocks in the object storage backup files to verify their integrity. For more information, see the Performing Health Check and Repair for Unstructured Data Backups section.

To configure the health-check settings for backup job:

1. At the **Backup Repository** step of the wizard, click **Advanced**.
2. On the **Maintenance** tab, select **Perform backup files health check** to enable the health check option. It ensures that all data and metadata is backed up correctly.

3. Click **Configure** and specify the time schedule for the health check.

4. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

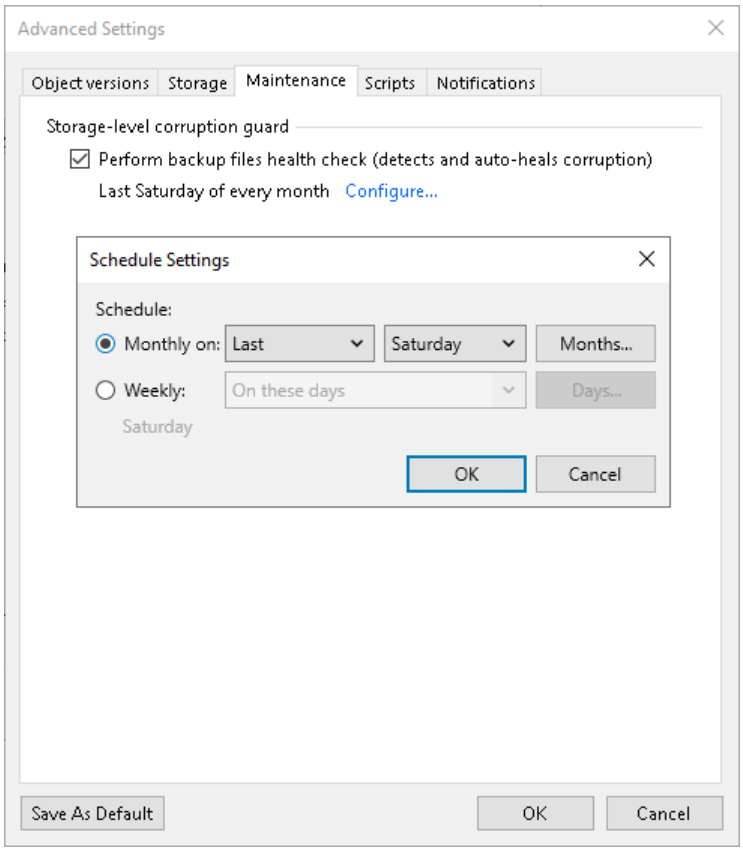


## Script Settings

You can configure custom scripts to run before or after the object storage backup job.

To specify script settings for the the object storage backup job:

1. At the **Backup Repository** step of the wizard, click **Advanced**.
2. Click the **Scripts** tab.
3. If you want to execute custom scripts, select the **Run the following script before the job** and **Run the following script after the job** check boxes and click **Browse** to choose executable files from a local folder on the backup server. The scripts are executed on the backup server.

   You can select to execute pre- and post-backup actions after a number of backup sessions or on specific week days.

   - If you select the **Run scripts every <N> backup session** option, specify the number of the backup job sessions after which the scripts must be executed.
   - If you select the **Run scripts on the selected days only** option, click **Days** and specify week days on which the scripts must be executed.

4. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

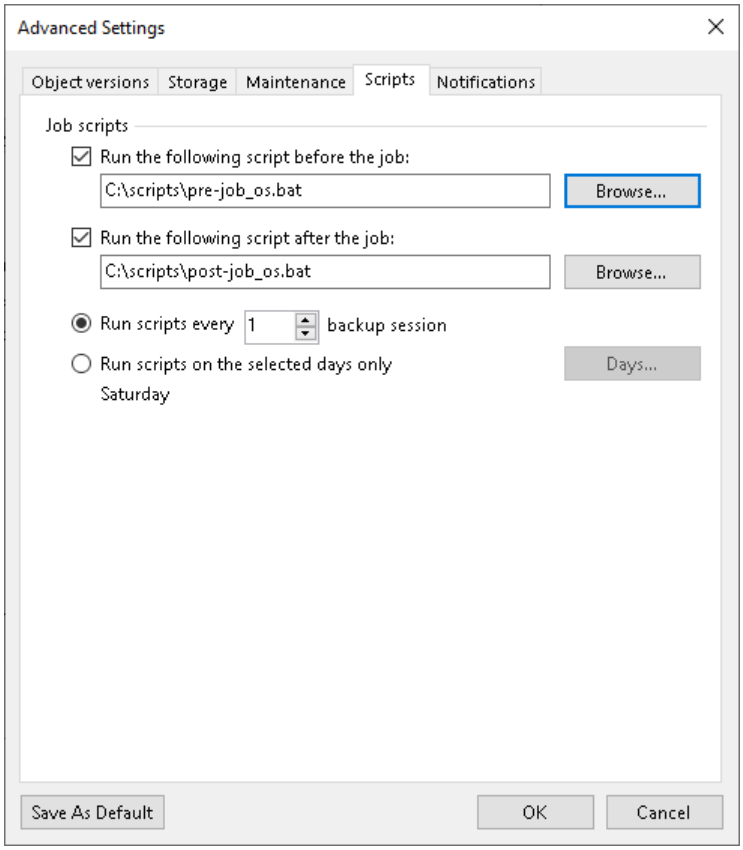


## Notification Settings

To specify notification settings for the the object storage backup job:

1. At the **Backup Repository** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully.

   SNMP traps will be sent if you specify global SNMP settings in Veeam Backup & Replication and configure software on recipient's machine to receive SNMP traps. For more information, see the Specifying SNMP Settings section.
4. Select the **Send email notifications to the following recipients** check box if you want to receive notifications about the job completion status by email. In the field under the check box, specify the recipient email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication. For more information, see the Configuring Global Email Notification Settings section.
5. You can choose to use global notification settings or specify custom notification settings:
   - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see the Configuring Global Email Notification Settings section.

- To configure a custom notification for the job, select **Use custom notification settings specified below**. You can specify the following notification settings:

  i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%Time%* (completion time), *%JobName%*, *%JobResult%*, *%ObjectCount%* (number of objects in the job) and *%Issues%* (number of objects in the job that have finished with the *Warning* or *Failed* status).

  ii. Select the **Notify on success**, **Notify on warning**, and **Notify on error** check boxes to receive email notification if the job completes successfully, completes with a warning or fails.

  iii. Select the **Suppress notifications until the last retry** check box to receive a notification about the final job status. If you do not enable this option, Veeam Backup & Replication will send one notification per every job retry.

6. Specify severity options if required:

   - Select the **Make object processing issues result in a job warning** to receive a warning at the end of the job processing session if any issues with object processing occur.

   - Select the **Make object attribute processing issues result in a job warning** to receive a warning at the end of the job processing session if any issues with processing of object attributes occur.

7. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

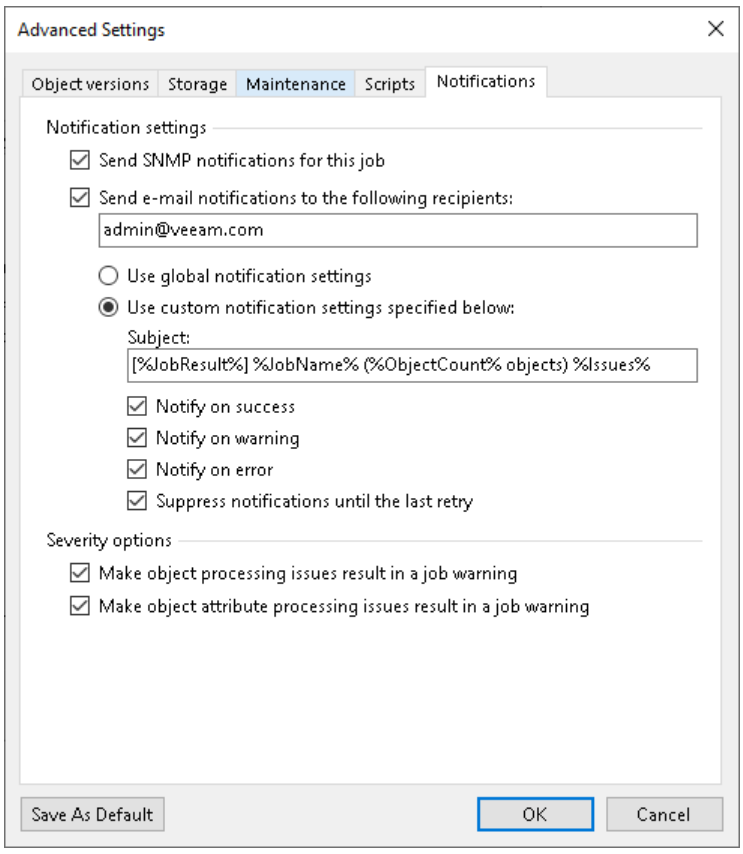

# Step 6. Specify Archive Repository Settings

At the **Archive Repository** step of the wizard, define the archive repository, where the object storage backup job must store backup files, and settings for moving files and folders to this repository. To learn what storage types you can assign the role of the archive repository to, see Storage Repositories in the **Backup Infrastructure for Unstructured Data Backup** section.

1. If you need to keep versions of some objects for a longer time after they are moved from the backup repository, you can configure archiving options to move object versions to the archive repository. You can also copy the recent object versions to the archive repository to store them according to the backup repository retention policy.

   To use the archive repository, select the **Archive object versions to the following repository** check box.

2. From the drop-down list under the **Archive object versions to the following repository** check box, select the storage to be used as a repository to store archived objects.

   By default, all objects deleted from the backup repository will be moved to the archive repository. If you do not need all the objects in the archive, you can choose what objects to keep.

3. If you need to keep the copy of the data stored in the backup repository also in the archive repository, select the **Archive recent object versions** check box.

   When you archive the recent object versions, Veeam Backup & Replication immediately copies all backed up objects from backup repository to the archive repository and stores them according to the backup repository retention policy. For example, if backup repository is configured to store object versions for 20 days, archive repository will also store these object versions for 20 days.

   > **NOTE**
   >
   > The copy mode in object storage backup requires a license. Thus, this feature is not supported in the Veeam Backup & Replication Community (free) Edition. For details, see Veeam Editions Comparison.

4. To specify the number of months or years during which backup files must be retained, select the **Archive previous object versions for** check box and specify the period.

   The period setting denotes a time period starting from the creation of the backup files in the backup repository, not from the moment when the object versions are moved from the backup repository to the archive repository.

   When you archive the previous object versions, Veeam Backup & Replication moves backed up object versions after their retention period is over from backup repository to the archive repository and stores them according to the archive repository retention policy. For example, if backup repository is configured to store object versions for 20 days, Veeam Backup & Replication moves object versions to the archive repository on the 21 day.

5. To specify what objects must be archived or excluded from the archive, do the following:

   a. Click **Choose** to open the **File Archive Settings** window.

   b. Under **Files to archive**, specify what objects must be archived:

      - **All files** – select this option to archive all objects moved from the backup repository to the archive repository.
      - **All files except the following extensions** – select this option to exclude objects with certain extensions from the selection to be archived. Specify extensions for objects to exclude from the selection. Objects with the specified extensions will not be copied from the backup repository to the archive repository.

- **Files with the following extensions only** – select this option to archive objects with certain extensions only. Specify extensions for objects to archive. Objects with these extensions will be copied from the backup repository to the archive repository.

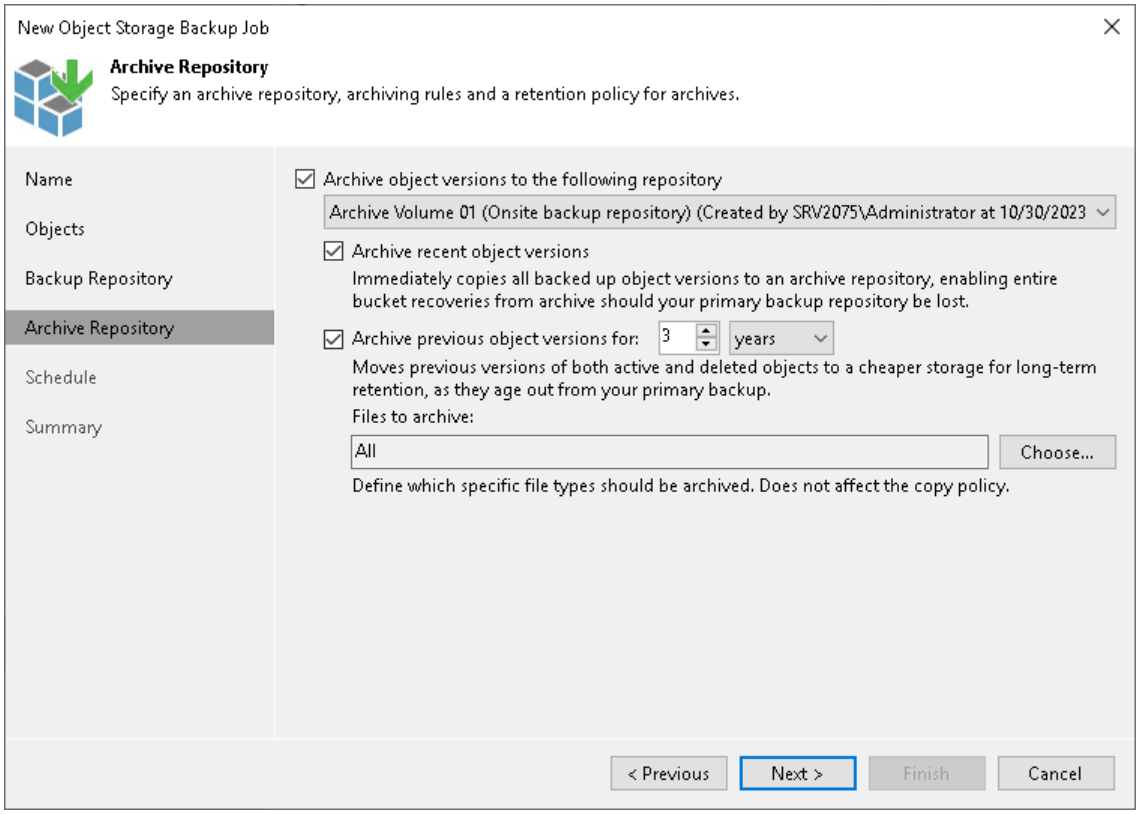

# Step 7. Specify Secondary Repository Settings

At the **Secondary Target** step of the wizard, you can specify a secondary repository that will be used to store additional copies of backup files from the backup storage for redundancy. To learn what storage types you can assign the role of the secondary repository to, see Storage Repositories in the **Backup Infrastructure for Unstructured Data Backup** section.

If you add a secondary repository, Veeam Backup & Replication will create a separate job for backup copy to it. The data copy process will start automatically after each primary object storage backup job runs.

> **NOTE**
>
> This step is available, if you select the **Configure secondary destinations for this job** check box at the Backup Repository step of the wizard.

To add a secondary repository:

1. Click **Add**.
2. From the list of existing repositories, select a repository that will keep additional copy of the backup files. You can add several secondary repositories for copying files of the primary backup job. To quickly find the repository, use the search field at the bottom of the wizard.

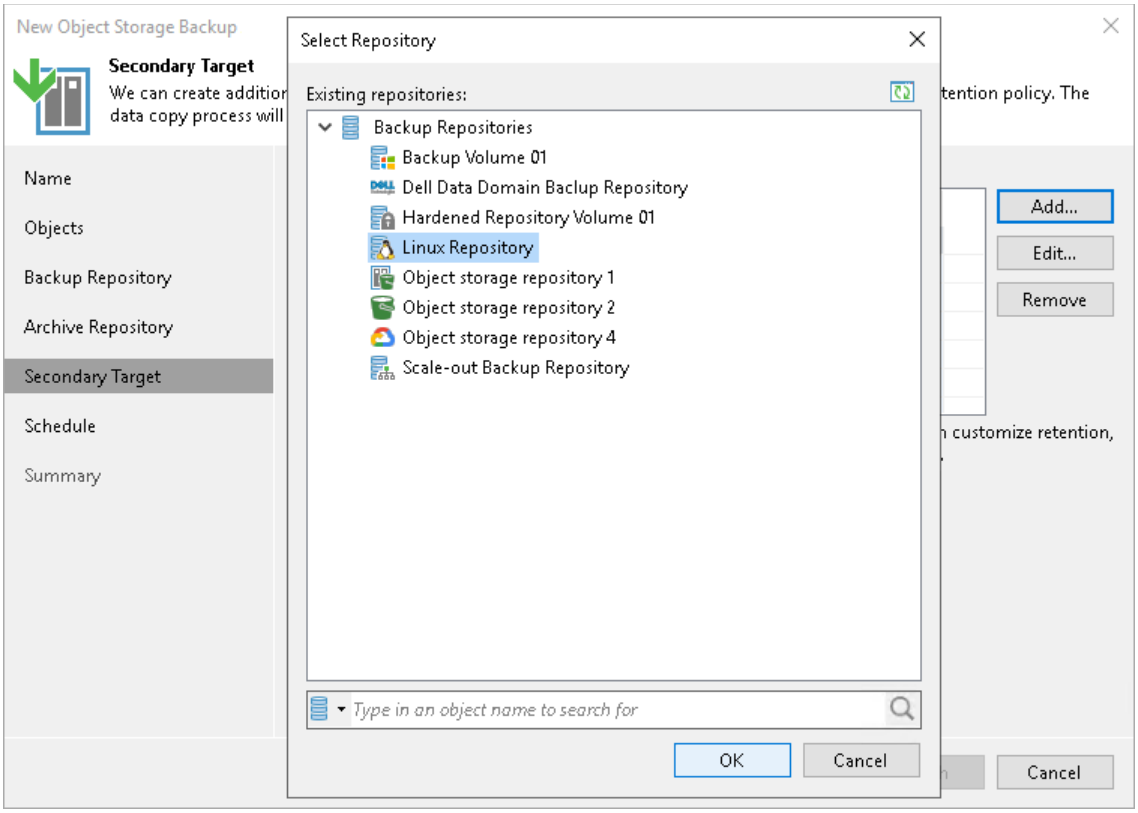


3. By default, retention and encryption settings for the secondary target repository are inherited from the primary job. To customize them, select the necessary repository in the **Secondary repositories** list and click **Edit**.
   - To enable custom retention settings:
     i. Select **Use custom retention policy**.
     ii. Specify how long all versions of each file will be kept in the secondary repository.

- To specify encryption settings different from those of the primary repository:

  i. Select **Use custom data encryption settings**.

  ii. In the **Password** field, select a password that you want to use for encryption. If you have not created the password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see the Password Manager section.

  If the backup server is not connected to Veeam Backup Enterprise Manager and does not have the Veeam Universal License or a legacy socket-based Enterprise or Enterprise Plus license installed, you will not be able to restore data from encrypted backups in case you lose the password. Veeam Backup & Replication will display a warning about it. For more information, see the Decrypting Data Without Password section.

- Configure time intervals at which the data can be copied to the secondary repository.

  - If you select the **Any time (continuously)** option, Veeam Backup & Replication will copy backup files to the secondary repository as soon as the primary object storage backup job completes.
  - If you want to specify time periods when it is permitted to start copying backup files to the secondary repository, select the **During the following time periods only** option and configure allowed and prohibited hours. These periods do not work as the backup window, so they will not cause the file backup copy to fail.

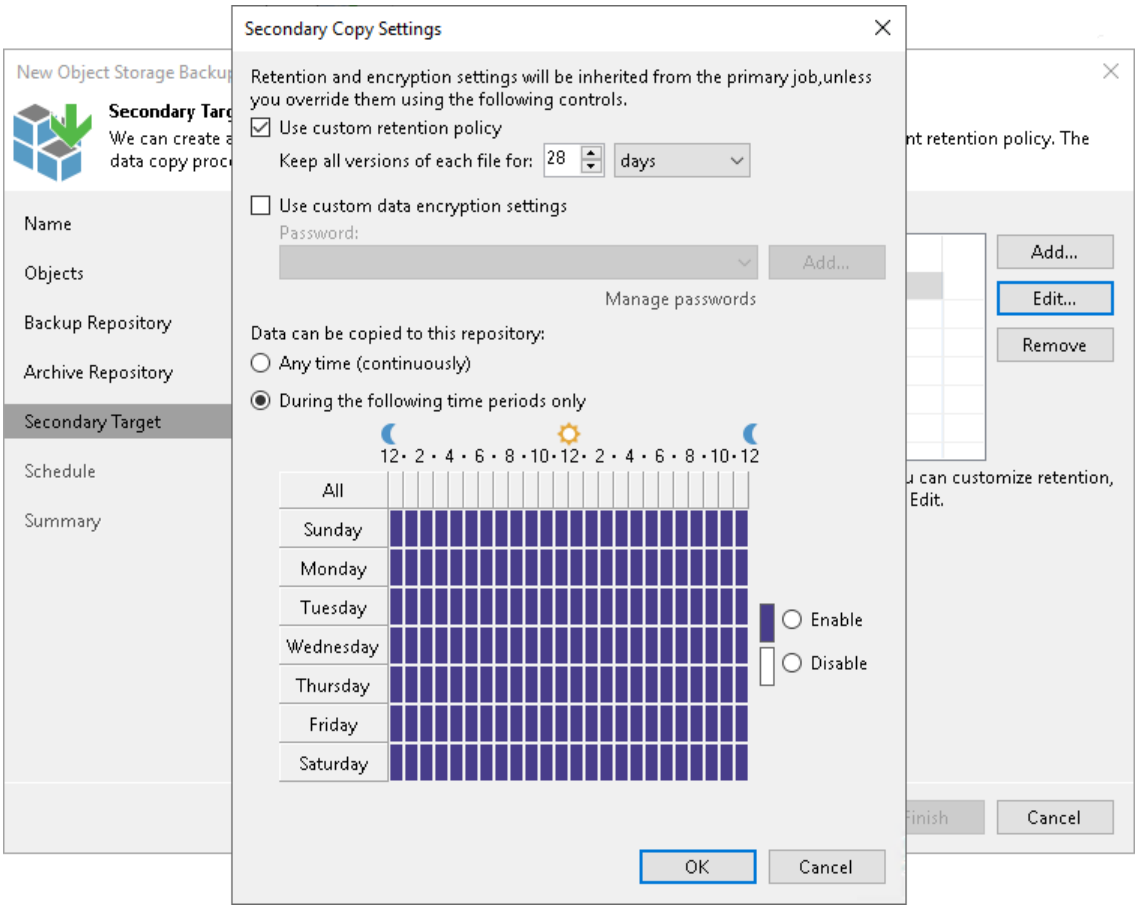

# Step 8. Define Job Schedule

At the **Schedule** step of the wizard, select to run the object storage backup job manually or schedule the object storage backup job to run on a regular basis.

To specify the object storage backup job schedule:

1. Select the **Run the job automatically** check box. If this check box is not selected, you will have to start the job manually to create the object storage backup.
2. Define scheduling settings for the job:
   - To run the job at specific time daily, on defined week days or with specific periodicity, select the **Daily at this time** option. Use the fields on the right to configure the necessary schedule.
   - To run the job once a month on specific days, select the **Monthly at this time** option. Use the fields on the right to configure the necessary schedule.

     > **NOTE**
     >
     > When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a specific time interval, select the **Periodically every** option. In the field on the right, select the necessary time unit: **Hours** or **Minutes**. Click **Schedule** and use the time table to define the permitted time window for the job. In the **Start time within an hour** field, specify the exact time when the job must start.

     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run the job by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.
   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right. A new backup job session will start as soon as the previous backup job session finishes.
   - To chain jobs, use the **After this job** field. In the common practice, jobs start one after another: when job *A* finishes, job *B* starts and so on. If you want to create a chain of jobs, you must define the time schedule for the first job in the chain. For the rest of the jobs in the chain, select the **After this job option** and choose the preceding job from the list.
3. In the **Automatic retry** section, define whether Veeam Backup & Replication must attempt to run the backup job again if the job fails for some reason. During a job retry, Veeam Backup & Replication processes failed object storage sources only. Enter the number of attempts to run the job and define time intervals between them. If you select continuous backup, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job runs.

4. In the **Backup window** section, define the time interval within which the backup job must complete. The backup window prevents the job from overlapping with production hours and ensures that the job does not provide unwanted overhead on the production environment. To set up a backup window for the job:
   a. Select the **Terminate job outside of the backup window** check box and click **Window**.
   b. In the **Window** window, define the allowed hours and prohibited hours for backup. If the job exceeds the allowed window, it will be automatically terminated.

> **NOTE**
>
> The **After this job** function will automatically start a job if the first job in the chain is started automatically by schedule. If you start the first job manually, Veeam Backup & Replication will display a notification. You will be able to choose whether Veeam Backup & Replication must start the chained job as well.

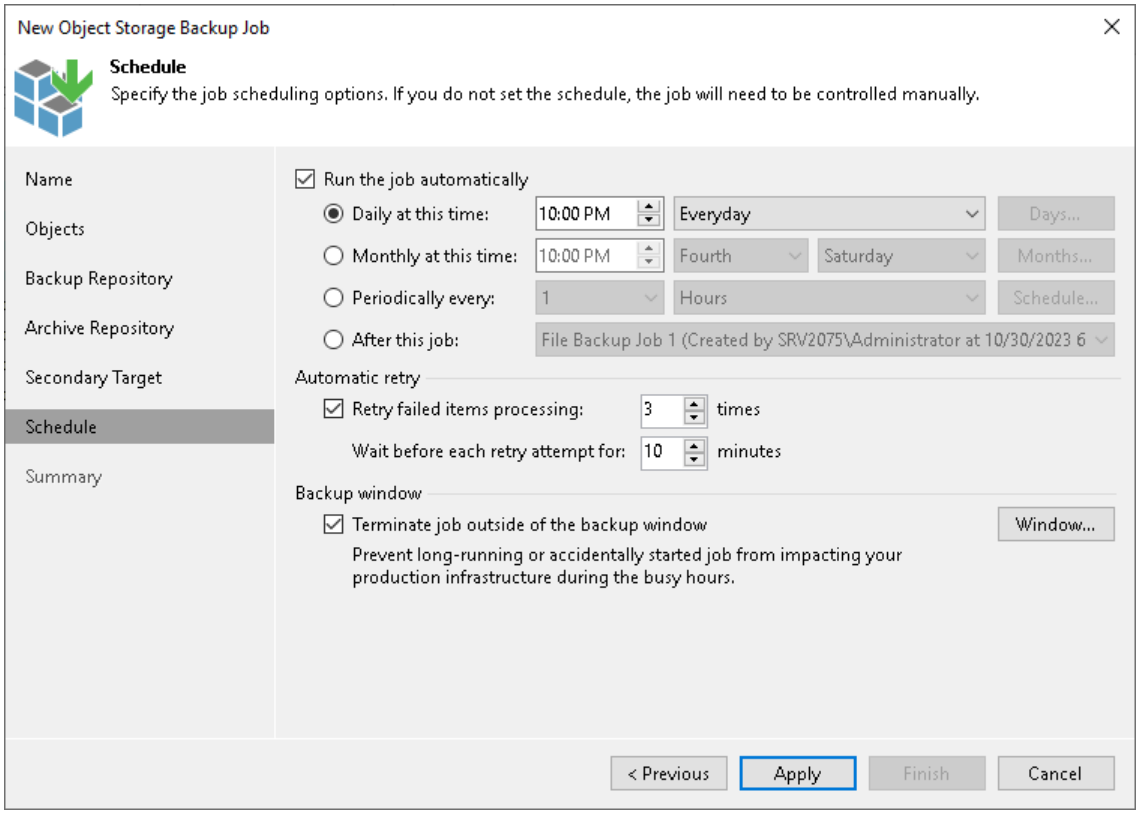

# Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of object storage backup job configuration.

1. Review details of the object storage backup job.
2. If you want to start the job right after you finish working with the wizard, select the **Run the job when I click Finish** check box.
3. Click **Finish** to close the wizard.

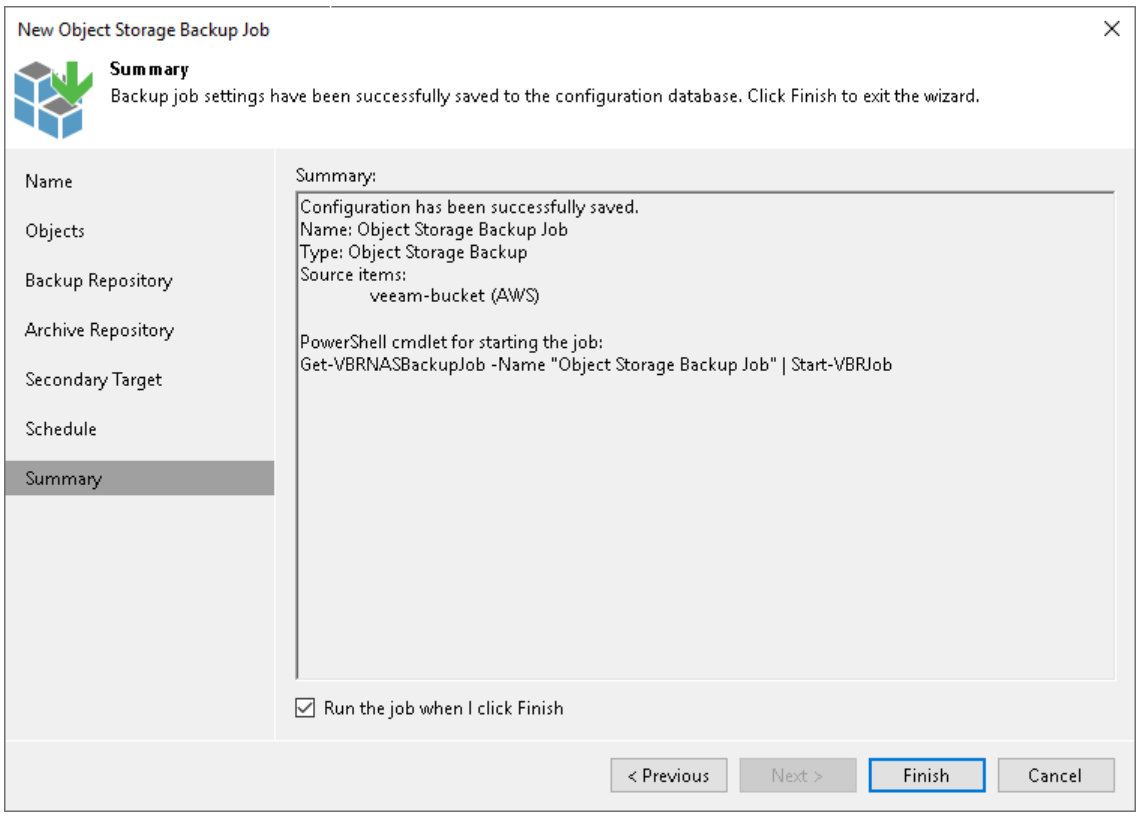

# Managing Unstructured Data Backups

You can perform the following operations with backups:

- View unstructured data backup properties.
- Copy unstructured data backups to another location.
- Start a new backup chain for an unstructured data backup job.
- Perform the health check and repair operations for unstructured data backups.
- Detaching Unstructured Data Backups from Jobs.
- Deleting Unstructured Data Backups from Disk.
- Convert backups from SMB or NFS shares to NAS filer shares.
- Convert backups from non-root to root shared folders.
- Update the source file share path for backup jobs with a secondary target.

# Viewing Unstructured Data Backup Properties

You can view summary information about the unstructured data backup. The summary information provides the following data:

- Name and path to the backup repository that stores backup files.
- Name and path to the archive repository that stores archived backup data.
- Path to the backup source and its original size.
- Available restore points: date of their creation, their type (**Backup** or **Archive**) and status.

  For the **Backup** type, the table shows all restore points stored in the backup repository. You can restore the unstructured data to the state as of any of these points. To learn how to restore the file share data, see File Share Data Recovery, to restore object storage data — Object Storage Data Recovery.

  For the **Archive** type, the table shows only a single record. The time stamp of this record denotes the date and time of the restore point, which was created in the backup repository and files of which were the first to be moved to the archive repository according to retention settings. To learn how to restore files from the archive, see the Restoring Backup Files from Archive Repository sections for the file share data recovery and the object storage data recovery.

To view summary information for backups:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click the backup and select **Properties**.
4. To see the list of available restore points, select the required object from the **Objects** list.

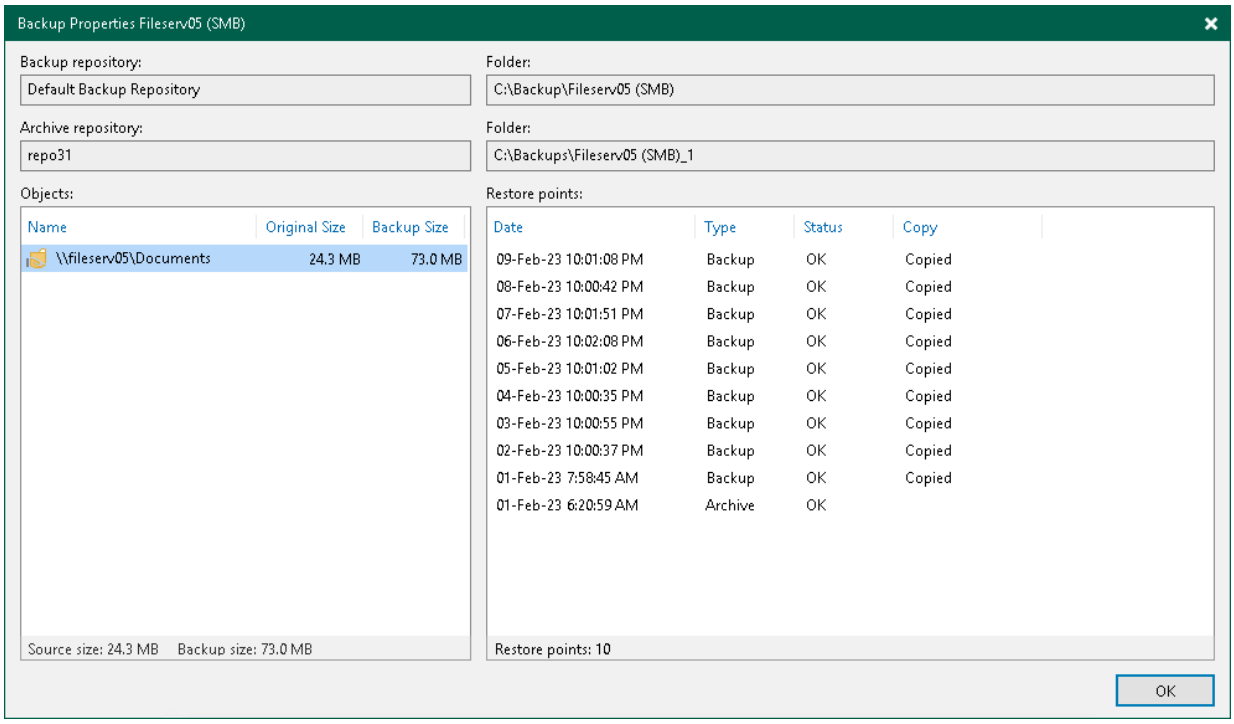

# Copying Unstructured Data Backups

Copying backups can be helpful if you want to copy file share backups to a repository or local or shared folder. Veeam Backup & Replication copies the whole backup chain.

When Veeam Backup & Replication performs the copy operation, it disables the job, copies files to the target location and then enables the job. After the copy operation finishes, the copied backups are shown in a node with the **(Exported)** postfix in the inventory pane.

> **NOTE**
>
> This section is about one-time copy operation. If you want to copy backups on a schedule, configure a secondary repository as described in the Creating File Backup Jobs and Creating Object Storage Backup Jobs sections.

## Copying Backups

To copy file share backups, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** node.
3. In the working area, select the necessary job.
4. Right-click the job and select **Copy backup**. Alternatively, click **Copy Backup** on the ribbon.
5. In the **Copy Backup to Another Location** window, choose a repository to which you want to copy backups.
6. [For archive backups] If you also want to copy archive backups to another archive repository, select the **Copy the associated archive** check box. From the drop-down list, select the necessary repository.
7. Click **OK**.

After the copy process finishes, the copied backups are shown in the **Disk (Exported)** node in the inventory pane.

> **NOTE**
>
> Consider the following:
>
> - If you copy backups from a scale-out backup repository and some backups are stored on extents in the Maintenance mode, such backups are not copied.
> - Veeam Backup & Replication copies backups only from the performance tier of the scale-out backup repository. If you want to copy data from the capacity tier, you first need to download it to the performance tier. For more information, see Downloading Data from Capacity Tier.

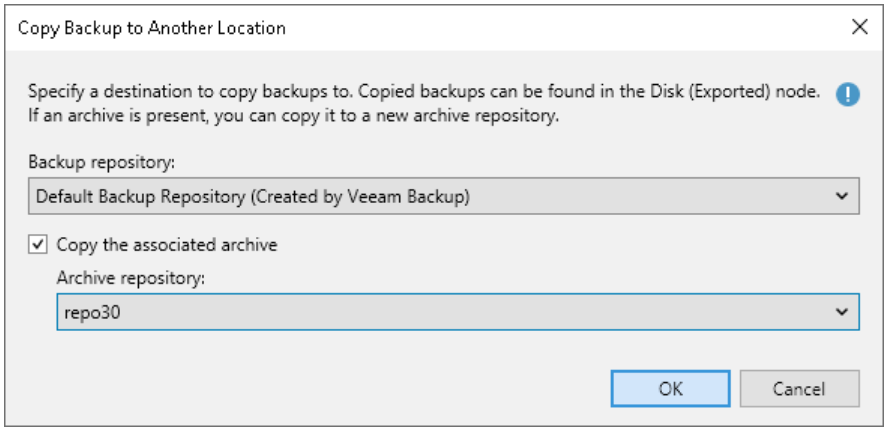
Copy Backup to Another Location
Specify a destination to copy backups to. Copied backups can be found in the Disk (Exported) node.
If an archive is present, you can copy it to a new archive repository.
Backup repository:
Default Backup Repository (Created by Veeam Backup)
Copy the associated archive
Archive repository:
repo30
OK
Cancel

# Starting New Backup Chain

You can start a new backup chain for the unstructured data backup job. Veeam Backup & Replication then creates a new active full backup that starts the new chain for the entire protected file share or bucket. All existing backup files are moved to the **Disk (Orphaned)** node under the **Backups** node in the Veeam Backup & Replication Console. Data files are stored to the same folder in the backup repository. The data files for the new backup chain are stored to a new separate folder in the backup repository.

To start a new backup chain:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs** > **Backup**.
3. In the working area, select a job of the **File Backup** or **Object Storage Backup** type, press and hold the [Ctrl] key, right-click the job and select **Start new backup**.

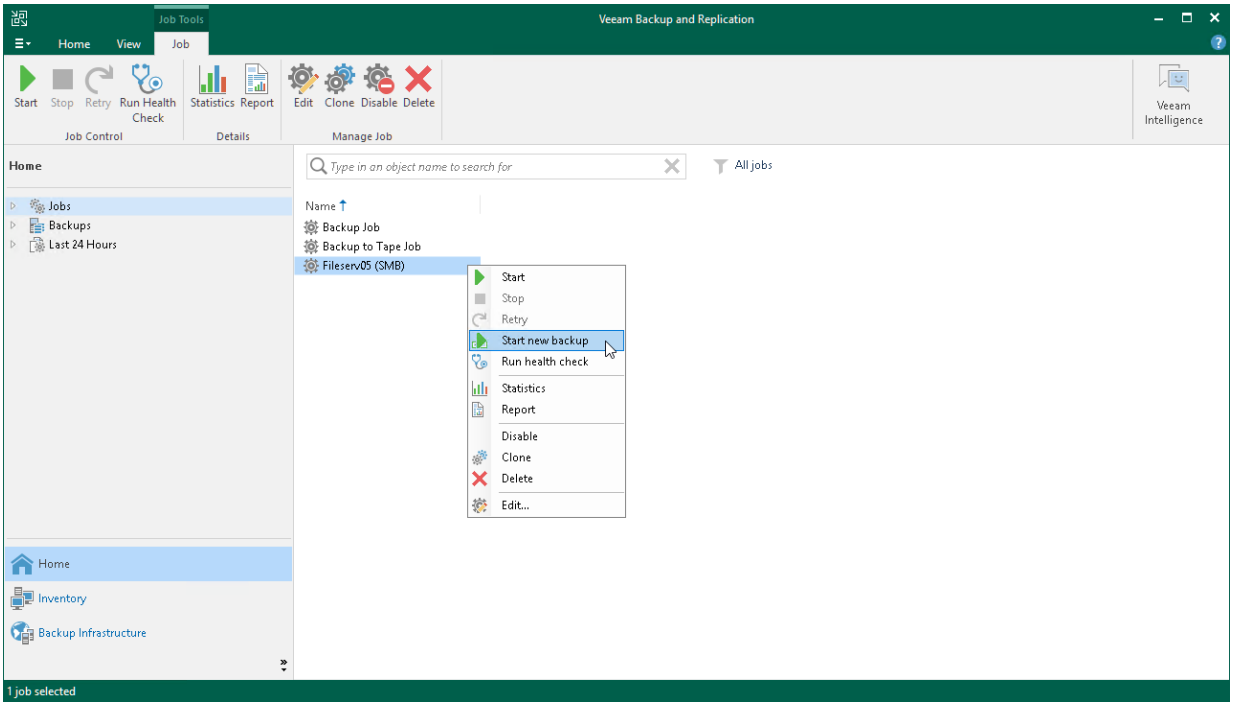

# Performing Health Check and Repair for Unstructured Data Backups

In this section you will learn how to perform:

- Health check for file share backup files
- Repair of file share backup files

## Health Check for File Share Backup Files

You can manually perform a health check for the backup chain. During the health check, Veeam Backup & Replication performs a CRC check for metadata and a hash check for data blocks in backup files to verify their integrity. The health check helps make sure that the restore point is consistent, and you will be able to restore data from this restore point.

To run the health check:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs** > **Backup**.
3. In the working area, select a job of the **File Backup** or **Object Storage Backup** type and click **Run Health Check** on the ribbon or right-click the job and select **Run health check**.

To run the health check periodically, you must enable the **Perform backup files health check** option in the backup job settings and define the health check schedule. By default, the health check is performed on the last Friday of every month. You can change the schedule and run the health check weekly or monthly on specific days. To learn how to configure periodic health check, see Maintenance Settings.

> **IMPORTANT**
>
> If you store your file backups on public cloud object storage repositories, running the health check operations may result in constantly downloading and uploading data to and from the storage, which may lead to higher costs. To avoid this, use helper appliances, configured for the repositories within the public clouds. For more information, see the Unstructured Data Backups in Object Storage Repositories section.

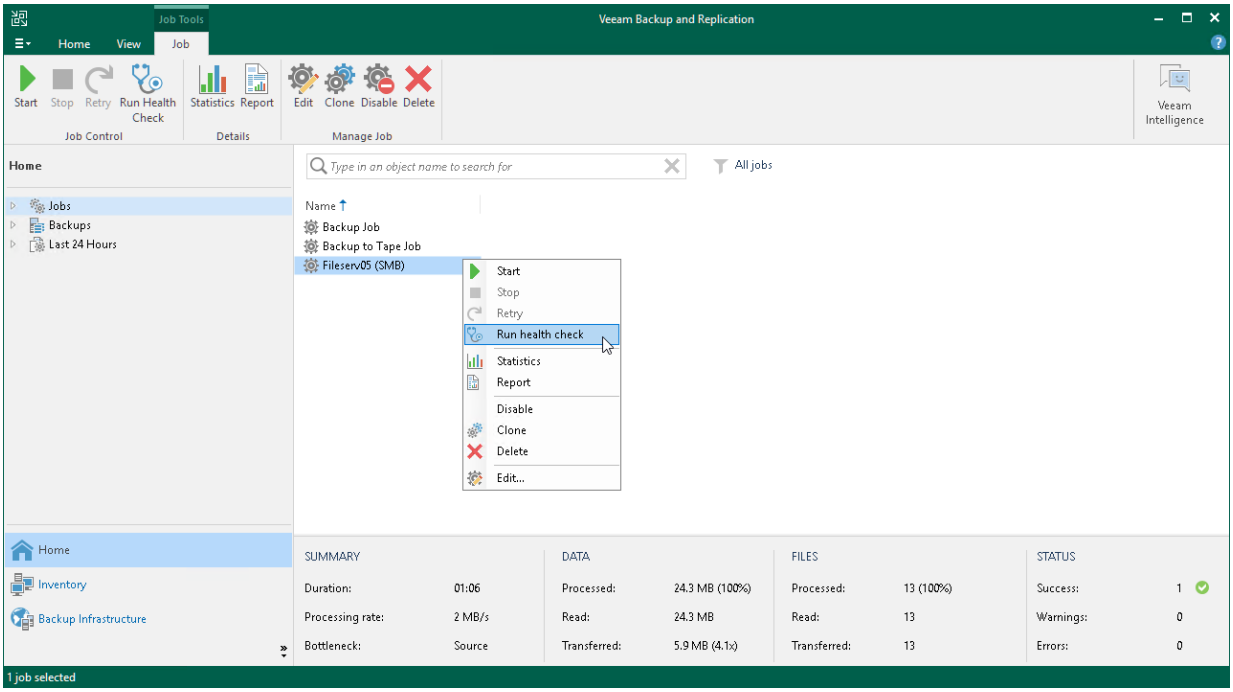


# Repair of File Share Backup Files

If during the health check Veeam Backup & Replication detects some inconsistency in the file share backup files, you can run the backup repair procedure to fix the issues.

To run the backup repair:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select a job of the **File Backup** type, right-click the job and select **Repair backup**.

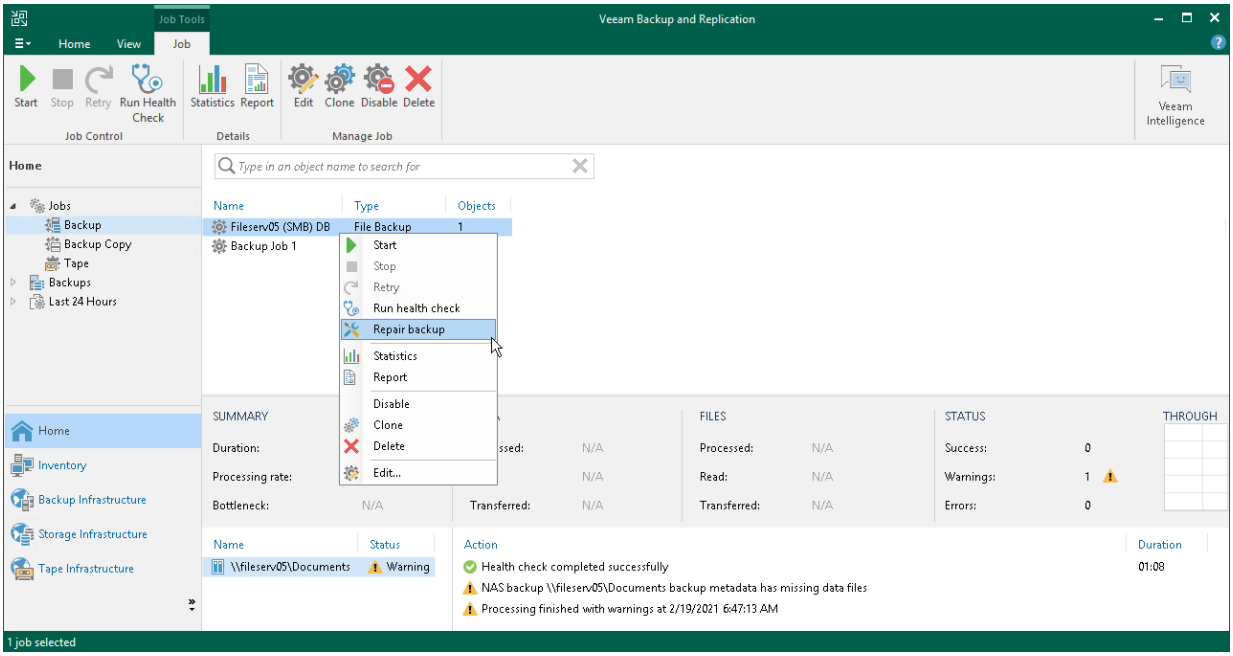

# Detaching Unstructured Data Backups from Jobs

If you want to detach backups from a job, you can use the **Detach from job** operation.

The detached backup files remain in the backup repository and the Veeam Backup & Replication console. Veeam Backup & Replication shows the detached backups in the inventory pane under the **Disk (Orphaned)** or **Object Storage (Orphaned)** node. These backups are retained according to the retention policy.

To detach backups from a job:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** node.
3. In the working area, right-click the necessary backup and select **Detach from job**. Alternatively, click **Delete from** > **Job** on the ribbon.

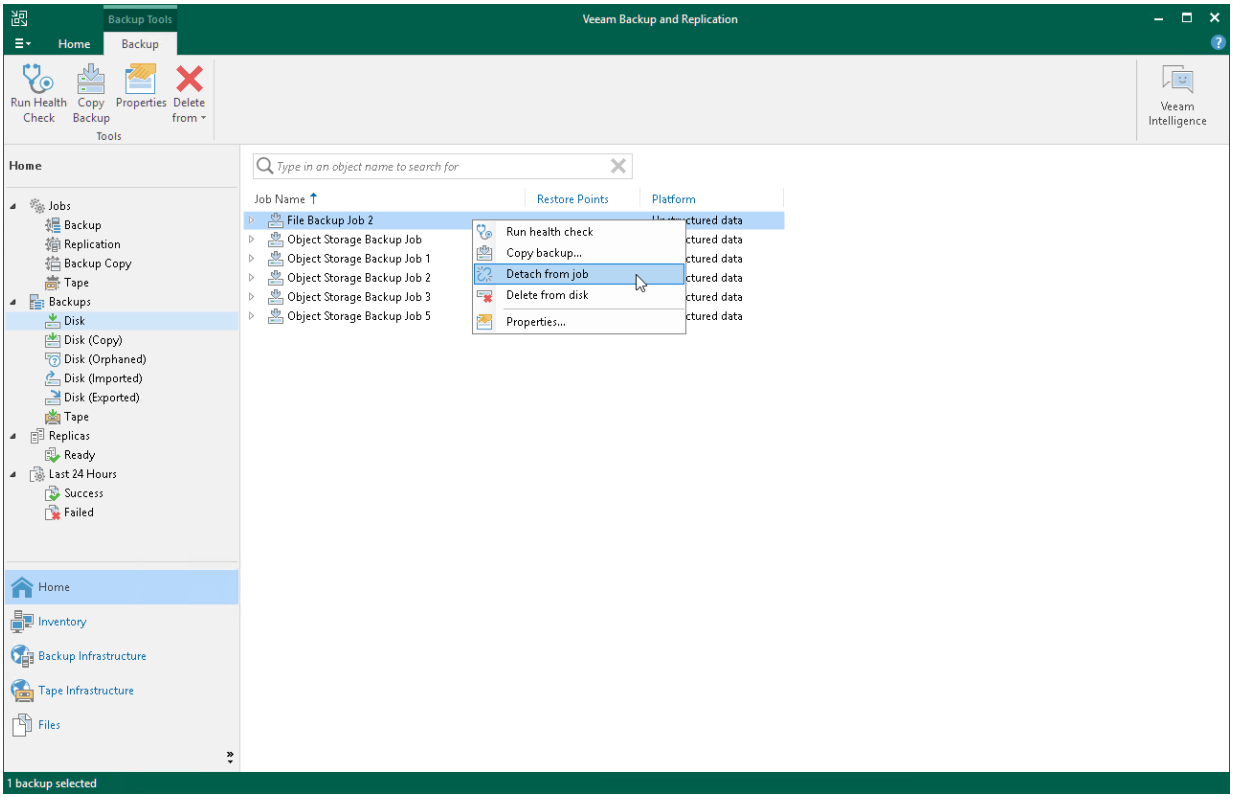

# Deleting Unstructured Data Backups from Disk

If you want to permanently remove unstructured data backup files from the backup repository and Veeam Backup & Replication console, you can use the **Delete from disk** operation.

Consider the following:

- Do not delete unstructured data backup files from the backup repository manually. Use the **Delete from disk** instead. If you delete unstructured data backup files manually, Veeam Backup & Replication database will not reflect these changes, causing subsequent backup job sessions to fail.
- When you delete unstructured data backup files from a disk, Veeam Backup & Replication deletes the whole chain from the backup repository. Thus, on the next backup job run, Veeam Backup & Replication will create a full backup for objects on the data source.
- **Delete from disk** operation does not remove unstructured data from processing with subsequent backup job sessions. If you do not want to back up this data anymore, remove or exclude it from processing in the backup job's settings.

To delete backups from disk:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** node.
3. In the working area, select the backup and click **Delete from** > **Disk** on the ribbon. You can also right-click the backup and select **Delete from disk**.

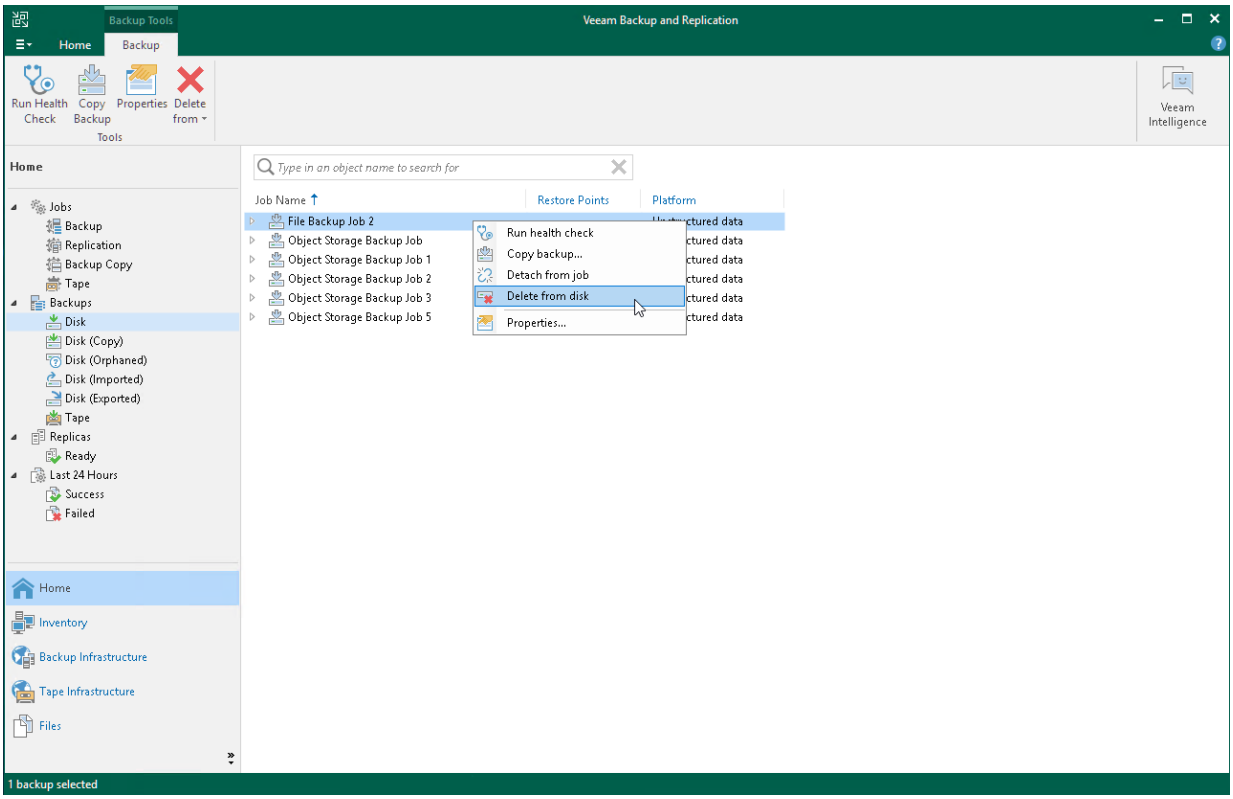

# Converting Backups from SMB or NFS Shares to NAS Filer Shares

You can use enterprise storage systems integrated with Veeam Backup & Replication both to host simple SMB or NFS shares and to act as NAS filer shares.

> **NOTE**
>
> You cannot convert backups stored on a root server or shares from the root server.

To use all the advantages of NAS filer shares, for example the native file change tracking technology, you can convert backups created for existing SMB or NFS shares into the format of NAS filer shares. After that you can continue to protect the file shares as NAS filer shares by running existing file backup jobs and by using existing backup files. Perform the conversion with extreme caution.

To convert SMB or NFS shares into NAS filer shares, do the following:

1. Disable file backup jobs protecting SMB or NFS shares, for which you want to convert backups. To do that, right-click the required job in the **Jobs** node of the inventory pane in the **Home** view and select **Disable**. Alternatively, you can click **Disable** on the ribbon.
2. Make sure that you have created NAS filer, which corresponds to existing SMB or NFS shares, added to the Veeam Backup & Replication inventory. The NAS filer and these shares must reside on the same storage system. The correspondence of the shares must be full except for the host name.
3. Run the `Convert-VBRNASBackupSANFormat` PowerShell cmdlet to convert the format of the file share backup to provide support of NAS filer shares.

   ```
   $nasbackup = Get-VBRNASBackup -name "File Backup Job 1"
   $netapp = Get-NetAppHost -name "pdc-ontap-1"
   $netapp_filer = Get-VBRNASServer -SANEntity $netapp
   Convert-VBRNASBackupSANFormat -Backup $nasbackup -Server $netapp_filer
   ```

   For more information, see the description of the `Convert-VBRNASBackupSANFormat` cmdlet in the Veeam PowerShell Reference.

   As a result, the backup will be moved from **Backups** > **Disk** node to **Backups** > **Disk (Orphaned)** node in the inventory pane of the **Home** view.

   At this step, you can check if the cmdlet has correctly converted the backup. To do that, check if backup object names in the **Disk (Orphaned)** node have changed and now show the path to the NAS filer share. If object names have not changed and show the path to the SMB or NFS share as before, continuing the conversion process can lead to the unwanted result. For example, when you enable the backup job for the converted backup, it will back up the NAS filer share not with an incremental run, but with a full run instead, which may lead to extra costs.
4. Use the Edit File Backup Job wizard to edit the file backup job that protects the file shares:
   a. At the Files and Folders step of the wizard, remove the existing SMB and NFS shares from the job and add NAS filer shares instead.
   b. At the Backup Repository step of the wizard, map the job to the backup that was converted at step 2.
5. Enable file backup jobs protecting file shares, for which you converted backups. To do that, right-click the required job from the **Jobs** node of the inventory pane in the **Home** view and clear selection of **Disable**. Alternatively, you can click **Disable** on the ribbon.

Veeam Backup & Replication supports conversion of backups created by backup copy jobs. To continue the old backup chain created by the backup copy job, do the following:

1. Disable file backup jobs protecting SMB or NFS shares, for which you want to convert backups. To do that, right-click the required job in the **Jobs** node of the inventory pane in the **Home** view and select **Disable**. Alternatively, you can click **Disable** on the ribbon.
2. Make sure that you have created NAS filer, which corresponds to existing SMB or NFS shares, added to the Veeam Backup & Replication inventory. The NAS filer and these shares must reside on the same storage system. The correspondence of the shares must be full except for the host name.
3. Run the `Convert-VBRNASBackupSANFormat` PowerShell cmdlet to convert the format of the file share backup to provide support of NAS filer shares.

   ```
   $backup_copy = Get-VBRNASBackup -name "File Backup Job 1 (Copy)"
   $netapp = Get-NetAppHost -name "pdc-ontap-1"
   $netapp_filer = Get-VBRNASServer -SANEntity $netapp
   Convert-VBRNASBackupSANFormat -Backup $backup_copy -Server $netapp_filer
   ```

   For more information, see the description of the `Convert-VBRNASBackupSANFormat` cmdlet in the Veeam PowerShell Reference.
4. Use the Edit File Backup Job wizard to remove the secondary target storage for the file backup job that protects the file shares:
   a. At the Secondary Target step of the wizard, remove the required repository selected as a secondary target for the file backup job.
   b. Go through all the wizard steps without running the job. Click **Finish**.
5. Use the Edit File Backup Job wizard to add the secondary target storage back for the file backup job that protects the file shares:
   a. At the Secondary Target step of the wizard, add the required repository as a secondary target for the file backup job.
   b. Go through all the wizard steps. Click **Finish**.
6. Enable file backup jobs protecting file shares, for which you converted backups. To do that, right-click the required job from the **Jobs** node of the inventory pane in the **Home** view and clear selection of **Disable**. Alternatively, you can click **Disable** on the ribbon.

   File share backup copy will automatically map to the file backup copy job. After that, the backup copy job will back up new points of the main file backup job if they were created.

# Converting Backups from Non-Root to Root Shared Folders

Veeam Backup & Replication allows adding a server root folder as a source for file backup jobs. In this case, all changes to separate shared folders residing on this server will be reflected in the file backup job where the root shared folder of this server is added. You can even add shared root folders using different protocols to one file backup job and thus protect all file shares that are or will be added on the server.

If you previously had several separate non-root shared folders residing on the same server and want to switch to using a single root shared folder to cover the same shares, you do not have to run full backups to update data of protected shares. Instead, you can convert existing backups and update existing file backup jobs to protect single root shared folders comprising all other non-root shared folders residing on the same server. Perform the conversion with extreme caution.

To convert backups from non-root to root shared folders, do the following:

1. Disable file backup jobs protecting file shares, for which you want to convert backups. To do that, right-click the required job in the **Jobs** node of the inventory pane in the **Home** view and select **Disable**. Alternatively, you can click **Disable** on the ribbon.
2. Make sure that your backup infrastructure has a root share (for example, NFS or SMB) added for the whole server or storage system where existing non-root shares reside. These shares must reside on the same server or storage system. The correspondence of the shares must be full except for the host name.
3. Run the `Convert-VBRNASBackupRootFormat` PowerShell cmdlet to convert backups created by one file backup job for separate non-root shared folders residing on the same server into the backup created for the server single root folder with all the non-root shared folders of the same type under it. For more information, see the description of the `Convert-VBRNASBackupRootFormat` cmdlet in the Veeam PowerShell Reference.

   As a result, the backup will be moved from **Backups** > **Disk** node to **Backups** > **Disk (Orphaned)** node in the inventory pane of the **Home** view.

   At this step, you can check if the cmdlet has correctly converted the backup. To do that, check if backup object names in the **Disk (Orphaned)** node have changed and now show the path to the server root folder. If object names have not changed and show the paths to multiple separate non-root shared folders as before, continuing the conversion process can lead to the unwanted result. For example, when you enable the backup job for the converted backup, it will back up all shared folders under root folder not with an incremental run, but with a full run instead, which may lead to extra costs.
4. Use the Edit File Backup Job wizard to edit the file backup job that protects the file shares:
   a. At the Files and Folders step of the wizard, remove the existing non-root shared folders from the job and add the server root folder instead.
   b. At the Backup Repository step of the wizard, map the job to the backup that was converted at step 2.
5. Enable file backup jobs protecting file shares, for which you converted backups. To do that, right-click the required job from the **Jobs** node of the inventory pane in the **Home** view and clear selection **Disable**. Alternatively, you can click **Disable** on the ribbon.

# Updating Source File Share Path for Backup Jobs with Secondary Target

Veeam Backup & Replication does not support auto mapping of the source file share path for the file backup copy when updating the source file share path for the main file backup job.

For example, you have a **File Backup** job protecting the file share located at `\\shared_server\documents`. It has an associated backup copy job - **File Backup (Copy) 1**. The file backups are stored in the **main_storage** repository, the file backup copies are stored in the **copy_storage** repository. Then, you decide to update the protected file share name to `\\share_server_new\documents`.

To correctly update the file backup job protecting this file share, do the following:

1. Add a new file share `\\share_server_new\documents` in the inventory, as described in section Adding File Share. Do not remove the old file share from the inventory yet as it is associated with the old file backup job.
2. Update the path to the source file share for the file backup:

```
$unstructuredbackup = Get-VBRUnstructuredBackup -Name "File Backup"
$sourceserver = Get-VBRUnstructuredServer -Backup $unstructuredbackup
$targetserver = Get-VBRUnstructuredServer -Name "\\share_server_new\docume
nts"
Update-VBRUnstructuredBackupPath -Backup $unstructuredbackup -SourceUnstru
cturedServer $sourceserver -TargetUnstructuredServer $targetserver
```

   As a result, the backup will be moved from the **Backups** > **Disk** to **Backups** > **Disk (Orphaned)**.
3. Update the path to the source file share for the file backup copy:

```
$unstructuredbackup = Get-VBRUnstructuredBackup -Name "File Backup (Copy)
1"
$sourceserver = Get-VBRUnstructuredServer -Backup $unstructuredbackup
$targetserver = Get-VBRUnstructuredServer -Name "\\share_server_new\docume
nts"
Update-VBRUnstructuredBackupPath -Backup $unstructuredbackup -SourceUnstru
cturedServer $sourceserver -TargetUnstructuredServer $targetserver
```

   As a result, the backup copy will be moved from **Backups** > **Disk (Copy)** to **Backups** > **Disk (Orphaned)**.
4. Edit settings of the **File Backup** job defining file backup properties:
   a. Remove the old protected file share at `\\shared_server\documents` and add the new file share at `\\share_server_new\documents`, as described in sections Select Files and Folders to Back Up of Creating File Backup Jobs.
   b. Map the file backup job to the **File Backup** backup converted at step 2, as described in step 4 of the Creating File Backup Jobs procedure.
   c. Remove the secondary target repository added at step 7 of Creating File Backup Jobs procedure.
   d. Ensure that the **Run the job when I click Finish** check box is not selected when you close the wizard, as described in step 9 of the Creating File Backup Jobs procedure.

5. Edit settings of the **File Backup** job defining file backup copy properties:

   a. Enable creation of the file backup copy by selecting the **Configure secondary destinations for this job** check box, as described in step 4 of the Creating File Backup Jobs procedure.

   b. Add the **copy_storage** repository for storing backup copies, as described in step 7 of the Creating File Backup Jobs procedure.

      Veeam Backup & Replication will automatically map the existing file backup copy, which was previously located in **Backups** > **Disk (Orphaned)**, to the file backup job.

   c. Ensure that the **Run the job when I click Finish** check box is selected when you close the wizard, as described in step 9 of the Creating File Backup Jobs procedure.

As a result, if the job and its backups are updated correctly, the first session of the updated job will create not a full backup, but an incremental one.

# Unstructured Data Recovery

Veeam Backup & Replication offers the following types of recovery:

- File share data recovery — to restore data previously backed up with file backup jobs.
- Object storage data recovery — to restore data previously backed up with object storage backup jobs.

## Related Topics

Data Recovery

# File Share Data Recovery

You can restore data previously backed up with file backup jobs. You can restore the following data:

- SMB file share files and folders
- NFS file share files and folders
- Files and folders of a managed Microsoft Windows server
- Files and folders of a managed Linux server

Veeam Backup & Replication offers several recovery options for different recovery scenarios:

- Instant file share recovery allows you to publish a point-in-time file share state to instantly access all protected files.
- Restore of the entire file share allows you to recover all files and folders of the file share to one of the restore points.
- Rollback to a point in time allows you to restore only changed files to one of the restore points.
- Restore of files and folders allows you to select files and folders to restore to one of the restore points.
- Restore of files from an archive repository allows you to select archived files to restore to one of the restore points.

# Instant File Share Recovery

The instant file share recovery feature works in 2 modes depending on the type of the file share:

- For NFS file shares, you can use this feature to publish a point-in-time file share state to enable users to instantly access all protected files in the read-only mode.
- For SMB file shares, you can use this feature to publish a point-in-time file share state to enable users to add, update, remove files in the mounted file share. The updated file share may be then migrated to the production server.

## Performing Instant File Share Recovery

Before you perform instant file share recovery, check prerequisites.

### Before You Begin

Before you perform instant file share recovery, consider the following:

- NFS file shares recovered with instant recovery are available in the read-only mode.
- Files larger than 1 GB are not available for update in file shares recovered with instant recovery.
- You cannot rename folders in file shares recovered with instant recovery.
- You cannot change folder attributes and apply folder security in published file shares recovered with instant recovery.

## Step 1. Launch Instant File Share Recovery Wizard

To launch the **Instant File Share Recovery** wizard, do one of the following:

- In the **Home** tab on the ribbon, click **Restore** > **File Share**. In the **Restore from File Backup** window, click **Instant file share recovery**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the file share that you want to restore. In the **Backup** tab on the ribbon, click **Instant Recovery**.
  - Right-click the file share that you want to restore and select **Instant file share recovery**.

  You can perform the instant file share recovery by using a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane. If the secondary repository is an object storage repository, backup copies created in it are represented in the **Backups** > **Object Storage (Copy)** node in the inventory pane.

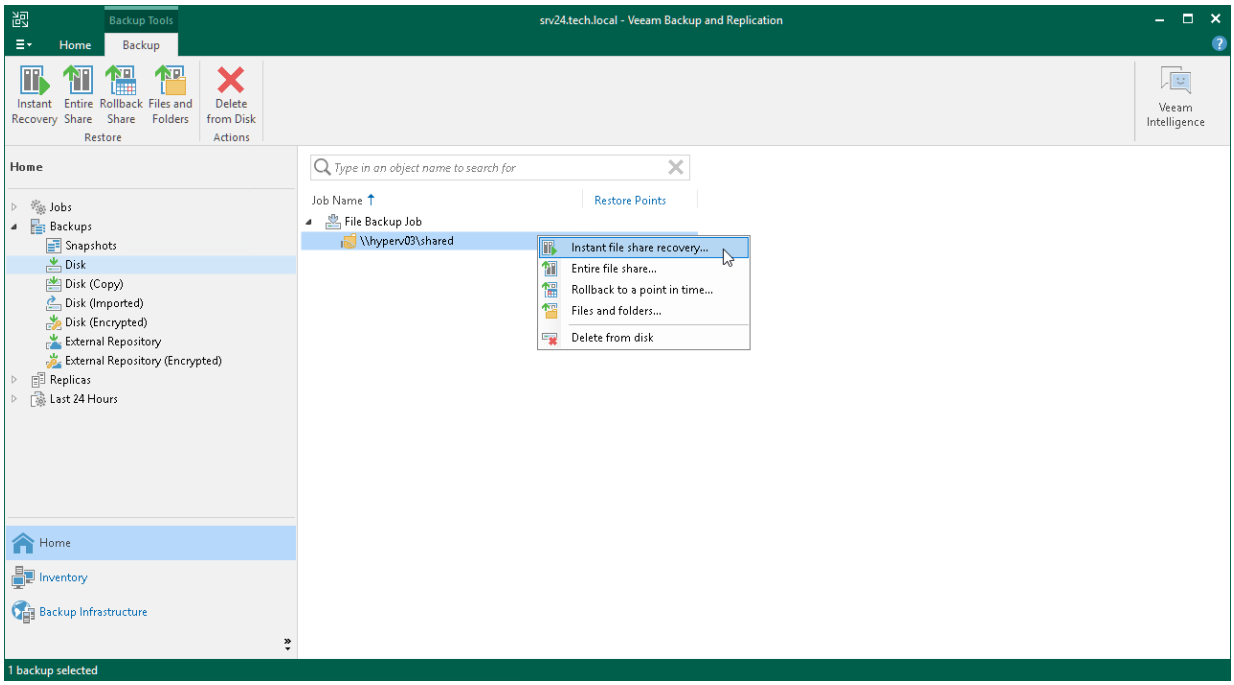

## Step 2. Select File Share to Restore

At the **File Shares** step of the wizard, select the file share you want to instantly recover:

1. Click **Add**.
2. In the **Backups Browser** window, expand the necessary backup job to select the required file share to restore.

To quickly find a file share, you can use the search field at the bottom of the window.

1. Enter a file share name or a part of it in the search field.
2. Press [Enter] to start the search.

Alternatively, you can use the **File shares to publish** search field to quickly search the required file share and add it to the list of file shares to publish.

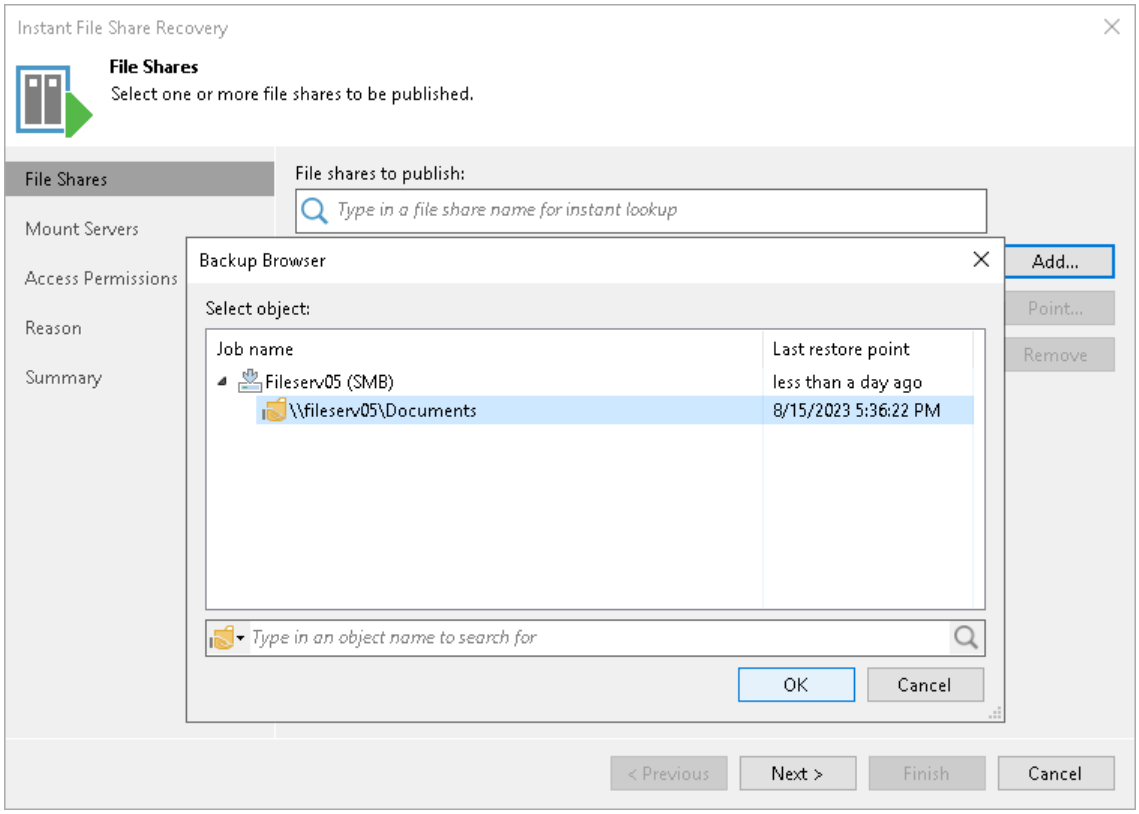

## Step 3. Specify Mount Server Settings

At the **Mount Servers** step of the wizard, specify mount server settings for published file shares.

- If you select the **Automatic selection** option, Veeam Backup & Replication will automatically choose the mount server where to restore file shares. The job will use the mount server from the repository where the backup files reside. For scale-out backup repositories, Veeam Backup & Replication will use the least occupied mount server.
- If you select the **Manual selection** option, you can specify which mount server to use to individually publish each file share:
  a. In the **Mount servers** list, select a file share for which you want to assign a mount server.
  b. Click **Edit**. Alternatively you can double-click the required file share in the list.
  c. In the **Select mount server** window, select a mount server to use to publish the chosen file share.

     To quickly find a mount server, you can use the search field at the bottom of the window.
  d. Click **OK** to confirm selection.

> **NOTE**
>
> Consider that data on the mounted file share may be available to the users added to the Administrators group on this mount server.

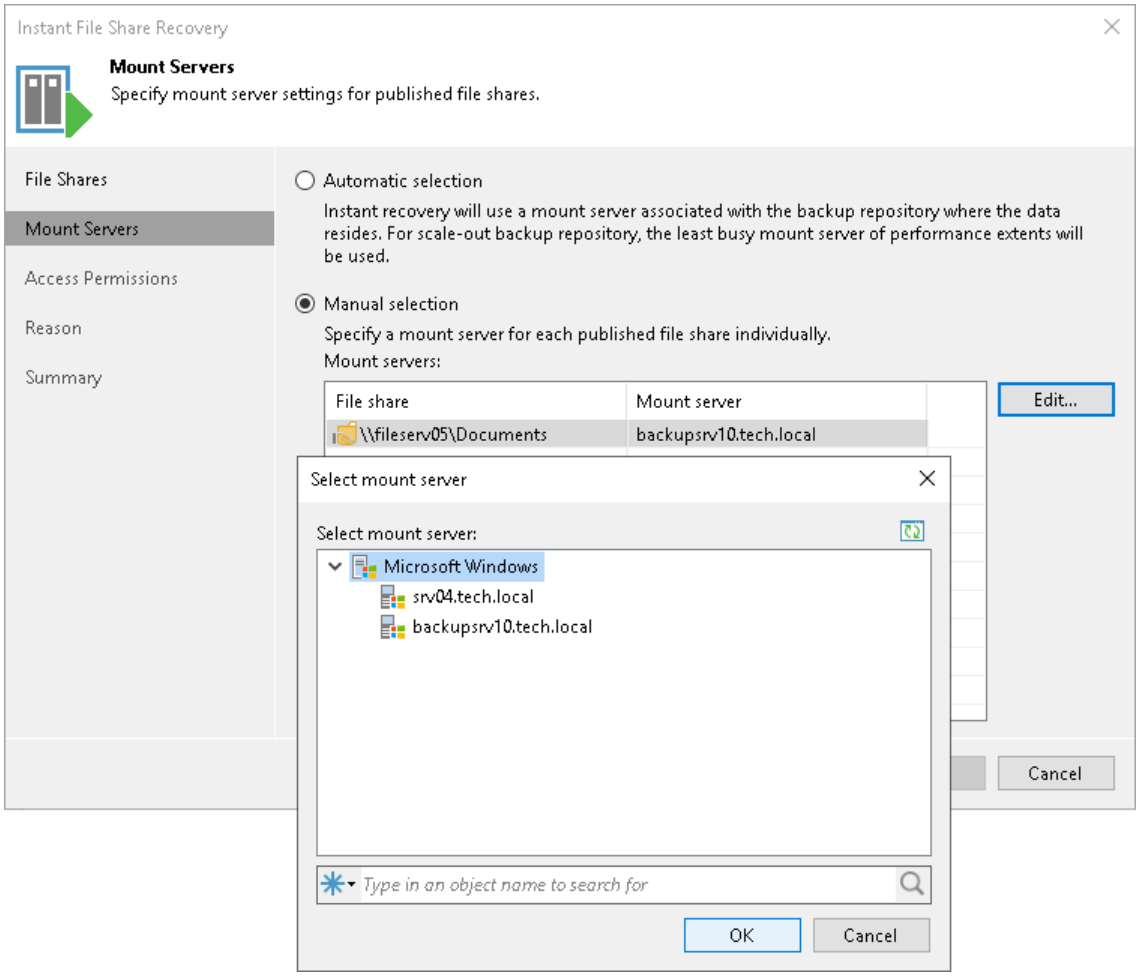

## Step 4. Specify Access Permissions

After you specify file shares and mount servers, Veeam Backup & Replication validates them. If Veeam Backup & Replication detects missing security descriptors on the file shares, it adds the **Access Permissions** step to the wizard. At this step you can specify the owner account and permissions for the file share.

1. From the **File shares** list, select a file share for which you want to specify an owner account and permissions.
2. Click **Set Owner** and specify the owner account for the file share.
3. Click **Permissions** and configure access permissions for the file share. The following options are available:
   - Deny to everyone
   - Allow to everyone
   - Allow to the following accounts or groups only

     Use **Add** and **Remove** buttons to configure accounts and groups to which you want to grant permissions for accessing the file share.

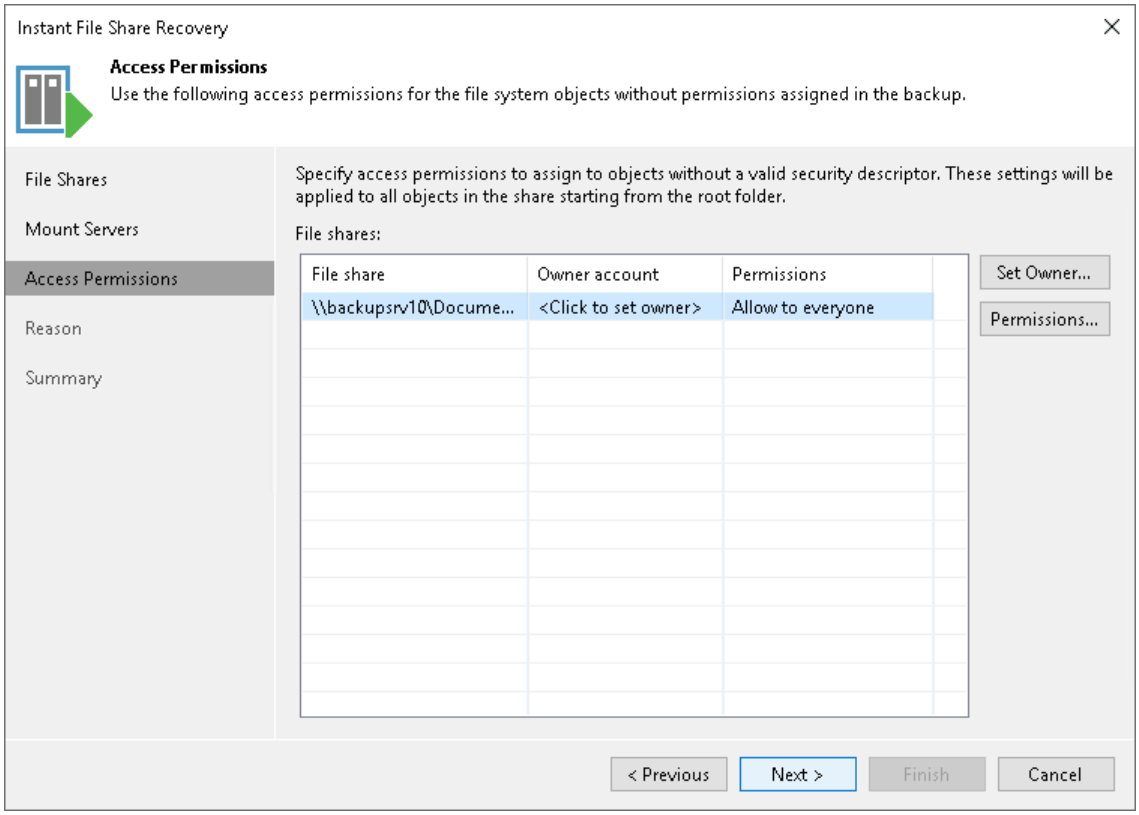

## Step 5. Specify Reason for Recovery

At the **Reason** step of the wizard, specify the reason for performing instant file share recovery. You can leave the field blank.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

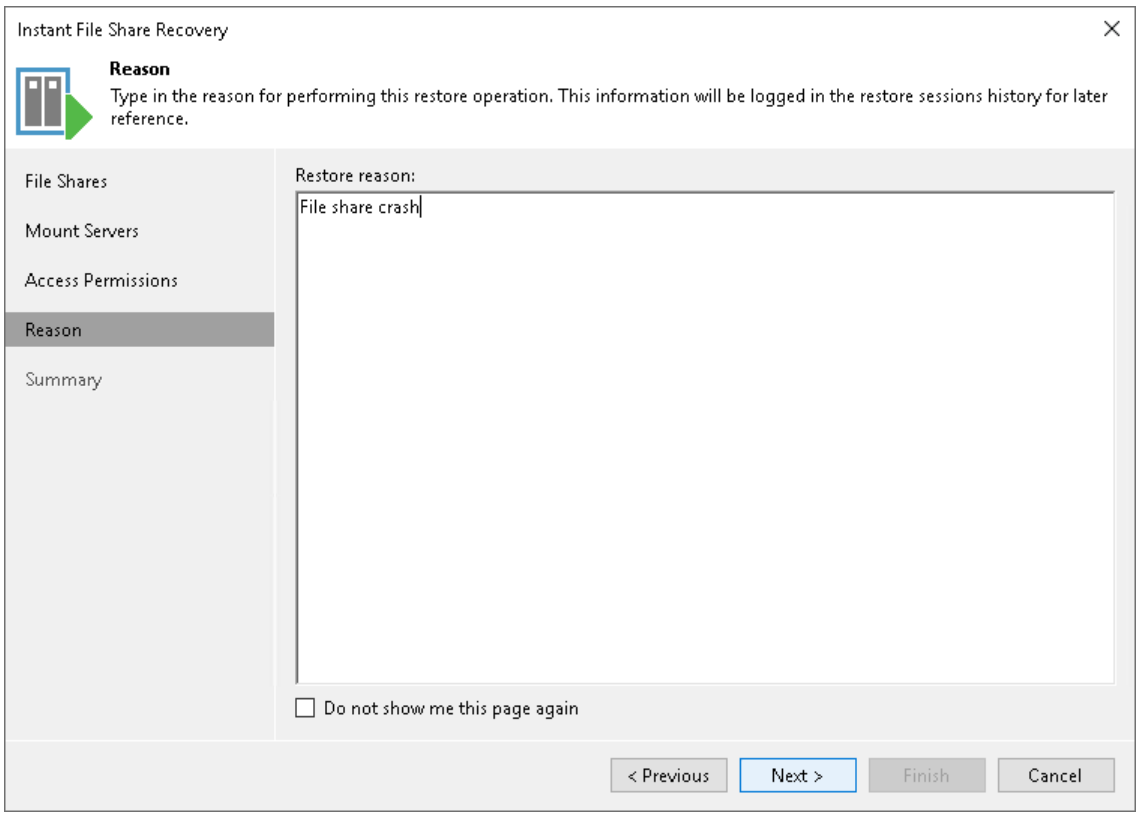

### Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review the instant file share recovery settings and click **Finish**. Veeam Backup & Replication will publish the file share to the specified mount servers.

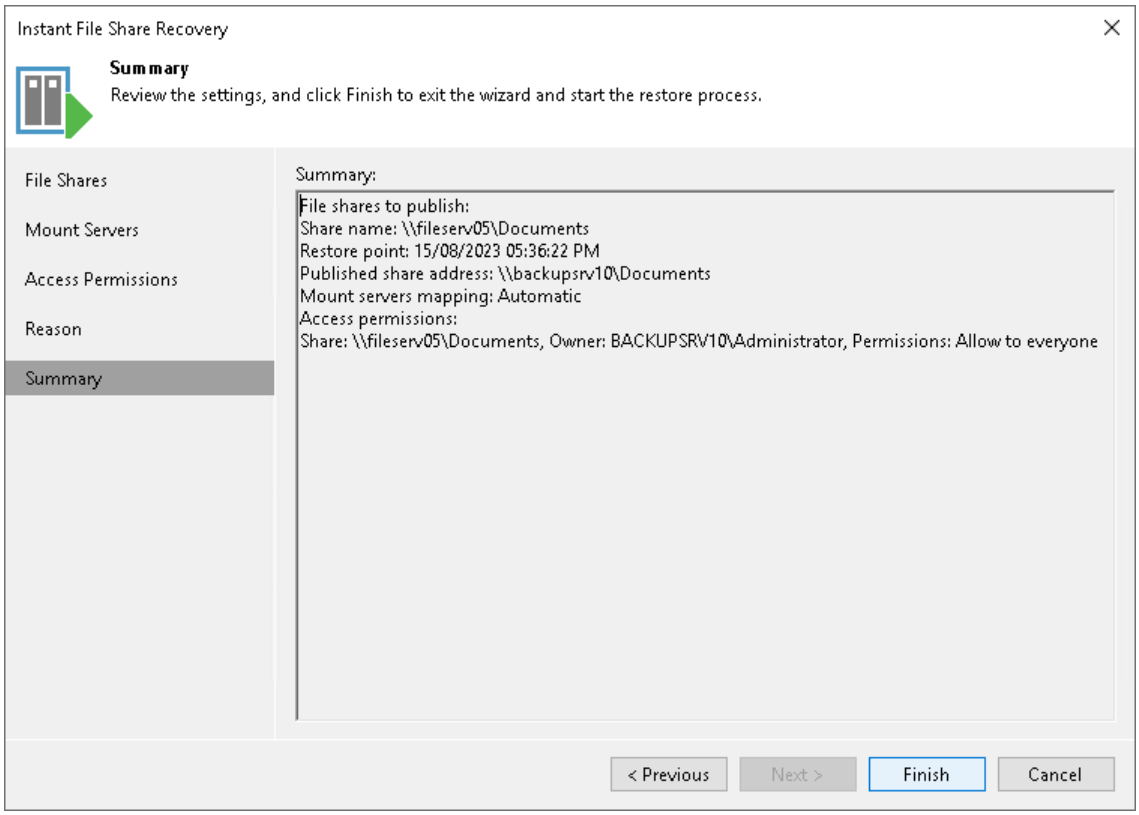


## Migrating File Share to Production

> **NOTE**
>
> Migration to production is supported for backups of SMB file shares only.

After you recover an SMB file share as described in section Performing Instant File Share Recovery, you can update the content of the mounted file share. After that, you can keep all the changes by migrating the mounted file share to production.

## Step 1. Launch Migrate to Production Wizard

To launch the **Migrate to Production** wizard, open the **Home** view. In the inventory pane, select **Instant Recovery**. In the working area, select the necessary published file share and do one of the following:

- In the **Instant Recovery** tab on the ribbon, click **Migrate to Production**.
- Right-click the file share that you want to migrate and select **Migrate to Production**.

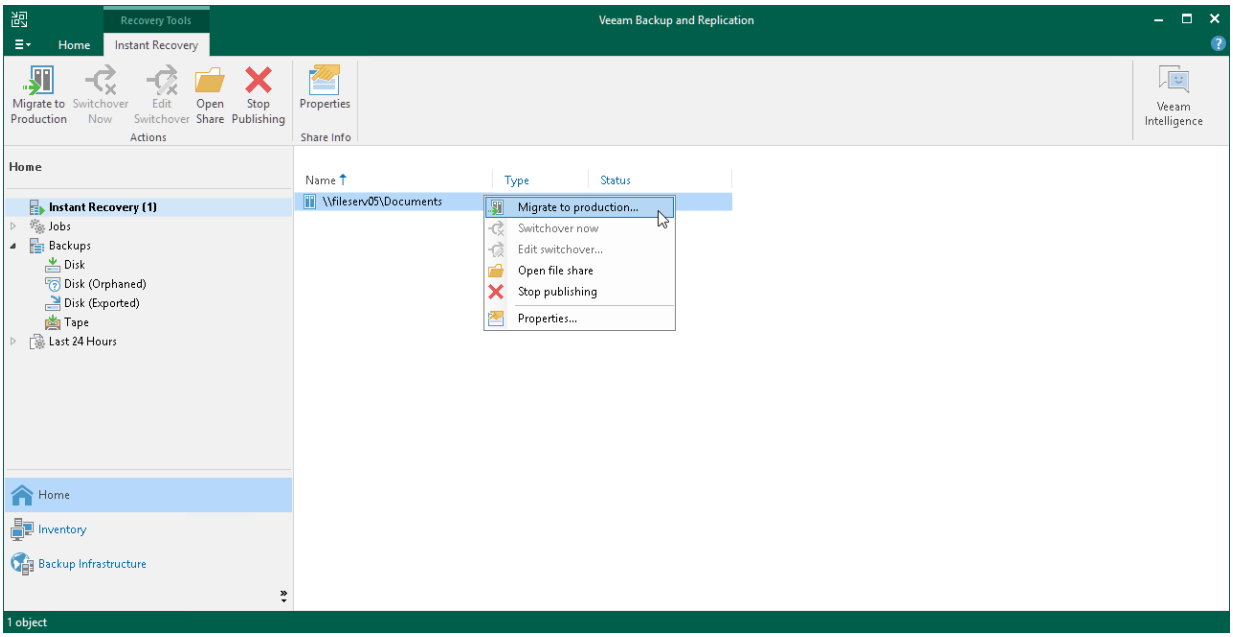

## Step 2. Define Destination for Migration

At the **Destination** step of the wizard, specify the location where you want to migrate the selected file share to.

- Select **Original location** to migrate data to the location where it resided originally. This type of migration is only possible if the original device is connected to Veeam Backup & Replication and powered on.
- Select **This server** to migrate data to another location:
  a. In the **This server** field, select a file share where files must be migrated to. You can select any file share added to the backup inventory. If the required file share is missing in the drop-down list, click **Add** and add a new file share to Veeam Backup & Replication. For more information on how to add a new file share, see Adding Unstructured Data Source.
  b. In the **Path to folder** field, specify a path to the folder on the selected file share where files must be migrated to.

     To select a specific folder on the file share to migrate files to, click **Browse**. In the **Select Folder** window, select the target location for the file share.

     If you want to restore the file share to a new folder, click **New Folder** at the bottom of the window, enter the folder name and click **OK** to confirm the new folder creation.

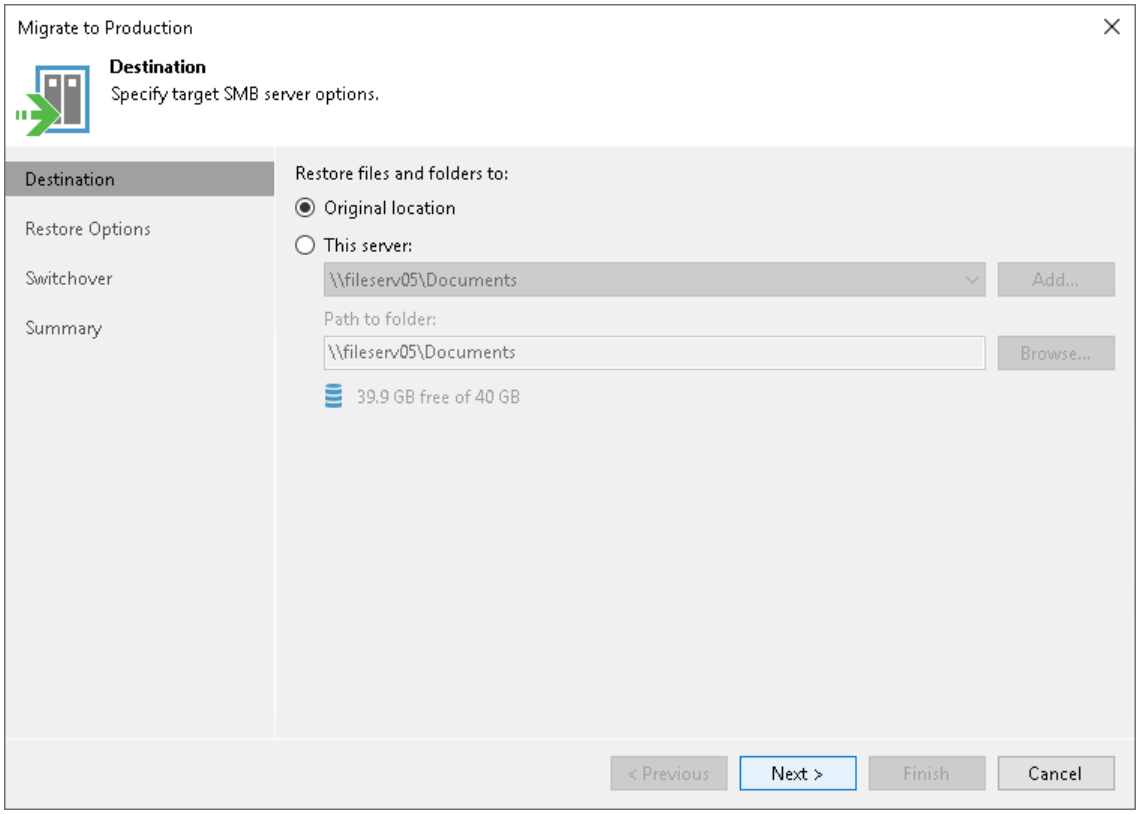

## Step 3. Specify Restore Options

At the **Restore Options** step of the wizard, specify overwrite options in case the file with the same name already exists in the destination:

- If you want to overwrite the existing files only if they are older than the restored files, select the **Replace older files only** option.
- If you want to restore the whole file share and overwrite the existing files with the restored files, select the **Restore anyway (overwrites the existing file)** option.

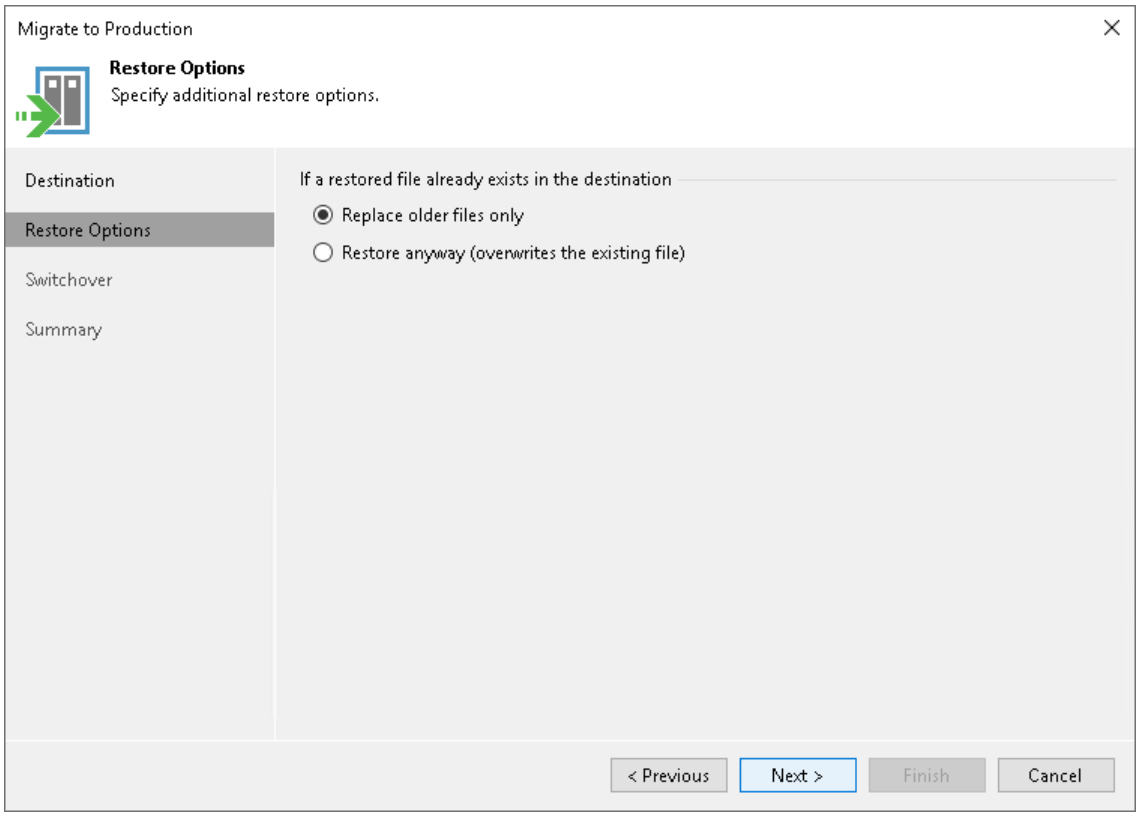

## Step 4. Specify Switchover Options

During migration to production, Veeam Backup & Replication moves to the production site not only content of the initial file share, but also incremental changes made by users in the mounted file share. When incremental changes are being moved, the mounted share is not available to users. We call this stage a switchover. The switchover may take some time, so ensure you properly plan when it is performed.

At the **Switchover** step of the wizard, specify file share switchover options:

- If you want Veeam Backup & Replication to perform the switchover automatically once the entire file share is migrated, select the **Automatic** option.
- If you want to perform the switchover manually after the entire file share is migrated, select the **Manual** option.

  To launch the switchover manually, open the **Home** view. In the inventory pane, select **Instant Recovery**. In the working area, select the necessary published file share and do one of the following:

  - In the **Instant Recovery** tab on the ribbon, click **Switchover Now**.
  - Right-click the file share that you want to switch over to production and select **Switchover now**.
- If you want Veeam Backup & Replication to perform the switchover at a certain moment, select the **Scheduled at** option and specify when you want it to be done.

  Before the scheduled switchover starts, you can edit the switchover time. To edit it, open the **Home** view. In the inventory pane, select **Instant Recovery**. In the working area, select the necessary published file share and do one of the following:

  - In the **Instant Recovery** tab on the ribbon, click **Edit Switchover**.
  - Right-click the file share that you want to update the switchover time for and select **Edit switchover**.

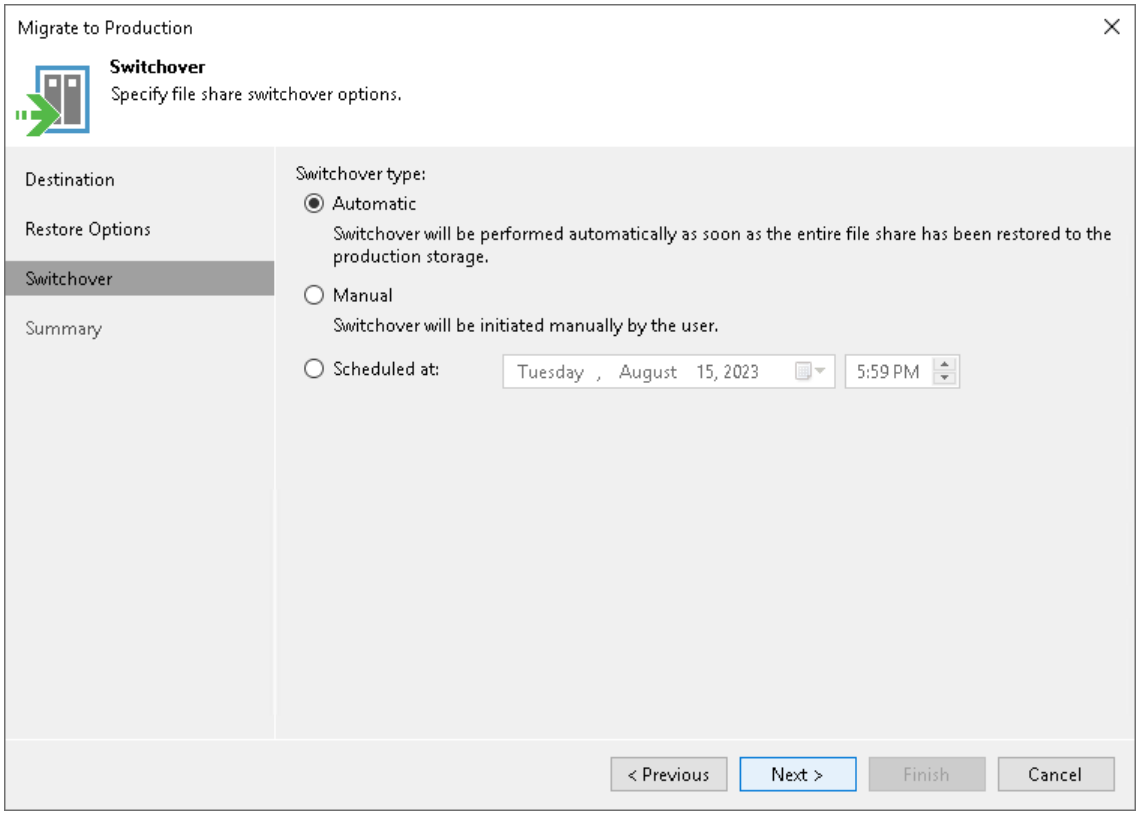

### Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review the migration to production settings and click **Finish**. Veeam Backup & Replication will migrate the selected file share to the specified production destination.

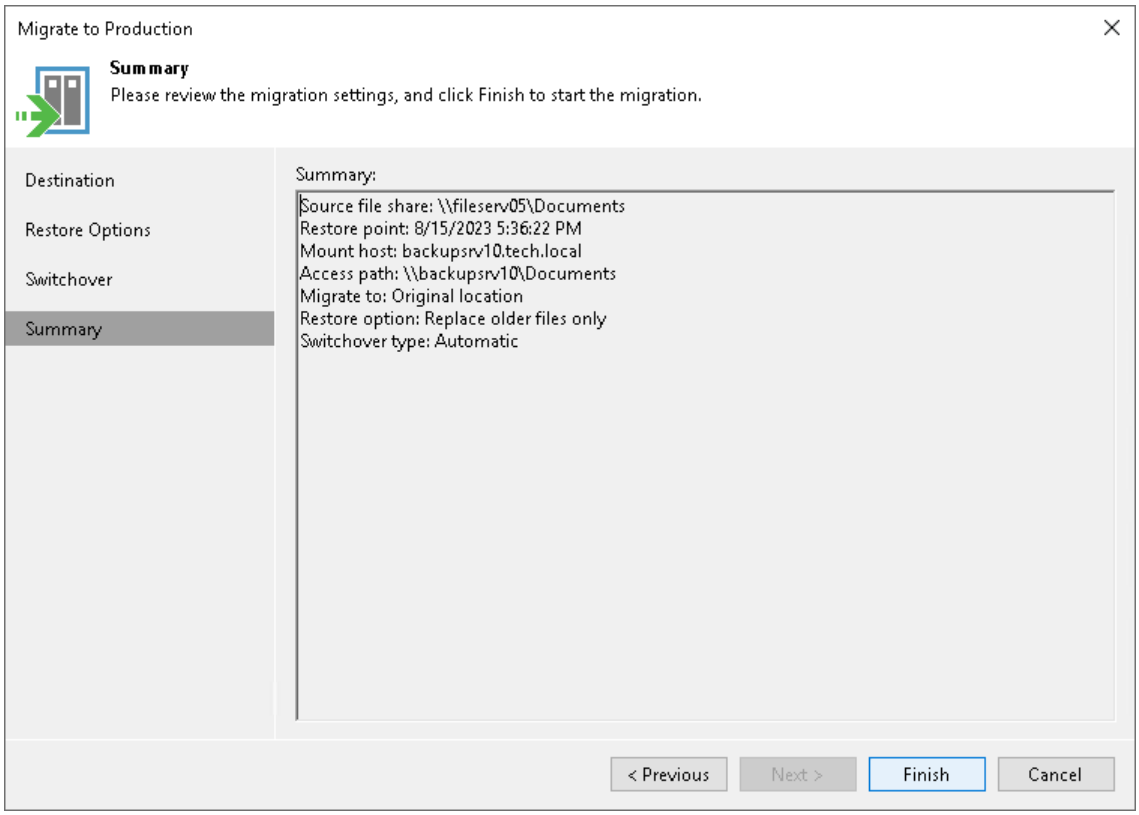


## Stopping File Share Publishing

If you do not need the published share any more, you can stop the file share publishing. To do that, open the **Home** view. In the inventory pane, select **Instant Recovery**. In the working area, select the necessary published file share and do one of the following:

- In the **Instant Recovery** tab on the ribbon, click **Stop Publishing**.
- Right-click the file share that you want to unmount and select **Stop publishing**.

# Restoring Entire File Share

You can restore the entire file share from the backup to a specific restore point. That can be helpful, for example, if your file share device gets out of order and you need to restore the entire file share to the original or other location.

> **NOTE**
>
> If the **Archive recent file versions** option is selected (the copy mode is enabled) at the Archive Repository step of the file backup job wizard, you may restore an entire file share from the archive repository. You can do that only for restore points that are stored in the archive repository and have the **Copied** label in backup properties.
>
> If this option is not selected (the copy mode is disabled), the restore of the entire file share or even whole folders from the archive repository is not supported.

Before you restore an entire file share, check prerequisites.

## Before You Begin

Before you restore the entire file share, consider the following:

- You can restore files and folders of the file share from a backup that has at least one successfully created restore point.
- The file share on which you plan to save restored files and folders must be added to the backup infrastructure.
- During the backup of an SMB file share, Veeam Backup & Replication does not collect ACL for the root shared folder. Therefore, if you restore the entire file share, ACL of the root folder is not restored. To solve this issue, prepare a folder with the required permissions before running the entire file share restore. During the restore, use this folder as a target root folder for the restored share.
- Regardless of restore options, if a backup has an item with the same name but a different type as one in the file share, this file will not be restored. For example, folder named `test.txt` and a text file named `text`.

# Step 1. Launch File Restore Wizard

To launch the **File Restore** wizard, do one of the following:

- In the **Home** tab on the ribbon, click **Restore** > **File Share**. In the **Restore from File Backup** window, click **Restore entire share**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the file share that you want to restore. In the **Backup** tab on the ribbon, click **Entire share**.
  - Right-click the file share that you want to restore and select **Entire file share**.

  You can restore the file share to the state as of a specific restore point by using a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane. If the secondary repository is an object storage repository, backup copies created in it are represented in the **Backups** > **Object Storage (Copy)** node in the inventory pane.

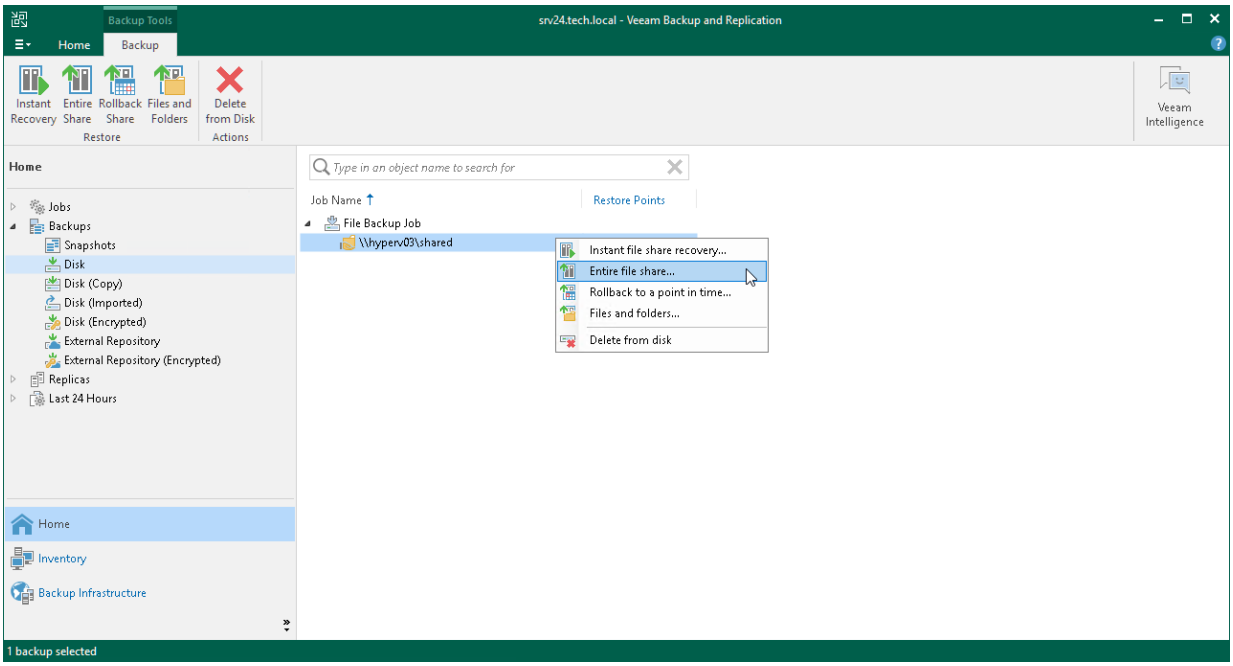

## Step 2. Select File Share to Restore

At the **File Shares** step of the wizard, select the file share that you want to restore:

1. Click **Add**.
2. In the **Backups Browser** window, select the file backup job and a file share in it that you want to restore. Click **OK**.
3. In the **File shares** table, choose the file share to select a restore point to restore to. Click **Point**.
4. In the **Select Restore Point** window, select the restore point to which you want to restore the file share. To select the required restore point, do one of the following:
   - Use the **Restore point** slider.
   - Click the date link under the **Restore point** slider. In the calendar in the left pane of the **Select Restore Point** window, select the date when the required restore point was created. The list of restore points in the right pane displays restore points created on the selected date. Select the point to which you want to restore the file share.

   In the **Files in backup** tree, you can see what folders and files are covered by the selected restore point and the date when each of them was modified.

   Click **OK**.

To quickly find a file share, you can use the search field at the top of the window. Enter a file share name or a part of it in the search field and press [Enter].

To exclude the file share from the restore process, select the file share in the table and click **Remove**.

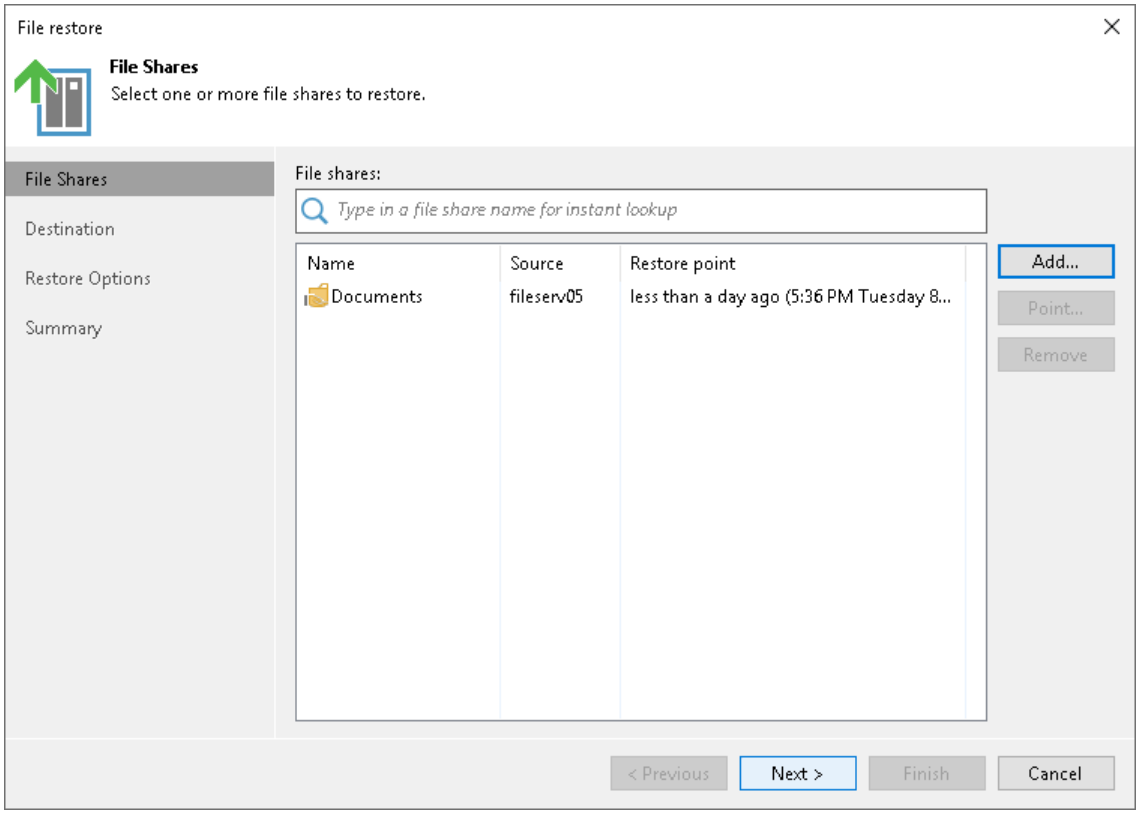

## Step 3. Specify Destination for Data Restore

At the **Destination** step of the wizard, specify the location where you want to restore the file share.

- Select **Original location** to restore data to the location where it resided originally. This type of restore is only possible if the original device is connected to Veeam Backup & Replication and powered on.
- Select **This server** to restore data to another location:
  a. In the **This server** field, select a file share to restore files to. You can select any file share added to the backup inventory. If the required file share is missing in the drop-down list, click **Add** and add a new file share to Veeam Backup & Replication. For more information on how to add a new file share, see Adding Unstructured Data Source.
  b. In the **Path to folder** field, specify a path to the folder on the selected file share to restore files to.

     To select a specific folder on the file share to restore files to, click **Browse**. In the **Select Folder** window, select the target location for the file share.

     If you want to restore the file share to a new folder, click **New Folder** at the bottom of the window, enter the folder name and click **OK** to confirm the new folder creation.
  c. Select **Preserve folder hierarchy** to keep the folder hierarchy of the original file share in the new location.

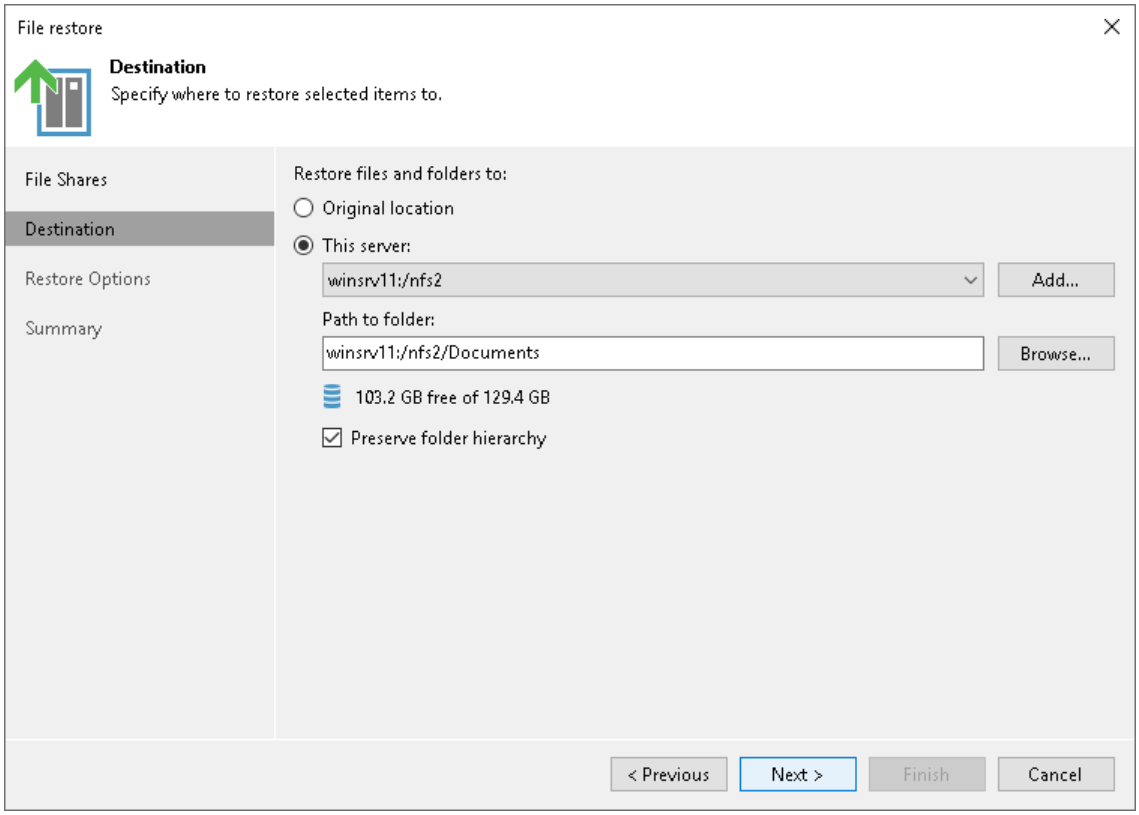

## Step 4. Specify Restore Options

At the **Restore Options** step of the wizard, specify overwrite options in case the file with the same name already exists in the target folder:

- **Skip restoring (keeps the existing file)**. Select this option if you do not want to overwrite the existing file with the restored one.
- **Replace older files only (use if a share was reverted to a storage snapshot)**. Select this option if you want to overwrite the existing file only if it is older than the restored file.
- **Replace newer files only (use to discard unwanted contents changes)**. Select this option if you want to overwrite the existing file only if the restored file is older than the source share file.
- **Restore anyway (overwrites the existing file)**. Select this option if you want to overwrite the existing file with the restored file in all cases.

Select the **Restore permissions and security attributes** check box if you want the restored files to keep their original ownership and security permissions. If you do not select this check box, Veeam Backup & Replication will change security settings. The user account under which the Veeam Backup Service runs will be set as the owner of the restored objects. Access permissions will be inherited from the target folder to which the objects are restored.

> **NOTE**
>
> Consider that Veeam Backup & Replication does not collect ACL handling settings of the source file share root folder, so you cannot restore them. Before restoring an entire file share, you will have to specify required ACL handling settings for the root folder of the target file share.

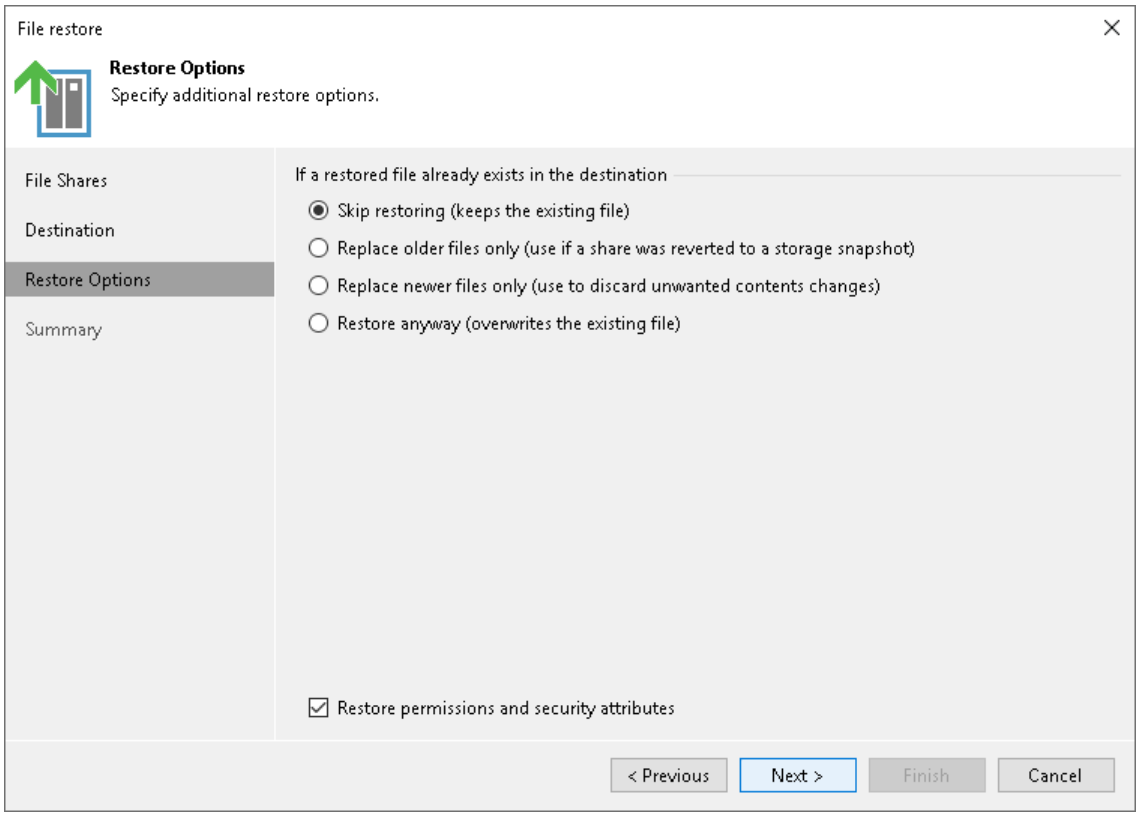

## Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review the file share restore settings and click **Finish**. Veeam Backup & Replication will restore the file share to the specified location.

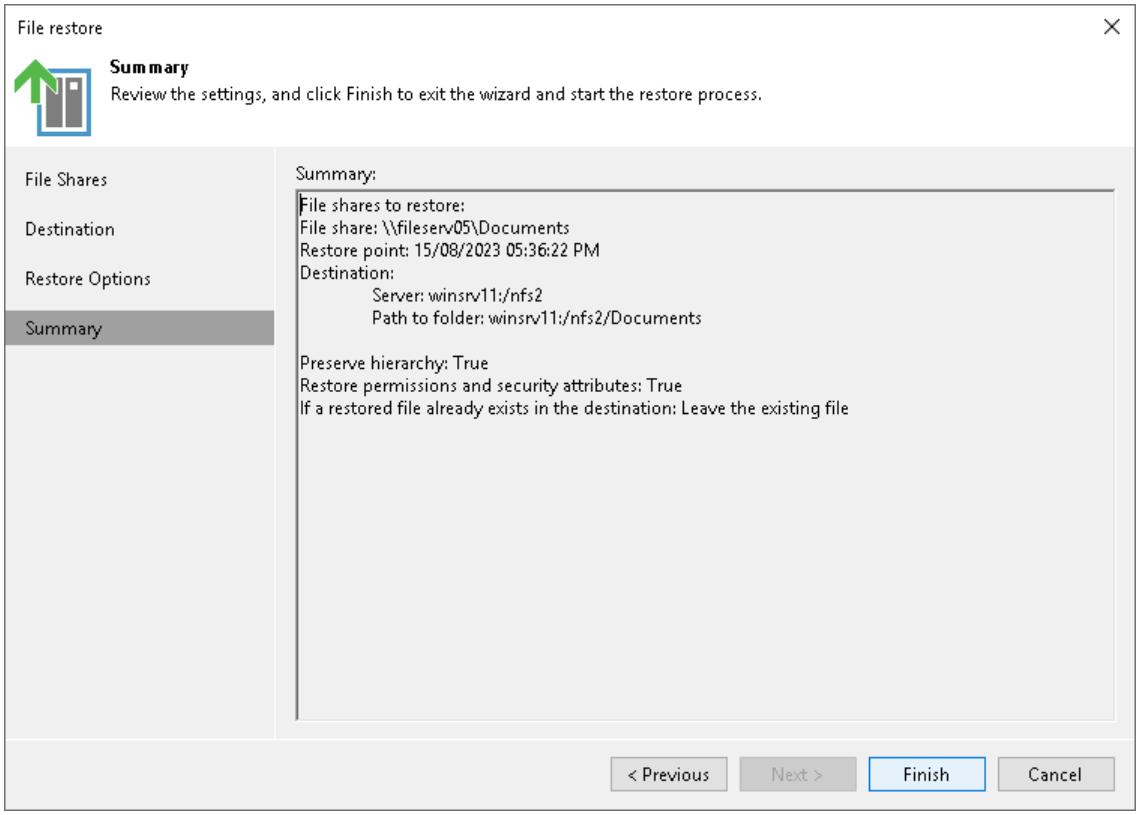


# Rolling Back to a Point in Time

You can roll back changes made to files and folders on the file share to a specific restore point that is older than the current file share state. This option can be useful, for example, when the original file share was attacked by ransomware. In this case you can roll back all the files that were changed by the ransomware to the state before the attack.

Before you roll back the file share to a point in time, check prerequisites.

## Before You Begin

Before you restore files and folders to a point in time, consider that you can restore files and folders of the file share from a backup that has at least one successfully created restore point.

# Step 1. Launch File Restore Wizard

To launch the **File Restore** wizard, do one of the following:

- In the **Home** tab on the ribbon, click **Restore** > **File Share**. In the **Restore from File Backup** window, click **Rollback to a point in time**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the file share backup whose files you want to restore. In the **Backup** tab on the ribbon, click **Rollback to a point in time**.
  - Right-click the file share backup whose files you want to restore and select **Restore** > **Rollback to a point in time**.

  You can roll back the file share to a point in time by using a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane. If the secondary repository is an object storage repository, backup copies created in it are represented in the **Backups** > **Object Storage (Copy)** node in the inventory pane.

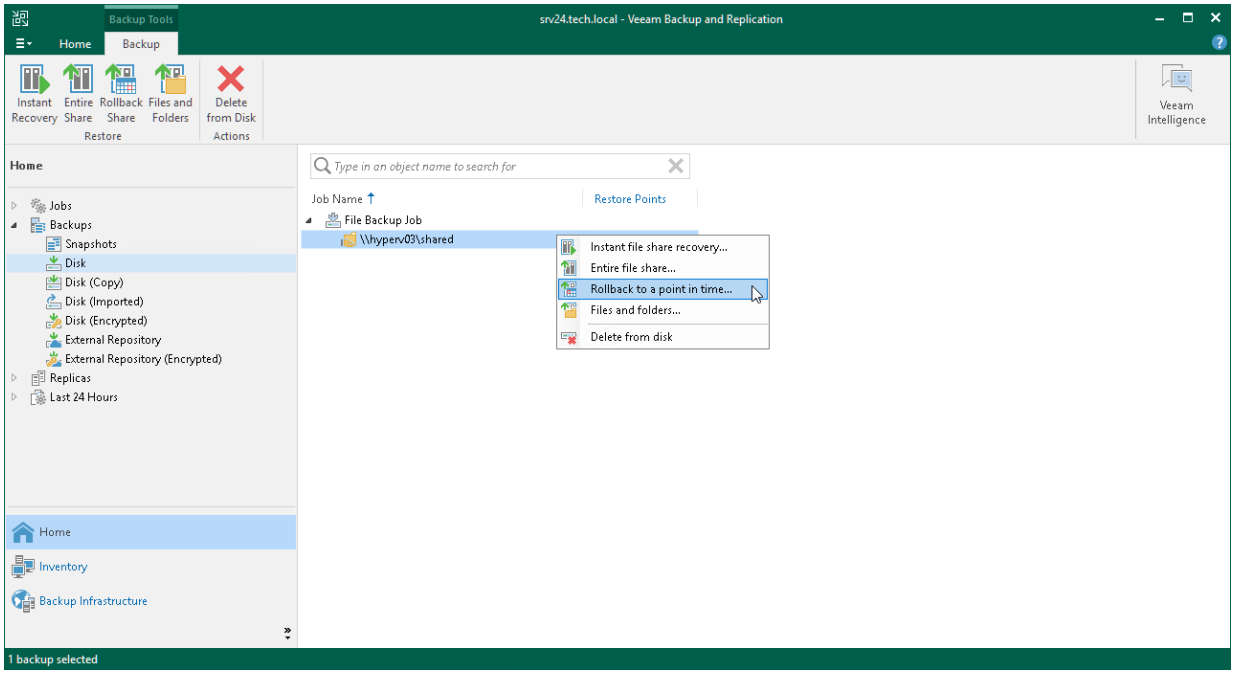

## Step 2. Select Object to Restore

At the **File Shares** step of the wizard, select the file share you want to roll back:

1. Click **Add**.
2. In the **Backups Browser** window, select the file backup job and then choose the file share you want to roll back. Click **OK**.
3. In the **File shares** table, choose the file share to select a restore point to roll back to. Click **Point**.
4. In the **Select Restore Point** window, select the restore point to which you want to roll back the files. To select the required restore point, do one of the following:
   - Use the **Restore point** slider.
   - Click the date link under the **Restore point** slider. In the calendar in the right pane of the **Restore points** window, select the date when the required restore point was created. The list of restore points on the left pane displays restore points created on the selected date. Select the point to which you want to roll back the files to.

   In the **Files in backup** tree, you can see what folders and files are covered by the selected restore point and the date when each of them was modified.

   Click **OK**.

To quickly find a file share, you can use the search field at the top of the window. Enter a file share name or a part of it in the search field and press [Enter].

To exclude the file share from the rollback process, select the file share in the table and click **Remove**.

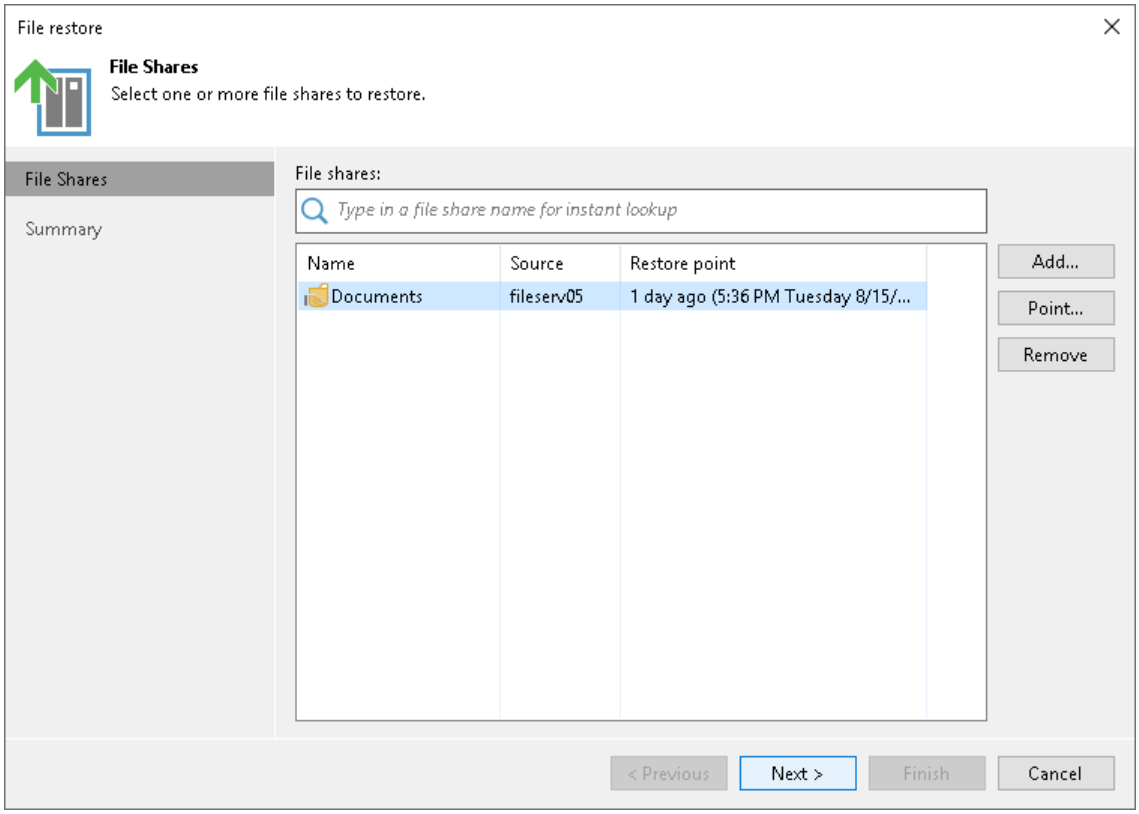

## Step 3. Finish Working with Wizard

At the **Summary** step of the wizard, review the file share restore settings and click **Finish**. Veeam Backup & Replication will restore the files to the specified point in time.

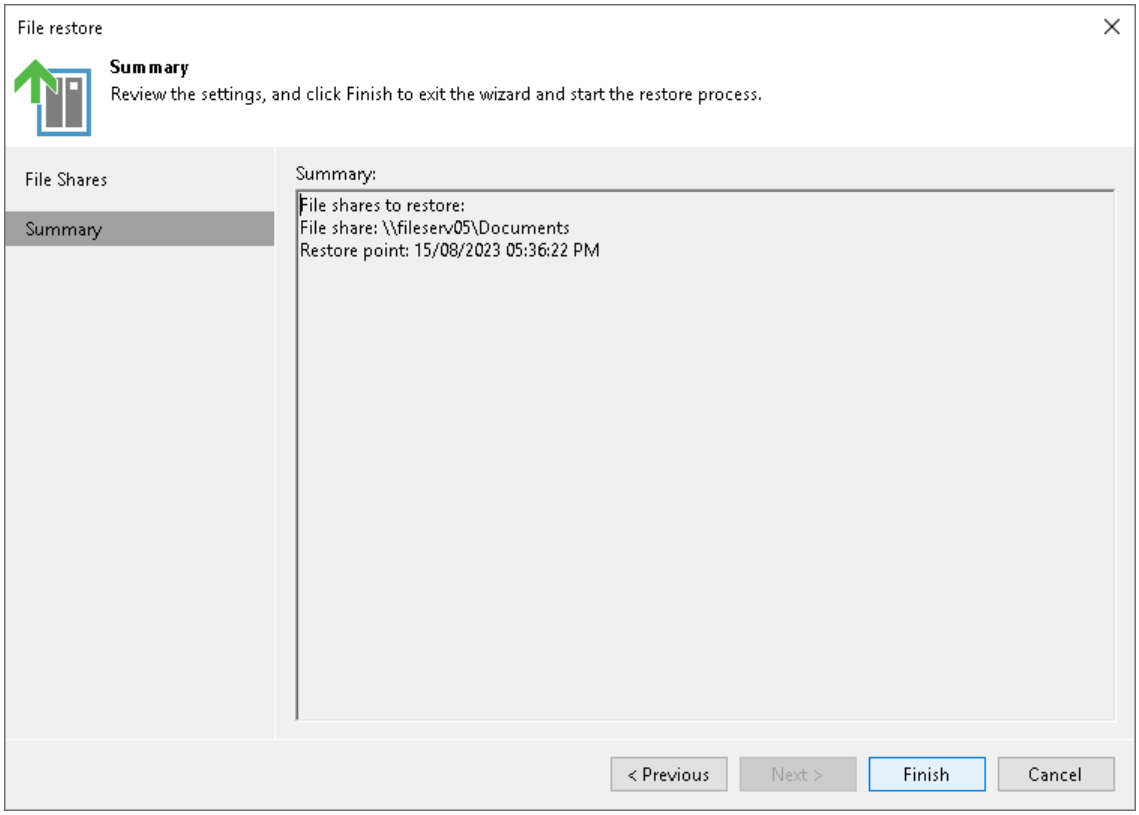


# Restoring Specific Files and Folders

You can restore specific files and folders to the original or a new location. This option can be useful, for example, if you need to get an older version of some files and folders from the backup.

When you restore specific files, you can extract file versions not only from the backup repository, but also from the archive repository. For more information, see Restoring Backup Files from Archive Repository.

> **NOTE**
>
> Consider that from the archive repository you can restore files only. Restore of whole folders from the long-term repository is not supported.

Besides, you can restore multiple versions of the same file.

Before you restore specific files and folders, check prerequisites.

## Before You Begin

Before you restore files and folders from the backup, consider the following:

- You can restore files and folders from a backup that has at least one created restore point, even if it is incomplete.

- The file share on which you plan to save restored files and folders must be added to the backup infrastructure.

# Step 1. Launch File Restore Wizard

To launch the **File Restore** wizard, do one of the following:

- In the **Home** tab on the ribbon, click **Restore** > **File Share**. In the **Restore from File Backup** window, click **Restore individual files and folders**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the file share backup whose files you want to restore. In the **Backup** tab on the ribbon, click **Files and folders**.
  - Right-click the file share backup whose files you want to restore and select **Restore** > **Files and folders**.

  In this case, you will pass directly to the Backup Browser.

  You can restore files and folders from a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane. If the secondary repository is an object storage repository, backup copies created in it are represented in the **Backups** > **Object Storage (Copy)** node in the inventory pane.

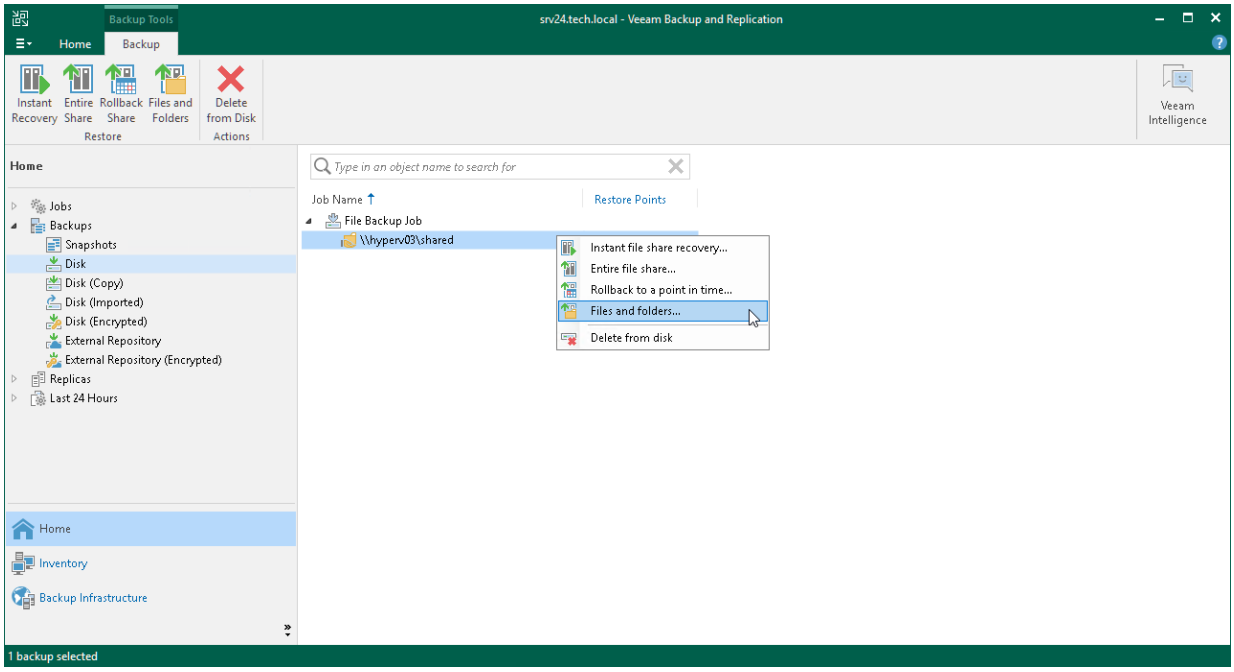

## Step 2. Select Object to Restore

At the **Backup** step of the wizard, select the file share backup you want to restore files from:

1. In the **Object to restore** list, expand the necessary backup job.
2. Select the file share.

To quickly find a file share, you can use the search field at the bottom of the window.

1. Enter a file share name or a part of it in the search field.
2. Click the **Start search** button on the right or press [Enter].

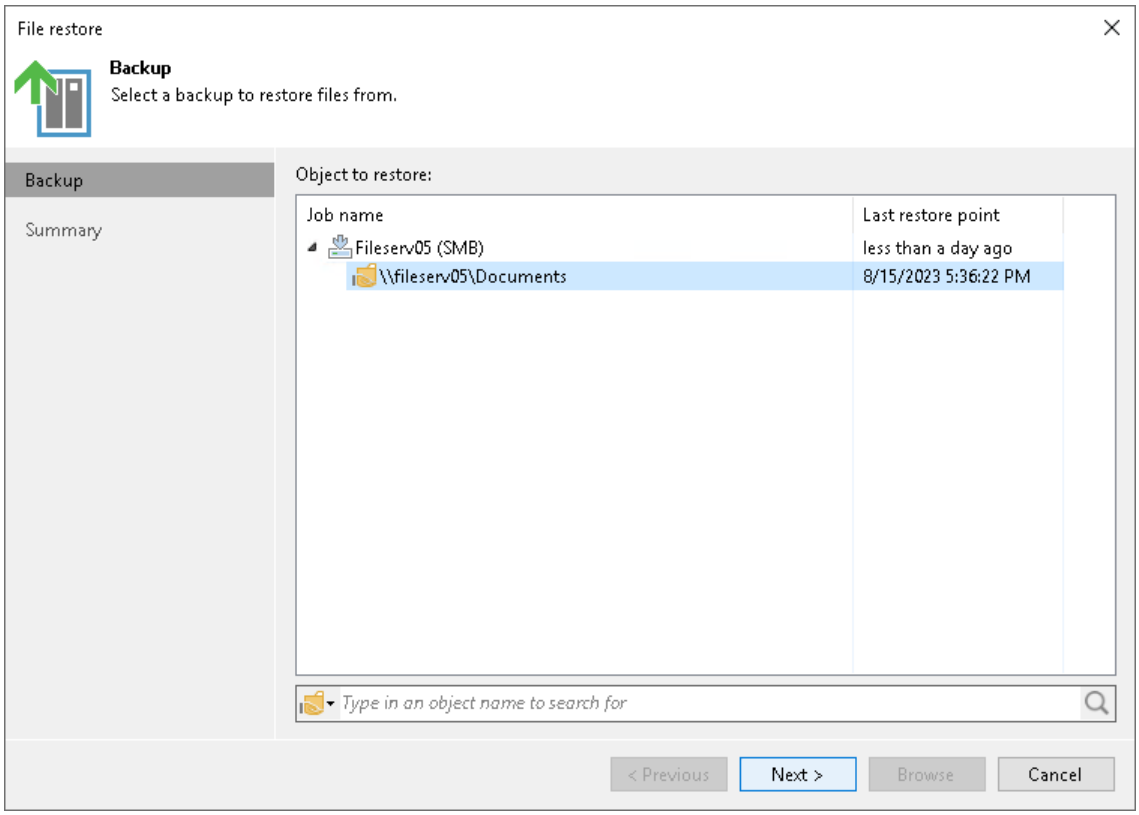

## Step 3. Verify Object Restore Settings

At the **Summary** step of the wizard, review selected restore object and click **Browse** to switch to the Backup Browser step and select files and folders to restore.

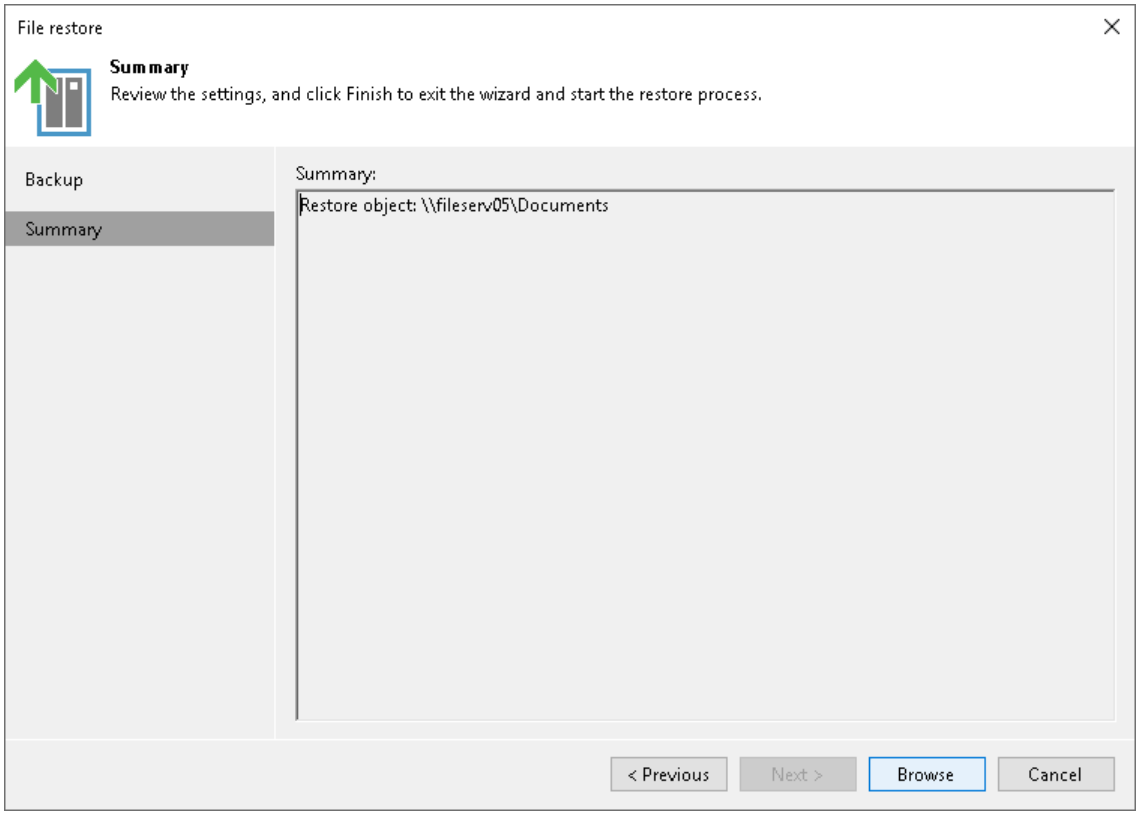

## Step 4. Select Files and Folders to Restore

In the **Backup Browser**, select files and folders to restore. Backup browser has three representations of restore points. You can select one of the following options in the **Restore Point** group on the ribbon to display file versions:

- **Latest** – the **Backup Browser** shows the latest versions of files and folders on the file share.
- **All Time** – the **Backup Browser** shows all files and folders ever backed up by the backup job. This option retrieves file versions stored both in the backup and archive repositories. This representation additionally shows how many file versions of each file are stored in the backup and the date when the latest file version was created.

  After you select this option:

  - If you restore a whole folder, you will be prompted to the Select Restore Point to Use step to select a restore point to restore files from.

    > **NOTE**
    >
    > Consider that during file-level restore you cannot restore whole folders from the archive restore points. If the files are already moved to the archive repository, you have to restore them one by one.

  - If you restore a single file, you will be prompted to the Select File Version to Restore step to select a file version to restore the file from.
- **Selected** – the **Backup Browser** shows versions of files and folders backed up as of the certain restore point. Select the restore point in the list on the right of the **Selected** option on the ribbon.

You can use the search field at the top of the working area to search for specific files and folders.

> **NOTE**
>
> To keep the operation of the Backup Server stable, the number of retrieved search result records is limited to 1000. Therefore, if you work with backup folders that store large volumes of data, it is recommended to narrow the search criteria to fit into the limitation.

You can restore files and folders to their original location or a new location.

# Restoring Files to Original Location

To restore files and folders to the original location, in the Veeam Backup browser right-click a file or folder and select one of the following commands:

- To overwrite the original file on the file share with the file restored from the backup, select **Restore** > **Overwrite**.
- To save the file restored from the backup next to the original file, select **Restore** > **Keep**.

  Veeam Backup & Replication will add the *_RESTORED_YYYYMMDD_HHMMSS* suffix to the original file name and store the restored file in the same folder where the original file resides.

Alternatively, you may use similar options under **Restore** in the ribbon.

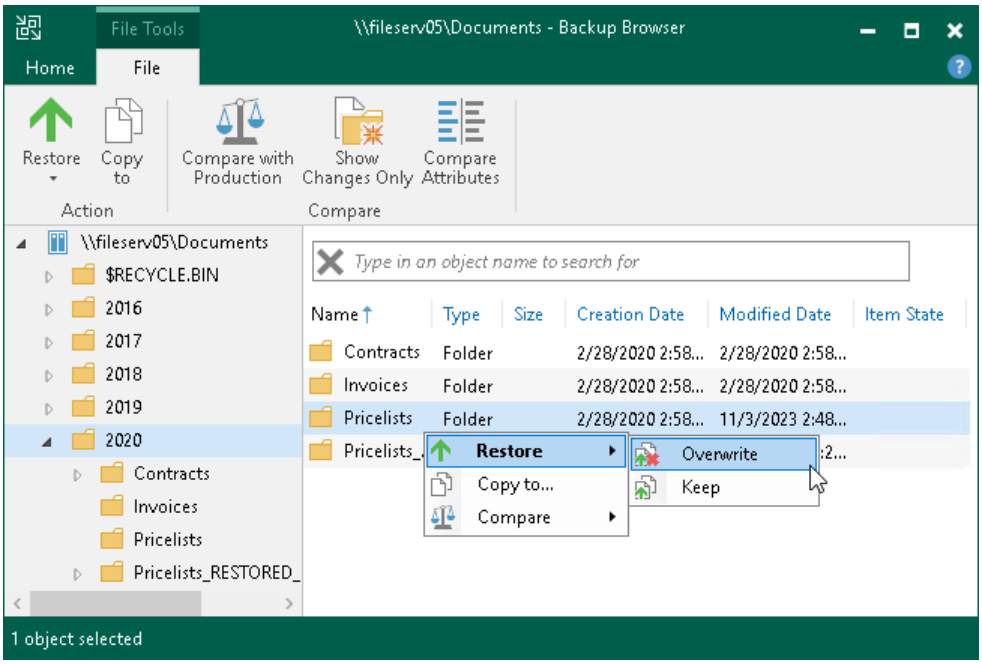


# Saving Files to New Location

To save files and folders to a new location, right-click the necessary file or folder and select **Copy to**. Alternatively, you may use the **Copy to** option in the ribbon.

If you restore files and folders to a new location, you will be prompted to the Specify Destination for File Restore step to specify a new destination.

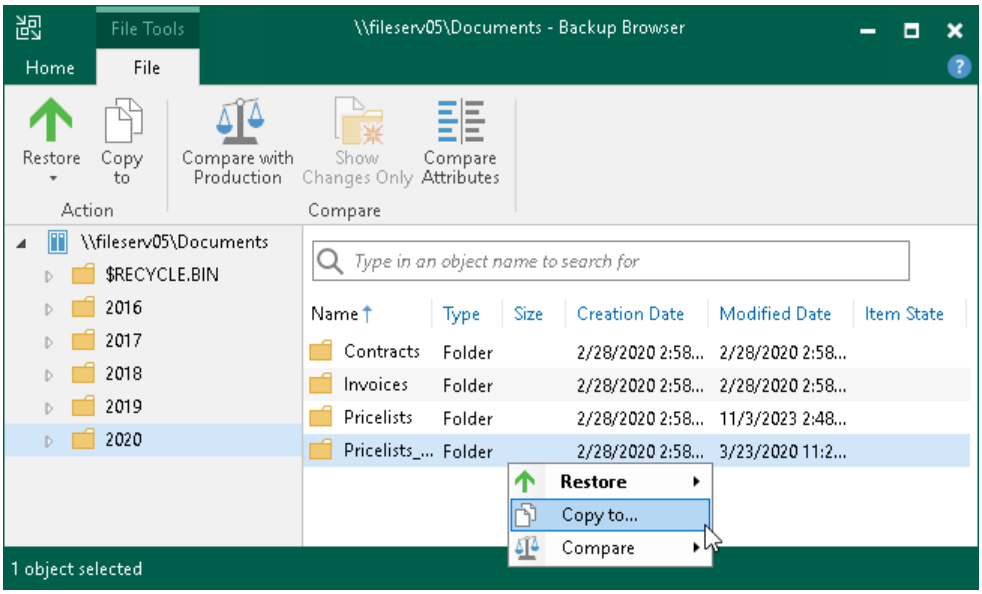

# Comparing Backup File and Folder Versions with Production Objects

> **NOTE**
>
> The compare with production feature does not work if you select the **All Time** option at the **Home** tab of the Backup Browser.

You can compare backup versions of specific files and folders with their production sources. To compare them, do the following:

1. Select a folder in the file tree in the left pane or a folder or file in the right pane. You can use [Ctrl] to select multiple objects in the right pane.
2. Right-click the selected objects and select **Compare** > **Compare**. Alternatively, click **Compare with Production** on the ribbon.

After the comparison, files and folders will have the following comparison states in the **Item State** column: *changed*, *unchanged*, *deleted*, *comparing*, or *failed to compare*. The states are updated when you turn off and then turn on the comparison mode, and when you start restoring changes of files and folders. Note that when comparing symbolic links, Veeam Backup & Replication compares attributes of the links, not the attributes of files and folders which the symbolic link points to.

> **TIP**
>
> To show only changed files and folders (in the *changed* and *deleted* states), perform the compare operation, right-click any area in the Veeam Backup browser and select **Compare** > **Show changes only** or click **Show Changes Only** on the ribbon. To show all files and folders, click the **Show changes only** option once again.
>
> To switch off the comparison states, select an item in the comparison state and click **Compare** > **Compare** or click **Compare with Production** on the ribbon. Note that if you switch off comparison for child files and folders, comparison for parent folders will also be switched off.

# Restoring Changed Files and Folders

After you run the compare with production session, you can restore the changed and deleted objects. To restore them to the original location, in the Veeam Backup browser right-click a file or folder with **Changed** or **Deleted** item state and select one of the following commands:

- To overwrite the original object on the object storage with the object restored from the backup, select **Restore changed and deleted** > **Overwrite**.
- To save the object restored from the backup next to the original object, select **Restore changed and deleted** > **Keep**.

  Veeam Backup & Replication will add the *_RESTORED_YYYYMMDD_HHMMSS* suffix to the original object name and store the restored object in the same folder where the original object resides.

Alternatively, you may use similar options under **Restore Changes** in the ribbon.

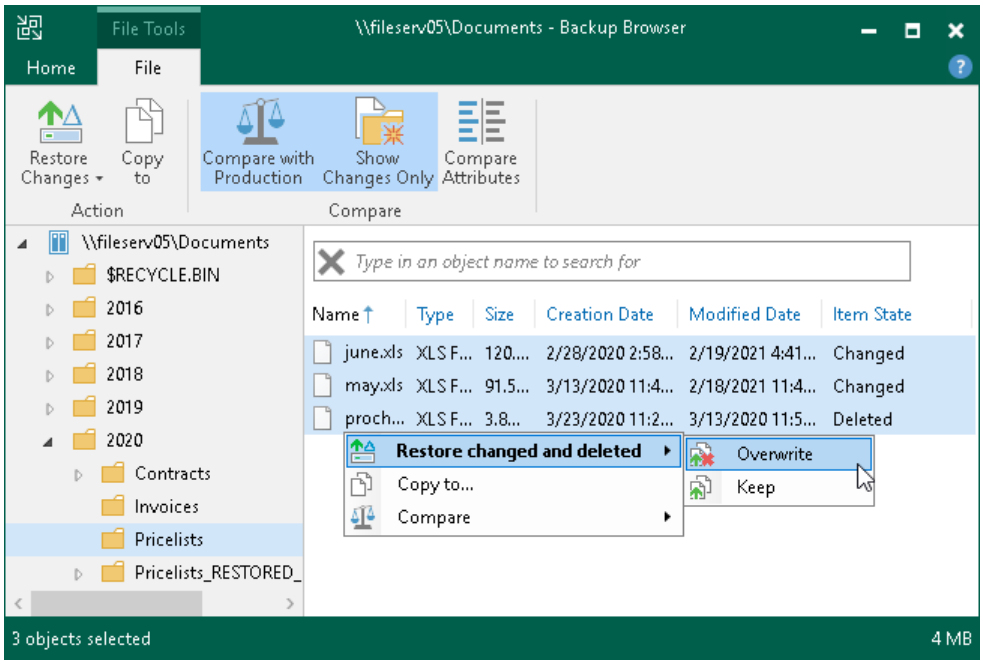


You can view which attributes were changed for files and folders:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected items and select **Compare** > **Compare attributes** or click Compare Attributes on the ribbon.

In the **Compare Attributes** window, Veeam Backup & Replication shows changed attributes. If you want to show all attributes, click the **Show unchanged attributes** check box at the top right corner. Note that Veeam Backup & Replication shows attributes maximum for 500 files and folders and shows attributes for the selected files and folders, not for the nested files.

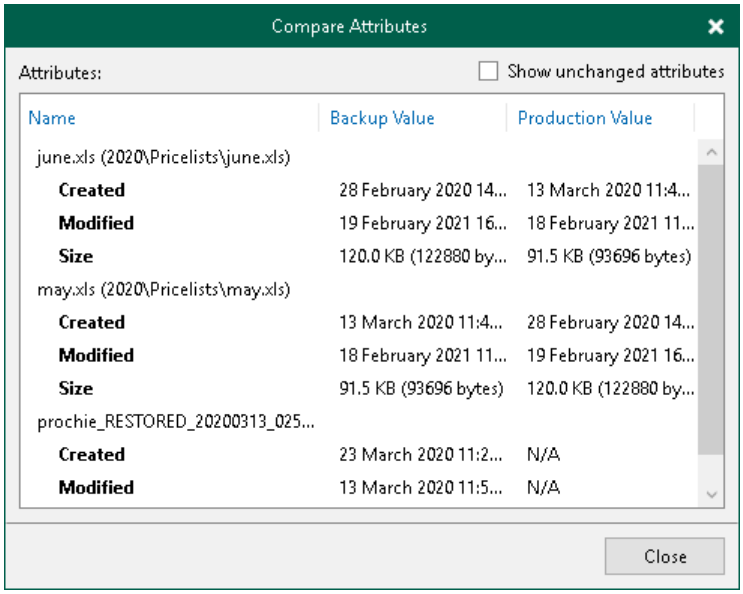

# Step 5. Select Restore Mode

The **Restore Mode** step is required if you use the **All Time** option at the Select Files and Folders to Restore step and the selected folders have more than one restore point.

Choose to what point you want to restore files and folders:

- To restore the folder to the latest available restore point, select **Latest state**.
- To select a specific restore point, select **Earlier restore point**.

  Choosing this option will open the Restore Point step.

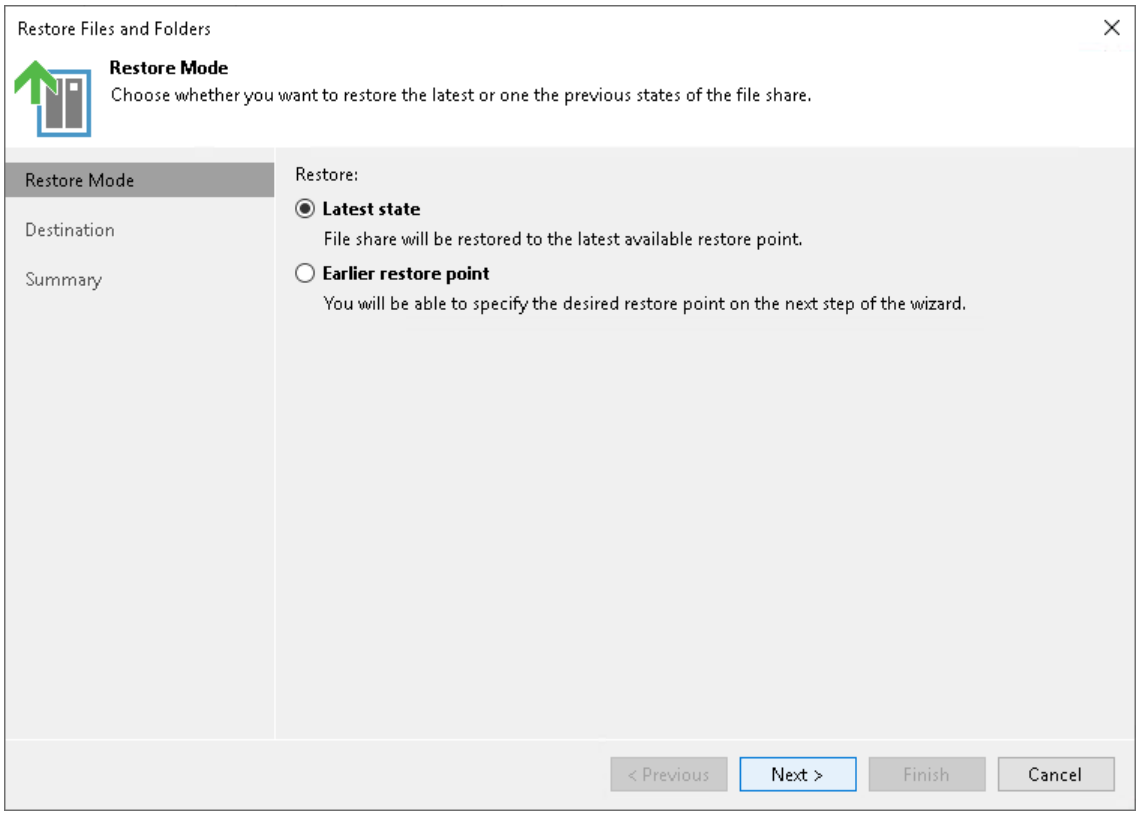

# Step 6. Select Restore Point

The **Restore Point** step is required if you use the **All Time** option at the Select Files and Folders to Restore step, the selected folders have more than one restore point, and you select the **Earlier restore point** option at the Restore Mode step.

At the **Restore Point** step of the wizard, select the point in time to restore folders to. To select the required restore point, do one of the following:

- Use the **Restore point** slider.
- Click the date link under the **Restore point** slider. In the calendar in the left pane of the **Restore points** window, select the date when the required restore point was created. The list of restore points in the right pane displays restore points created on the selected date. Select the point to which you want to restore the files and folders.

In the **Files in backup** tree, you can see what folders and files are covered by the selected restore point and the date when files and folders were modified.

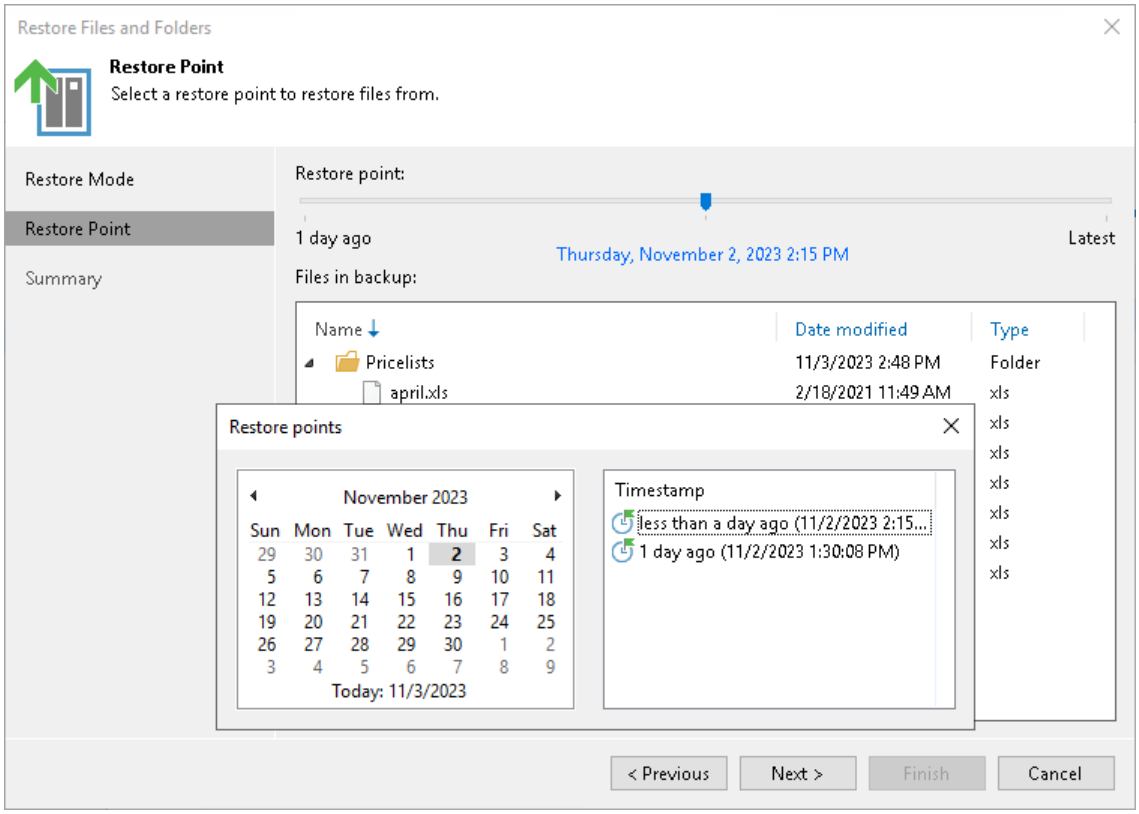

# Step 7. Select File Version to Restore

The **File Version** step is required if you use the **All Time** option at the Select Files and Folders to Restore step and the selected files have more than one file version.

If at the Backup Browser step you have selected to keep original objects, select one or more versions to restore. You can restore files both from the backup repository and archive repository. To select several file versions, hold [Ctrl] and select multiple records in the table. Restore of multiple file versions can be helpful, for example, when you need to search for a specific version of the file, but you do not know for sure which one contains required changes.

If at the Backup Browser step you have selected to overwrite original objects, you can select only one version to restore.

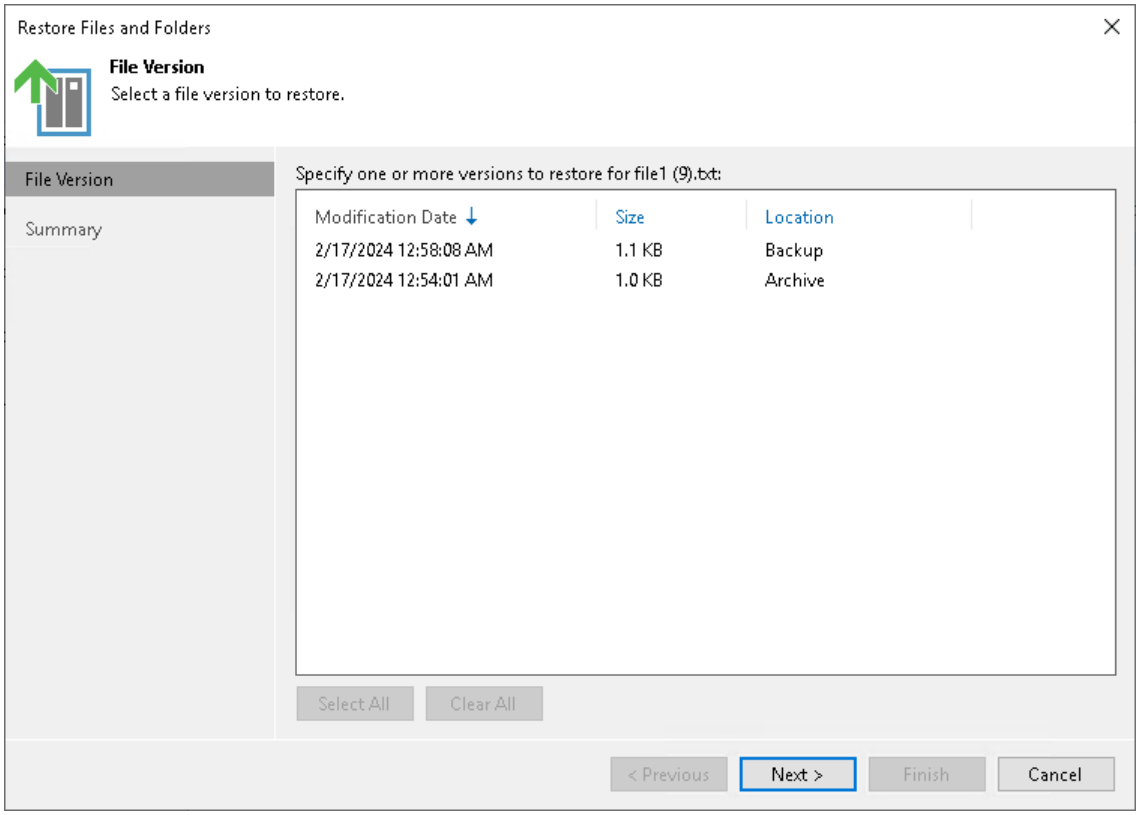

# Step 8. Specify Destination for File Restore

The **Destination** step is required if you choose the **Copy To** option at the Select Files and Folders to Restore step. Specify the destination where the restored files must be stored:

1. In the **Restore files and folders to** field, select a file share to which the files must be restored. All file shares added to the inventory of Veeam Backup & Replication are available. If the required file share is missing in the drop-down list, click **Add** and add a new file share to Veeam Backup & Replication.

   For more information on how to add a new file share, see Adding Unstructured Data Source.

2. In the **Path to folder** field, specify a path to the folder on the selected file share where files must be restored.

   To create a dedicated folder for restored files, click **Browse**. In the **Select Folder** window, select the target location for the file share.

   If you want to restore the file share to a new folder, click **New Folder** at the bottom of the window. Confirm the new folder creation.

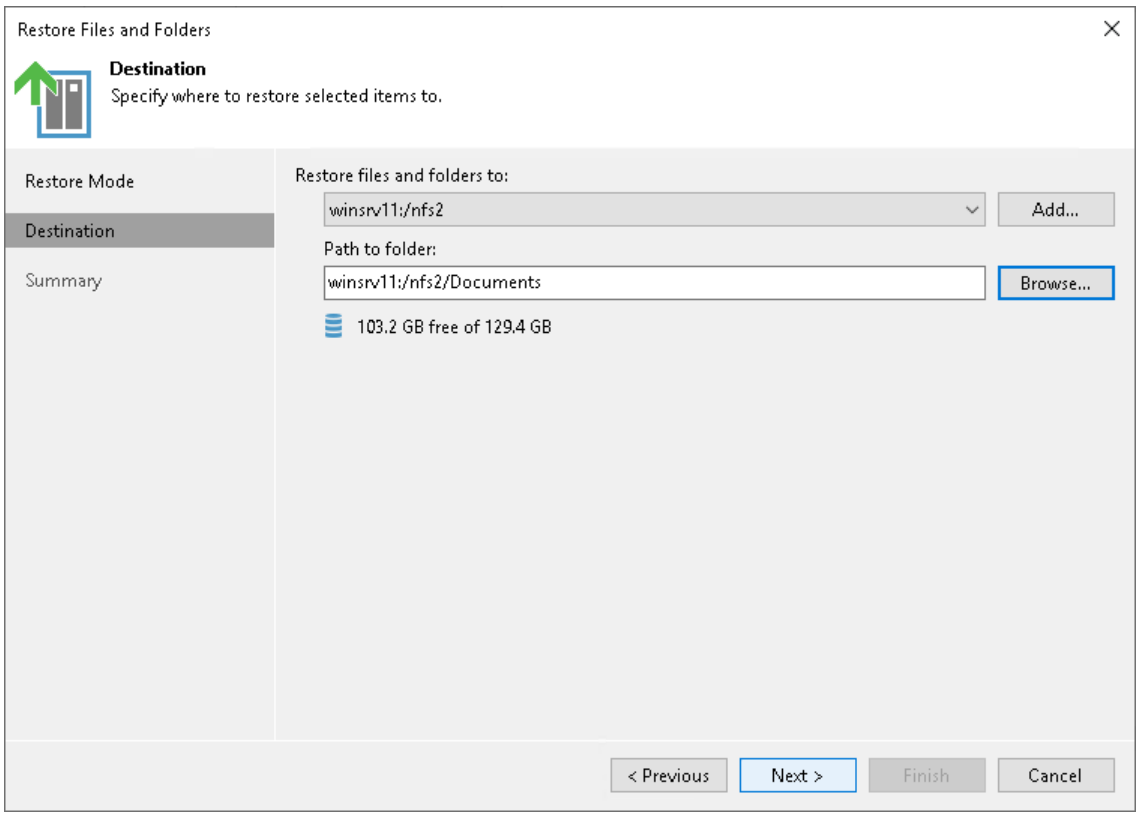

## Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review the file restore settings and click **Finish**.

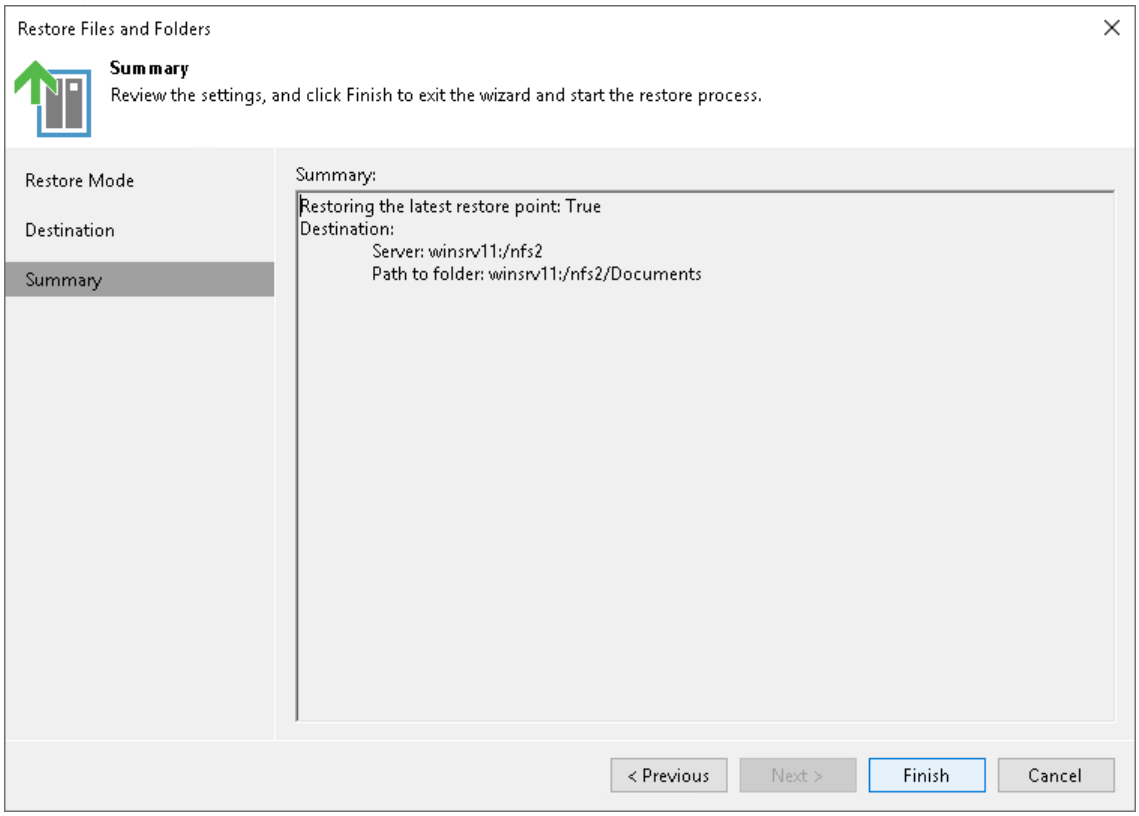


# Restoring Backup Files from Archive Repository

You can restore any file from the archive repository to the state of any file version stored in the archive. Depending on the circumstances, such a restore can require different actions.

> **NOTE**
>
> Consider that from the archive repository you can restore files only. Restore of whole folders from the long-term repository is not supported.

# Regular Restore from Archive

To perform a regular restore from the archive repository when you have all required backup data stored both in the backup repository and archive repository, follow the instructions given in Restoring Specific Files and Folders. Consider that to restore data from the archive repository, you must select the **All Time** option for selecting files and folders to restore.

# Emergency Restore from Archive

There can be different situations where backups are not available in the backup repository: for example, backup files are broken or deleted from the backup repository, backup configuration files are removed from the configuration, archive repository is added into backup infrastructure of another backup server, or backup archive is moved from one archive repository to another one. In these cases, you can restore specific files from the backup archive.

To restore files from the archive repository, for example, from an object storage:

1. If necessary, add the storage that keeps the required archive to the backup infrastructure as described in section Adding Backup Repositories.
2. Rescan the added archive repository as described in section Rescanning Backup Repositories.
3. Download metadata for the archive backup:
   a. Locate the required file backup archive under **Backups** > **Object Storage (Imported)** node in the **Home** view.
   b. Right-click the file backup and select **Download backup metadata**.

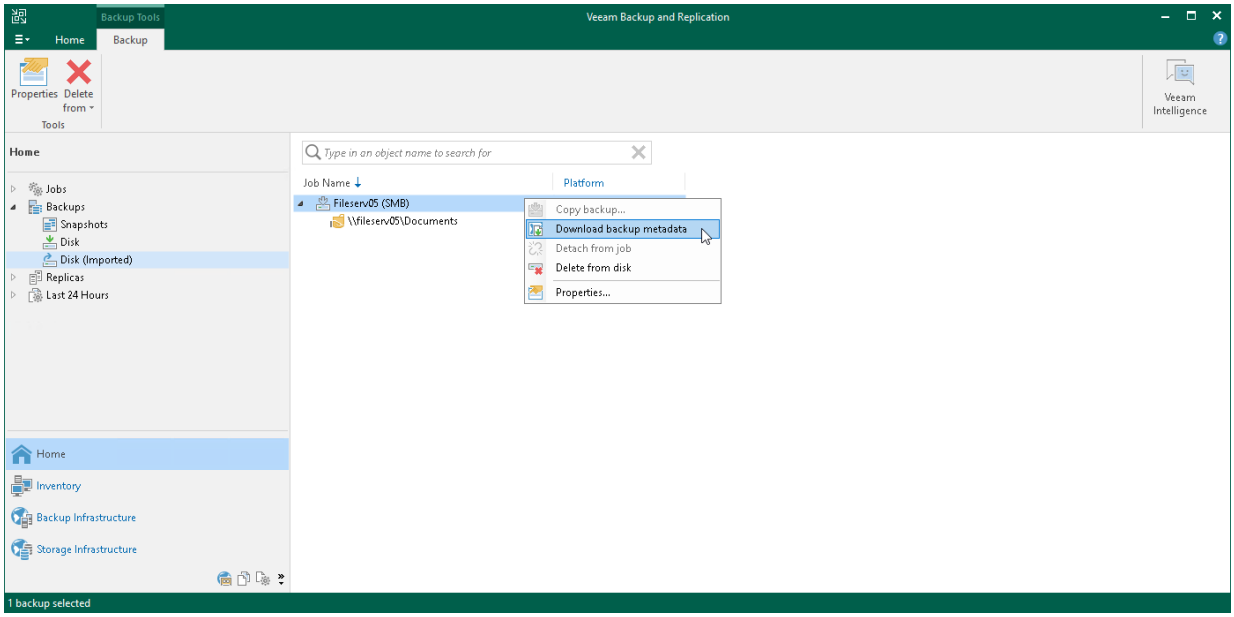


   c. From the drop-down list, select a backup repository to store the downloaded metadata of files archive and click **OK**.

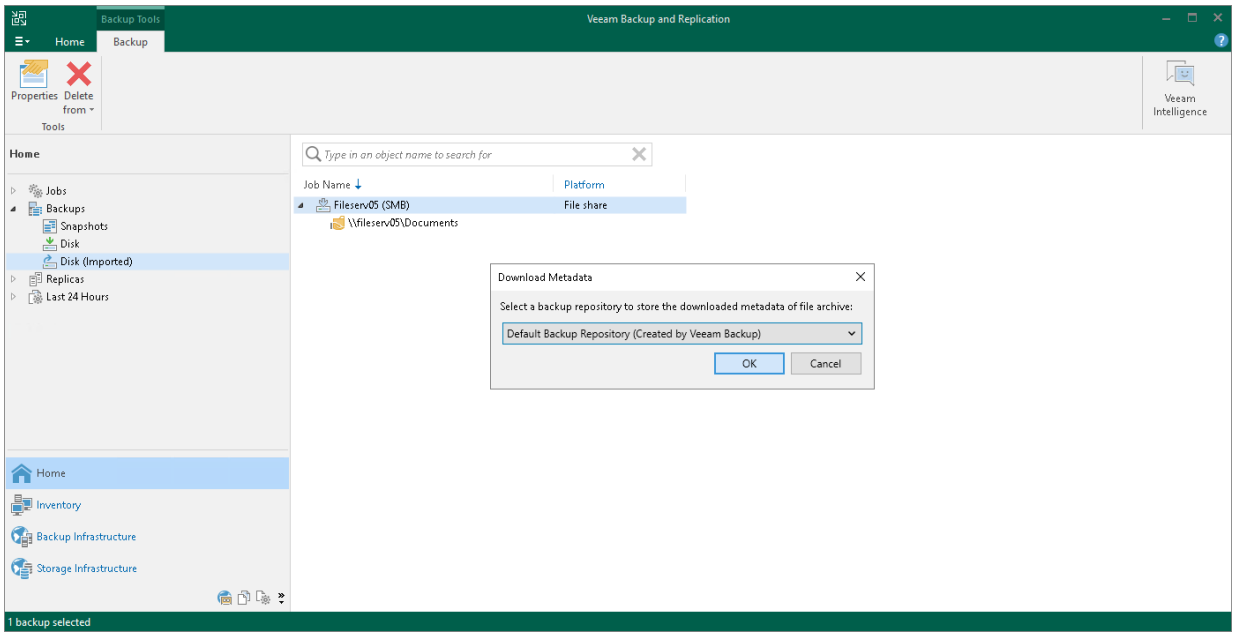


4. Restore files from the archive backup as described in section Restoring Specific Files and Folders.

# Object Storage Data Recovery

You can restore data previously backed up with object storage backup jobs. You can restore the following data:

- Whole buckets or containers of object storage repositories.
- Objects of the certain restore point.
- Multiple objects versions.

Veeam Backup & Replication offers several recovery options for different recovery scenarios:

- Restore of the entire bucket or container allows you to restore a whole object storage bucket or container to one of the restore points.
- Rollback to a point in time allows you to restore only changed objects to one of the restore points.
- Restore of individual objects allows you to select objects to restore to one of the restore points.
- Restore of objects from the archive repository allows you to select archived objects to restore to one of the restore points.

## Restoring Entire Bucket or Container

You can restore an entire object storage bucket or container from the backup to a specific restore point. That can be helpful, for example, if your object storage data gets corrupted or unavailable and you need to restore the entire bucket or container to the original or other location.

> **NOTE**
>
> If the **Archive recent object versions** option is selected (the copy mode is enabled) at the Archive Repository step of the object storage backup job wizard, you may restore an entire bucket or container from the archive repository. You can do that only for restore points that are stored in the archive repository and have the **Copied** label in backup properties.
>
> If this option is not selected (the copy mode is disabled), the restore of the entire bucket or container from the archive repository is not supported.

Before you restore a bucket or container, check prerequisites.

### Before You Begin

Before you restore the entire bucket or container, consider the following:

- You can restore the bucket or container from a backup that has at least one successfully created restore point.
- The object storage where you plan to save the restored bucket or container must be added to the backup infrastructure.
- Veeam Backup & Replication does not support restore of tags if you restore data from Microsoft Azure Blob storage to Amazon S3 object storage and S3 compatible object storage.
- Veeam Backup & Replication does not support Instant recovery of data previously backed up with object storage backup jobs.

- [For Microsoft Azure Blob storage] Veeam Backup & Replication does not support the following restore scenarios if you change the Azure Storage access tier of blocks to the *archive* access tier:
  - Restore data to the original location or to another location with the overwrite option.
  - Restore individual objects or versions with the overwrite option.
- The target location where you want to restore the bucket or container has the following limitations:
  - You can restore backups of Amazon S3 and s3 compatible only to a new bucket located either in Amazon S3 or S3 compatible object storage.
  - You can restore backups of Microsoft Azure Blob Storage only to a new container in Microsoft Azure Blob Storage.

> **TIP**
>
> If you want to restore these objects, you must change the access tier to *hot*, *cool* or *cold* access tiers. If you do not want to restore these objects, you need to remove them from your Azure container.

# Step 1. Launch Bucket Restore Wizard

To launch the **Bucket Restore** wizard, do one of the following:

- In the **Home** tab on the ribbon, click **Restore** > **Object Storage**. In the **Restore from Object Storage Backup** window, click **Restore entire buckets**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the buckets or containers that you want to restore. In the **Backup** tab on the ribbon, click **Entire Bucket**.
  - Right-click the the buckets or containers that you want to restore and select **Restore entire bucket**.

  You can restore the bucket or container to the state as of a specific restore point by using a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane. If the secondary repository is an object storage repository, backup copies created in it are represented in the **Backups** > **Object Storage (Copy)** node in the inventory pane.

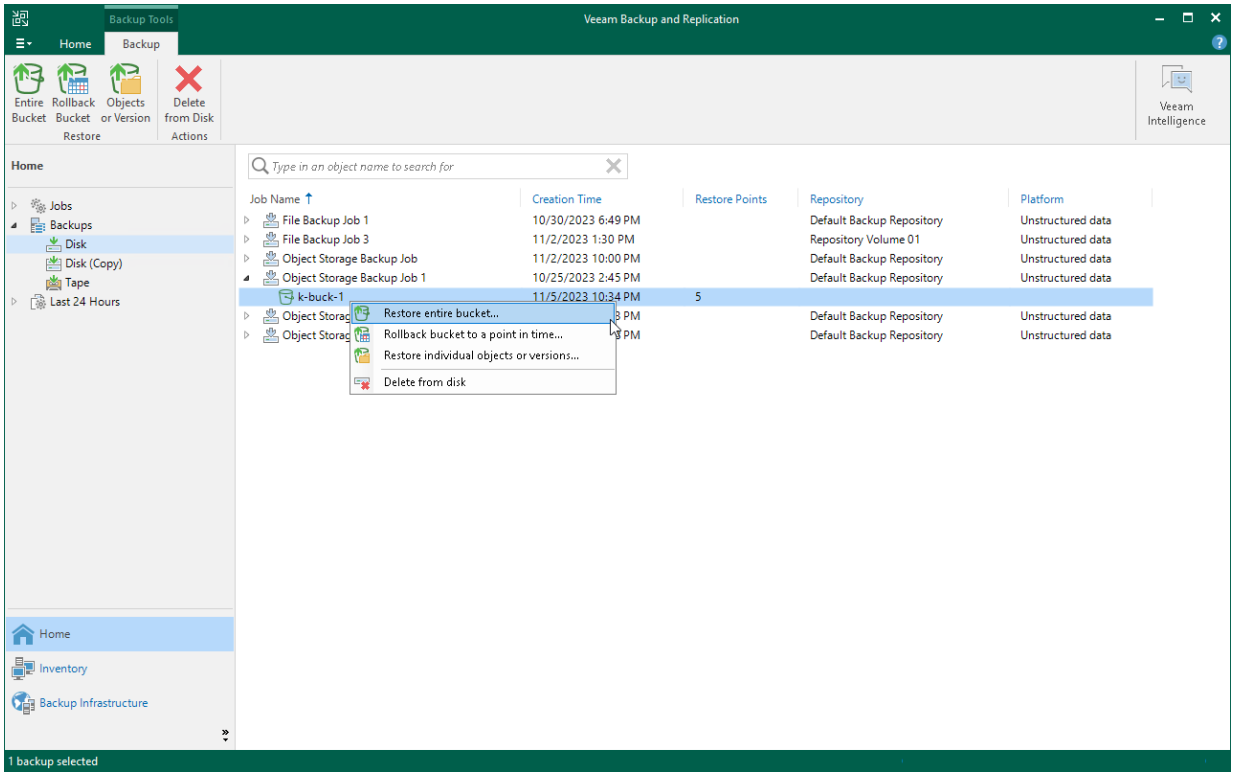

## Step 2. Select Buckets to Restore

At the **Buckets** step of the wizard, select the buckets or containers that you want to restore:

1. Click **Add**.
2. In the **Backups Browser** window, select the object storage backup job and a bucket or container in it that you want to restore. You can select multiple buckets or containers by holding [Ctrl] and clicking the required buckets or containers. Click **OK**.
3. In the **Buckets** table, choose the bucket or container to select a point to restore to. Click **Point**.
4. In the **Select Restore Point** window, select the restore point to which you want to restore the bucket or container. To select the required restore point, do one of the following:
   - Use the **Restore point** slider.
   - Click the date link under the **Restore point** slider. In the calendar in the left pane of the **Select Restore Point** window, select the date when the required restore point was created. The list of restore points in the right pane displays restore points created on the selected date. Select the point to which you want to restore the bucket or container.

   In the **Files in backup** tree, you can see what prefixes and objects are covered by the selected restore point and the date when each of them was modified.

   Click **OK**.

To quickly find a bucket or container, you can use the search field at the top of the window. Enter a bucket or container name or a part of it in the search field and press [Enter].

To exclude the bucket or container from the restore process, select the bucket or container in the table and click **Remove**.

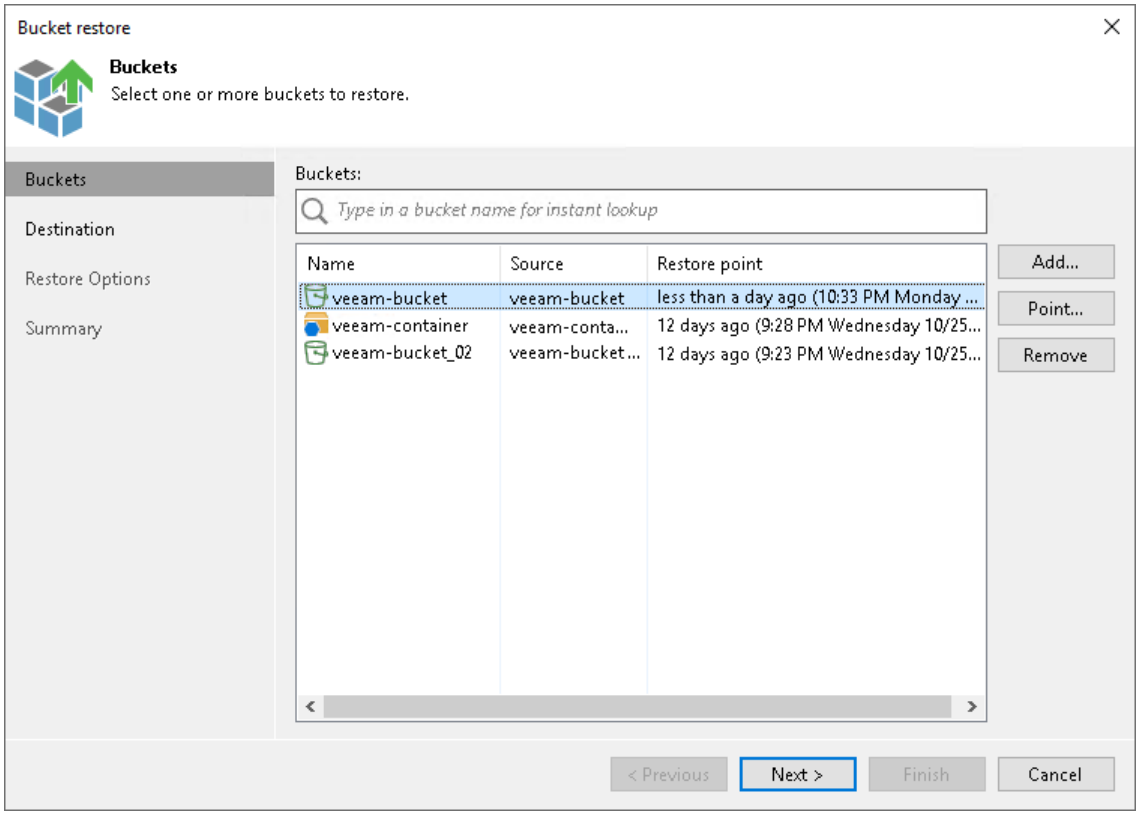

## Step 3. Specify Destination for Bucket or Container Restore

At the **Destination** step of the wizard, specify the location where you want to restore the bucket or container.

- Select **Original location** to restore data to the location where it resided originally. This type of restore is only possible if the original device is connected to Veeam Backup & Replication and powered on.
- Select **New bucket** to create a new bucket or container and restore the backed-up bucket or container content to it. You can restore only a single bucket or container to a new bucket or container.

  > **NOTE**
  >
  > The target location where you want to restore the bucket or container has the following limitations:
  >
  > - You can restore backups of Amazon S3 and s3 compatible only to a new bucket located either in Amazon S3 or S3 compatible object storage.
  > - You can restore backups of Microsoft Azure Blob Storage only to a new container in Microsoft Azure Blob Storage.

- Select **This location** to restore data to another location:

  a. In the **This location** field, select a location to restore objects to. You can select any object storage added to the backup inventory. If the required location is missing in the drop-down list, click **Add** and add a new location to Veeam Backup & Replication. For more information, see the Adding Unstructured Data Source section.

  b. In the **Path** field, specify a path to the prefix in the selected location to restore objects to.

  To select a specific prefix, click **Browse**. In the **Select Folder** window, select the target location for the bucket or container content.

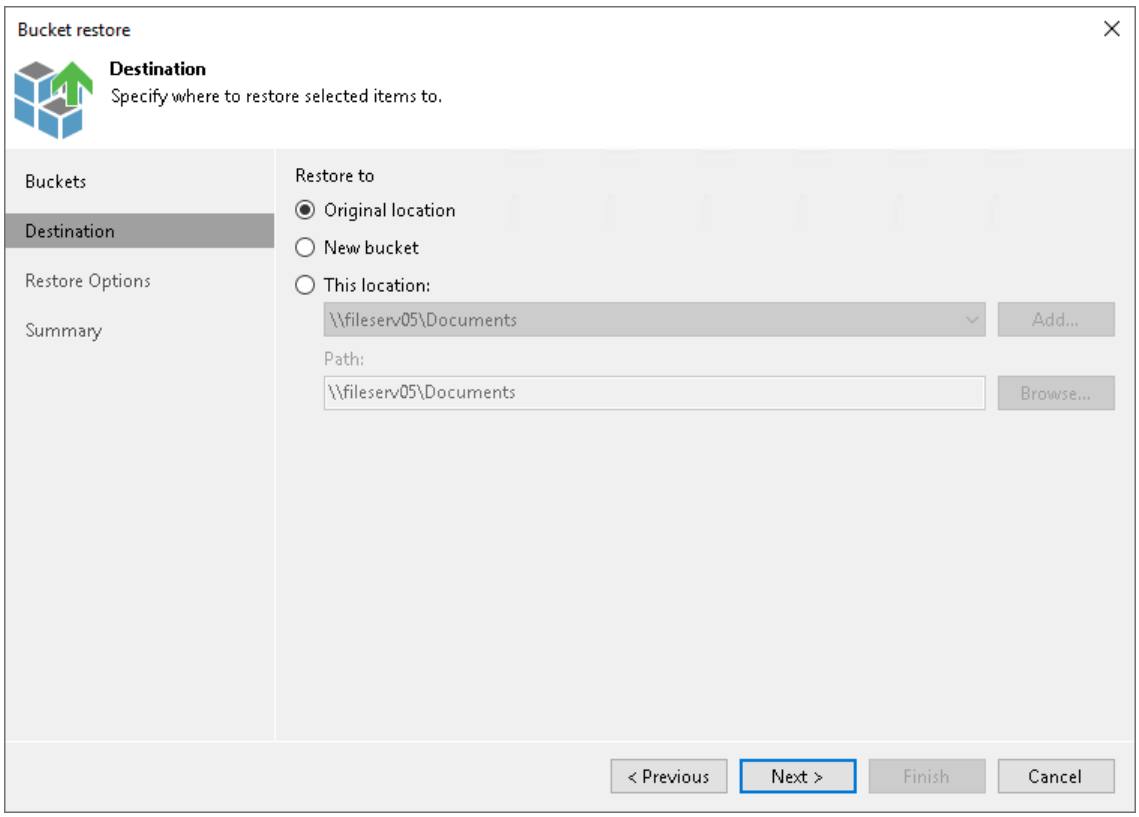

# Step 4. Create New Bucket or Container

The **New Bucket** step is available if you select the **New bucket** option at the Destination step.

At the **New Bucket** step of the wizard, create a new bucket or container where you want to restore the bucket or container.

1. From the **Object storage** drop-down list, select an object storage where you want to create a new bucket or container. If the required object storage is missing in the drop-down list, click **Add** and add a new object storage to the Veeam Backup & Replication inventory. For more information, see the Adding Object Storage section.

   > **NOTE**
   >
   > The target bucket or container where you want to restore the data has the following limitations:
   >
   > - You can restore backups of Amazon S3 and s3 compatible only to a new bucket located either in Amazon S3 or S3 compatible object storage.
   > - You can restore backups of Microsoft Azure Blob Storage only to a new container in Microsoft Azure Blob Storage.

2. In the **Bucket name**, specify the name for the bucket or container following bucket or container naming rules for the selected object storage.

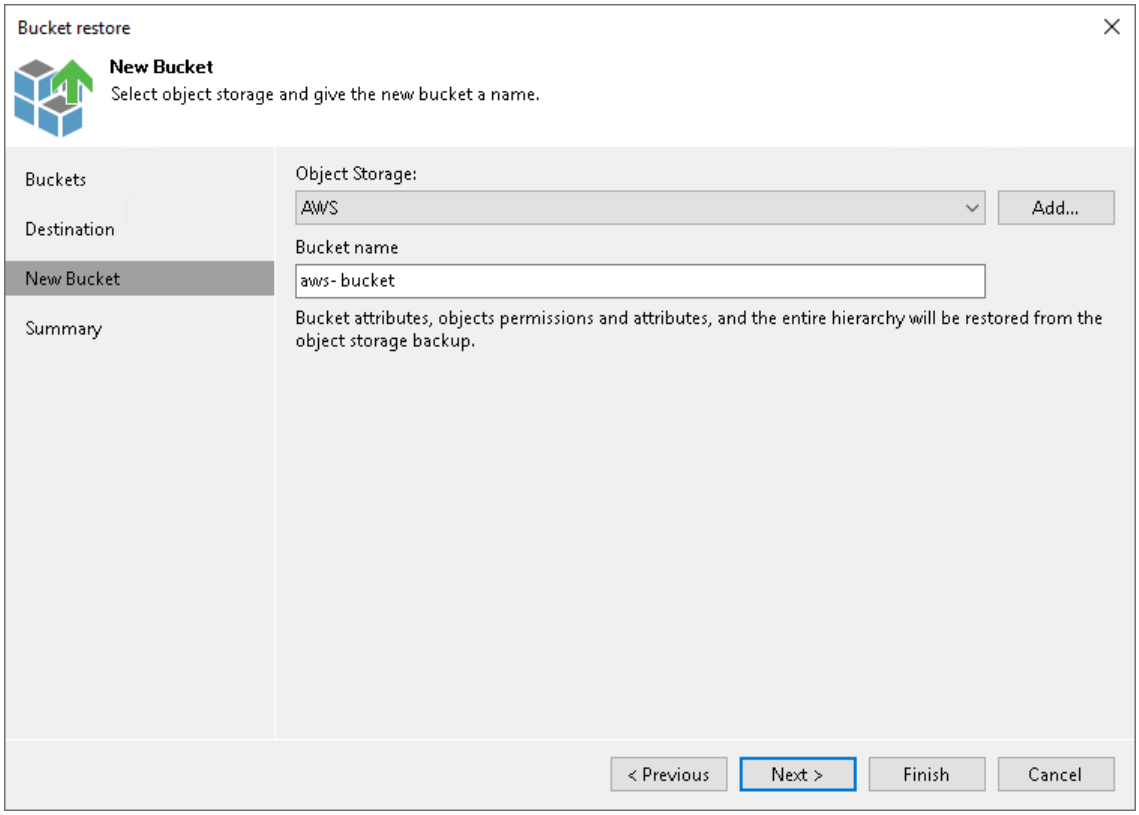

# Step 5. Specify Restore Options

The **Restore Options** step is available if you select either the **Original location** option or the **This location** option at the Destination step.

At the **Restore Options** step of the wizard, specify overwrite options in case the bucket or container with the same name already exists in the target location:

- **Skip restoring (keeps the existing object)**. Select this option if you do not want to overwrite the existing objects with the restored objects.
- **Replace older objects only (use if a bucket was reverted to a snapshot)**. Select this option if you want to overwrite the existing objects only if they are older than the restored objects.
- **Replace newer objects only (use to discard unwanted contents changes)**. Select this option if you want to overwrite the existing objects only if the restored objects are older than the source objects.
- **Restore anyway (overwrites existing objects)**. Select this option if you want to overwrite the existing objects with the restored objects in all cases.

Select the **Overwrite current bucket attributes with attributes from the backup** check box if you want the target bucket or container to inherit attributes from the restored bucket or container.

> **NOTE**
>
> Note that you can select this check box only when restoring to original location.

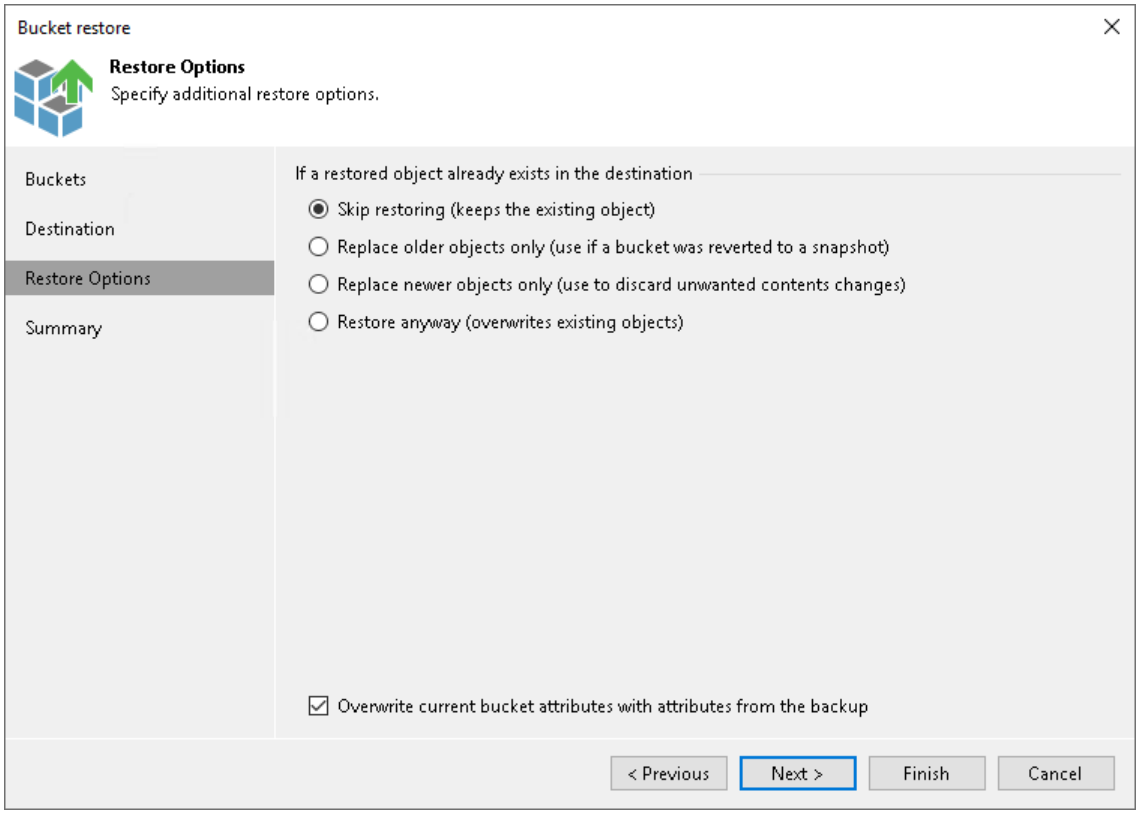

## Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review the bucket or container restore settings and click **Finish**. Veeam Backup & Replication will restore the bucket or container to the specified location.

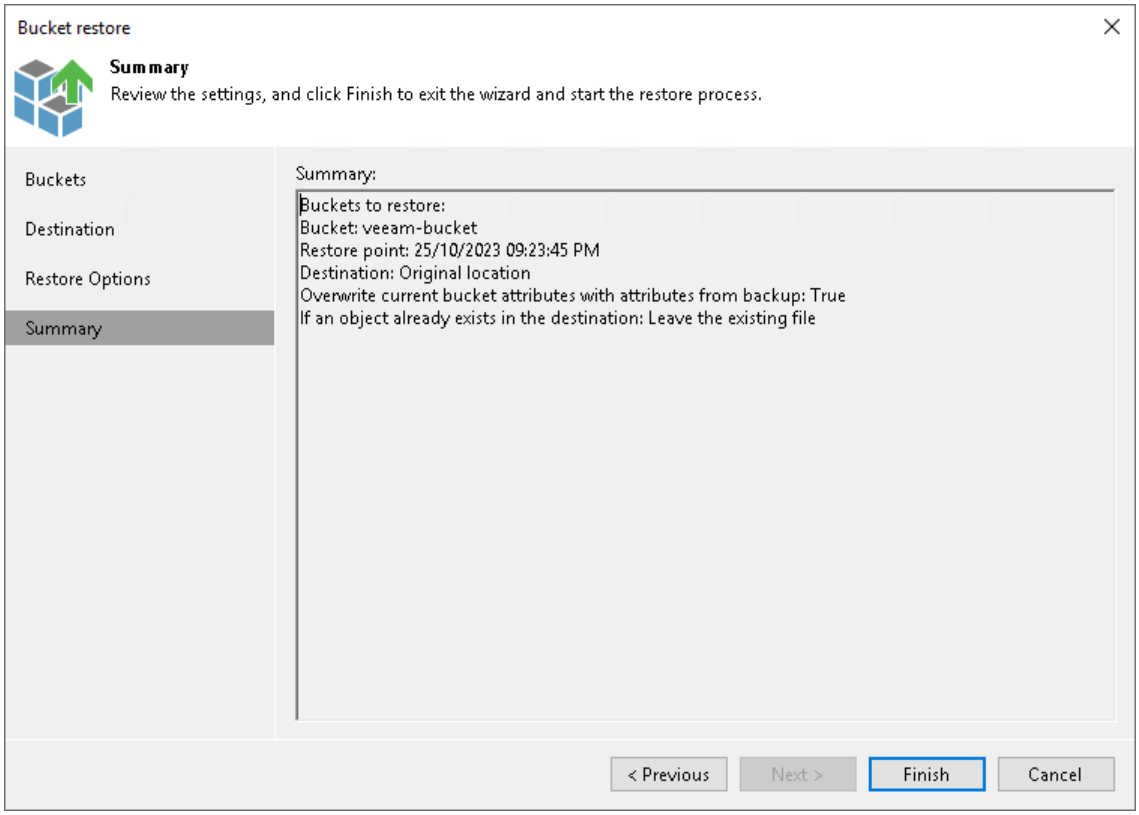


# Rolling Back to a Point in Time

You can roll back changes made to data in object storage buckets or containers to a specific restore point that is older than the current buckets or container state. This option can be useful, for example, when the original bucket or container was attacked by ransomware. In this case you can roll back all the data that was changed by the ransomware to the state before the attack.

Before you roll back the buckets or containers to a point in time, check prerequisites.

## Before You Begin

Before you restore buckets or containers to a point in time, consider that you can restore buckets or containers from a backup that has at least one successfully created restore point.

# Step 1. Launch Bucket Rollback to a Point in Time Wizard

To launch the **Bucket Rollback to a Point in Time** wizard, do one of the following:

- In the **Home** tab on the ribbon, click **Restore** > **Object Storage**. In the **Restore from Object Storage Backup** window, click **Rollback buckets to a point in time**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the bucket or container that you want to restore. In the **Backup** tab on the ribbon, click **Rollback Bucket**.
  - Right-click the object storage that you want to restore and select **Rollback bucket to a point in time**.

  You can roll back the bucket or container to the state as of a specific restore point by using a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane. If the secondary repository is an object storage repository, backup copies created in it are represented in the **Backups** > **Object Storage (Copy)** node in the inventory pane.

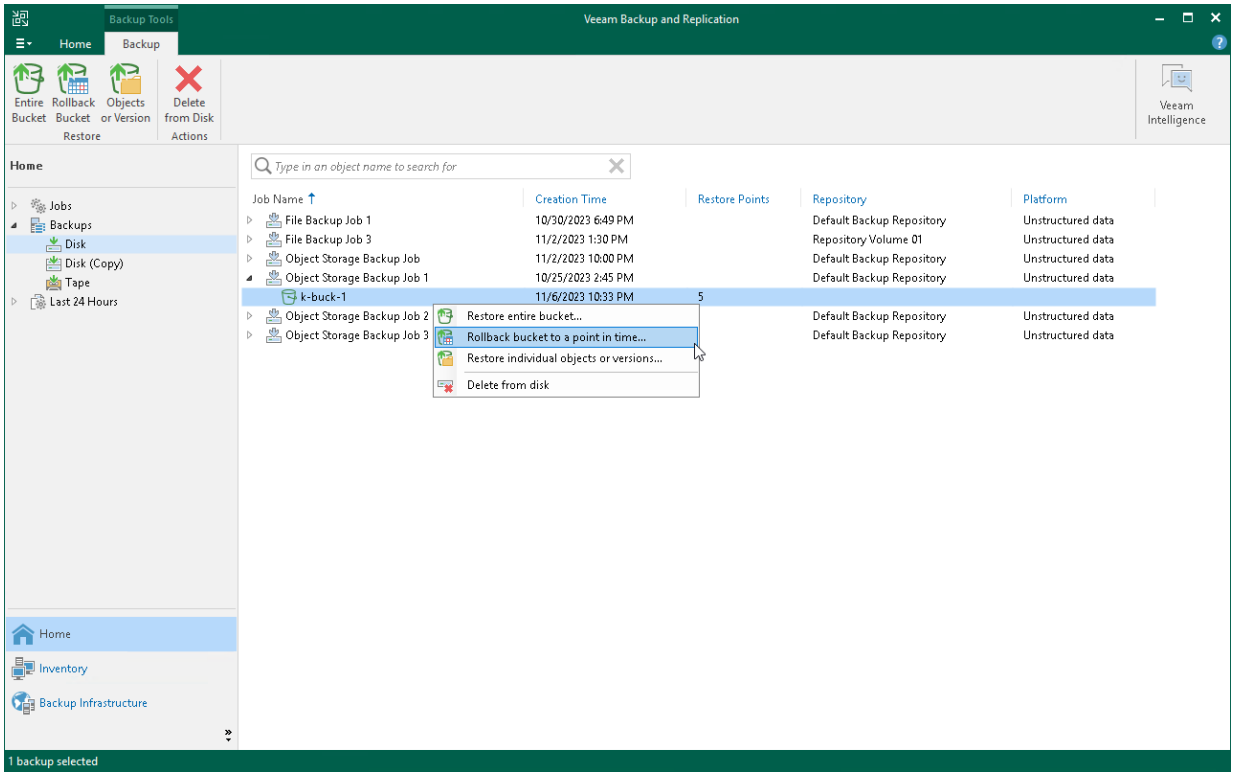

## Step 2. Select Buckets to Restore

At the **Buckets** step of the wizard, select the buckets or containers that you want to roll back:

1. Click **Add**.
2. In the **Backups Browser** window, select the object storage backup job and a bucket or container in it that you want to restore. You can select multiple buckets or containers by holding [Ctrl] and clicking the required buckets or containers. Click **OK**.
3. In the **File shares** table, select the bucket or container to select a restore point to rollback to. Click **Point**.
4. In the **Select Restore Point** window, choose the restore point to roll back the bucket or container to. To select the required restore point, do one of the following:
   - Use the **Restore point** slider.
   - Click the date link under the **Restore point** slider. In the calendar in the left pane of the **Restore points** window, select the date when the required restore point was created. The list of restore points on the right pane displays restore points created on the selected date. Select the point to which you want to roll back the bucket or container.

   In the **Files in backup** tree, you can see what prefixes and objects are covered by the selected restore point and the date when each of them was modified.

   Click **OK**.

To quickly find a bucket or container, you can use the search field at the top of the window. Enter a bucket or container name or a part of it in the search field and press [Enter].

To exclude the bucket or container from the restore process, select the bucket or container in the table and click **Remove**.

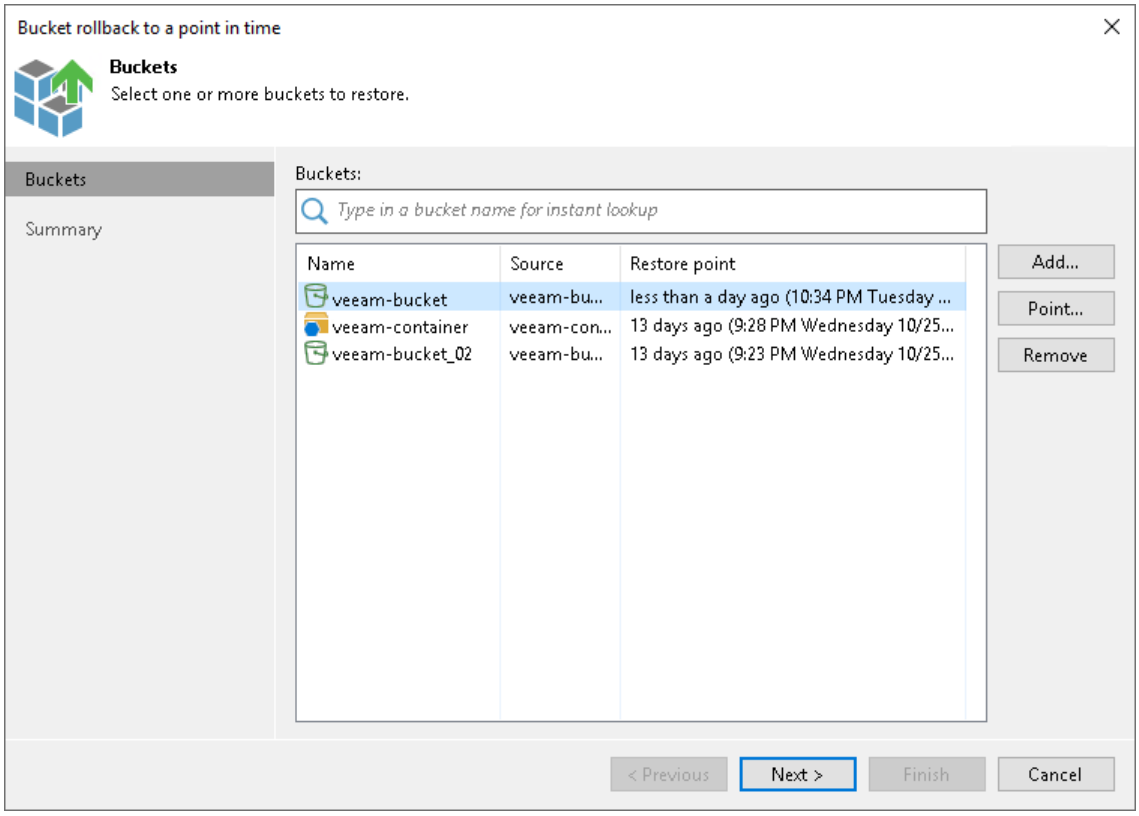

## Step 3. Finish Working with Wizard

At the **Summary** step of the wizard, review the bucket or container rollback settings and click **Finish**. Veeam Backup & Replication will roll back the bucket to the restore point.

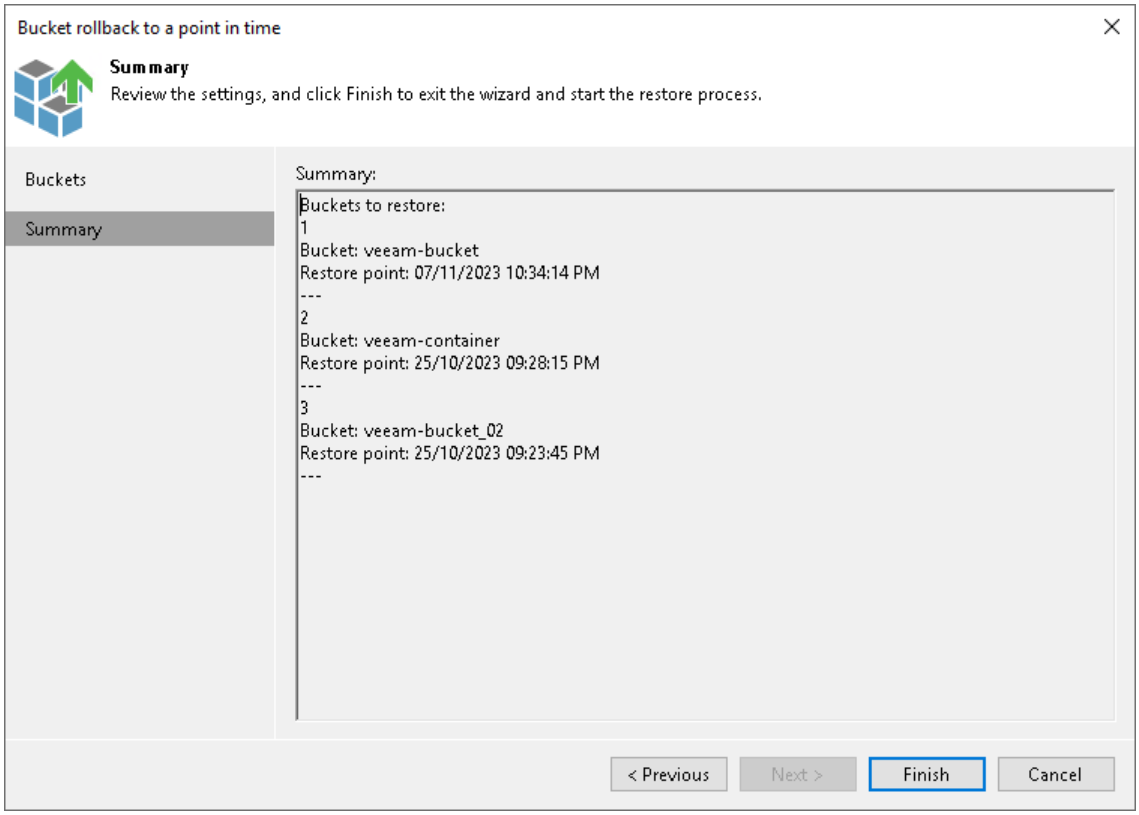


# Restoring Individual Objects or Versions

You can restore individual objects from the bucket or container backups to the original or a new location. This option can be useful, for example, if you need to get an older version of some objects from the backup.

When you restore specific objects, you can extract object versions not only from the backup repository, but also from the archive repository. For more information, see the Restoring Objects from Archive Repository section.

> **NOTE**
>
> Consider that from the archive repository you can restore objects only. Restore of prefixes from the long-term repository is not supported.

Besides, you can restore multiple versions of the same object.

Before you restore specific objects, check prerequisites. Then use the **Objects or Versions Restore** wizard to restore objects or their specific versions.

1. Launch the Objects or Versions Restore wizard.
2. Select an object to restore.
3. Verify object restore settings.
4. Select objects to restore.
5. Select a restore mode.

6. Select a restore point.
7. Select an object version to restore.
8. Specify the destination for object restore.
9. Finish working with the wizard.

## Before You Begin

Before you restore objects from the backup, consider the following:

- You can restore objects from a backup that has at least one created restore point, even if it is incomplete.
- The object storage on which you plan to save restored objects must be added to the backup infrastructure.

# Step 1. Launch Object Restore Wizard

To launch the **Objects or Versions Restore** wizard, do one of the following:

- In the **Home** tab on the ribbon, click **Restore** > **Object Storage**. In the **Restore from Object Storage Backup** window, click **Restore individual objects or their specific versions**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Click the bucket or container that you want to restore. In the **Backup** tab on the ribbon, click **Objects or Versions**.
  - Right-click the object storage that you want to restore and select **Restore individual objects or versions**.

  You can restore objects from a backup copy. Backup copies created in the secondary repositories are represented in the **Backups** > **Disk (Copy)** node in the inventory pane. If the secondary repository is an object storage repository, backup copies created in it are represented in the **Backups** > **Object Storage (Copy)** node in the inventory pane.

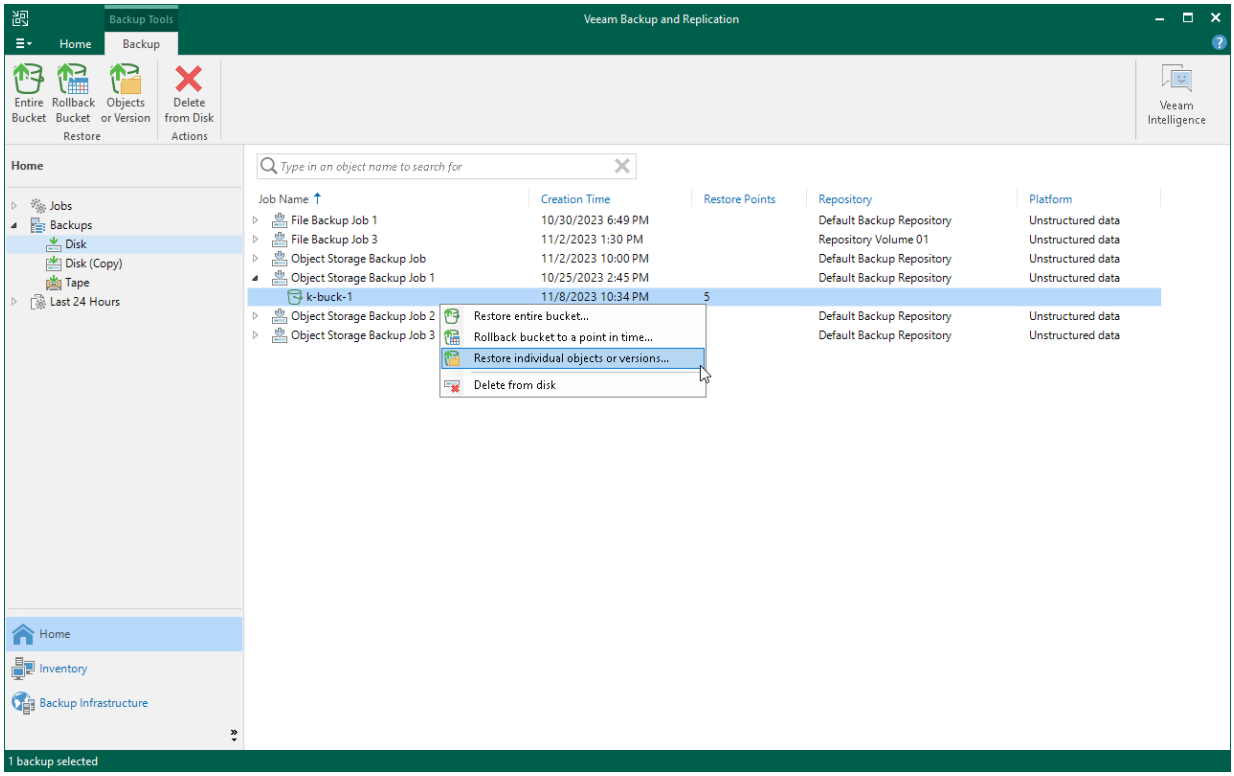

## Step 2. Select Object to Restore

At the **Backup** step of the wizard, select the bucket or container you want to restore:

1. In the **Object to restore** list, expand the necessary backup job.
2. Select the bucket or container.

To quickly find a bucket or container, you can use the search field at the bottom of the window.

1. Enter a bucket or container name or a part of it in the search field.
2. Click the **Start search** button on the right or press [Enter].

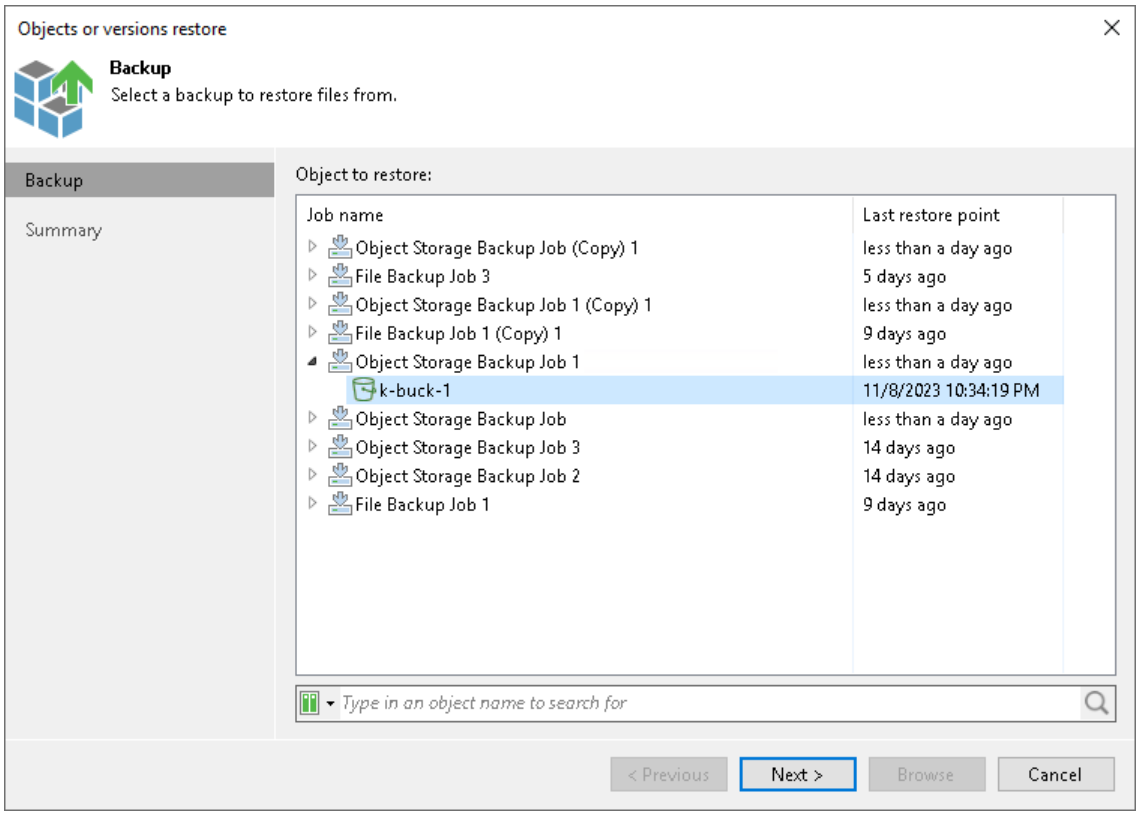

## Step 3. Verify Object Restore Settings

At the **Summary** step of the wizard, review the selected bucket or container and click **Browse** to switch to the Backup Browser step and select objects to restore.

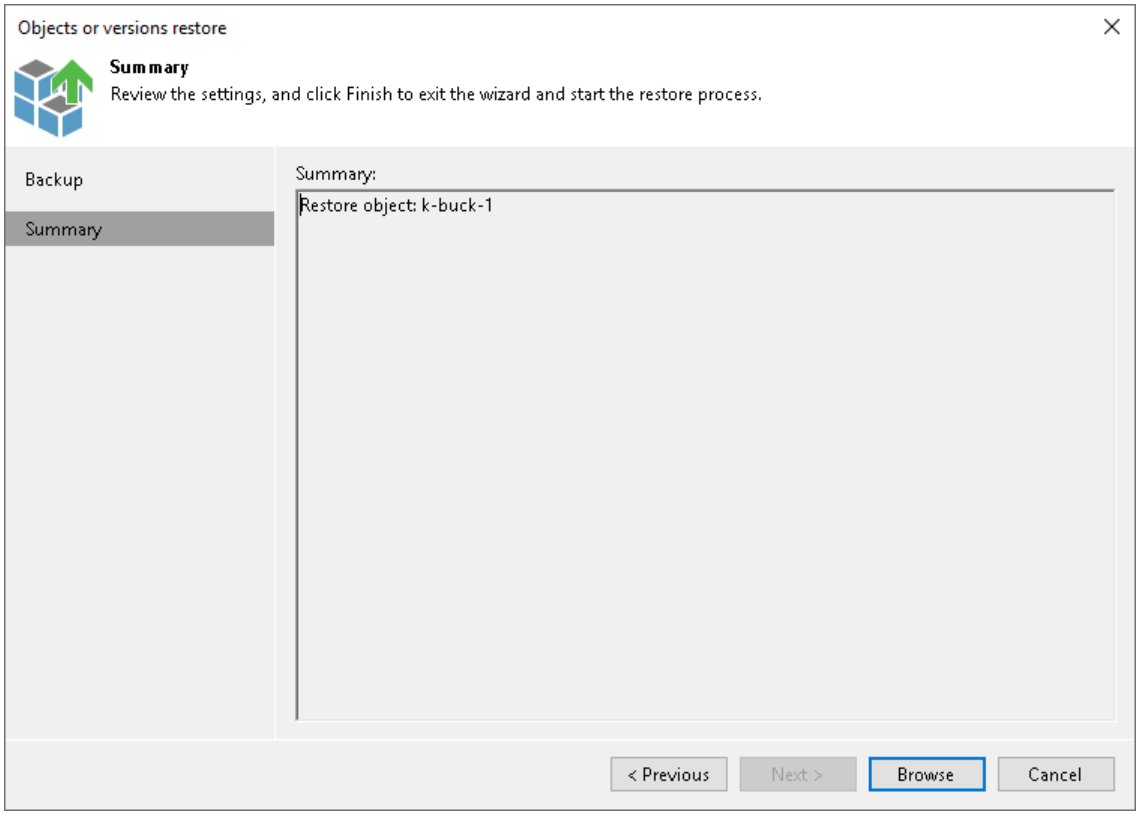

## Step 4. Select Objects to Restore

In the **Backup Browser**, select objects and prefixes to restore. Backup browser has three representations of restore points. You can select one of the following options in the **Restore point** group on the ribbon to display object versions:

- **Latest** – the **Backup Browser** shows the latest versions of objects and prefixes in the bucket or container.
- **All Time** – the **Backup Browser** shows all objects and prefixes ever backed up by the backup job. This option retrieves object versions stored both in the backup and archive repositories. This representation additionally shows how many object versions of each object are stored in the backup and the date when the latest object version was created.

  After you select this option:

  - If you restore a whole prefix, you will be prompted to the Select Restore Mode step to select a restore mode and restore point to restore objects from.

    > **NOTE**
    >
    > Consider that during object-level restore you cannot restore whole prefixes from the archive restore points. If objects and prefixes are already moved to the archive repository, you have to restore only objects one by one.

  - If you restore a single object, you will be prompted to the Select Object Version to Restore step to select a object version to restore the object from.
- **Selected** – the **Backup Browser** shows versions of objects and prefixes backed up as of the certain restore point. Select the restore point in the list on the right of the **Selected** option on the ribbon.

You can use the search field at the top of the working area to search for specific objects and prefixes.

> **NOTE**
>
> To keep the operation of the Backup Server stable, the number of retrieved search result records is limited to 1000. Therefore, if you work with backup prefixes that store large volumes of data, it is recommended to narrow the search criteria to fit into the limitation.

You can restore objects and prefixes to their original location or a new location.

# Restoring Objects and Prefixes to Original Location

To restore objects and prefixes to the original location, in the Veeam Backup browser right-click an object or prefix and select one of the following commands:

- To overwrite the original object in the object storage with the object restored from the backup, select **Restore** > **Overwrite**.
- To save the object restored from the backup next to the original object, select **Restore** > **Keep**.

  Veeam Backup & Replication will add the *_RESTORED_YYYYMMDD_HHMMSS* suffix to the original object name and store the restored object in the same prefix where the original object resides.

Alternatively, you may use similar options under **Restore** in the ribbon.

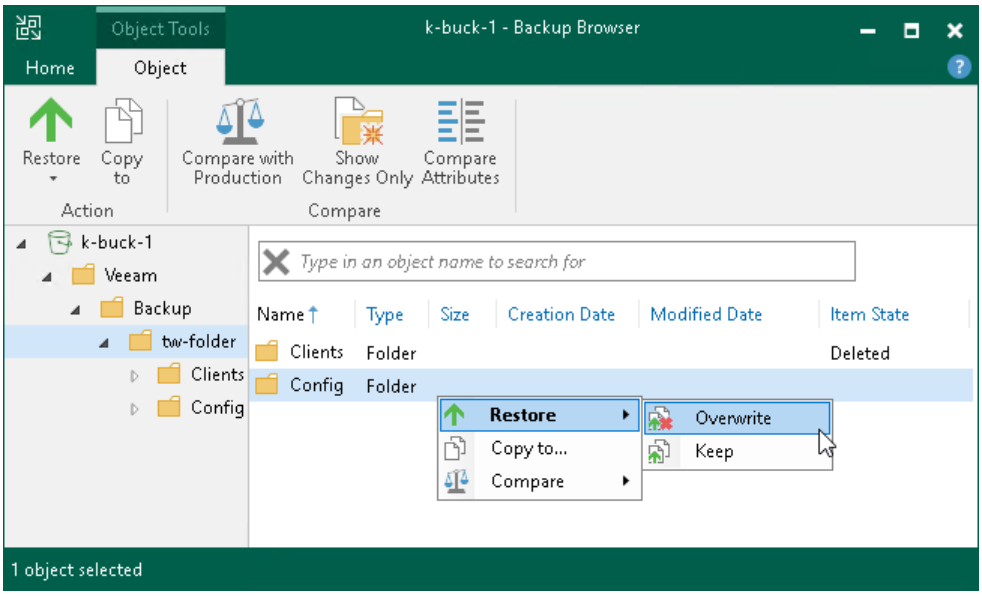


# Saving Objects and Prefixes to New Location

To save objects and prefixes to a new location, right-click the necessary object or prefix and select **Copy to**. Alternatively, you may use the **Copy to** option in the ribbon.

If you restore objects and prefixes to a new location, you will be prompted to the Specify Destination for Object Restore step to specify a new destination.

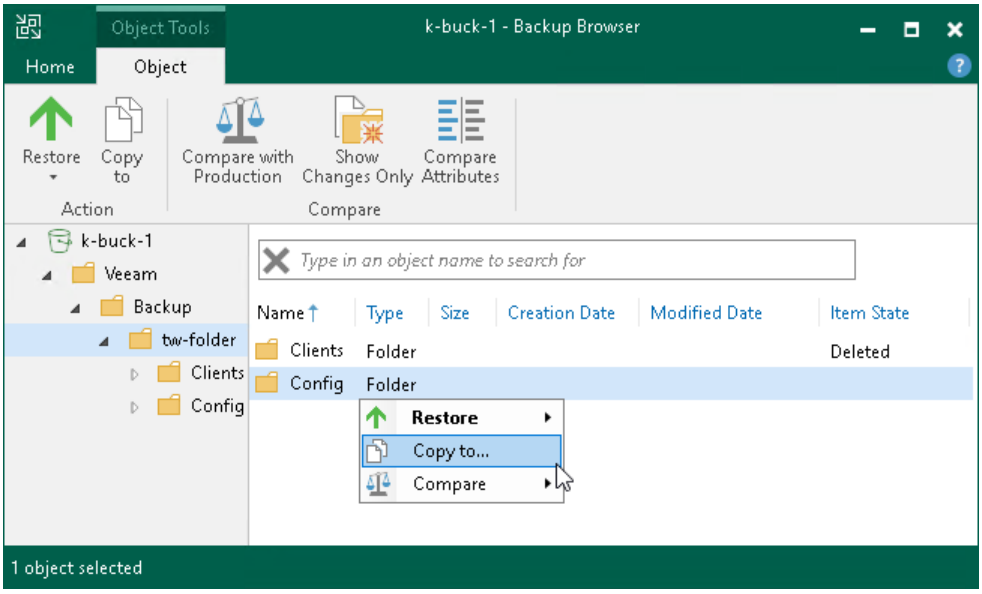


# Comparing Backup Object and Prefix Versions with Production Objects and Prefixes

> **IMPORTANT**
>
> Running the operations of the compare with production feature for backups of object storage requires sending `HEAD` and `GET` requests to the storage. That may lead to additional costs.

> **NOTE**
>
> The compare with production feature does not work if you select the **All Time** option at the **Home** tab of the Backup Browser.

You can compare backup versions of specific objects and prefixes with their production sources. To compare them, do the following:

1. Select an object or prefix in the items tree in the left pane or object in the right pane. You can use [Ctrl] to select multiple objects and prefixes in the right pane.
2. Click **Compare with Production** on the ribbon. Alternatively, you can right-click the necessary object or prefix and select **Compare** > **Compare**.

After the comparison, objects will have the following comparison states in the **Item State** column: *changed*, *unchanged*, *deleted*, *comparing* or *failed to compare*. The states are updated when you turn off and then turn on the comparison mode, and when you start restoring changes of objects. Note that when comparing symbolic links, Veeam Backup & Replication compares attributes of the links, not the attributes of objects which the symbolic link points to.

> **TIP**
>
> Consider the following:
>
> - To show only changed objects (in the *changed* and *deleted* states), perform the compare operation, right-click any area in the Veeam Backup browser and select **Compare** > **Show changes only** or click **Show Changes Only** on the ribbon.
> - To switch off the comparison states, select an object in the comparison state and click **Compare** > **Compare** or click **Compare with Production** on the ribbon. Note that if you switch off comparison for child objects, comparison for parent objects will also be switched off.

# Restoring Changed Objects and Prefixes

After you filter the changed objects and prefixes, you can restore them.

To restore changed and deleted objects and prefixes to the original location, do the following:

1. In the Veeam Backup browser right-click an object or prefix with **Changed** or **Deleted** item state.
2. Select one of the following commands:
   - To overwrite the original object or prefix on the object storage with the object restored from the backup, select **Restore changed and deleted** > **Overwrite**.
   - To save the object or prefix restored from the backup next to the original object, select **Restore changed and deleted** > **Keep**.

     Veeam Backup & Replication will add the *_RESTORED_YYYYMMDD_HHMMSS* suffix to the original object and prefix name and store the restored object and prefix where the original object and prefix resides.

Alternatively, you may use similar options under the **Restore changes** group in the ribbon.

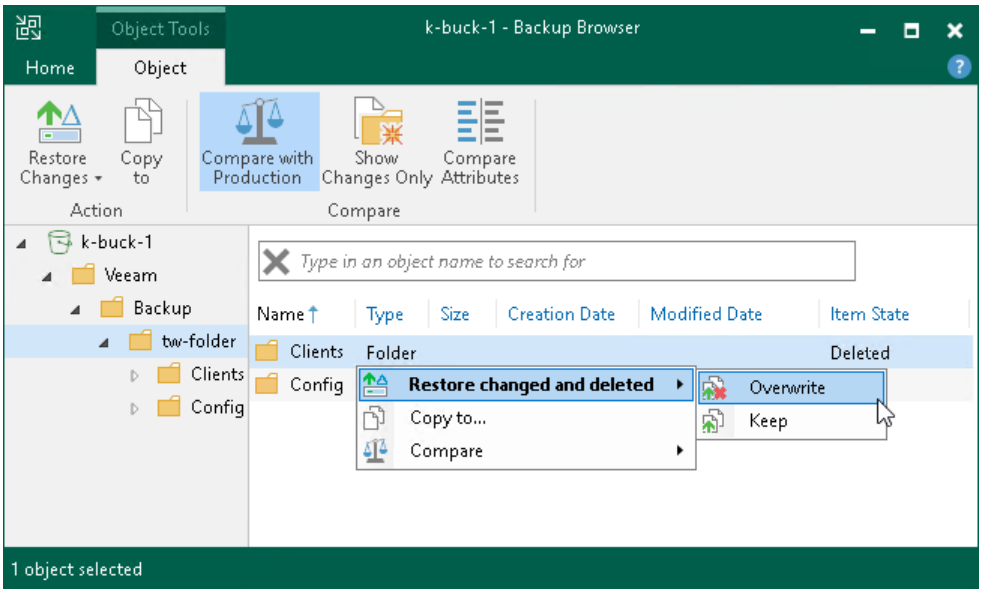


You can also compare attributes of the changed and deleted objects and prefixes against objects and prefixes in the source object storage. To compare them, do the following:

1. Select a prefix in the items tree in the left pane or an object and the prefix in the right pane. You can use [Ctrl] to select multiple objects and prefixes in the right pane.
2. Right-click one of the selected items and select **Compare** > **Compare attributes** or click **Compare Attributes** on the ribbon. Alternatively, you can right-click one of the selected items and select **Compare** > **Compare** or click **Compare with Production** on the ribbon.

In the **Compare Attributes** window, Veeam Backup & Replication shows changed attributes. If you want to show all attributes, click the **Show unchanged attributes** check box at the top right corner. Note that Veeam Backup & Replication shows attributes maximum for 500 objects and prefixes and shows attributes for the selected objects, not for the nested objects.

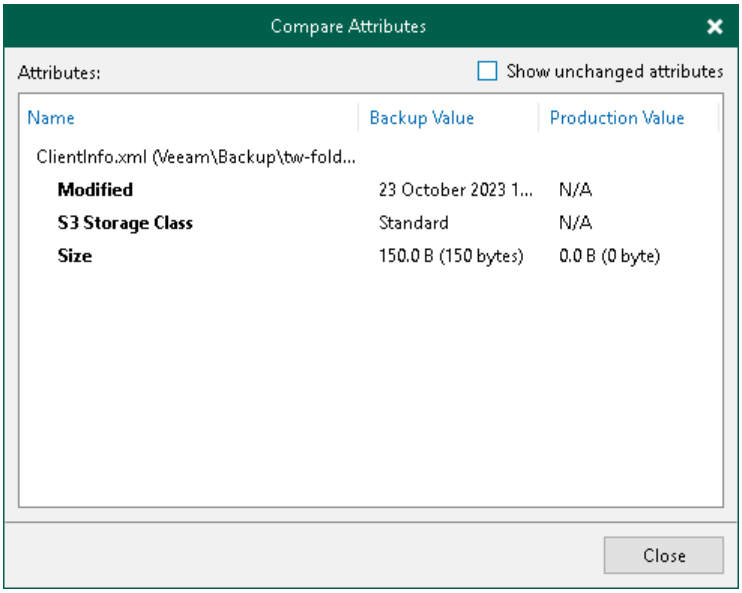

# Step 5. Select Restore Mode

The **Restore Mode** step is required if you use the **All Time** option at the Select Objects to Restore step and the selected prefixes have more than one restore point.

Choose to what point you want to restore prefixes:

- To restore the prefixes to the latest available restore point, select **Latest state**.
- To select a specific restore point, select **Earlier restore point**.

  Choosing this option will open the Restore Point step.

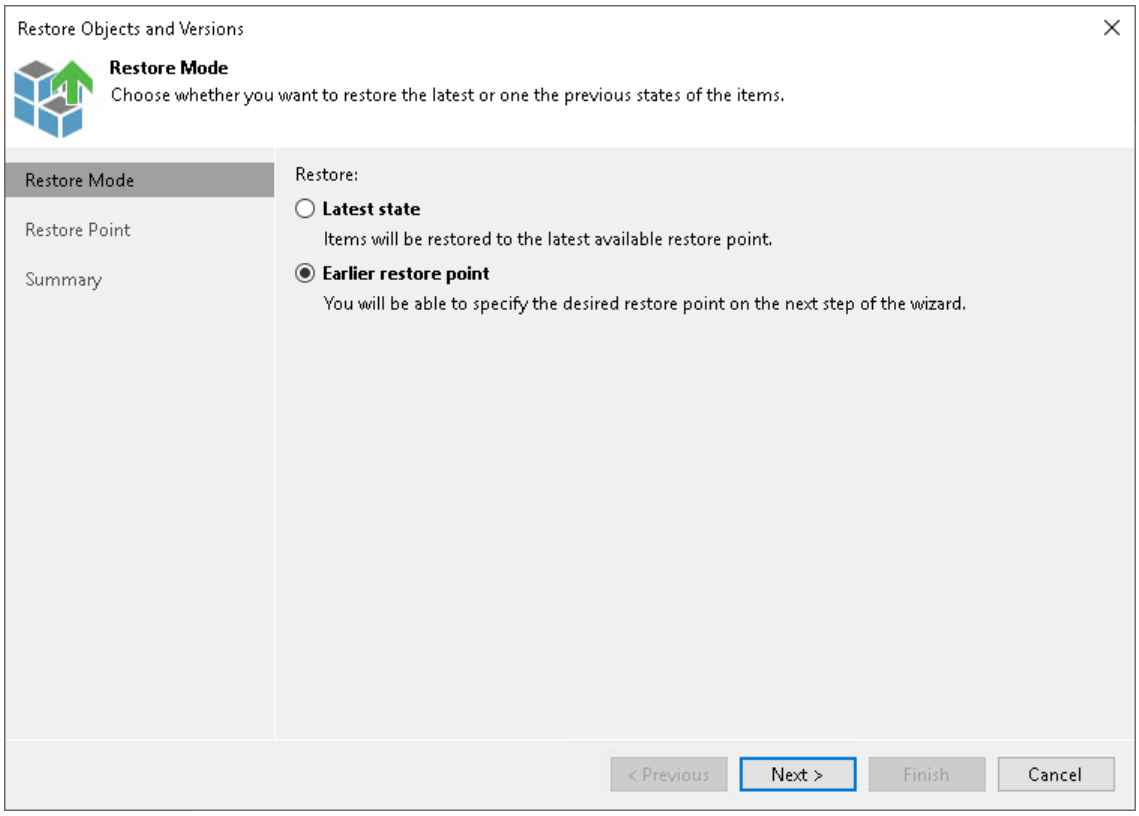

# Step 6. Select Restore Point

The **Restore Point** step is required if you use the **All Time** option at the Select Objects to Restore step, the selected prefixes have more than one restore point, and you select the **Earlier restore point** option at the Restore Mode step.

At the **Restore Point** step of the wizard, select the point in time to restore prefixes to. To select the required restore point, do one of the following:

- Use the **Restore point** slider.
- Click the date link under the **Restore point** slider. In the calendar in the left pane of the **Restore points** window, select the date when the required restore point was created. The list of restore points in the right pane displays restore points created on the selected date. Select the point to which you want to restore the objects and prefixes.

In the **Files in backup** tree, you can see what objects and prefixes are covered by the selected restore point and the date when object and prefixes were modified.

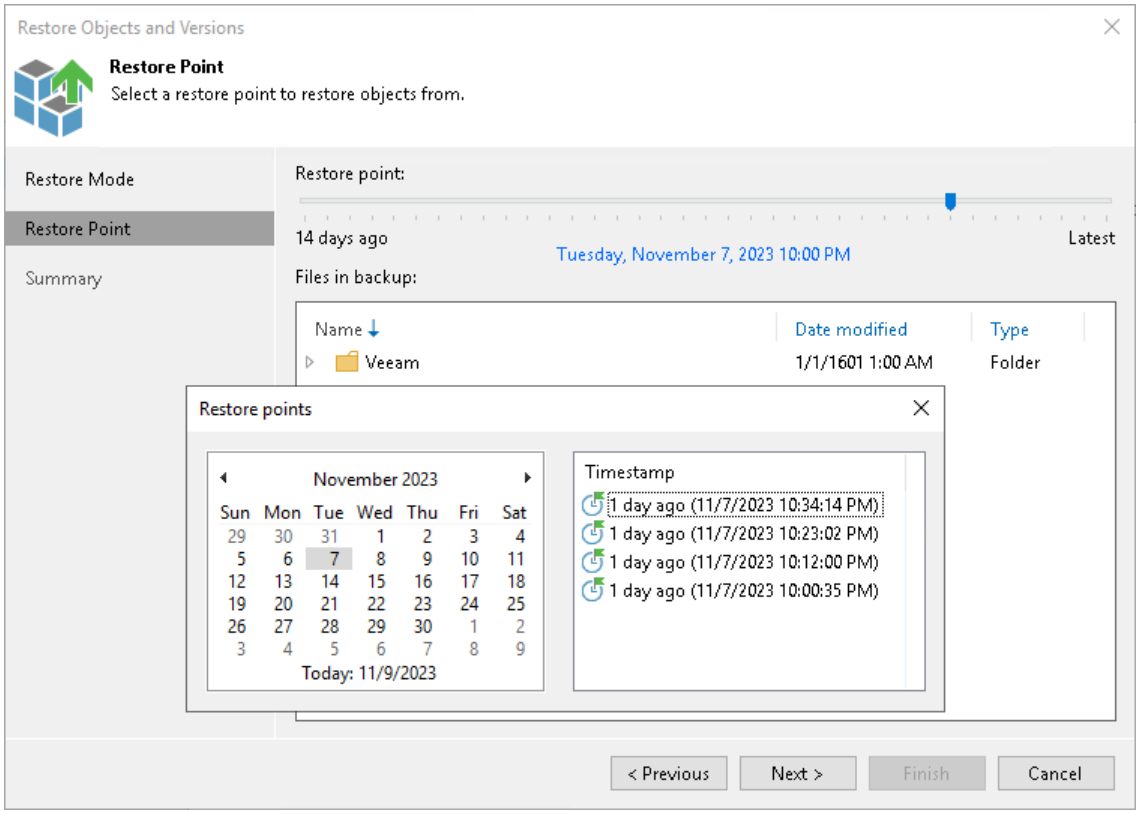

# Step 7. Select Object Version to Restore

The **Object Version** step is required if you use the **All Time** option at the Select Objects to Restore step and the selected objects have more than one object version.

Select one or more versions to restore. You can restore objects both from the backup repository and archive repository. To select several object versions, hold [Ctrl] and select multiple records in the table. Restore of multiple object versions can be helpful, for example, when you need to search for a specific version of the object, but you do not know for sure which one contains required changes.

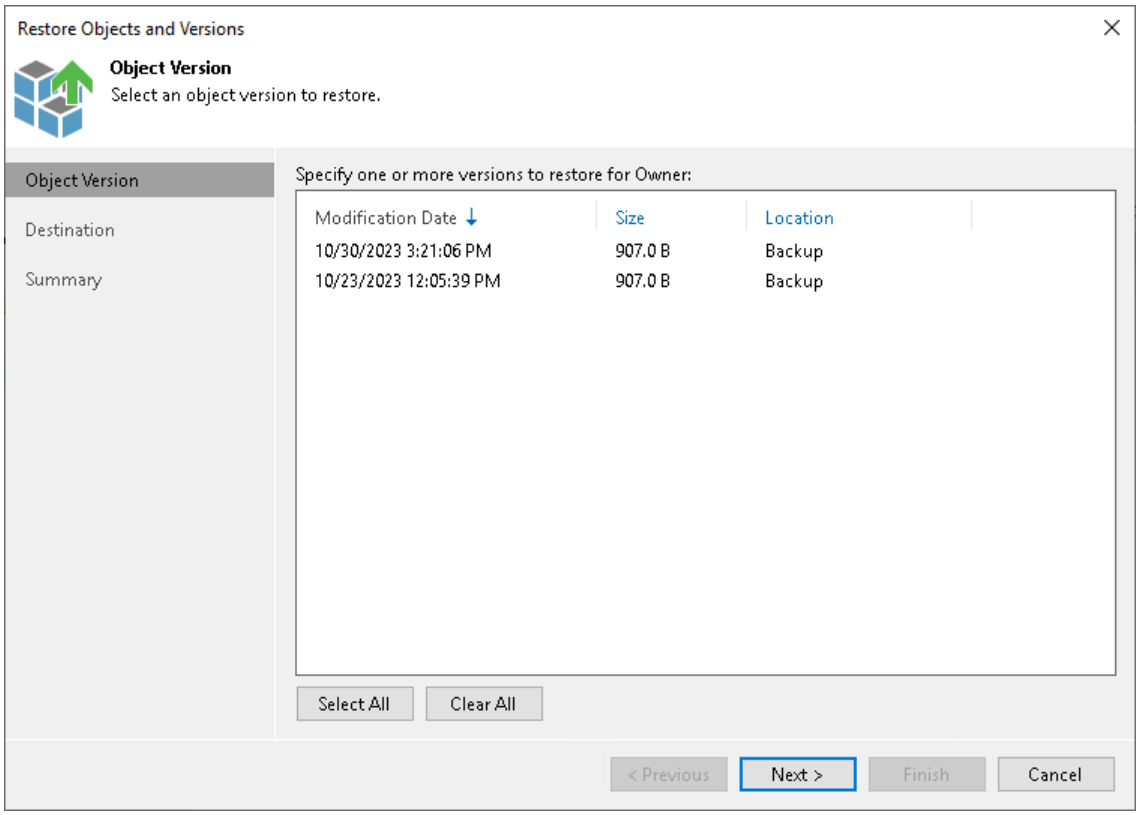

# Step 8. Specify Destination for Object Restore

The **Destination** step is required if you choose the **Copy To** option at the Select Objects to Restore step. Specify the destination where the restored objects must be stored:

1. In the **Restore to** field, select an unstructured data storage to which the objects must be restored. All unstructured data sources added to the inventory of Veeam Backup & Replication are available. If the required object storage is missing in the drop-down list, click **Add** and add the necessary object storage to Veeam Backup & Replication, as described in the Adding Unstructured Data Source section.
2. In the **Path** field, specify a bucket or container in the selected object storage or a path to the prefix in the selected object storage where objects must be restored.

   To select a dedicated prefix for restored objects, click **Browse**. In the **Select Folder** window, select the target location.

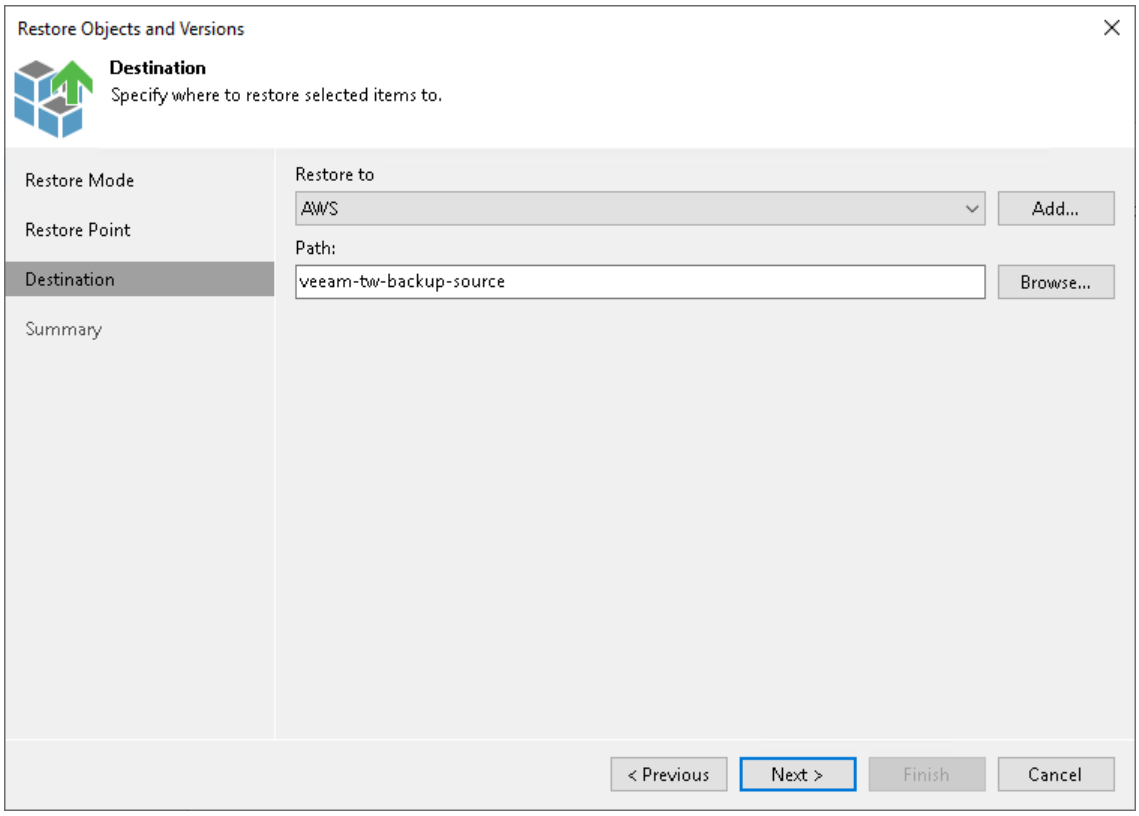

### Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review the object restore settings and click **Finish**.

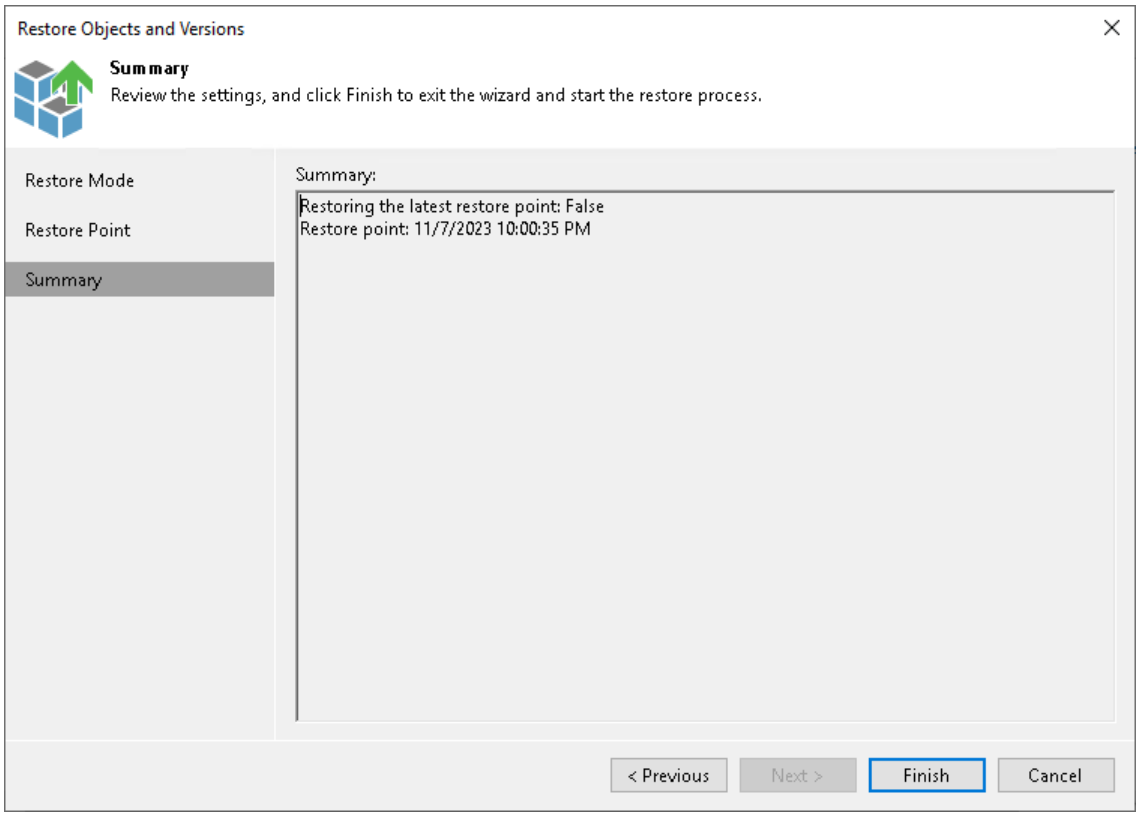


## Restoring Objects from Archive Repository

You can restore any object from the archive repository to the state of any object version stored in the archive. Depending on the circumstances, such a restore can require different actions.

> **NOTE**
>
> Consider that from the archive repository you can restore objects only. Restore of whole prefixes from the long-term repository is not supported.

## Regular Restore from Archive

To perform a regular restore from the archive repository when you have all required backup data stored both in the backup repository and archive repository, follow the instructions given in the Restoring Individual Objects or Versions section. Consider that to restore data from the archive repository, you must select the **All Time** option for selecting objects to restore.

## Emergency Restore from Archive

There can be different situations where backups are not available in the backup repository: for example, backed-up objects are corrupted or deleted from the backup repository, backup configuration files are removed from the configuration, archive repository is added into backup infrastructure of another backup server, or backup archive is moved from one archive repository to another one. In these cases, you can restore specific objects from the backup archive.

To restore objects from the archive repository, for example, from an object storage:

1. If necessary, add the storage that keeps the required archive to the backup infrastructure as described in section Adding Backup Repositories.
2. Rescan the added archive repository as described in section Rescanning Backup Repositories.
3. Download metadata for the archive backup:
   a. Locate the required object backup archive under **Backups** > **Object Storage (Imported)** node in the **Home** view.
   b. Right-click the object backup and select **Download backup metadata**.

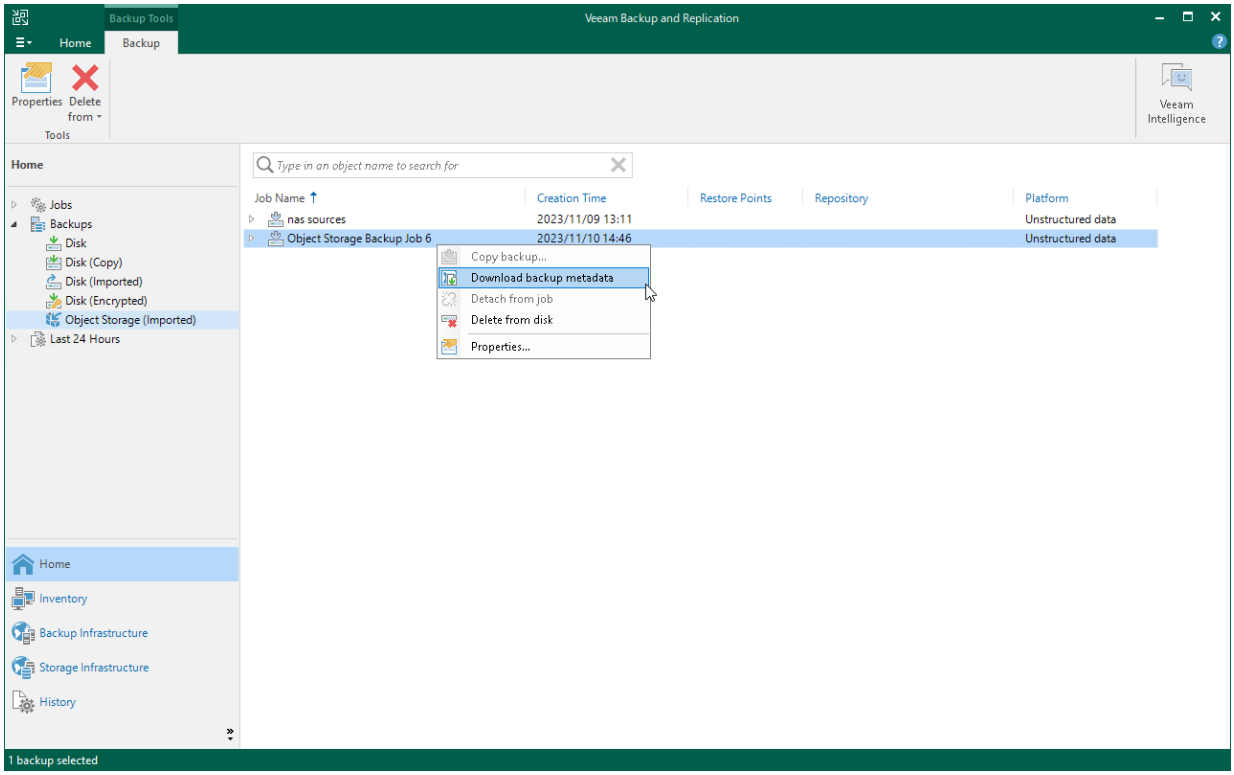

c. From the drop-down list, select a backup repository to store the downloaded metadata of objects archive and click **OK**.

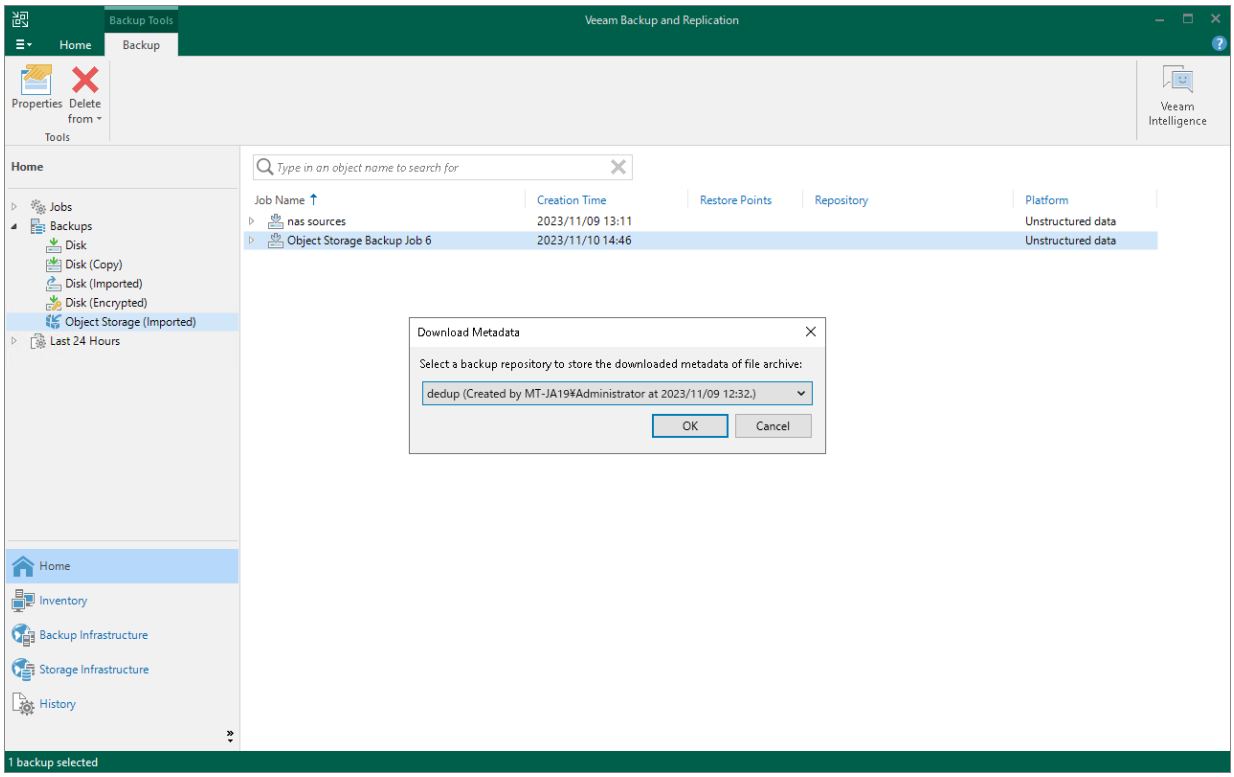


4. Restore objects from the archive backup as described in the Restoring Individual Objects or Versions section.

# VeeamZIP

With Veeam Backup & Replication, you can quickly perform backup of one or several VMs with VeeamZIP.

VeeamZIP is similar to a full VM backup. The VeeamZIP job always produces a full backup file (VBK) that acts as an independent restore point. You can store the backup file in a backup repository, in a local folder on the backup server or in a network share.

> **IMPORTANT**
>
> Consider the following:
>
> - Veeam Backup & Replication does not enforce backup repository throttling rules during VeeamZIP jobs.
> - You cannot use a Veeam Cloud Connect repository as a target for VeeamZIP jobs.

Backup files produced with VeeamZIP jobs are displayed in the **Home** view, under the **Backups** > **Disk (Exported)** node.

When you perform backup with VeeamZIP, you do not have to configure a backup job and schedule it. Instead, you can start the backup process for selected VMs immediately. This type of a backup requires minimum settings — you should only select the backup destination, choose the necessary compression level and enable or disable encryption and application-aware processing if necessary. For more information, see Creating VeeamZIP Backups.

To view the progress or results of the VeeamZIP job session, you can use the **History** view. For more information, see Viewing Real-Time Statistics.

To restore VM data from VeeamZIP backups, you can right-click it in the **Home** view and select the necessary restore option. You can also double-click the necessary VeeamZIP backup file on the machine where Veeam Backup & Replication is installed.

# Creating VeeamZIP Backups

You can quickly back up running and powered off VMs with VeeamZIP. VeeamZIP can be helpful if you want to create an ad-hoc backup for VMs, archive VMs before decommissioning and so on. You can create VeeamZIP backups for one or more VMs.

> **TIP**
>
> Veeam Backup & Replication keeps settings of the latest VeeamZIP task. To quickly create VeeamZIP backups with the same settings and store backups in the same location, right-click the necessary VM and select **VeeamZIP to**.

To create VeeamZIP backups:

1. Open the **Inventory** view. In the infrastructure tree, select a host or VM container (host, cluster, folder, resource pool, VirtualApp, datastore or tag) in which the VMs that you want to back up reside.
2. In the working area, select the VMs and click **VeeamZIP > VeeamZIP** on the ribbon or right-click the VMs and select **VeeamZIP**.

   To quickly find the necessary VMs, type the VM name or a part of it in the search field at the top of the working area and click the **Start search** button on the right or press [Enter] on the keyboard.

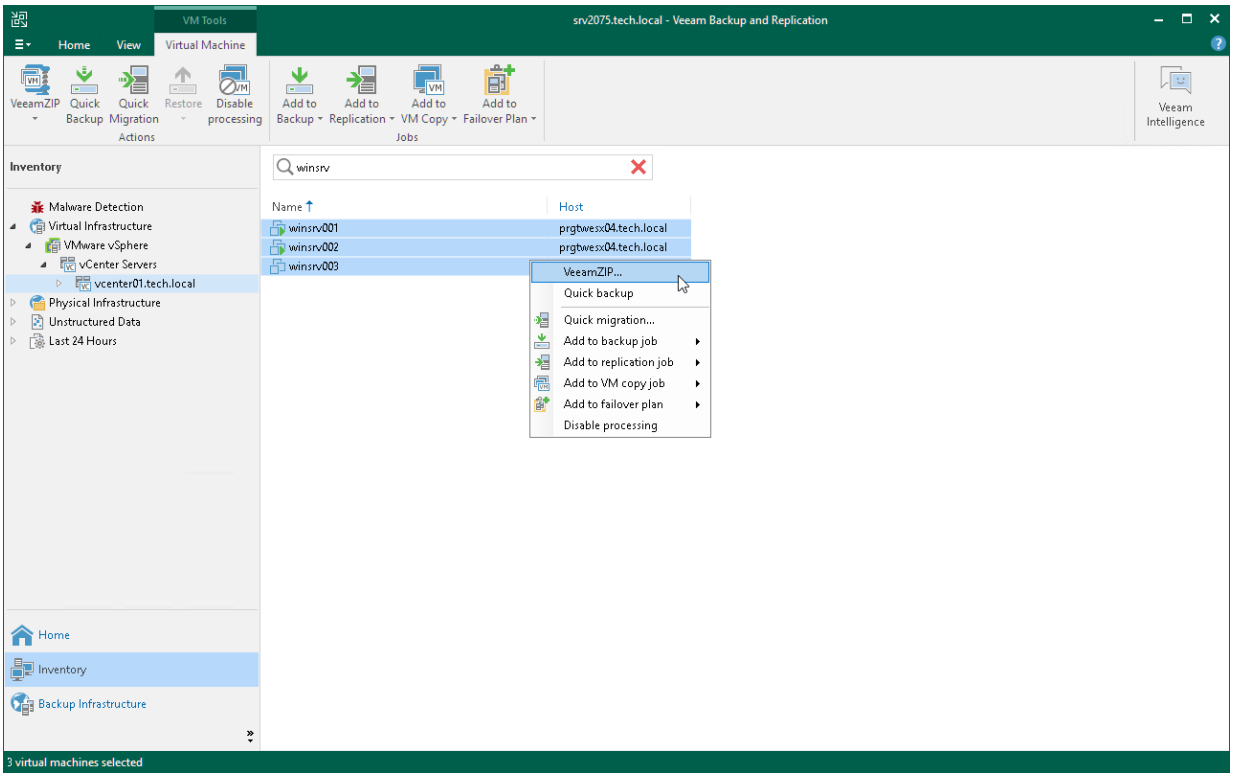


3. In the opened **VeeamZIP <N> VM** window:
   a. In the **Destination** section, specify a location in which you want to store VeeamZIP backups.
      - To store VeeamZIP backups in a backup repository, select **Backup repository** and choose the necessary backup repository from the list. In this case, VeeamZIP backups will be saved to the VeeamZIP subfolder of the folder where the backup repository stores backups. You can check this folder at the Configure Backup Repository Settings step of the backup repository wizard.

- To store VeeamZIP backups in a local folder on the backup server, select **Local or shared folder**, click **Browse** on the right and select a folder in which VeeamZIP backups must be stored.
- To store VeeamZIP backups in a shared folder, select **Local or shared folder** and type in the UNC name of the shared folder in the field under it. Keep in mind that the UNC name always starts with two back slashes (\\).

  If the shared folder requires authentication, select the necessary credentials from the **Credentials** list. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add necessary credentials. For more information, see Credentials Manager.

b. If you want to specify retention settings for the created VeeamZIP backups, select the **Delete this backup automatically** check box. From the drop-down list, select the retention period. The VeeamZIP backup file will be removed at 12:00:00 AM on the next day after the retention period ends.

   [For hardened repository] Veeam Backup & Replication sets an immutability period for backup files with retention period as equal to the longest of these periods. For more information, see How Immutability Works.

   If you do not want to delete VeeamZIP backups, leave the **Delete this backup automatically** check box unselected.

   > **TIP**
   >
   > You can customize retention period values in the drop-down list as described in this Veeam KB article.

c. To encrypt VeeamZIP backups, select the **Enable backup file encryption** check box. In the **Password** field, select the encryption method you want to use — a password that will generate secret keys or an external KMS server that will generate KMS keys. For more information about VeeamZIP encryption, see Encrypting VeeamZIP Backups.

   If you have not created the password beforehand, click **Add** or use the *Manage passwords* link to specify a new password. For more information, see Password Manager.

   If you do not use password loss protection, Veeam Backup & Replication will display a warning about it. For more information, see Password Loss Protection.

d. From the **Compression level** list, select a compression level for created backups.

e. By default, Veeam Backup & Replication uses VMware Tools quiescence to create a transactionally consistent image of VMs. You can disable VM quiescence. To do this, select the **Disable guest quiescence** check box. In this case, Veeam Backup & Replication will create a crash-consistent VM backup.

f. Click **OK**. The VeeamZIP task will start immediately. Veeam Backup & Replication will create a full backup file (VBK) and store it in the specified location. The VM name, date and time of the file creation are appended to the file name so you can easily find the necessary backups afterwards.

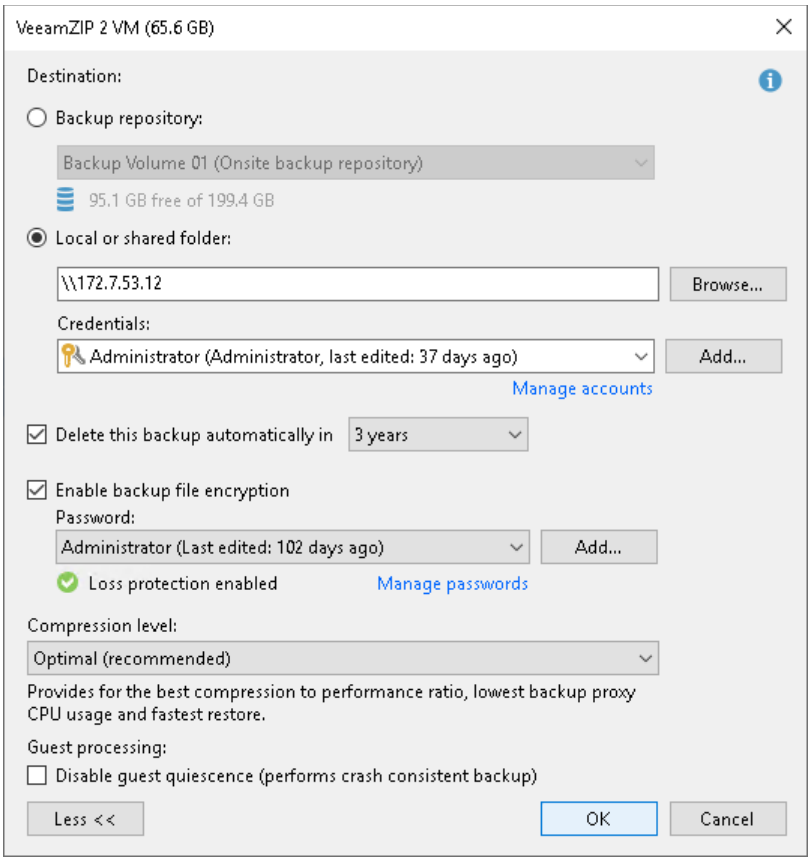


4. As the job runs, you can track the job performance in the real-time mode. To see the job results once it completes, open the **History** view, expand the **Jobs** node and click **Backup**. Then double-click the job session in the list.

# Encrypting VeeamZIP Backups

The encryption process for a VeeamZIP session includes the following steps:

1. Encryption is configured in the VeeamZIP session options. When you enable the encryption option, you need to enter a password or select a KMS server.
2. If you enter a password, Veeam Backup & Replication generates a secret key. If you select a KMS server, Veeam Backup & Replication asks the KMS server to generate an asymmetric KMS key. A secret key or a KMS key will be used to encrypt data encryption keys during the job. For more information, see How Data Encryption Works.
3. When you start a job, Veeam Backup & Replication encrypts backup files with data encryption keys. By default, this operation is performed on the backup proxy.
4. Encrypted data is passed to the target backup repository and stored to a resulting backup file.

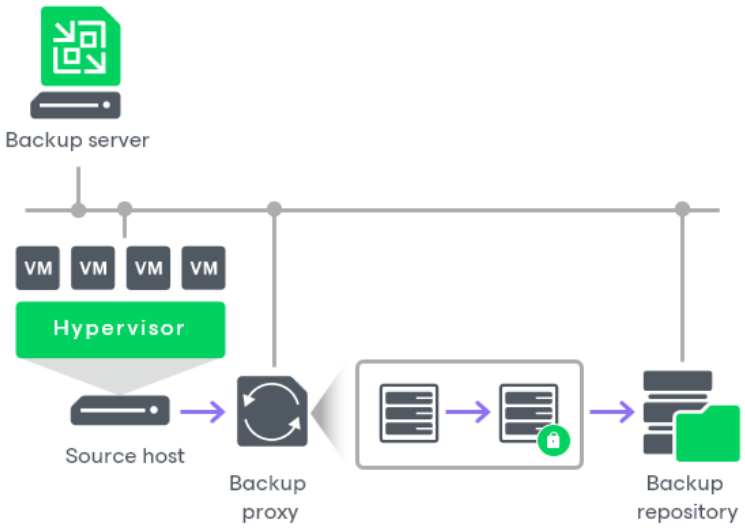


**NOTE**

If you plan to store backups in Veeam Data Cloud Vault, you must enable encryption for VeeamZIP backup.

# Managing and Restoring VeeamZIP Backups

Managing for VeeamZIP backups is practically the same as for regular backups. You can view backup properties, copy, move or delete backups from disk. For more information, see the following sections:

- Viewing Backup Properties
- Moving Backups
- Copying Backups
- Deleting Backups from Disk

You can restore your data directly from the VeeamZIP backups back to production servers. VeeamZIP data recovery does not differ from that of a standard backup data recovery and can be performed by using any of the following methods:

- Instant Recovery to VMware vSphere.
- Instant Recovery to Microsoft Hyper-V.
- Instant Disk Recovery.
- Entire VM Restore.
- Virtual Disk Restore.
- VM Files Restore.
- Guest OS File Restore.
- Application Item Restore.
- Restore to Amazon EC2.
- Restore to Microsoft Azure.
- Restore to Google Compute Engine.

# Backup Copy

The main backup purpose is to protect your data against disasters and virtual or physical machine failures. However, having just one backup does not provide the necessary level of safety. The primary backup may get destroyed together with production data, and you will have no backups from which you can restore data.

To build a successful data protection and disaster recovery plan, it is recommended that you follow the 3-2-1 rule:

- 3: You must have at least three copies of your data: the original production data and two backups.
- 2: You must use at least two different types of media to store the copies of your data, for example, local disk and cloud.
- 1: You must keep at least one backup off-site, for example, in the cloud or in a remote site.

Thus, you must have at least two backups and they must be in different locations. If a disaster takes out your production data and local backup, you can still recover from your off-site backup.

To help you adopt the 3-2-1 rule, Veeam Backup & Replication offers backup copy capabilities. Backup copy allows you to create several instances of the same backup data in different locations, whether on-site or off-site. Backup copies have the same format as those created by backup jobs and you can recover your data from them when you need it.

Backup copy is a job-driven process. Veeam Backup & Replication fully automates the backup copy process and lets you specify retention settings to maintain the desired number of restore points, as well as full backups for archival purposes.

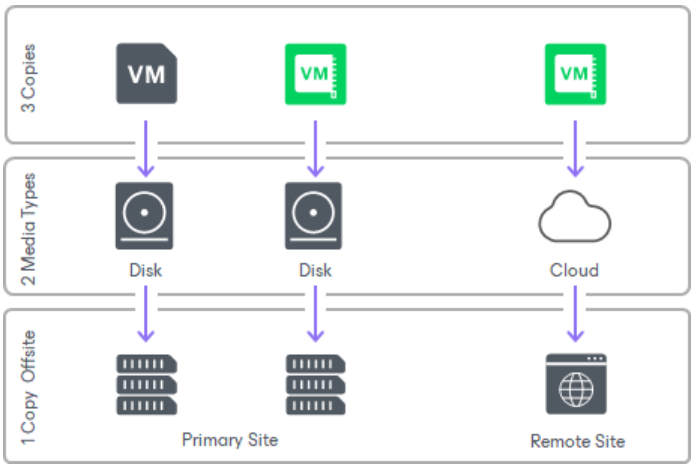

# About Backup Copy

With backup copy, you can create several instances of the same backup file and copy them to secondary (target) backup repositories for long-term storage. Target backup repositories can be located in the same site as the source backup repository or can be deployed off-site. The backup copy file has the same format as the primary backup, so you can restore necessary data directly from it in case of a disaster.

Veeam Backup & Replication supports backup copy for the following types of backups:

- Backups of VMware vSphere or VMware Cloud Director virtual machines created by Veeam Backup & Replication
- Backups of Microsoft Hyper-V virtual machines created by Veeam Backup & Replication
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for Mac, Veeam Agent for Oracle Solaris or Veeam Agent for IBM AIX
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV
- Backups of Oracle, SAP HANA and Microsoft SQL Server databases created by Veeam Plug-ins for Enterprise Applications
- Backups stored in an HPE StoreOnce backup repository
- File share backups created by Veeam Backup & Replication
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE
- Backups of Amazon EC2 instances created by Veeam Backup for AWS
- Backups of Microsoft Azure virtual machines created by Veeam Backup for Microsoft Azure
- Backups of VM instances created by Veeam Backup for Google Cloud
- Backups created by Veeam Backup for OLVM and RHV
- Backups exported by Kasten policies

> **IMPORTANT**
>
> Consider the following for copying backups created by Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud:
>
> - You can copy such backups from external repositories but not to them.
> - When copying backups from external repositories, consider that there can be a situation when new restore points were not copied. In this case, make sure the source chain for the backup copy job was not recreated. If it was recreated (all restore points have been deleted and created anew), you must create a new backup copy job with a new chain of backups.

When the backup copying process starts, Veeam Backup & Replication accesses backup files in the source backup repository, retrieves data blocks for a specific machine from the backup file, copies them to the target backup repository, and composes copied blocks into a backup file in the target backup repository. The backup copying process does not affect virtual and physical infrastructure resources, does not require creation of additional VM snapshots or VSS snapshots and does not produce load on machines whose backups are copied.

Backup copy is a job-driven process. To copy backups, you need to configure backup copy jobs. The backup copy job defines when, what, how and where to copy. For more information on how to create backup copy jobs, see Creating Backup Copy Jobs for VMs and Physical Machines. Note that to copy file share backups, you need to configure a file backup job, not the backup copy job. For more information, see Creating File Backup Jobs.

# How Backup Copy Works

Veeam Backup & Replication performs backup copy in the following way:

1. [For VM backup copy jobs only] Veeam Backup & Replication connects to vCenter Servers and ESXi hosts to gather information about VMs whose restore points you want to copy.
2. For backup copying process, Veeam Backup & Replication starts two Veeam Data Movers — source Veeam Data Mover and target Veeam Data Mover. Veeam Data Movers location depends on the backup repository type and data transport path. For more information, see Backup Copy Architecture.
3. The first backup copy run always produces a full backup file. Veeam Backup & Replication copies data blocks that are necessary to build a full backup of a machine as of the most recent state.

   Veeam Backup & Replication can copy data blocks from one or more backup files in the backup chain in the source backup repository.

   - If the backup chain is created in the reverse incremental backup method, Veeam Backup & Replication copies data blocks of the latest full backup.
   - If the backup chain is created in the forward or forever forward incremental backup method, Veeam Backup & Replication copies data blocks from the first full backup and a set of incremental backups.

   To minimize the amount of traffic going over the network, you can set Veeam Backup & Replication to use WAN accelerator, data compression and deduplication technologies.
4. Veeam Backup & Replication transfers copied data to the target backup repository and writes all copied data blocks to the full backup file.

   - New backup copy jobs always work in per-machine mode. In per-machine mode, data of every machine in the job is stored to separate backup files in the target backup repository.
   - Backup copy jobs created in the previous versions of Veeam Backup & Replication continue to create backup files in the specified formats: single-file backups if the **Use per-machine backup files** option was disabled or per-machine backups with single metadata file if the option was enabled. You can upgrade the backup chain format as described in section Upgrading Backup Chain Formats.
5. During every next backup copy run, when a new restore point appears in the source backup repository, Veeam Backup & Replication copies incremental changes from this most recent restore point and transfers them to the target backup repository. Veeam Backup & Replication writes the copied data blocks to the incremental backup file in the target backup repository, that is, Veeam Backup & Replication creates a new restore point in the forever forward incremental backup chain.

   To retain the desired number of restore points, Veeam Backup & Replication uses a retention policy. For more information, see Short-Term Retention Policy.

If you want to store some restore points for longer periods (for weeks, months or years) and enable long-term retention policy (GFS retention policy), Veeam Backup & Replication creates a forward incremental backup chain. For more information on GFS retention policy, see Long-Term Retention Policy (GFS).

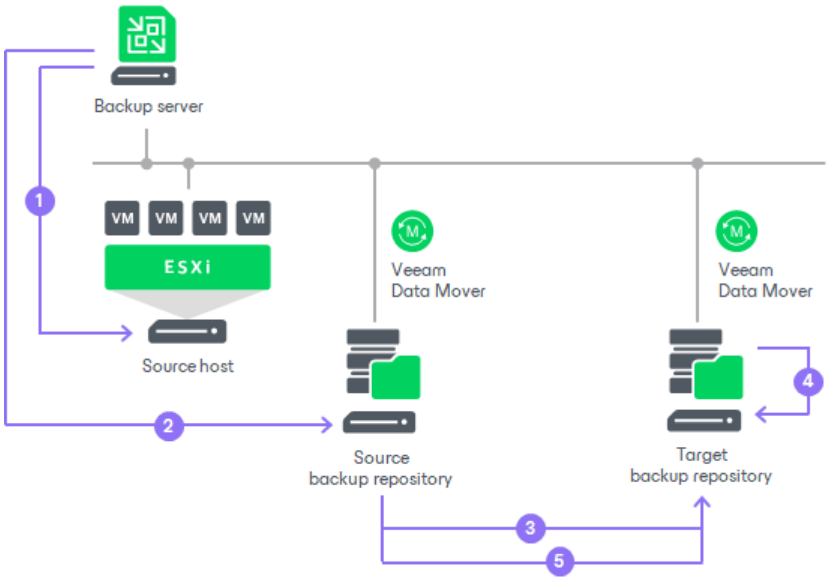


In some cases, the source backup job and backup copy job may overlap. This situation can occur if the source backup job needs to transform the source backup chain.

If a specific task in the backup copy job locks the source backup chain to read data from it, and the source backup job that needs to write data to this backup chain starts at this moment (for example, for reverse incremental backup), the task in the backup copy job is put on hold. The backup copy job can continue processing other tasks that use other sources (for example, backup files created by other backup jobs). After the source backup job releases the backup chain, the backup copy job resumes processing machines in this backup chain.

# Backup Copy Architecture

To transport data from the source backup repository to the target backup repository, the backup copy job uses one of the following paths:

- Direct transport path
- Transport path over WAN accelerators

## Direct Transport Path

Veeam Backup & Replication transports data directly from the source backup repository to the target backup repository. This type of data transport is recommended for copying backups to on-site backup repositories or off-site backup repositories over fast connections.

When Veeam Backup & Replication transports data over the direct data path, it uses Veeam Data Movers on the following backup infrastructure components:

- **Microsoft Windows and Linux repositories**. Veeam Backup & Replication uses the source Veeam Data Mover on the source backup repository and target Veeam Data Mover on the target backup repository.

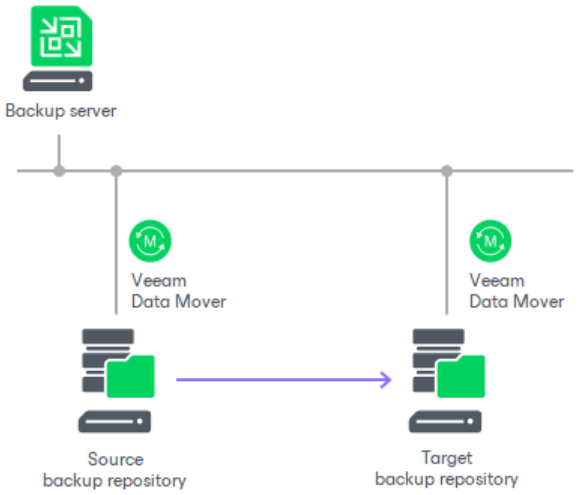

- **Shared folder backup repository**. If you have instructed Veeam Backup & Replication to automatically select the gateway server, Veeam Backup & Replication will use Veeam Data Movers deployed on mount servers associated with backup repositories. In case mount servers cannot be used for some reason, Veeam Backup & Replication will fail over to the backup server.

  If you have explicitly defined the gateway server, Veeam Backup & Replication will use the source Veeam Data Mover on the gateway server in the source site and target Veeam Data Mover on the gateway server on the target site.

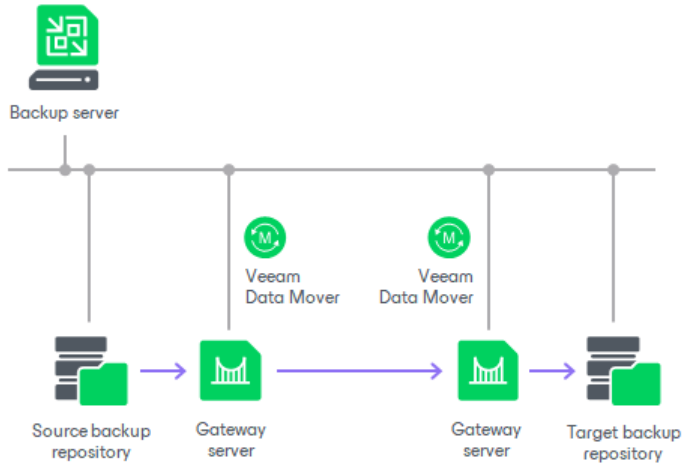


- **Object storage repository**. If you have not explicitly defined the gateway server, Veeam Backup & Replication will use both source and target Veeam Data Movers on the source backup repository site.

  If you have explicitly defined the gateway server, Veeam Backup & Replication will use the source Veeam Data Mover on the gateway server in the source site and target Veeam Data Mover on the gateway server on the target site.

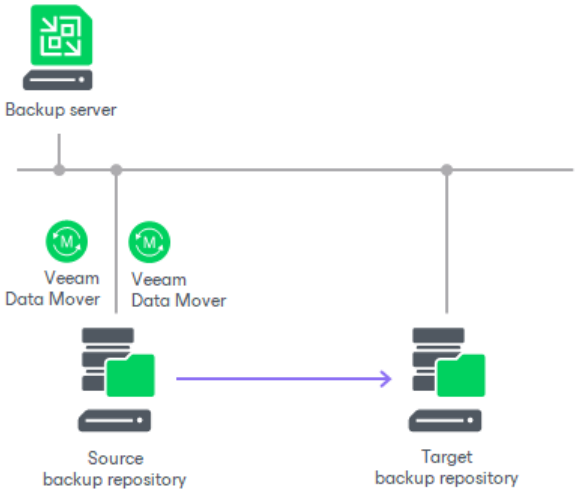


# Transport Path over WAN Accelerators

Veeam Backup & Replication transports data through a pair of WAN accelerators: one deployed on the source side and the other one deployed on the target side. WAN accelerators remove redundant blocks before transferring data and thus significantly reduce the amount of traffic going over the network. This type of data transport is recommended for copying backups off-site over slow connections or WAN.

> **IMPORTANT**
>
> The WAN acceleration technology is included in the Veeam Universal License. When using a legacy socket-based license, the Enterprise Plus edition is required. For more information, see WAN Acceleration.

When Veeam Backup & Replication transports data using WAN accelerators, it uses Veeam Data Movers on the following backup infrastructure components:

- **Microsoft Windows and Linux repositories**. Veeam Backup & Replication uses the source Veeam Data Mover on the source backup repository and target Veeam Data Mover on the target backup repository.
- **Shared folder backup repository**. If you have instructed Veeam Backup & Replication to automatically select the gateway server, Veeam Backup & Replication will use Veeam Data Mover deployed on the source and target WAN accelerator. If you have explicitly defined the gateway server, Veeam Backup & Replication will use the source Veeam Data Mover on the gateway server in the source site and target Veeam Data Mover on the gateway server on the target site.

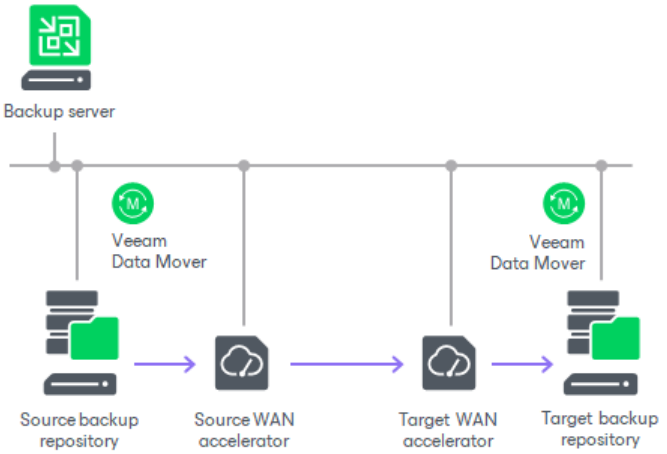

# Backup Copy Modes

Veeam Backup & Replication offers two backup copy modes:

- **Immediate copy**

  In the immediate copy mode, Veeam Backup & Replication copies restore points as soon as they appear in a source backup repository. Veeam Backup & Replication copies only restore points created by source backup jobs (backup jobs that you select when configuring a backup copy job).
  Veeam Backup & Replication can also copy transaction log backups if you enable this capability in job settings.

  The immediate copy mode is supported for the following backup types:

  - Backups of VMware vSphere and Microsoft Hyper-V VMs created by Veeam Backup & Replication.
  - Backups created by Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for IBM AIX, Veeam Agent for Oracle Solaris and Veeam Agent for Mac operating in the standalone or managed mode.
  - Backups created by Veeam Plug-ins for Enterprise Applications (Oracle RMAN, SAP HANA, SAP on Oracle).
  - Backups created by Veeam Backup for Proxmox VE.
  - Backups created by Veeam Backup for Nutanix AHV.
  - Backups created by Veeam Backup for OLVM and RHV.
  - Backups created by Veeam Backup for AWS.
  - Backups created by Veeam Backup for Microsoft Azure.
  - Backups created by Google Cloud.
  - Backups exported by Kasten policies.

- **Periodic copy**

  In the periodic copy mode, Veeam Backup & Replication copies the latest source restore point according to schedule specified in backup copy job settings. Veeam Backup & Replication copies only restore points created by source backup jobs (backup jobs that you select when configuring a backup copy job).

  The periodic copy mode is supported for the following backup types:

  - Backups of VMware vSphere and Microsoft Hyper-V VMs created by Veeam Backup & Replication.
  - Backups created by Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for IBM AIX, Veeam Agent for Oracle Solaris and Veeam Agent for Mac operating in the standalone or managed mode.
  - Backups created by Veeam Backup for Proxmox VE.
  - Backups created by Veeam Backup for Nutanix AHV.
  - Backups created by Veeam Backup for OLVM and RHV.
  - Backups created by Veeam Backup for AWS.
  - Backups created by Veeam Backup for Microsoft Azure.
  - Backups created by Google Cloud.

- Backups exported by Kasten policies.

> **NOTE**
>
> Periodic backup copy jobs created in Veeam Backup & Replication 11 or earlier became legacy periodic backup copy jobs in Veeam Backup & Replication 12. Legacy periodic backup copy jobs are fully operational, but you can only edit them and cannot create new ones. You can upgrade backup files of a legacy periodic backup copy job to the per-machine backup files using mapping. For more information, see Upgrading Backup Chain Formats.

# Changing Backup Copy Modes

Veeam Backup & Replication allows you to change the selected backup copy mode by editing backup copy job settings.

> **IMPORTANT**
>
> The periodic copy mode does not support processing of transaction log backups. Processing of transaction log backups must be turned off before changing the immediate copy mode to the periodic copy mode.

If you want to change the selected backup copy mode for a backup copy job created in earlier versions of Veeam Backup & Replication, it must have the per-machine backup with separate metadata files format. If a backup copy job has the per-machine backup with single metadata file format, you must upgrade its backup chain format to per-machine with separate metadata files or detach backups to start a new backup chain. For more information, see Upgrading Backup Chain Formats.

# Restore Point Selection

Veeam Backup & Replication always copies the most recent restore points, even if a backup copy job runs for the first time and source backup repositories already contain chains of restore points.

In the immediate copy mode, backup copy job copies the recent complete restore point created by a source backup job on the first run. On the next runs, Veeam Backup & Replication copies the oldest source restore points that were created after the point initially copied until there are no restore points left.

In the periodic copy mode, backup copy job starts according to schedule. Backup copy job copies the recent complete restore point created by a source backup job on the first run. On the next runs backup copy job continues to copy the recent complete restore points.

If there are no restore points considered as recent, Veeam Backup & Replication does not copy data from source backup repositories. Instead, it waits for new restore points to appear. Only after that, Veeam Backup & Replication copies the most recent data blocks to the target repository.

In the periodic copy mode, you can also specify the search scope for restore points. For more information, see Select Machines to Process.

## Limitations for Restore Points Selection

The following limitations apply when Veeam Backup & Replication selects restore points that must be copied to the target repository:

- Veeam Backup & Replication does not copy restore points from the target backup repository.
- Veeam Backup & Replication does not copy restore points from imported backups.
- Veeam Backup & Replication does not copy restore points that have already been copied by the same backup copy job to the target backup repository.
- Veeam Backup & Replication does not copy incomplete restore points.
- Veeam Backup & Replication does not copy restore points that are locked by the backup transformation process (merge, transform).
- A backup copy job does not copy a restore point if its data block size differs from the data block size of restore points that the job has already copied to the target backup repository. To copy restore points with the changed block size, you need to create active full backups. For details, see Change Storage Optimization Settings for Backup Copy Job.

  For example, if you have changed the block size for restore points in the source backup job (the **Storage optimization** option in the Storage Settings), Veeam Backup & Replication will not copy newly created restore points and will display the *Restore point is located in backup file with different block size* message.

# Transformation Processes

Veeam Backup & Replication can perform additional transformations in the target backup repository after the backup copying task or at the end of the backup copy interval. Transformation processes are the following:

- **Backup chain transformation**

  When a new restore point is copied to the target backup repository, Veeam Backup & Replication checks the retention policy settings for the backup copy job. If the limit in restore points is exceeded, Veeam Backup & Replication does the following:

  - If only short-term retention policy is enabled, Veeam Backup & Replication transforms the backup chain to make room for a new restore point. For more information, see Short-Term Retention Policy.
  - If long-term retention policy is enabled, Veeam Backup & Replication removes unnecessary restore points. Veeam Backup & Replication removes restore points in a way similar to the one described in section Forward Incremental Backup Retention Policy.

  For more information on retention policies, see Retention Policy for Backup Copy Jobs.

  After the transformation process, Veeam Backup & Replication can perform additional operations: remove data of deleted workloads from the backup chain and compact a full backup file.

- **Removal of deleted items**

  In the backup copy job settings, you can specify after which period you want to delete data of deleted workloads from backups created by backup copy jobs. After the period ends, Veeam Backup & Replication checks the list of workloads included in the job and removes data of deleted workloads from the backup chain in the target backup repository. For more information on how data is deleted and which limitations apply, see Deleted Items Retention. For more information on how to configure the deleted items retention, see Specifying Advanced Settings.

- **Full backup file compact**

  In the backup copy job settings, you can select to periodically compact a full backup file to reduce its size and increase the speed of read and write operations. For more information, see Compact of Full Backup File.

> **NOTE**
>
> If backup copy job processes the per-machine backup files, transformation processes will be performed for each object individually.

# Backup Copy Window

A backup copy window is a period of time when a backup copy job is allowed to transport data between source and target repositories.

By default, the backup copy window is configured to allow data transfer at any time. If you do not want the backup copy job to overlap the production hours, you can reduce the backup copy window and specify "prohibited" hours.

During the prohibited hours Veeam Backup & Replication cannot transfer data between source and target repositories. Other aspects of how Veeam Backup & Replication behaves during the prohibited hours and backup copy window depend on the selected backup copy mode. For more information, see Backup Copy Window and Prohibited Hours in Immediate Copy Mode and Backup Copy Window and Prohibited Hours in Periodic Copy Mode.

> **NOTE**
>
> During the prohibited hours, Veeam Backup & Replication stops only data transferring operations. Transformation processes in the target repository are still performed. For more information, see Transformation Processes.

## Backup Copy Window and Prohibited Hours in Immediate Copy Mode

When the prohibited hours start, Veeam Backup & Replication stops backup copy job sessions during which data transfer was not finished and assigns Failed to the session statuses. During the prohibited hours, Veeam Backup & Replication does not start new backup copy job sessions.

When the backup copy window starts, Veeam Backup & Replication transfers all restore points that were not transferred and creates as many incremental backup files as were omitted.

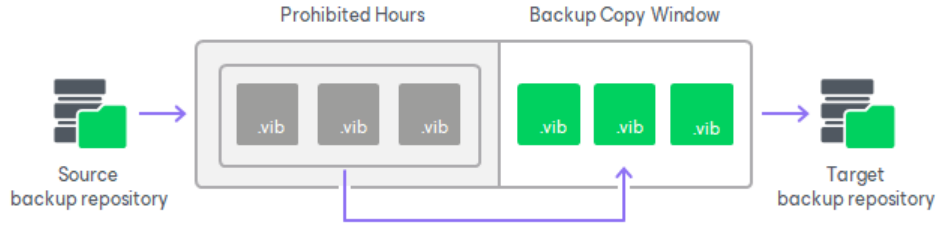


## Backup Copy Window and Prohibited Hours in Periodic Copy Mode

Backup copy job can not start during prohibited hours. If a prohibited time interval starts and backup copy job is still processing, it will cause backup copy job to stop.

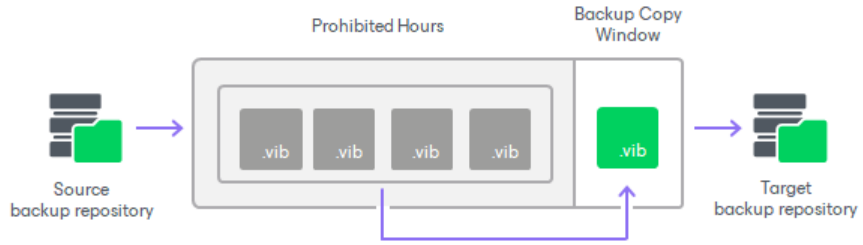

# Retention Policy for Backup Copy Jobs

The retention policy of a backup copy job defines for how long Veeam Backup & Replication must retain copied restore points in the target backup repository. The retention policy of a backup copy job does not depend on retention policy of the source backup job. The backup copy job has its own retention policy settings.

Veeam Backup & Replication offers two retention policy schemes for backup copy jobs:

- Short-Term Retention Policy
- GFS Retention Policy (Weekly, Monthly, Yearly)

Also, there is a separate retention policy for machines that has been removed from the infrastructure. For details, see Deleted Items Retention.

## Short-Term Retention Policy

The short-term retention policy allows retaining restore points created by backup copy jobs for a specified number of days or until the number of restore points reaches the specified number in the retention settings.

During the first backup copy session, Veeam Backup & Replication creates the first restore point — a full backup. The next backup copy sessions add incremental backups to the backup chain. As a result, the regular backup cycle produces a chain of a full backup and set of incremental backups in the target backup repository. When the retention policy is exceeded, Veeam Backup & Replication removes the earliest restore points from backup chains in the target backup repositories.

Since Veeam Backup & Replication creates forever forward incremental backup chains while backup copy jobs run, Veeam Backup & Replication applies the forever forward incremental retention policy to remove restore points and maintain the desired number of restore points.

When configuring short-term retention policy settings for a backup copy job, you have two options:

- Specify the number of restore points.

  Veeam Backup & Replication keeps the last *N* restore points, where *N* is the number of restore points that you specify in the settings. The minimum number that you can specify is 2.

- Specify the number of days.

  Veeam Backup & Replication keeps restore points created during the last *N* days, where *N* is the number of days that you specify in the settings. The minimum number that you can specify is 2.

  Consider the following for the daily retention policy:

  - The minimum number of retained restore points is 3. This number does not depend on the number of days set in the retention policy. For example, the retention policy is set to 5 days. You launch the job after it was stopped for 10 days. Normally, Veeam Backup & Replication deletes all previous restore points. However, due to the minimum number of retained restore points, you will still have at least 3 restore points: the newly created restore point and the two previous ones.

    You can change the minimum number of retained restore points with a registry value. For more information, contact Veeam Customer Support.

  - If the backup job starts at the end of the day and finishes the next day, Veeam Backup & Replication assumes that the restore point is created at the moment when the backup job started. However, Veeam Backup & Replication starts counting retention policy days only after the backup job finishes processing workloads.

- When determining whether the number of allowed days is exceeded, Veeam Backup & Replication ignores the day when the daily retention policy runs. In fact, Veeam Backup & Replication keeps restore points for the *N + 1* days, where *N* is the number of days that you specify in the settings.
- When determining whether the number of allowed days is exceeded, Veeam Backup & Replication also counts days when the backup job did not create any backups.

**NOTE**

If you want to create full backups periodically (weekly, monthly, yearly), enable the GFS retention policy. If you do not enable the GFS retention, the regular backup copy cycles will create only incremental backups.

**IMPORTANT**

This section describes how the short-term retention policy functions when the GFS retention policy is disabled. If you enable the GFS retention policy, consider the following:

- The backup copy chain will contain more restore points than you have specified in the short-term retention policy.
- With enabled GFS retention policy, Veeam Backup & Replication applies the forward-incremental retention policy to the backup copy chain.

For details, see GFS Retention Policy.

# Example

The regular backup cycle is based on the short-term retention policy scheme. When you specify retention policy settings, you define how many restore points you want to retain in the backup chain in the target backup repository.

For example, you have selected to retain 7 restore points. The backup copy runs once a day and starts on Sunday.

1. Veeam Backup & Replication creates a full backup on Sunday and add 6 incremental backups Monday through Saturday.

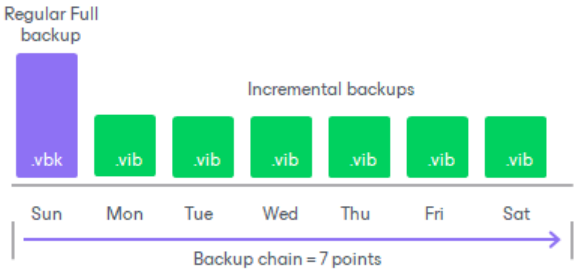

2. On Sunday, Veeam Backup & Replication creates another increment. As a result, there will be 8 restore points, which exceeds the retention policy. Thus, the oldest increment is merged to the full backup.

   After the oldest increment is merged to the full backup, Veeam Backup & Replication removes the increment as it is no longer needed.

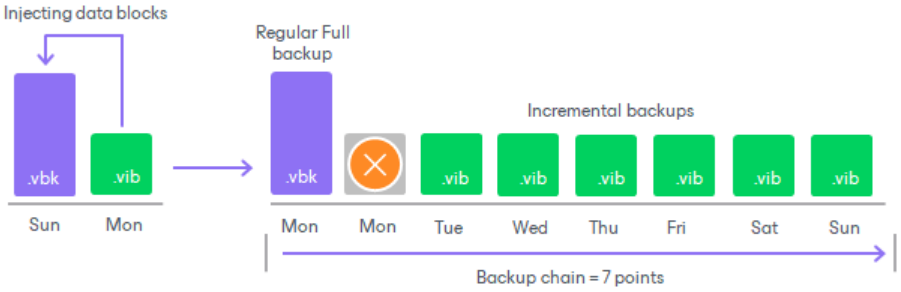


# Related Topics

Long-Term Retention Policy (GFS)

# Long-Term Retention Policy (GFS)

The long-term or Grandfather-Father-Son (GFS) retention policy allows you to store VM backups for long periods of time — for weeks, months and years. For this purpose, Veeam Backup & Replication creates synthetic or active full backup files and marks them with GFS flags. These GFS flags can be of three types: weekly, monthly or yearly. Depending on which flag is assigned to the full backup, it will be stored for specified number of weeks, months or years.

The GFS retention also helps you to mitigate risks that the short-term retention policy has, such as large number of subsequent incremental backups. Large number of subsequent incremental backups can increase recovery time, because Veeam Backup & Replication has to read data through the whole backup chain. Also, one corrupted increment can make the whole chain useless. When you configure the GFS retention, Veeam Backup & Replication creates weekly/monthly/yearly full backups, so instead of one backup chain consisting of one full backup and incremental backups, you will have several backup chains.

GFS backups are always full backup files that contain data of the whole machine image as of a specific date. GFS is a tiered retention policy and it uses a number of cycles to retain backups for different periods of time:

- Weekly backup cycle
- Monthly backup cycle
- Yearly backup cycle

In the GFS retention policy, weekly backups are known as 'sons', monthly backups are known as 'fathers' and yearly backups are known as 'grandfathers'. Weekly, monthly and yearly backups are also called archive backups.

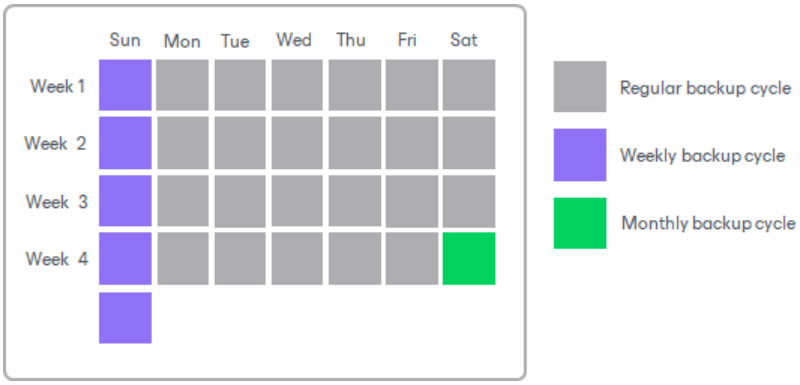


**NOTE**

GFS retention policy functions in combination with short-term retention policy. After you enable the GFS retention, the backup chain switches from the forever-forward incremental policy to forward incremental policy. Thus, the increments are no longer merged to the full backup file.

If you enable only yearly full backups without monthly and weekly backups, this can result in a large number of increments in a backup chain. To avoid this, it is recommended to enable an additional weekly GFS cycle. Weekly GFS cycle will update the backup chain every week which will allow Veeam Backup & Replication to remove excessive increment files.

# Related Topics

- How GFS Retention Works
- Limitations and Considerations for GFS Retention Policy

## How GFS Retention Works

To understand how GFS retention works, see the following sections:

- Backup Copy GFS Methods
- Backup Copy GFS Cycles
- Backup Chain for GFS Backups
- Limitations and Considerations for GFS Retention Policy

### Backup Copy GFS Methods

You can instruct Veeam Backup & Replication to create archive full backups with the following methods:

- Synthetic full method — Veeam Backup & Replication synthesizes archive full backups using restore points in the target backup repository.
- Active full method — Veeam Backup & Replication copies data for archive full backups from the source backup repository.

# Synthetic Full Method for Archive Backups

The synthetic full backup is the default method to create archive full backups. Veeam Backup & Replication does not copy data for archive full backups from the source backup repository. It synthesizes archive full backups from backup files that are already stored in the target backup repository. This approach helps reduce load on the network and production environment.

> **NOTE**
>
> The synthetic full method is not recommended if you use a deduplication storage appliance as a target backup repository. Performing a synthetic full backup in such repositories requires additional time and resources to download and decompress backup data blocks.
>
> This recommendation does not apply to HPE StoreOnce, Dell Data Domain and ExaGrid:
>
> - HPE StoreOnce and Dell Data Domain use virtual synthetics. Veeam Backup & Replication creates archive full backups by virtually synthesizing data blocks from existing backup files.
> - ExaGrid uses adaptive deduplication. Veeam Backup & Replication creates archive full backups from existing backup files that are stored in complete form in ExaGrid high-speed cache.

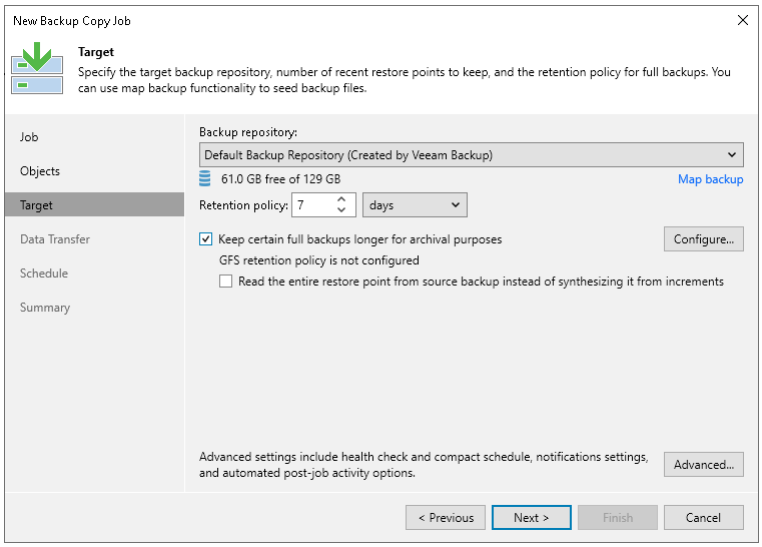


# Active Full Method for Archive Backups

You can instruct Veeam Backup & Replication to create archive full backups (backups retained by the GFS scheme) with the active full backup method. The active full backup method is recommended if you use a deduplicating storage appliance as the target backup repository. Active full backup helps improve the backup copy job performance and reduce the load on the target backup repository.

By default, Veeam Backup & Replication uses the synthetic backup method to create archive full backups. However, synthesizing archive full backups can cause problems with storage performance on deduplicating storage appliances. Deduplicating storage appliances are optimized for sequential data access. The synthetic backup creation, however, takes random I/O operations — Veeam Backup & Replication reads data from existing backup files and writes data to the synthesized archive full backup file. As a result, the storage performance can degrade.

In addition, backups reside in the target backup repository in the deduplicated and compressed state. Before creating synthetic full backups, Veeam Backup & Replication needs to download and decompress data blocks of backups, which requires additional time and resources.

**NOTE**

Consider the following:

- The active full backup method does not always copy the most recent restore point from the source backup repository. If the recent restore point is not created by the time the GFS task must start, Veeam Backup & Replication copies the latest available restore point from the source backup repository.
- If Veeam Backup & Replication does not manage to transfer the restore point according to backup copy schedule, Veeam Backup & Replication will finalize the transfer anyway.

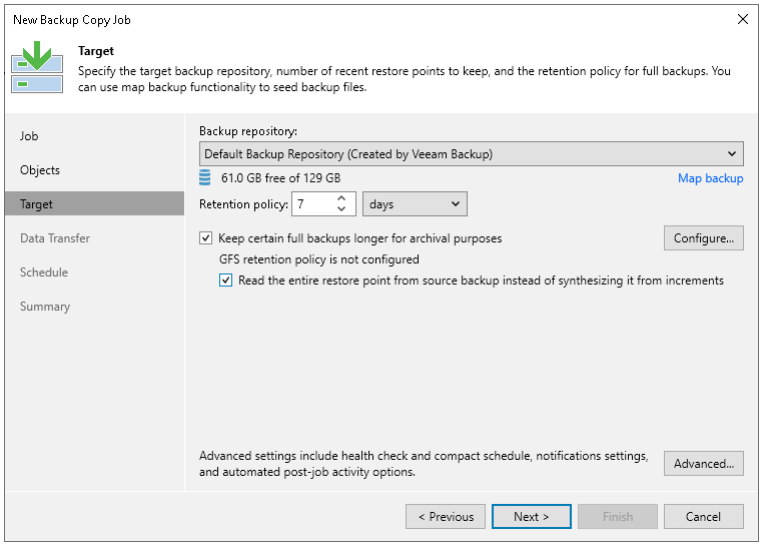


## Backup Copy GFS Cycles

When you configure the GFS schedule in the backup copy job settings, you specify for how long archive backups must be stored and on which day certain GFS backup must be created. You can configure the GFS schedule in the settings of a backup copy job. For details, see Backup Copy Job: Define Backup Copy Target.

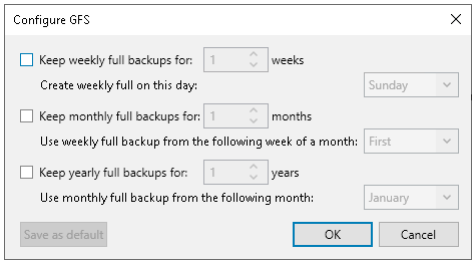

There are separate schedules for weekly, monthly and yearly full cycles. For details on settings of the GFS retention, see the following table.

| GFS Backup Option | Description |
|---|---|
| **Weekly GFS cycle** | If you want to create weekly full backups, select the **Keep weekly full backups for** check box. Then, specify the number of weeks during which the weekly backup must be stored on the target repository. During this period the weekly backup cannot be deleted or modified.<br><br>Veeam Backup & Replication creates a weekly full backup on the specified day of the week. On this day, creation of a weekly full backup starts as soon as the backup copy interval starts. |
| **Monthly GFS cycle** | To create monthly restore points, you can select the **Keep monthly full backups for** check box. Then, specify the number of months during which the monthly backup must be stored on the target repository. During this period the monthly backup cannot be deleted or modified.<br><br>Veeam Backup & Replication creates monthly full backups according to a schedule that depends on whether the weekly cycle is enabled or disabled:<br>• If weekly backups are enabled, Veeam Backup & Replication uses the weekly backup schedule and adds a monthly flag to the weekly backup.<br>• If weekly backups are disabled, Veeam Backup & Replication creates monthly full backups on the first day of the selected week. You can select **First**, **Second**, **Third**, **Fourth** or **Last** week of a month. If you select **First**, monthly backups are created on the first day of each month. If you select **Last**, monthly backups are created on different dates depending on the number of days in a month:<br>o For 31 days: 25th of the month.<br>o For 30 days: 24th of the month.<br>o For 29 days: 23rd of the month.<br>o For 28 days: 22nd of the month.<br><br>If the first day of the current week has already passed and you select the first week for the monthly full backup cycle, Veeam Backup & Replication creates a monthly full backup even if it is not the first day of the current week.<br><br>If you enable the monthly GFS cycle and select the first week that has already passed, Veeam Backup & Replication will create a monthly full backup only in the next month. Monthly full backup for the current month will not be created. |

| GFS Backup Option | Description |
| --- | --- |
| **Yearly GFS cycle** | If you want to create yearly restore points, select the **Keep yearly full backups for** check box. Then, specify the number of years during which the yearly backup must be stored on the target repository. During this period the yearly backup cannot be deleted or modified.<br><br>Veeam Backup & Replication creates yearly full backups according to a schedule that depends on whether the monthly cycle is enabled or disabled:<br>• If monthly backups are enabled, Veeam Backup & Replication uses the monthly backup schedule and adds a yearly flag to the monthly backup.<br>• If monthly backups are disabled, Veeam Backup & Replication creates yearly full backups on the first day of the selected month.<br><br>If the first day of the current month has already passed and you select the current month for the yearly full backup cycle, Veeam Backup & Replication creates a yearly full backup even if it is not the first day of the current month.<br><br>If you enable the yearly GFS cycle and select the month that has already passed, Veeam Backup & Replication will create a yearly full backup only in the next year. Yearly full backup for the current year will not be created.<br><br>**TIP**<br><br>If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server. |

## Backup Chain for GFS Backups

GFS retention creates yearly, monthly and weekly full backups and functions in a combination with the short-term retention. Short-term retention policy cannot delete or merge these GFS full backups.
Veeam Backup & Replication removes GFS backups only after the specified retention period for yearly/monthly/weekly backup is exceeded. Thus, the backup chain may contain more restore points than specified in the short-term retention policy.

When you enable the GFS retention, Veeam Backup & Replication no longer merges increments to full backups because GFS full backups cannot be modified. Thus, the short-term retention policy counts retention points only in the active backup chain not in the whole combination of backup chains.

**NOTE**

Veeam Backup & Replication removes GFS backup files only during running backup copy job sessions. This means that if the backup copy job does not run on the expected retention date,
Veeam Backup & Replication will remove the GFS backup file later during the next job session.

# Multiple GFS Flags

If you schedule a monthly or yearly full backup on the same day when the weekly full backup is scheduled, Veeam Backup & Replication creates only one archive full backup. The created backup will be marked at the same time as weekly, monthly and yearly GFS backup. In the Veeam Backup & Replication console, you will see all GFS flags assigned to the backup.

The full backup can be marked as weekly, monthly and yearly. When transforming weekly, monthly and yearly backup chains, Veeam Backup & Replication checks flags set for the full backup file. If the full backup file belongs to some other retention policy tier and must be retained in the target backup repository, such backup file will not be removed.

# Checking Which Restore Point Has GFS Flag

To check whether a restore point has a GFS flag, you can open the backup properties in the Veeam Backup & Replication console. Weekly, monthly, yearly backups have "*W*", "*M*" and "*Y*" flag in the *Retention* column. For instructions, see Viewing Backup Properties.

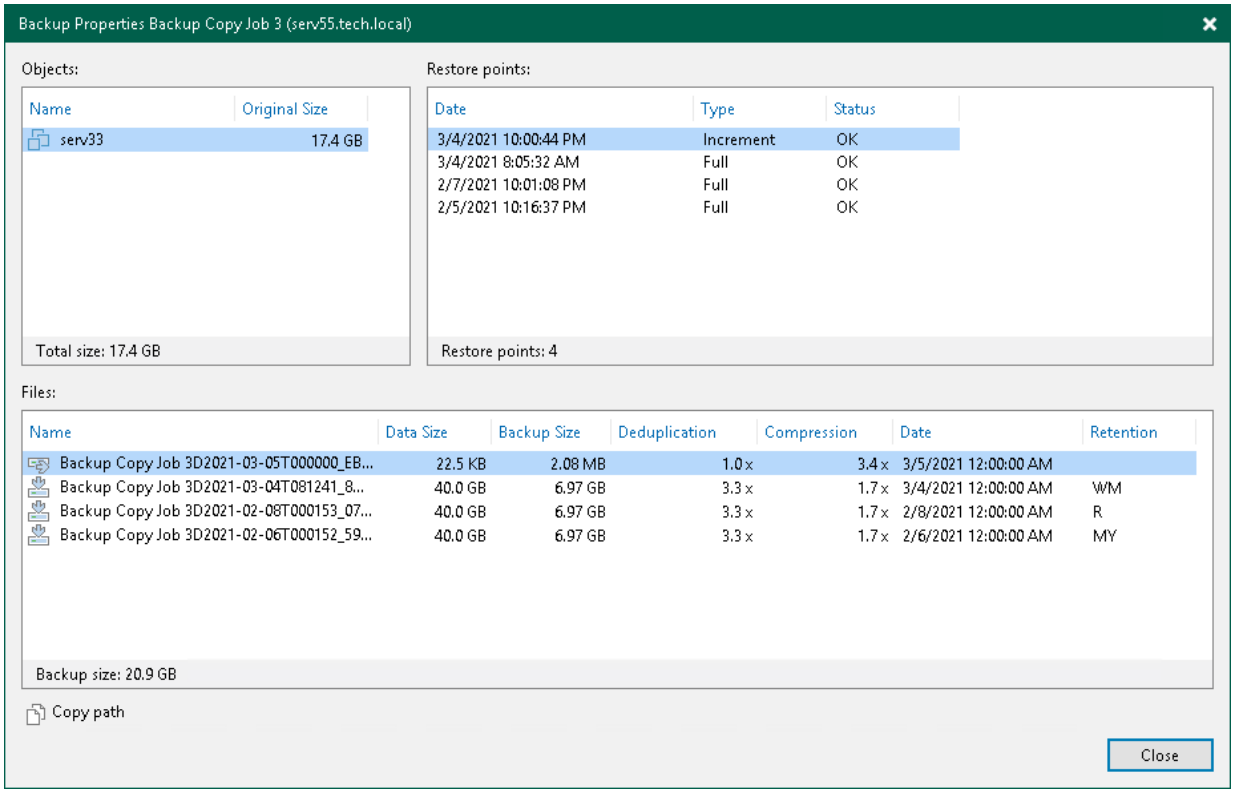


## Limitations and Considerations for GFS Retention Policy

Before configuring the GFS retention policy for a backup copy job, consider the following limitations and considerations:

- General Settings
- Periodic Copy Mode
- Changes in GFS Retention After Upgrading from Veeam Backup & Replication 10 to version 12

# General Settings

Consider the following for general settings of a backup copy job:

- You cannot enable GFS retention settings if you use a backup repository with rotated drives as the target backup repository.
- [For yearly GFS cycle] If you enable only the yearly GFS cycle, you can encounter the case when there is one full backup and a large number of increments for the whole year. To avoid this case, it is recommended to enable an additional weekly GFS cycle. Weekly GFS cycle will update the backup chain every week which will allow removing excessive increment files.
- If for some reason the GFS synthetic full was not created in the scheduled day, Veeam Backup & Replication will create the synthetic GFS full after the next run of the backup copy job.
- If it is the day when the GFS full backup must be created and there were no new backup files since the last run of the backup copy, Veeam Backup & Replication will create the GFS full backup from the latest available backup chain.
- GFS full backups cannot be merged or deleted by short-term retention. However, regular (R) full backups can be merged and removed by short-term retention if GFS is disabled.
- [For immediate copy mode] If the backup copy job was not run when the GFS full backup must be created, the GFS full backup will not be created on this day.

# Periodic Copy Mode

If you want to use the periodic copy mode, consider the following:

- If Veeam Backup & Replication does not manage to transfer the restore point according to backup copy schedule, Veeam Backup & Replication will finalize the transfer anyway.
- Veeam Backup & Replication creates a GFS full backup even if the GFS full backup creation is scheduled when the backup copy scheduled run is not finished. On the day when the GFS full backup must be created, Veeam Backup & Replication shows a warning that the current backup copy scheduled run will be completed. The way Veeam Backup & Replication behaves further depends on the selected backup copy GFS method:
  - In case of the synthetic full method, Veeam Backup & Replication first copies data for an incremental backup from the source backup repository and then, on the target backup repository, synthesizes the GFS full backup using this data and data of the already stored backup files.
  - In case of the active full method, Veeam Backup & Replication copies data for the GFS full backup from the source backup repository and creates the GFS full backup on the target backup repository.

# Changes in GFS Retention After Upgrading from Veeam Backup & Replication 10 to version 12

In Veeam Backup & Replication version 12, the backup copy GFS retention settings are different from the settings in version 10.

> **IMPORTANT**
>
> If you run an upgrade from version 10 to version 12 when the GFS retention policy is disabled, Veeam Backup & Replication will delete all GFS storage repositories created before.

If you upgrade to Veeam Backup & Replication 12, all GFS retention settings of existing backup copy jobs are automatically switched to the new format with minimal changes:

- **Weekly GFS retention**: If in version 10 the GFS setting was to keep *5* weekly backups, in version 12 the setting is changed to keep weekly backups for *5* weeks.
- **Monthly GFS retention**: If in version 10 the monthly GFS schedule was set to a period between the 1st Monday and 2nd Sunday, in version 12 the monthly GFS settings is changed to *First week*.

  If the monthly GFS schedule was set to a period between the 3rd Monday and last Sunday, in version 12 the monthly GFS settings is changed to *Last week*.
- **Yearly GFS retention**: If the yearly GFS schedule was set to a certain day of the month, in version 12 the schedule is set to the first day of the specified month.

  If in version 10 the yearly GFS schedule was set to the first-fourth Monday-Sunday of the year, in version 12 the schedule is changed to the first day of January.

  If the yearly GFS schedule was set to the last Monday-Sunday of the year, in version 12 the schedule is set to the first day of December.
- **Quarterly GFS retention:** Since Veeam Backup & Replication version 11, quarterly GFS retention policy option is deprecated.

  If in version 10 the quarterly GFS policy was enabled, in version 12 three additional months are added to the monthly GFS policy to compensate the quarterly full backups.

  If in version 10, the GFS policy was set to X monthly backups and Y quarterly backups. Then, in version 12, the retention policy is switched to store monthly backups for (X + 3Y) months.

> **IMPORTANT**
>
> [For synthetic method] Before the upgrade to Veeam Backup & Replication 12 or later, make sure that all GFS candidates (incremental restore points created on days when GFS was scheduled and that are expected to be transformed into full GFS restore points) are already transformed into GFS restore points. To force the backup copy job to transform all GFS candidates, you can temporarily decrease the short-term retention to a value less than the number of restore points between the latest restore point and the most recent GFS candidate and then wait till all the candidates are transformed.
>
> Before Veeam Backup & Replication version 11, Veeam Backup & Replication created GFS candidates on days when GFS was scheduled and only then transformed them into full GFS restore points according to the short-term retention. For more information on how restore points were transformed, see Synthetic Weekly Full Backups. Starting from Veeam Backup & Replication version 11, Veeam Backup & Replication creates GFS restore points according to a new schedule and creates them right on the scheduled days. After the upgrade, Veeam Backup & Replication no longer transforms previous GFS candidates into full GFS restore points. This means, that all GFS candidates lose their GFS status, they become regular incremental restore points and are deleted according to the short-term retention policy.

# Deleted Items Retention

After you configure a backup copy job, you may want to change something in the virtual infrastructure. For example, you may remove some virtual or physical machines or move machines to another location. You may also exclude machines from the backup copy job that has already run for some time.

By default, when you remove a machine protected by Veeam Backup & Replication from the virtual infrastructure, exclude a machine from the backup copy job or stop protecting a machine with Veeam Agent, the copied data still remains in backup files in the target backup repository. To avoid keeping redundant data on disk, you can enable the **Remove deleted items data after** option in the backup copy job settings. With this option enabled, at the end of every synchronization cycle Veeam Backup & Replication will remove data for deleted machines from backup files in the target backup repository.

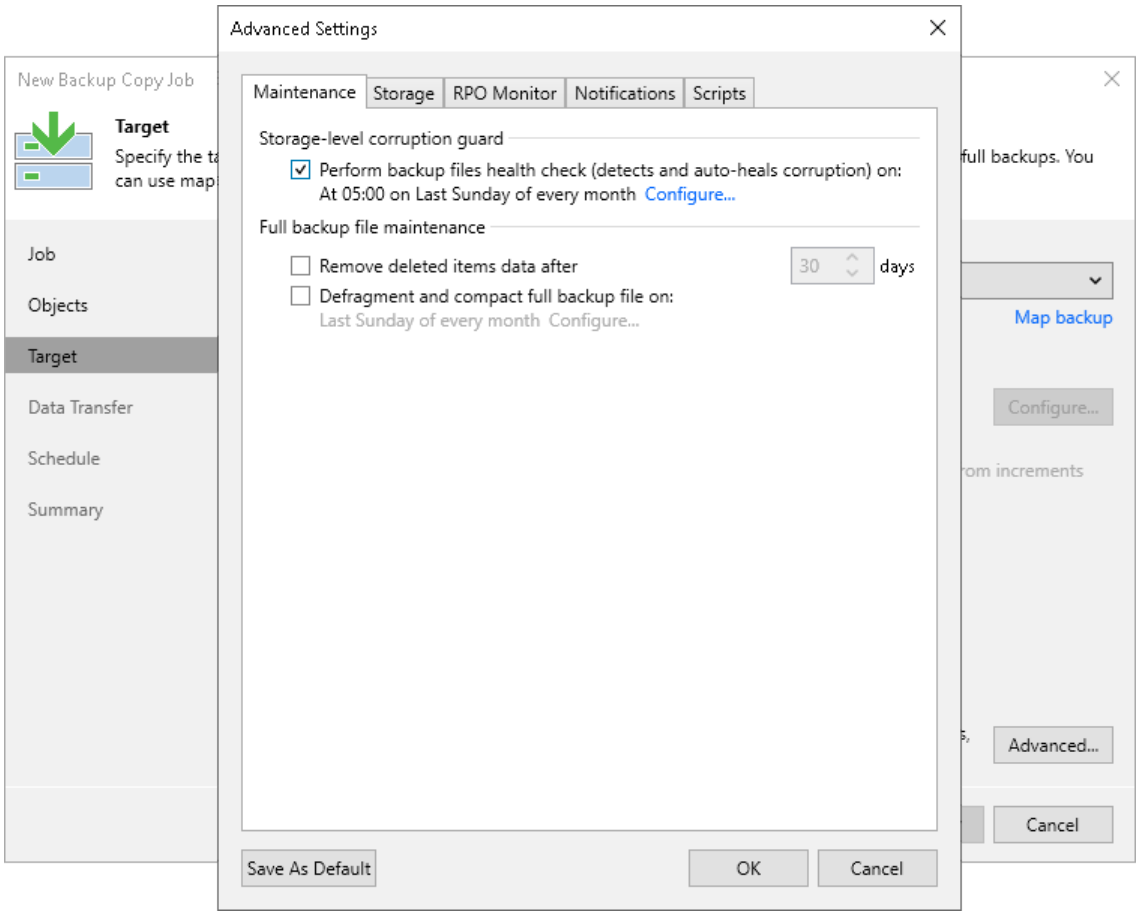


Veeam Backup & Replication removes data for the deleted machine only if two conditions are met:

1. Veeam Backup & Replication has not created valid restore points for the machine for the number of days specified in the **Remove deleted items data after** field.
2. The machine is not processed by the backup copy job.

This approach helps ensure that data for deleted machines can be saved by the GFS retention.

For example:

- The retention for the backup copy job is set to 7.
- The retention period for deleted machine is set to 3 days.

The backup copy job has created 3 successful restore points — a full backup and two incremental backups. During the next 4 days, no successful restore points were created. At the next synchronization cycle, Veeam Backup & Replication will not remove data for the deleted machine from the target backup repository as the backup chain contains successful incremental restore points for this machine.

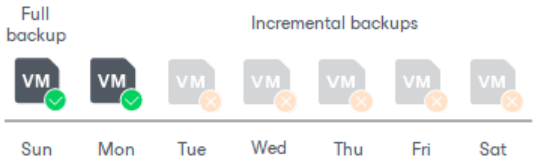


**IMPORTANT**

Consider the following:

- The deleted items retention applies only to regular backup chains. Veeam Backup & Replication does not remove data for deleted machines from weekly, monthly and yearly backups.
- [For single-file backups] When Veeam Backup & Replication removes data for deleted machines from regular backup chains, it does not free up space in the backup repository. It marks the space as available to be overwritten, and this space is overwritten during subsequent job sessions or the backup file compact operation.
- When Veeam Backup & Replication removes data for deleted machines from per-machine backup chains, it does not mark the space as available but deletes backup files since they contain data for 1 machine only.

Veeam Backup & Replication does not analyze the reason for which the machine has not been processed during the backup copy session. For example, the machine can be regarded as deleted if Veeam Backup & Replication has failed to obtain data for the machine from the virtual infrastructure, the machine has failed to be processed in time during the backup copy session and so on.

For this reason, you must be careful when specifying the retention period for deleted machines. If the retention period is too short, Veeam Backup & Replication may remove from the backup chain restore points that you still require.

For example, a backup copy job is configured to process 2 machines and has the following settings:

- The backup copy starts once a day.
- The retention for the backup copy job is set to 2.
- The retention period for deleted machines is set to 1 day.

The backup copy job runs in the following way:

1. On Sunday, the backup copy job creates a full backup for 2 machines - Machine1 and Machine2.
2. On Monday, the backup copy job creates an incremental backup for Machine1. The backup copy job does not manage to process Machine2 in time.
3. On Tuesday, the backup copy job creates an incremental backup for Machine1. The backup copy job does not manage to process Machine2 in time.
4. At the end of the backup copy job session on Tuesday, Veeam Backup & Replication transforms the backup chain and detects deleted machines. Veeam Backup & Replication regards Machine2 as a deleted machine — the deleted machines retention is set to 1 day, and after the transformation, there are no valid restore points for this machine in the backup chain.

As a result, after the backup copy session on Tuesday, backup files in the target backup repository will not contain data for VM2.

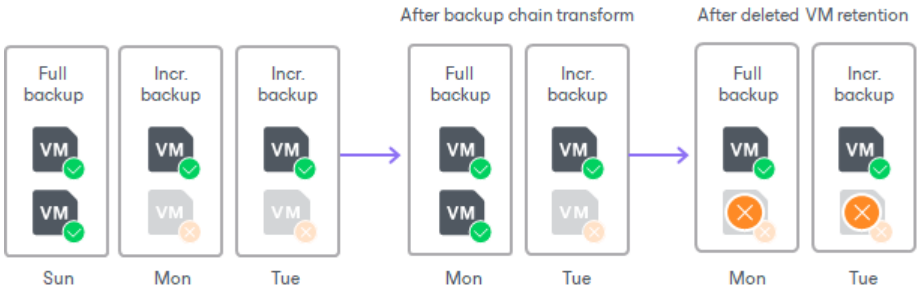

# Backup Copy Space Requirements

If the target backup repository is insufficiently provisioned, or the retention policy is configured to cause the job to retain more points than the target backup repository can contain, enforcement of retention points can fail. This can lead to the repository having not enough free space and restore points being unrestorable.

To calculate the space needed for backup copy job:

1. Find the total combined size of a full backup from each of the backup jobs that will be included in the backup copy job. (X)
2. Determine how many full backups you will have (Y):
   - One full backup for the short-term retention policy.
   - One full backup for every weekly, monthly and yearly backup cycle for the long-term retention policy (GFS).
3. If you have long-term retention policy configured, add one more full backup (Y+1).

   When GFS point creation takes place, there will be an extra full backup file on the disk until the operation is complete and retention enforcement deletes the oldest GFS restore point.
4. Multiply the total (Y+1) by the size estimated for a full backup (X). Make sure your repository has the free space for at least this much data and a little bit more for incrementals/variance in data.

# Example

The backup copy job is created to copy data from two backup jobs. A full backup from the first backup job is 750 GB and a full backup from the second backup job is 500 GB. Combining them, 1.25 TB should be accounted for the combined full backups created by the backup copy job, and a minimum of another 1.25TB should be allowed for the merge.

The backup jobs have a combined daily rate of change of 170GB (backup 1 has increments sized 100, 120, 100, 150, 125 and 175 GB for an average of 128.334 GB and backup 2 has increments of 50, 40, 55, 30, 45 and 30 GB for an average of 41.667 GB). If the backup copy job will retain 14 restore points, allow at minimum 14 * 170 GB or 2.38 TB for increments. Given that it is impossible to consistently predict rates of change, it is best to plan for the largest restore point.

Depending on your retention policy, the target backup repository should have the following free space available:

- For short-term retention policy: [1.25 TB full backup] + [1.25 TB merge] + [2.38 TB incremental points*14] = 4.88 TB.
- For long-term retention policy (GFS): [1.25 TB full backup] + [5.0 TB GFS points] + [1.25 TB merge] + [2.38 TB incremental points] = 9.88TB.

# Health Check for Backup Files

You can instruct Veeam Backup & Replication to periodically perform a health check for the latest restore point or backup file in the backup chain. The health check helps Veeam Backup & Replication make sure that further restore will be possible.

The health check starts according to the schedule. By default, the health check is scheduled to start at 5:00 on last Sunday of every month. The health check verifies restore points (full backup files or related full and incremental backup files). Only the latest restore points are verified.

**NOTE**

Consider the following:

- The health check process differs for backup files stored in the HPE StoreOnce repository. For more information, see Health Check for Backup Files Stored on HPE StoreOnce.
- If you perform the health check for the encrypted backup files, Veeam Backup & Replication will pass encryption keys to the regular backup repository or cloud repository. For more information on encryption, see Data Encryption.
- If the backup copy job is still running when the health check is scheduled to run, the health check will start after the backup copy job finishes.

## How Health Check Works

Veeam Backup & Replication performs the health check in the following way:

1. Veeam Backup & Replication calculates CRC values for backup metadata and hash values for data blocks of a disk in the backup file and saves these values in the metadata of the backup file, together with copied data.
2. On the day when the health check is scheduled, Veeam Backup & Replication performs the following actions:
    a. Veeam Backup & Replication performs the health check for the latest restore point in the backup chain. If the latest restore point in the backup chain is incomplete, Veeam Backup & Replication checks the restore point preceding the latest one.

    Veeam Backup & Replication calculates CRC values for backup metadata and hash values for disks data blocks in the backup file and compares them with the CRC and hash values that are already stored in the backup file.

b. If the health check detects corrupted data blocks, together with data blocks for the new restore point, Veeam Backup & Replication transports valid data blocks for the corrupted restore point. The valid data blocks are stored to the new incremental restore point created by this backup copy session. As a result, the backup chain gets "fixed", and you get a possibility to restore data from restore points following the corrupted restore point.

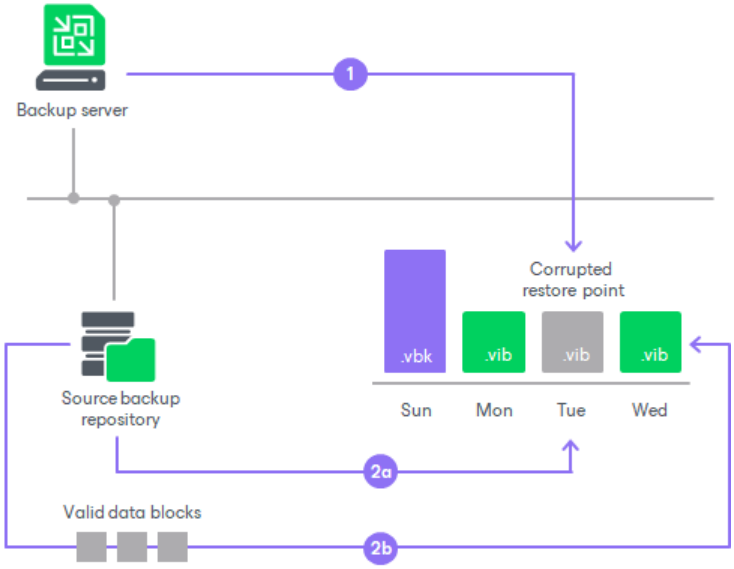


# Health Check for Backup Files Stored on HPE StoreOnce

The health check starts as soon as a backup copy job transfers backup files to the target repository. The health check verifies backup files, not restore points, and only those backup files transferred during the current job session.

Veeam Backup & Replication performs the health check in the following way:

1. As soon as a backup file is transferred to the target repository, Veeam Backup & Replication calls the HPE StoreOnce internal method that calculates checksums for data blocks of this backup file.

   The checksum calculation requires the rehydration of the verified data. This may lower the performance on the target repository especially during the first backup copy session. During this session, the health check verifies all backup files in source backup chains.

2. Veeam Backup & Replication compares the calculated checksums and the checksums already stored in the backup file.

3. If the health check detects corrupted data blocks, Veeam Backup & Replication deletes the backup file from the target repository. On the next backup copy job session, Veeam Backup & Replication transfers and performs the health check for this file again.

# Compact of Full Backup File

The backup copy job constantly transforms the full backup file in the backup chain to meet retention policy settings. The transformation process, however, has a side effect. In the long run, the full backup file grows large and gets fragmented. The file data occurs to be written to non-contiguous clusters on disk, and operations of reading and writing data from and to the backup file slow down.

To resolve the fragmentation problem, you can instruct Veeam Backup & Replication to compact the full backup file periodically. During the file compact operation, Veeam Backup & Replication creates a new full backup file in the target repository: it copies existing data blocks from the old backup file, rearranges and stores them close to each other. As a result, the full backup file gets defragmented, its size reduces and the speed of reading and writing from and to the file increases.

To compact the full backup file periodically, you must enable the **Defragment and compact full backup file** option in the backup copy job settings and define the compact operation schedule. By default, the compact operation is performed on the last Sunday of every month. You can change the compact operation schedule and instruct Veeam Backup & Replication to perform it weekly or monthly on specific days.

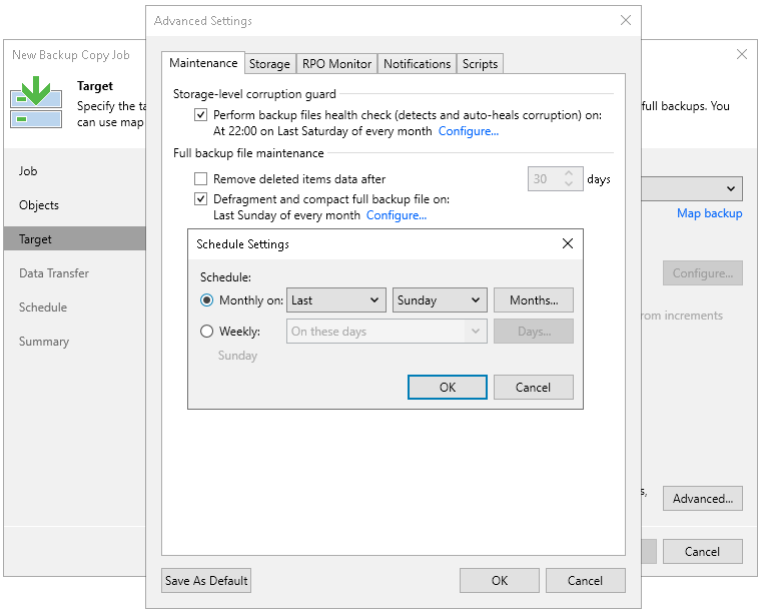

# Backup Copy Job Mapping

Backup copy job mapping helps you reduce the amount of data transferred over network and decrease the load on WAN accelerators or slow connections. You can also use mapping to upgrade from legacy backup chain formats to per-machine backup with separate metadata files format. For more information, see Backup Chain Formats.

[For legacy periodic backup copy job] If you use the target backup repository also as a target for other backup copy or backup jobs, you can already have a backup of machines that you want to copy. In this case, you can map the backup copy job to this backup.

A backup copy job mapped to a backup is performed in the following way:

1. Veeam Backup & Replication accesses a backup to which you map the backup copy job. The backup may have any number of restore points in the chain. This backup chain will be used as a seed for the further backup copying process.
2. During subsequent backup copy sessions, Veeam Backup & Replication copies restore points in a regular manner. It copies only incremental changes and stores them as new restore points next to the seed backup chain.

A mapped backup copy job does not store copied restore points in a dedicated folder in the target backup repository. Instead, it stores restore points to the same folder where the "seed" backup chain resides.

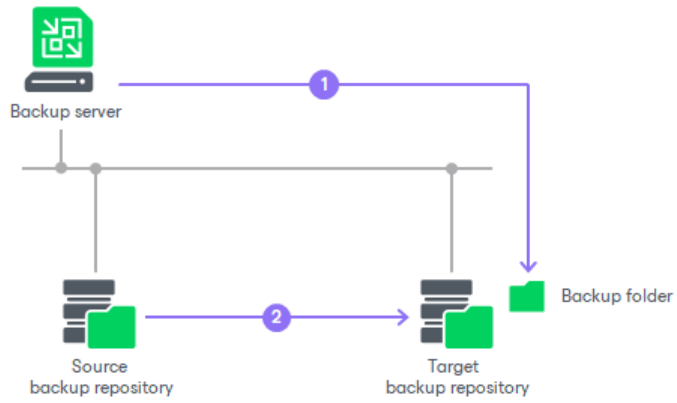


# Creating Seed for Backup Copy Job

Backup copy jobs have limitations for backups that can be used as seeds. The limitations are listed in section Map Backup File.

To create a seed for the primary backup copy job, do the following:

1. Create a backup copy job. Add jobs or repositories whose restore points you want to copy to this backup copy job. Target the backup copy job to some backup repository on the source side. This backup repository will be used as an intermediate one.
2. Run the backup copy job to create a full backup file (VBK) in the intermediate backup repository.
3. Transfer the created VBK file and VBM file from the intermediate backup repository to the target backup repository.
4. Perform repository rescan to populate the target backup repository.

   If the initial backup file was encrypted, you must enter a password to unlock the full backup file. Otherwise, Veeam Backup & Replication will not display the full backup file in the list of backups in the backup repository. For more information, see Importing Encrypted Backups.

5. Remap the backup copy job to the full backup file that you have created and transferred to the target backup repository.

As a result, Veeam Backup & Replication will use the full backup file as a seed. When a new restore point for the machine is available in the source backup repository, Veeam Backup & Replication will copy the restore point to the target backup repository and store it next to the full backup seed.

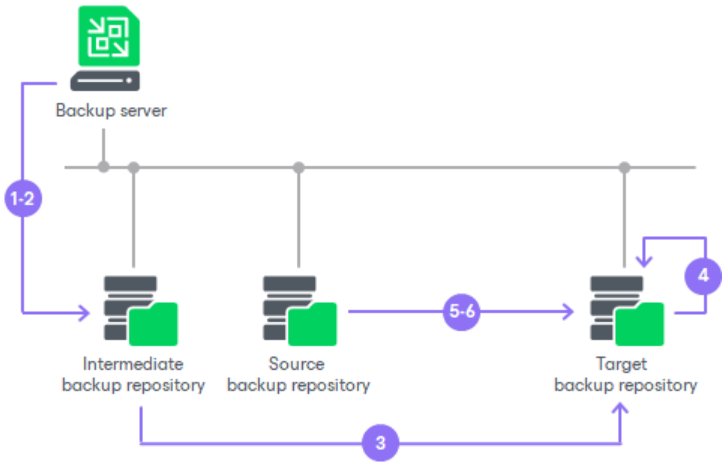


# Specifying Backup Copy Interval for Periodic Copy Mode

[For backup copy jobs created in Veeam Backup & Replication 11 or earlier] When you configure a backup copy job, make sure that its backup copy interval covers at least the latest restore point in the backup repository from which you plan to copy backups. The length of the backup copy interval has an impact on the algorithm of restore point selection. Veeam Backup & Replication copies only restore points that match the following criterion:

```
Time of restore point creation >= current time – backup copy interval
```

That is, if you have a backup chain whose latest restore point is 1 week old, you need to set the backup copy interval to 1 week. If you set the backup copy interval to a smaller time interval, for example, 1 day, the restore points that are older than 1 day will fall out of the search scope, and Veeam Backup & Replication will not transfer any restore points. For more information, see Restore Point Selection.

# Active Full Backup Copies

You can manually create an ad-hoc full backup for the backup copy job — active full backup copy, and add it to the backup chain in the target backup repository. To do this, you can use the **Active Full** button on the ribbon or the **Active Full** command from the shortcut menu.

Active full backup copy can be helpful if you want to change backup copy job settings, for example, enable or disable encryption. Veeam Backup & Replication will apply new settings starting from this full backup.

Veeam Backup & Replication treats archive full backups created with the active full backup method as regular backups and applies regular retention policy rules to maintain the necessary number of restore points.

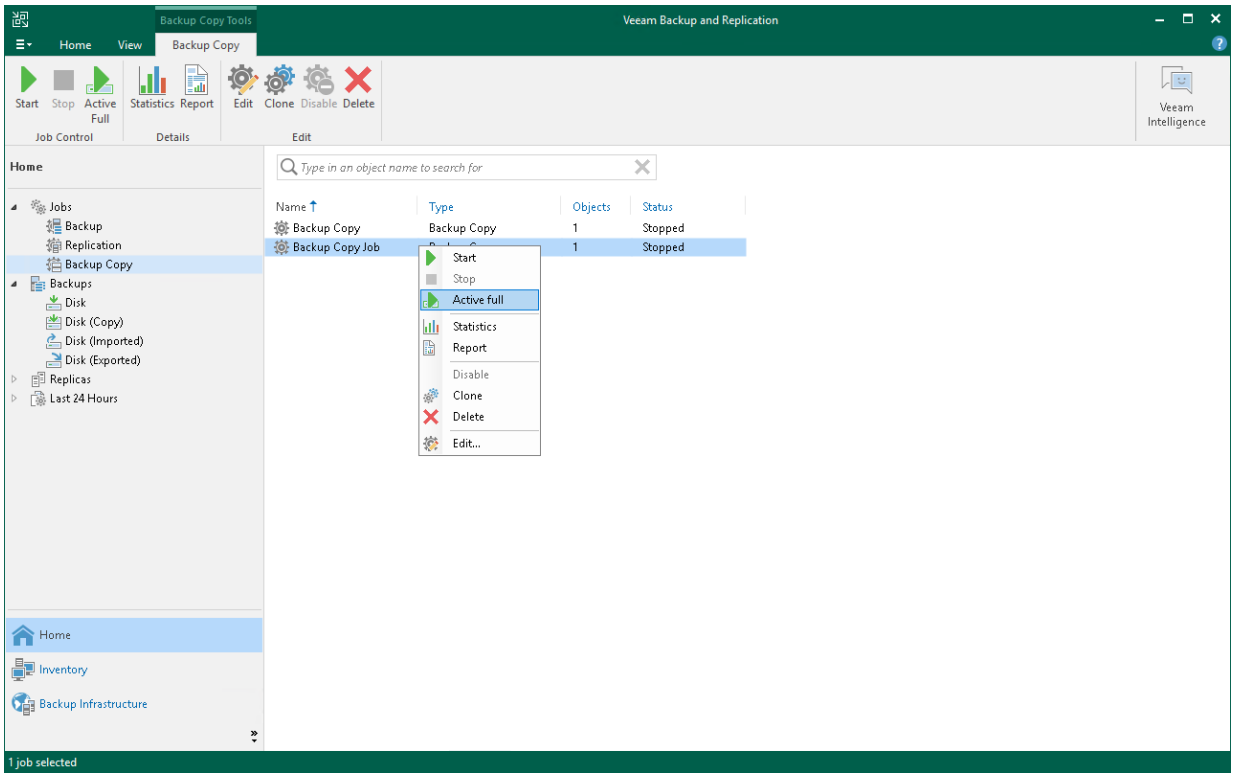


# Retention Policy for Active Full Backups

If you create active full backups for backup copy jobs, Veeam Backup & Replication applies to the backup chain retention rules of the forward incremental backup method. Veeam Backup & Replication waits until the number of restore points in the new backup chain is equal to the retention policy setting, and then removes the previous backup chain on the whole. For more information, see Forward Incremental Backup Retention Policy.

# Automatic Job Retries

Veeam Backup & Replication automatically retries several operations that are performed within a backup copy job session.

## Periodic Backup Copy Job Tasks Retry

By default, in the periodic copy mode, Veeam Backup & Replication automatically retries a failed backup copy task 5 times within one backup copy job session. A new task is started immediately after the previous one.

The backup copy task is retried only if the previous task has failed and a restore point has not been copied to the target backup repository. Veeam Backup & Replication does not perform a retry if a task has finished with the *Success* or the *Warning* status.

The backup copy task is retried during the same backup copy session only. If a restore point fails to be copied during all retries in the current backup copy session, Veeam Backup & Replication marks the current task as failed. [For the legacy periodic mode] After that, Veeam Backup & Replication performs the necessary transformation processes and starts a new backup copy session.

A backup copy job can process several machines. If only some machines are successfully processed by the backup copy task, Veeam Backup & Replication creates a restore point holding data for these machines in the target backup repository. Veeam Backup & Replication will attempt to process restore points for all machines during the next backup copy session.

> **NOTE**
>
> Some errors from WAN accelerators can block backup copy job retries. For example, if there is no space in the global cache on the target WAN accelerator, Veeam Backup & Replication puts backup copy operations on hold and waits for the expiration of the backup copy session.

## Immediate Backup Copy Job Retry

By default, in the immediate copy mode, Veeam Backup & Replication automatically retries a failed backup copy job 3 times. Instead of retrying each task within one backup copy job session, Veeam Backup & Replication performs each retry as a separate session run.

If no new restore points are available, Veeam Backup & Replication performs retries at 1 hour interval. If new restore points are available, Veeam Backup & Replication performs retries immediately.

## Transformation Retry

After the backup copying task, Veeam Backup & Replication may perform a number of additional transformation processes in the target backup repository. These processes include the backup chain transformation, removing of deleted machines from restore points and compacting a full backup file. For more information, see Transformation Processes.

Veeam Backup & Replication may fail to perform transformation: for example, if the backup file in the target backup repository is locked by the file-level restore session. By default, Veeam Backup & Replication automatically retries transformation processes for 5 times. The first interval between retries is 1 minute; the interval doubles with every new attempt. If all retries of transformation processes fail, Veeam Backup & Replication does the following:

- [For the immediate copy mode] Stops the job with the *Fail* status and waits for the new job session.

- [For the periodic copy mode] Stops the job with the *Fail* status and waits for the job to retry according to schedule.
- [For the legacy periodic copy mode] Puts the job to the idle state and waits for the new backup copy interval to begin.

# Virtual Infrastructure Access Retry

At the beginning of every backup copy session, Veeam Backup & Replication accesses the virtual infrastructure to make up a list of machines processed by the job.

Veeam Backup & Replication may fail to access the virtual infrastructure for some reason: for example, in case the vCenter Server is not responding. By default, Veeam Backup & Replication automatically retries access operations for 5 times with a 5 minute interval.

# Creating Backup Copy Jobs for VMs and Physical Machines

To copy backups to a secondary location, you must configure a backup copy job. The backup copy job defines how, where and when to copy backups. One backup copy job can be used to process one or multiple machines. Machines included in the job are processed in parallel. If a machine included in the backup copy job has multiple disks, disks are processed sequentially, one after another.

> **NOTE**
>
> If you want to copy backups between HPE StoreOnce repositories, follow the instructions listed in section Creating Backup Copy Jobs for HPE StoreOnce Repositories.
>
> If you want to copy file share backups, follow the instructions listed in section Creating File Backup Jobs.

Before you create a job, check prerequisites. Then use the **New Backup Copy Job** wizard to configure the backup copy job.

# Before You Begin

Before you create a backup copy job, check the following prerequisites:

- Backup infrastructure components that will take part in the backup copy process must be added to the backup infrastructure. This include target backup repository to which backups must be copied. For more information on adding components, see Backup Infrastructure Components.
- If you plan to use pre-job and post-job scripts, you must create scripts before you configure the backup copy job.
- If you plan to copy backups to an HPE StoreOnce repository, check limitations and requirements for it.
- If you plan to copy Veeam Agents backups, see Veeam Agent Management Guide for limitations and requirements.
- If you plan to copy backups to an immutable repository, you must enable the GFS retention policy. For more information, see Long-Term Retention Policy (GFS).
- If you plan to use WAN accelerators, check that you use the Enterprise Plus edition of Veeam Backup & Replication and that target and source WAN accelerators are added to the backup infrastructure. For more information, see Adding WAN Accelerators.
- If you delete the source backup job after creating the backup copy job, backup files will become orphaned. The orphaned backup files are displayed under the **Backups > Disk (Orphaned)** node. If the orphaned backup files were also stored in capacity tier or archive tier, they will also be displayed under the **Backups > Object Storage (Orphaned)** or **Backups > Archive (Orphaned)** nodes. The orphaned backup files can not be processed by any job.
- Even if **Use per-machine backup files** option is disabled in a repository that you are planning to use as the target, backup copy job will always create per-machine backup files in the backup repository.

# Step 1. Launch New Backup Copy Job Wizard

To run the **New Backup Copy Job** wizard:

- If you don't have application plug-ins backups, do one of the following:
  - On the **Home** tab click **Backup Copy**.
  - Open the **Home** view, in the inventory pane right-click **Jobs** or right-click anywhere in the working area, and click **Backup Copy**.
- If you have application plug-ins backups, do one of the following:
  - On the **Home** tab, click **Backup Copy** and click **Image-level backup copy** or **Application backup copy**.
  - Open the **Home** view, in the inventory pane right-click **Jobs** or right-click anywhere in the working area, click **Backup Copy** and click **Image-level backup copy** or **Application backup copy**.

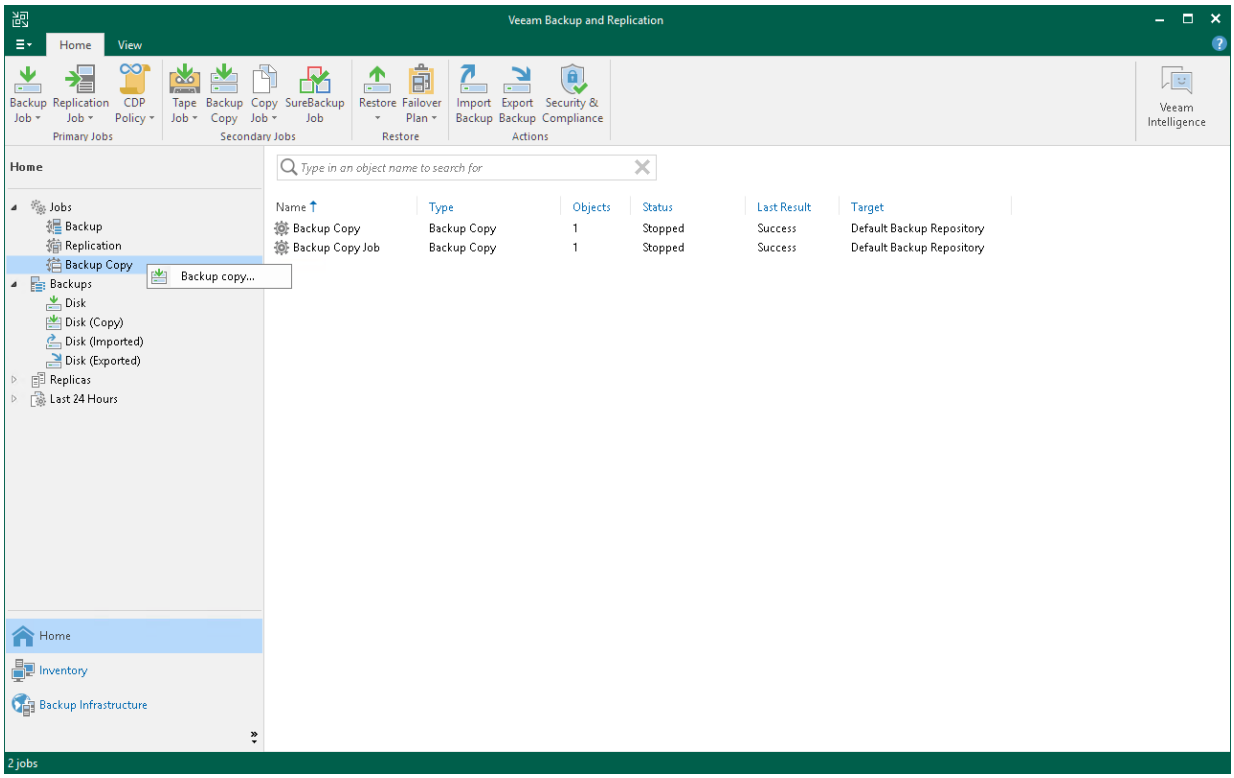

# Step 2. Specify Job Name and Copy Mode

At the **Job** step of the wizard, specify basic settings for the backup copy job:

1. In the **Name** field, specify a name for the job.
2. In the **Description** field, provide a description for the job.

   The default description contains information on a user who created the job, date and time when the job was created.
3. Select a backup copy mode. For more information on copy modes and backup types supported in each mode, see Backup Copy Modes.
   - Select **Immediate copy** to copy new restore points and, if required, log backups as soon as they appear.
   - Select **Periodic copy** to copy the most recent restore points.

> **NOTE**
>
> Backup copy job in the immediate copy mode processes only the latest backup file chain. To make such jobs copy all existing backups, create `BackupCopyMirrorAll (REG_MULTI_SZ)` registry value under the `HKLM\SOFTWARE\Veeam\Veeam Backup and Replication` key on the backup server. Populate this value with the names of backup copy jobs. Ensure that the registry value is created before the specified jobs are run for the first time.

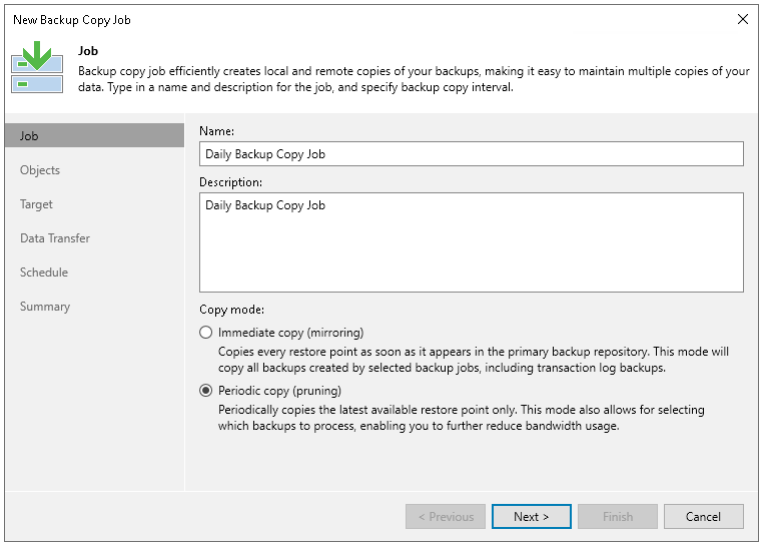

# Step 3. Select Workloads to Process

At the **Objects** step of the wizard, select workloads whose restore points you want to copy to the target backup repository:

1. Click **Add**.
2. Select a type of a source from which you want to copy restore points:
   - **From jobs**. You will see existing backup jobs. Veeam Backup & Replication will copy restore points created by the selected jobs.

     You can use another backup copy job as a source to process VMware, Cloud Director, and Hyper-V workloads, including workloads located on backup repositories with rotated drives.

     [For the periodic copy mode] If multiple jobs process one workload, Veeam Backup & Replication copies only restore points created by the first job in the **Objects to process** list.
   - **From repositories**. You will see all backup repositories in the backup infrastructure. Veeam Backup & Replication will copy restore points stored on the selected backup repositories.

     If you select repositories as sources, and target new jobs to the repositories in future, Veeam Backup & Replication will update backup copy job settings automatically to include these jobs to be copied.
   - **From backups**. You will browse for workloads in existing backups. Veeam Backup & Replication will search for restore points of the selected workloads in all backups of backup jobs created on the external repositories and will copy the most recent restore points. You can limit the search scope by selecting only specific backup policies for the backup copy job.

     This source is only available for backup copy jobs that process backups of Amazon EC2 instances, Microsoft Azure and Google Cloud.

     > **NOTE**
     >
     > You can use **From backups** source only if you have an external repository added to the backup infrastructure and have at least one backup file stored in it.
3. In the **Add Objects** window, select the necessary sources or workloads.
4. Click **Add**.
5. [For the immediate copy mode] If you have configured processing of transaction log backups in the source backup jobs, and want to copy these log backups to the target repository, select the **Include database transaction log backups** check box.

As an alternative, you can create a backup copy job with an empty source — that is, do not add any workloads at this step of the wizard. In this case, you need to configure a secondary destination for the source backup job and link it to the created backup copy job. For more information, see Linking Backup Jobs to Backup Copy Jobs.

# Limitations for Workload Selection

If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.

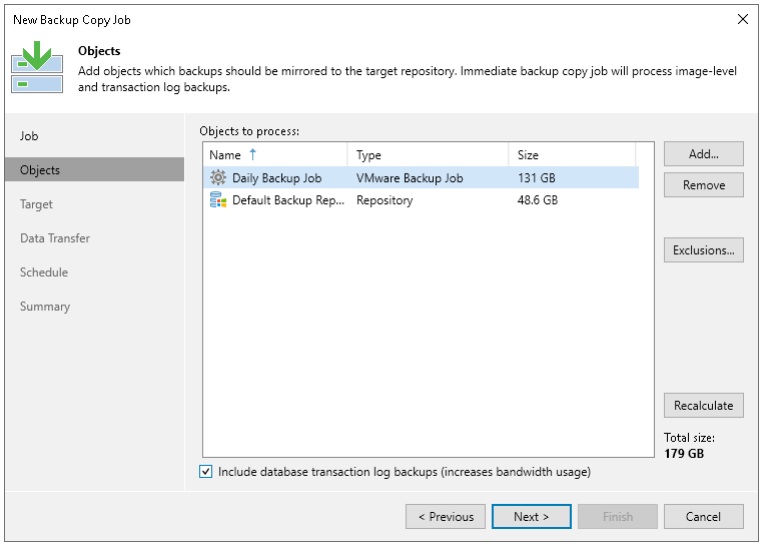

# Step 4. Exclude Objects from Backup Copy Job

This option is available only for virtual machines.

To specify which objects you want to exclude from the backup copy job:

1. At the **Objects** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, click **Add**.

   If you have added repositories as sources of restore points, you can exclude individual VMs or backup jobs from processing. To exclude individual VMs, click **Add** > **VMs**. To exclude individual jobs, click **Add** > **Jobs**.
3. In the **Add Objects** window, select objects that you want to exclude.

   When you exclude VMs, you can use the **Show full hierarchy** check box to display the hierarchy of all hosts added to Veeam Backup & Replication.
4. Click **Add**.
5. Click **OK**.

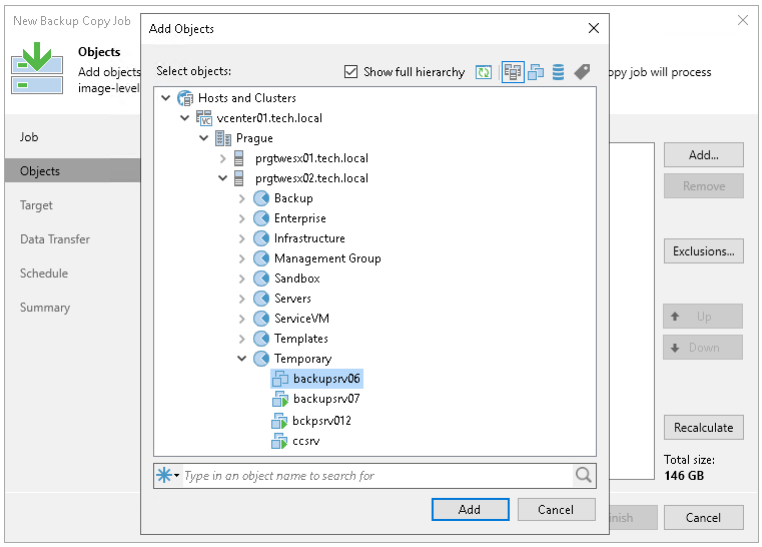

# Step 5. Define Processing Order

> **IMPORTANT**
>
> This functionality is available only when editing legacy periodic backup copy job created in Veeam Backup & Replication 11 or earlier.

You can define the order in which the backup copy job must process workloads. Configuring workload order can be helpful, if you want the backup copy job to process mission-critical workloads first. For this, put these workloads higher in the list to ensure that their processing fits the backup window.

Workloads inside a container are processed at random. To ensure that workloads are processed in the defined order, you must add them as standalone workloads, not as part of the container.

To define workload processing order:

1. At the **Objects** step of the wizard, select a workload whose order you want to change.
2. Use the **Up** and **Down** buttons on the right to move the workload up or down in the list.

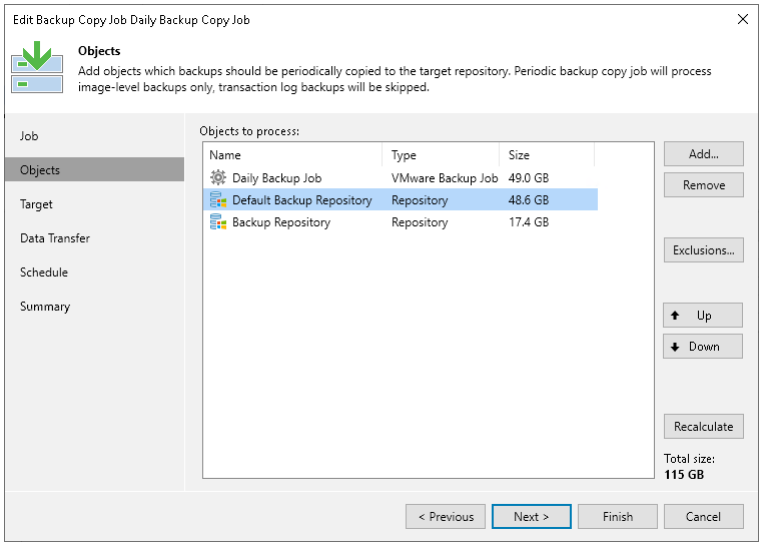

# Step 6. Specify Target Repository and Retention Settings

At the **Target** step of the wizard, define a target backup repository and configure retention policy:

1. From the **Backup repository** list, select a backup repository where copied backups must be stored.
2. In the **Retention Policy** field, configure the short-term retention policy for restore points:
   - If you want to keep the last <N> restore points, select *restore points* from the drop-down list and specify the number of restore points.
   - If you want to keep all restore points created during the last <N> days, select *days* from the drop-down list and specify the number of days.

   When the specified number is exceeded, the earliest restore point will be removed from the backup chain or will be merged with the next closest restore point. For more information on how Veeam Backup & Replication retains the desired number of restore points, see Short-Term Retention Policy.

   > **NOTE**
   >
   > If you enable the GFS retention, the short-term retention policy will not be able to delete and merge the GFS backup files. Thus, the backup copy chain will have more restore points than specified in the short-term retention policy.

3. If you want to create weekly, monthly and yearly full backups, you can configure long-term retention policy (GFS retention policy). GFS full backups will not be deleted or modified until the specified retention period expires. For more information on the GFS retention policy and its limitations, see Long-Term Retention Policy (GFS).

   To configure GFS retention policy, do the following:

   a. Select the **Keep certain full backups longer for archival purposes** check box.

   b. Click **Configure**.

   c. In the **Configure GFS** window, select the necessary GFS backup options. You can configure Veeam Backup & Replication to create weekly, monthly and yearly restore points. For details on settings of the GFS retention, see Backup Copy GFS Cycles.

   > **NOTE**
   >
   > Before you implement the GFS retention policy, see Limitations and Considerations.

4. You can define a way to create weekly, monthly and yearly full backups:
   - Synthetic Full Method: With this method, during the GFS backup copy creation, Veeam Backup & Replication does not copy data from the source backup repository but synthesizes full backups from backup files that are already stored in the target backup repository. This approach helps to reduce load on the network and production environment.

     The synthetic full method is used by default. To use this method, leave the **Read the entire restore point from source instead of synthesizing it from increments** option unselected.

- Active Full Method: With this method, Veeam Backup & Replication copies data for archive full backups from the source backup repository. This method decreases load on the target repository but increases load on the network and production environment.

  To use this method, select the **Read the entire restore point from source instead of synthesizing it from increments** option.

**NOTE**

Backup copy job with restore point-based retention policy does not support partial vApp active full backups. For more information, see Performing Partial Active Full Backup.

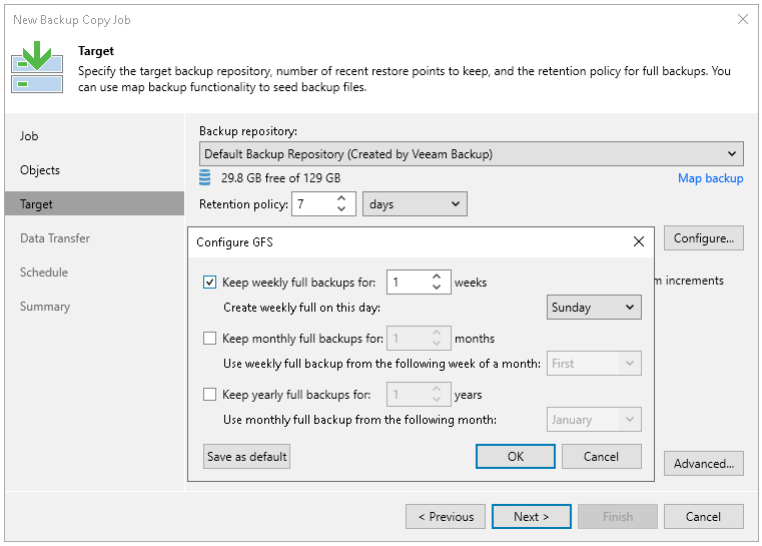

# Step 7. Map Backup File

If the target backup repository already stores a backup of workloads that you want to copy, you can map the backup copy job to this backup.

The backup copy job will use the backup as a seed. As a result, Veeam Backup & Replication will transfer less data over network. For more information, see Backup Copy Job Mapping.

To map the backup copy job to a backup:

1. At the **Target** step of the wizard, click the **Map backup** link.
2. In the **Select Backup** window, select a backup that contains restore points of workloads that you want to copy.

## Limitations for Mapping

When mapping a backup copy job, consider the following limitations:

- If a backup that you plan to use as a seed is encrypted, you must enable encryption for the backup copy job. The password that you use for the backup copy job can differ from the password used for the initial job.
- The following limitations apply to a backup that can be used as a seed:
  - You can map the backup copy job created in Veeam Backup & Replication 11 only to a backup file created in Veeam Backup & Replication 11. In this case, a backup file will not be upgraded to the per-machine backup file.
  - You can map the backup copy job created in Veeam Backup & Replication 12 only to a backup created by a backup copy job. Legacy backup copy job can be mapped to backups created by both backup job and backup copy job.
  - You can map the backup copy job created in Veeam Backup & Replication 12 only to the per-machine backup file. If you map it to the non per-machine backup file, this file will be upgraded to the per-machine backup file. Old version of the backup file will become orphaned.
  - [For the legacy periodic copy mode] If you map the backup copy job to a backup created by a backup job, this backup must be created with the incremental backup method only, that is, forever forward or forward incremental.

- [For the legacy periodic copy mode] You can map a Veeam Agent backup copy job only to a backup created by backup copy job that processes backups created by Veeam Agent operating in the standalone mode.

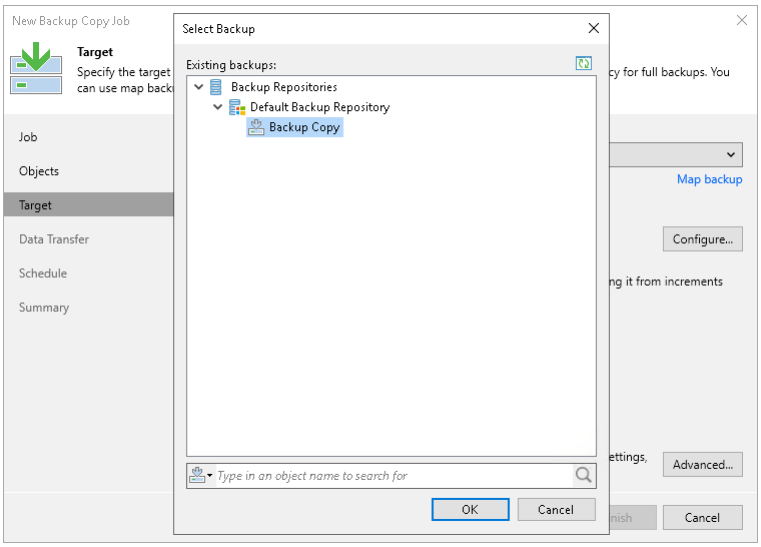

# Step 8. Specify Advanced Settings

At the **Target** step of the wizard, you can specify the following settings for the backup copy job:

- Maintenance settings
- Storage settings
- RPO Monitor settings
- Notification settings
- Script settings
- Advanced settings

> **TIP**
>
> After you specify advanced settings for the backup copy job, you can save them as default settings. For this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new backup copy job, Veeam Backup & Replication will automatically apply the default settings to the new job.

## Maintenance Settings

To specify settings for backup files stored in the target backup repository:

1. At the **Target** step of the wizard, click **Advanced**.
2. If you want to periodically perform a health check of the most recent restore point in the backup chain, select the **Perform backup files health check** check box. To specify the schedule for the health check, click **Configure**.

   By default, the health check is scheduled to start at 5:00 on last Sunday of every month. For more information on the health check, see Health Check for Backup Files.
3. Select the **Remove deleted items data after** check box and specify the retention policy settings for deleted workloads.

   By default, the deleted item retention period is 30 days. It is recommended that you set the retention period to 3 days or more to prevent unwanted data loss. For more information on the retention policy and its limitations, see Deleted Items Retention.
4. To periodically compact a full backup, select the **Defragment and compact full backup file** check box. To specify the schedule for the compacting operation, click **Configure**.

   By default, the compact operation is disabled. For more information on compact of full backup files, see Compact of Full Backup File.

5. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

**IMPORTANT**

When enabling the **Defragment and compact full backup file** option, consider the following:

- The **Defragment and compact full backup file** option can be enabled only if GFS retention policy is disabled.
- The target backup repository must have enough space to store a file of the full backup size. During the compact process, Veeam Backup & Replication creates an auxiliary VBK file that exists in the backup repository until the end of the compact operation.
- [For the legacy periodic backup copy job] If you do not want to copy data for workloads that have only one restore point in the full backup file and this restore point is older than 7 days, check that the following conditions are met: **Remove deleted items data** is disabled; **Use per-machine backup files** is disabled in the settings of the target backup repository.

  Veeam Backup & Replication will extract data for such workloads from the full backup file and write this data to a separate backup file. The file will be displayed under the **Backups > Disk (Imported)** node in the **Home** view.

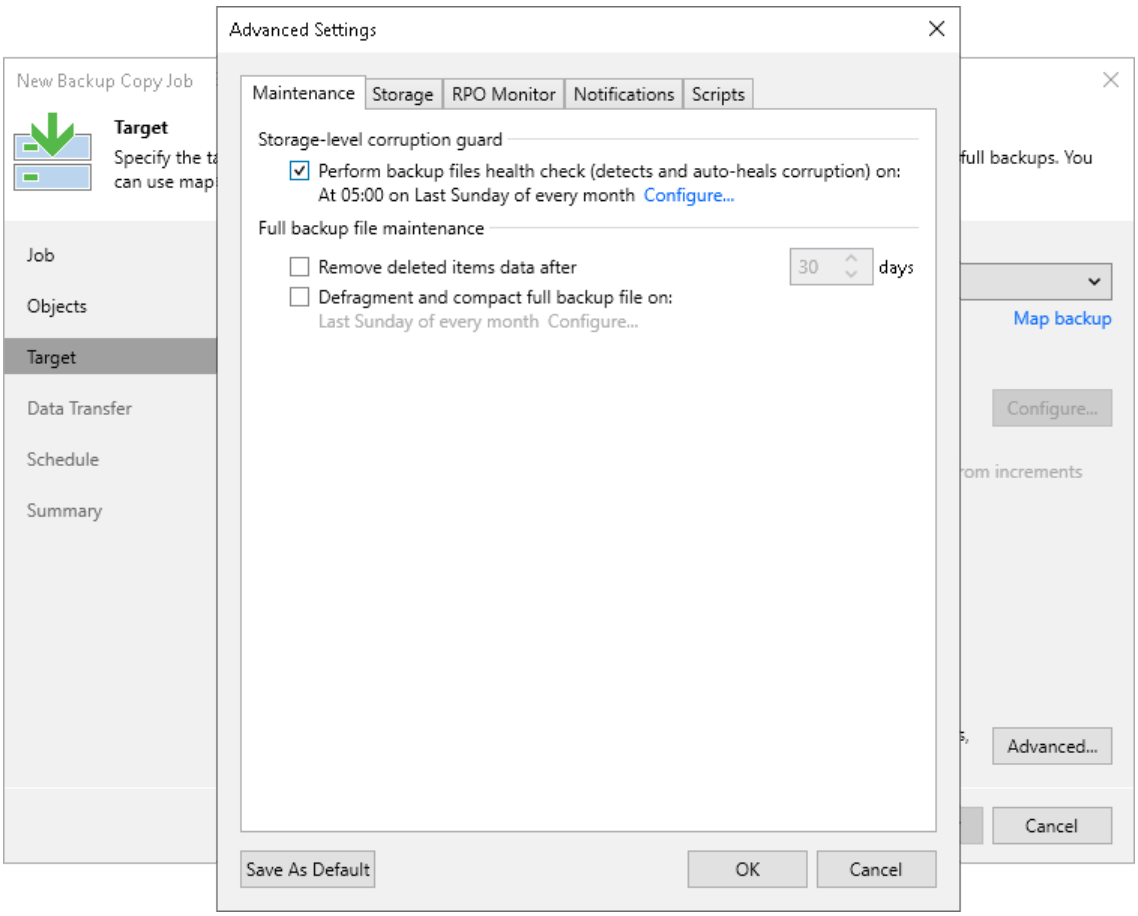


# Storage Settings

To specify compression, deduplication and encryption settings for backup files stored in the target backup repository, do the following:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **Storage** tab.

3. In the **Data reduction** section, specify data compression and deduplication settings:
   - By default, Veeam Backup & Replication performs deduplication before storing copied data in the target backup repository. To disable data deduplication, clear the **Enable inline data deduplication** check box.

     For more information on deduplication, see Deduplication.
   - From the **Compression level** list, select a compression level.

     For more information on data compression levels, see Data Compression.
4. To encrypt backup files, select the **Enable backup file encryption** check box. In the **Password** field, select the encryption method you want to use — a password that will generate secret keys or an external KMS server that will generate KMS keys. For more information about backup copy job encryption, see Encrypting Backup Copy Jobs.

   If you have not created the password beforehand, click **Add** or use the *Manage passwords* link to specify a new password. For more information, see Password Manager.

   If you do not use password loss protection, Veeam Backup & Replication will display a warning about it. For more information, see Password Loss Protection.

5. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

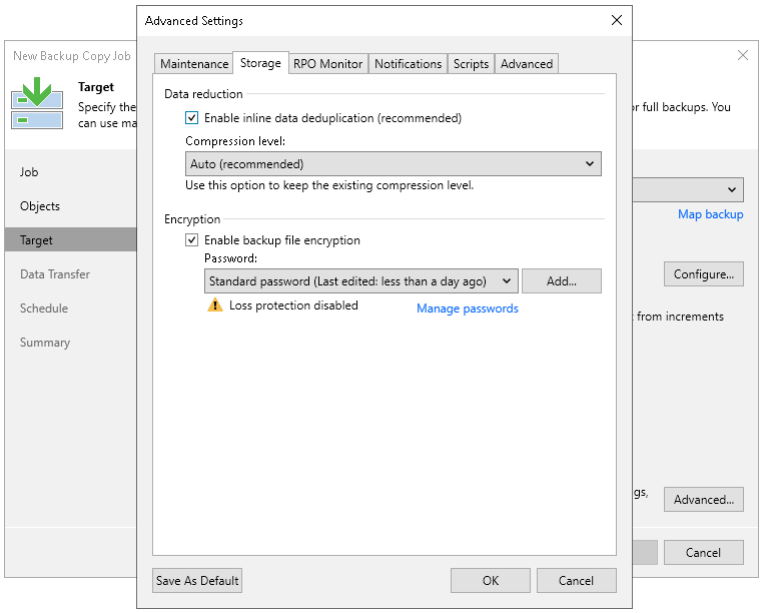


# RPO Monitor Settings

You can instruct a backup copy job to display a warning if a newly created restore point or transaction log is not copied within the desired recovery point objective (RPO). The RPO is counted down from the moment when the source backup job finishes and is ready to be copied.

To mark a job with the *Warning* status when the RPO is exceeded, do the following:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **RPO Monitor** tab.
3. Select the **When new backup is not copied within** check box.

4. In the fields on the right, specify the desired RPO in minutes, hours or days. If you specify days, RPO monitor will consider calendar days instead of the 24 hours period.
5. If you have enabled copying of log backups, select the **When new log backup is not copied within** check box.
6. In the fields on the right, specify the desired RPO in minutes, hours or days.
7. Click **Save as default** if you want to save this set of RPO settings as the default one. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

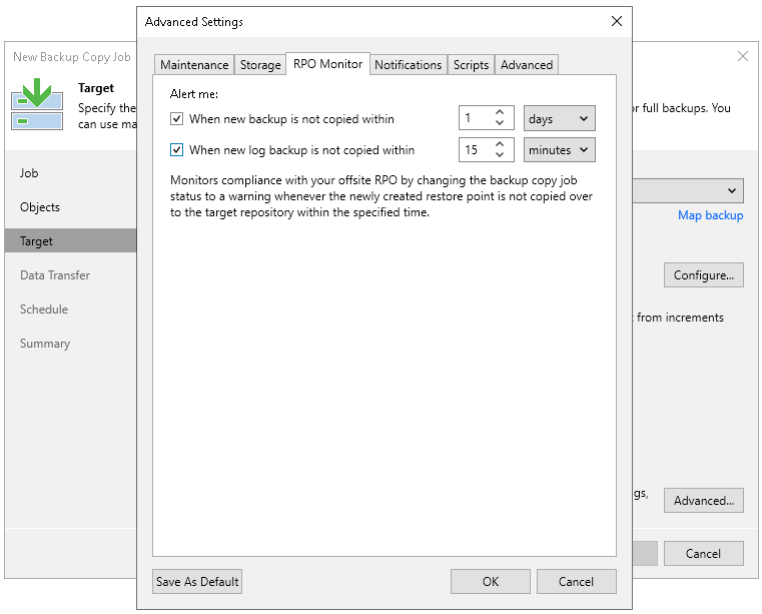


# Notification Settings

To specify notification settings for the backup copy job:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully. In the periodic copy mode, you will receive notifications when the entire backup copy job finishes. In the immediate copy mode — when copying of each source backup job finishes.

   SNMP traps will be sent if you specify global SNMP settings in Veeam Backup & Replication and configure software on the recipient workload to receive SNMP traps. For more information, see Specifying SNMP Settings.
4. Select the **Send email notifications to the following recipients** check box if you want to receive notifications by email in case of job failure or success. In the field under the check box, specify the recipient email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication.

   - Veeam Backup & Replication sends notifications when copying of each source backup job finishes. For example, if your backup copy job contains two source backup jobs, you will receive two emails.

- [For the legacy periodic copy mode] Veeam Backup & Replication sends a consolidated email notification once for the specified backup copy interval. Even if the synchronization process is started several times within the interval, for example, due to job retries, only one email notification will be sent.

For more information on how to configure global notification settings, see Configuring Global Email Notification Settings.

5. In the **Send daily summary at** field, specify when you want to send notifications about backup copy jobs that process log backups. Veeam Backup & Replication sends a consolidated report once a day at the specified time.
6. You can choose to use global notification settings or specify custom notification settings.
   - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see Configuring Global Email Notification Settings.
   - To configure a custom notification for a job, select **Use custom notification settings specified below**. You can specify the following notification settings:
     i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%Time%* (completion time), *%JobName%* (the backup copy job name and source backup job name in the *CopyJobName|SourceJobName* format), *%JobResult%*, *%ObjectCount%* (number of workloads in the job) and *%Issues%* (number of workloads in the job that have been processed with the *Warning* or *Failed* status).
     ii. Select the **Notify on success**, **Notify on warning** and **Notify on error** check boxes to receive email notification if data processing completes successfully, completes with a warning or fails.
7. [For the immediate copy mode] Select when you want to receive notifications, **Immediately after each copied backup** option or **Daily as a summary**.

8. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

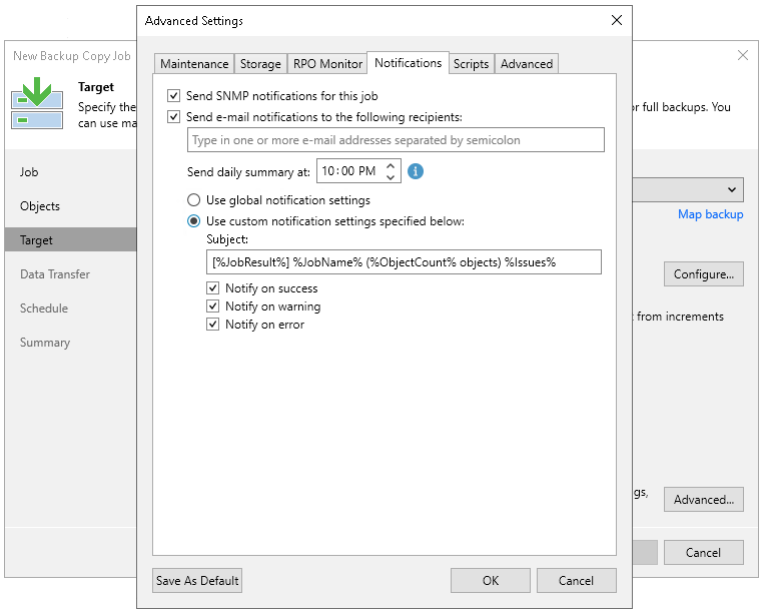

# Scripts Settings

To specify script settings for the backup copy job:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **Scripts** tab.
3. Select the **Run the following script before the job** and **Run the following script after the job** check boxes to execute custom scripts before and after the backup copy job. Note that in the immediate copy mode, scripts are executed for every source backup job.

   Then click **Browse** and select executable files from a local folder on the backup server. The scripts will be executed on the backup server after the transformation processes are completed on the target repository
4. You can change how often the scripts must be executed:
   - To run the scripts after a specific number of backup copy sessions, select **Run scripts every... backup session** option and specify the number of sessions.
   - To run the scripts on specific days, select the **Run scripts on selected days only** option and click the **Days** button to specify week days.

   > **NOTE**
   >
   > If you select the **Run scripts on the selected days only** option, Veeam Backup & Replication executes scripts only once on each selected day — when the job runs for the first time. During subsequent job runs, scripts are not executed.

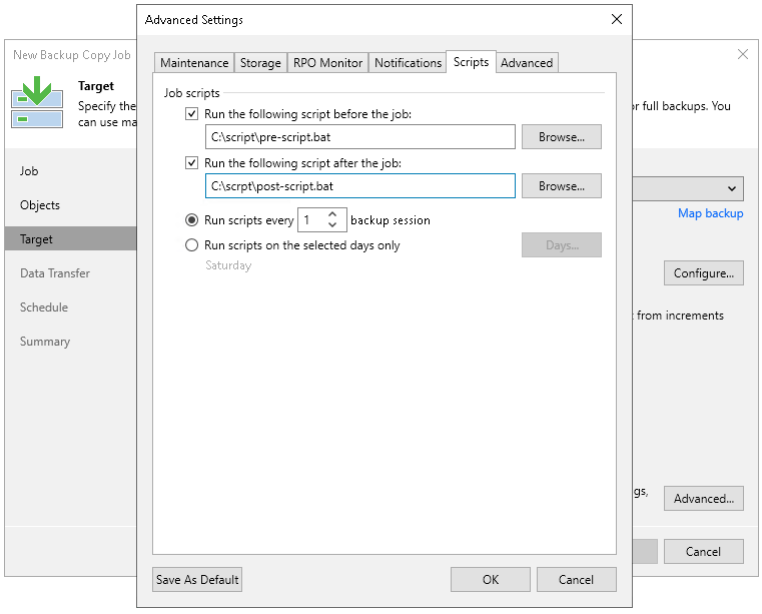


# Advanced

If a new restore point is not yet created, you can process the most recent restore point instead of waiting. To do so, select the **Process the most recent restore point instead of waiting** check box.

> **NOTE**
>
> This option is available only in the periodic backup copy job.

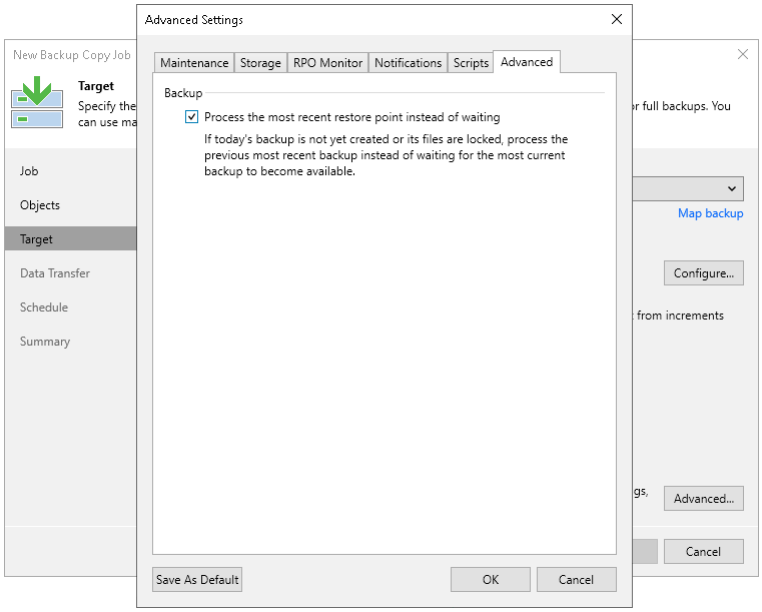
Advanced Settings
Maintenance
Storage
RPO Monitor
Notifications
Scripts
Advanced
Backup
Process the most recent restore point instead of waiting
If today's backup is not yet created or its files are locked, process the previous most recent backup instead of waiting for the most current backup to become available.
Save As Default
OK
Cancel
New Backup Copy Job
Target
Job
Objects
Target
Data Transfer
Schedule
Summary
Map backup
Configure...
Advanced...
Cancel

# Step 9. Specify Data Path Settings

The **Data Transfer** step of the wizard is available only if you copy backups of virtual or physical machines created by Veeam Backup & Replication or Veeam Agents.

At this step of the wizard, you can select how Veeam Backup & Replication will transport backed up data — directly or through WAN accelerators. By default, during the backup copy job Veeam Backup & Replication transports data directly from the source backup repository to target backup repository. This type of transport is recommended if you plan to copy backup files over high-speed connections.

If you plan to copy backup files over WAN or slow connections, it is recommended that you configure source and target WAN accelerators in the backup infrastructure and copy backups through these WAN accelerators. For more information, see WAN Acceleration.

To use WAN acceleration for the backup copy job:

1. At the **Data Transfer** step of the wizard, select the **Through built-in WAN accelerators** option.
2. From the **Source WAN accelerator** list, select a WAN accelerator configured in the source site.
3. From the **Target WAN accelerator** list, select a WAN accelerator configured in the target site.

## Requirements and Limitations for WAN Accelerators

- You must not assign one source WAN accelerator to several backup copy jobs that you plan to run simultaneously.

  The source WAN accelerator requires a lot of CPU and RAM resources and does not process multiple backup copy tasks in parallel. As an alternative, you can create one backup copy job for all work loads you plan to process over one source WAN accelerator. The target WAN accelerator, however, can be assigned to several backup copy jobs.

- [For WAN accelerators with the high bandwidth mode disabled] It is recommended that you pre-populate the global cache on the target WAN accelerator before you start the backup copy job. Global cache population helps reduce the amount of traffic transferred over WAN. For more information, see Manually Populating Global Cache.

- You cannot use WAN accelerators for backup copy jobs that copy backups of Amazon EC2 instances.

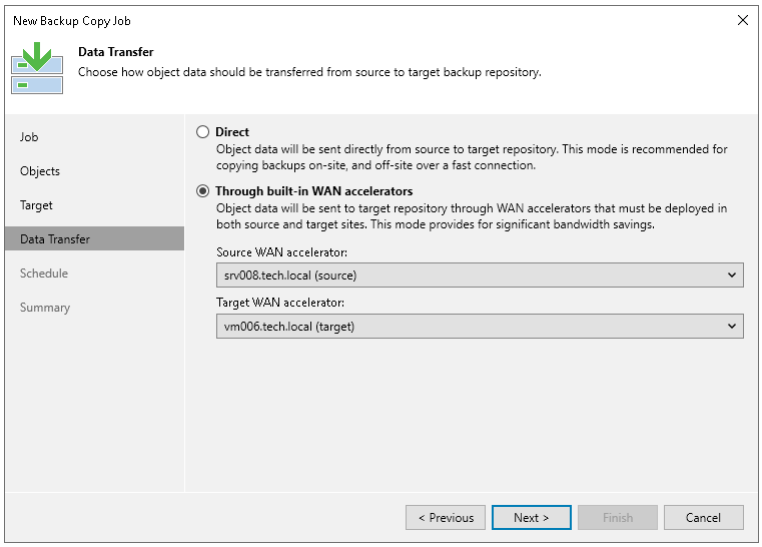

# Step 10. Define Backup Copy Window

At the **Schedule** step of the wizard, you can define a time span in which the backup copy job will transport data between source and target backup repositories. For more information, see Backup Copy Window.

To define a backup window for the periodic backup copy job:

1. Select the **Run the job automatically** check box.
2. Select the required schedule option:
   - **Daily at this time**. You can specify the time and select the days option from the drop-down list on which the backup copy job will run: **Everyday**, **On weekdays** or **On these days**. If you select the **On these days** option, click **Days** to specify them.
   - **Monthly at this time**. You can specify the time and select the day options from the drop-down lists on which the backup copy job will run. Click **Months** to specify months on which the backup copy job will run.
   - **Periodically every**. You can select the time options from the drop-down list or click **Schedule** to select the desired time area. Use the **Permitted** and the **Denied** options to mark the selected time segments. Use the **Start time within an hour** option to specify minutes.
   - **After this job**. If you choose this option, select the job from the drop-down list after which the backup copy job will start.
3. To configure the backup copy job automatic retries, select the **Retry failed items processing** check box and specify the amount of retries and time intervals between them.
4. Select the **Terminate job if it exceeds allowed backup window** check box if you want the job to terminate itself to prevent performance impact during production hours. Click **Window** to select the desired backup window time area. Use the **Permitted** and **Denied** options to mark the selected time segments.

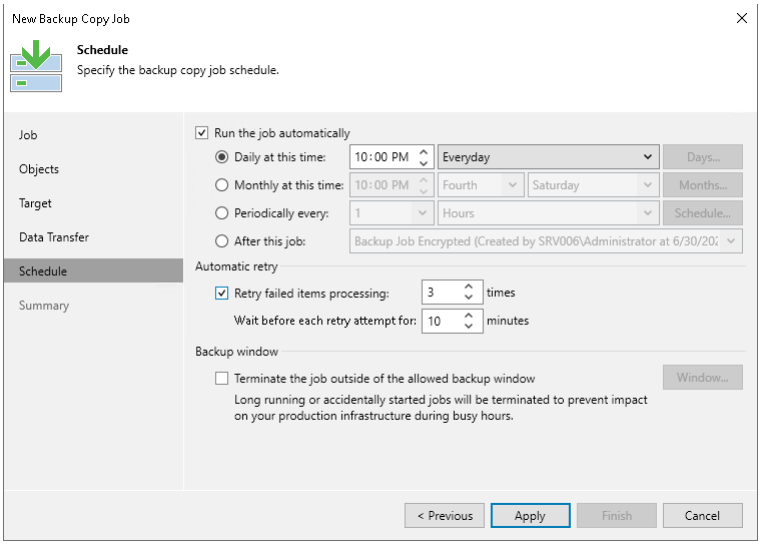


To define a backup window for the immediate backup copy job:

1. Select when the job can transfer data:
   - **Any time**
   - **During the following time periods only**

2. If you selected the **During the following time periods only** option, specify the required backup window option. Use the **Enable** and **Disable** options to mark the selected time segments as allowed or prohibited for the backup copy job.

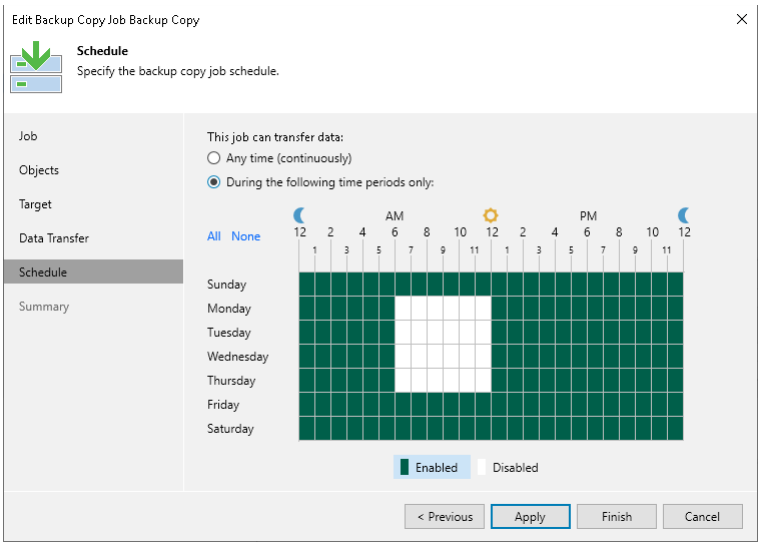

# Step 11. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of backup copy job configuration:

1. Review details of the backup copy job.
2. Select the **Enable the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

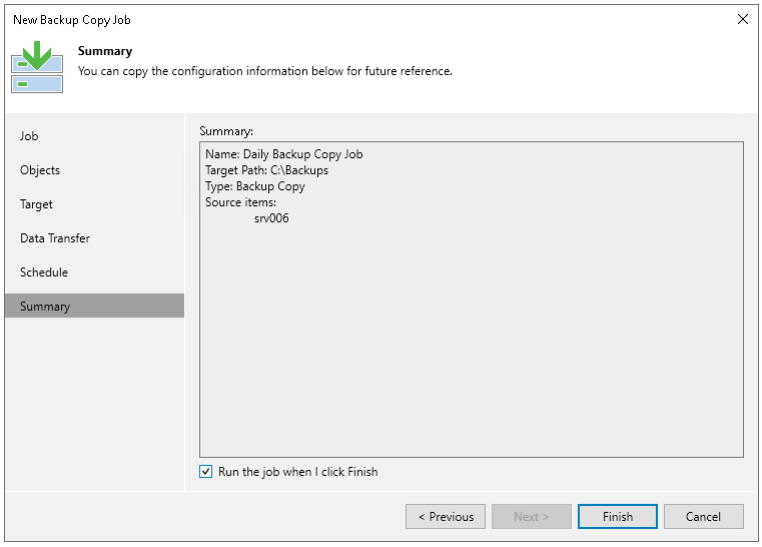

# Creating Backup Copy Jobs for HPE StoreOnce Repositories

To copy backup files between HPE StoreOnce backup repositories, you must configure a backup copy job for HPE StoreOnce repositories.

Unlike other backup copy jobs, the backup copy job for HPE StoreOnce mirrors data from the source repository. This backup copy job copies backup files as they are stored in the source repository, without any transformation. To copy the backup files, Veeam Backup & Replication uses the HPE StoreOnce Catalyst Copy technology.

The backup copy job copies only backup files created by backup jobs and other backup copy jobs. The backup files must be of the following types:

- Backup files of VMware vSphere and Microsoft Hyper-V VMs created by Veeam Backup & Replication. Log backup files are not copied.
- Physical machine backup files created by Veeam Agent backup jobs managed by the backup server.
- Backup files of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV.
- Backup files of oVirt VMs created by Veeam Backup for OLVM and RHV.
- Backups copied to an HPE StoreOnce repository by other backup copy jobs (regular backup copy jobs and backup copy jobs for HPE StoreOnce repositories).

> **IMPORTANT**
>
> Consider the following limitations:
>
> - The backup jobs must be configured on the same backup server where you configure the backup copy job for HPE StoreOnce repositories. Backups created by jobs configured on other backup servers are not copied.
> - The backup copy jobs for HPE StoreOnce repositories do not copy imported backups.
> - To copy backups created by other backup copy jobs (regular backup copy jobs), you must enable the GFS retention for these backup copy jobs. For more information on how to enable the GFS retention, see Specify Target Repository and Retention Settings.

When the backup copy job for HPE StoreOnce runs for the first time, it copies all existing backup files. Then the backup copy job starts each time a new backup file appears in the source repository. In case of a removed backup file, the backup copy job waits 21 days since the backup file creation and after removes the backup file from the target repository. If 21 days have already passed at the moment of removal, the backup copy job removes the backup file immediately. You can change this day limit in the backup copy job settings. For more information, see Maintenance Settings.

Before creating a job, check prerequisites and limitations. Then use the **New Backup Copy Job** wizard to configure the backup copy job.

# Before You Begin

Before you create a backup copy job for HPE StoreOnce backup repository, check the following requirements:

- The minimum supported software versions of HPE StoreOnce are the following:
  - For the third generation, the minimum version is 3.18.18.
  - For the fourth generation, the minimum version is 4.2.3.
- Make sure that all backup infrastructure components that take part in the backup copy process are added to the backup infrastructure. These components include the source and target repositories between which data is copied. For more information on how to add a backup repository, see Backup Repositories.
- Check that repositories between which you plan to copy data have a direct connection to each other. This is required because Veeam Backup & Replication uses the HPE StoreOnce Catalyst Copy technology to copy backup files.

  This direct connection must be of the same type as the connection that you select when adding the target HPE StoreOnce. For example, if you connected the target HPE StoreOnce repository over Fibre Channel, you must connect the source HPE StoreOnce to the target HPE StoreOnce over Fibre Channel.
- HPE StoreOnce repositories connected over Fibre Channel require the two-way connection. Zone the source initiator World Wide Names (WWNs) with the destination target WWNs, and zone the destination initiator WWNs with the source target WWNs.
- If you plan to use pre-job and post-job scripts, you must create scripts before you configure the backup copy job.
- If source backup repository has backup immutability disabled while the target repository is immutable, backup copy job will work in non-immutable mode.
- If source and target HPE StoreOnce backup repositories are configured with different block chunking algorithm, backup copy job will copy the data without changing the block sizes.
- If source backup job has GFS retention policy configured, its GFS immutability settings will be applied to the backup files copied to the target HPE StoreOnce backup repository.

# Step 1. Launch New Backup Copy Job Wizard

To run the **New Backup Copy Job** wizard, do one of the following:

- On the **Home** tab, click **Backup Copy** and click **Storage copy**.
- Open the **Home** view, in the inventory pane right-click **Jobs** or right-click anywhere in the working area, click **Backup Copy** and click **Storage copy**.

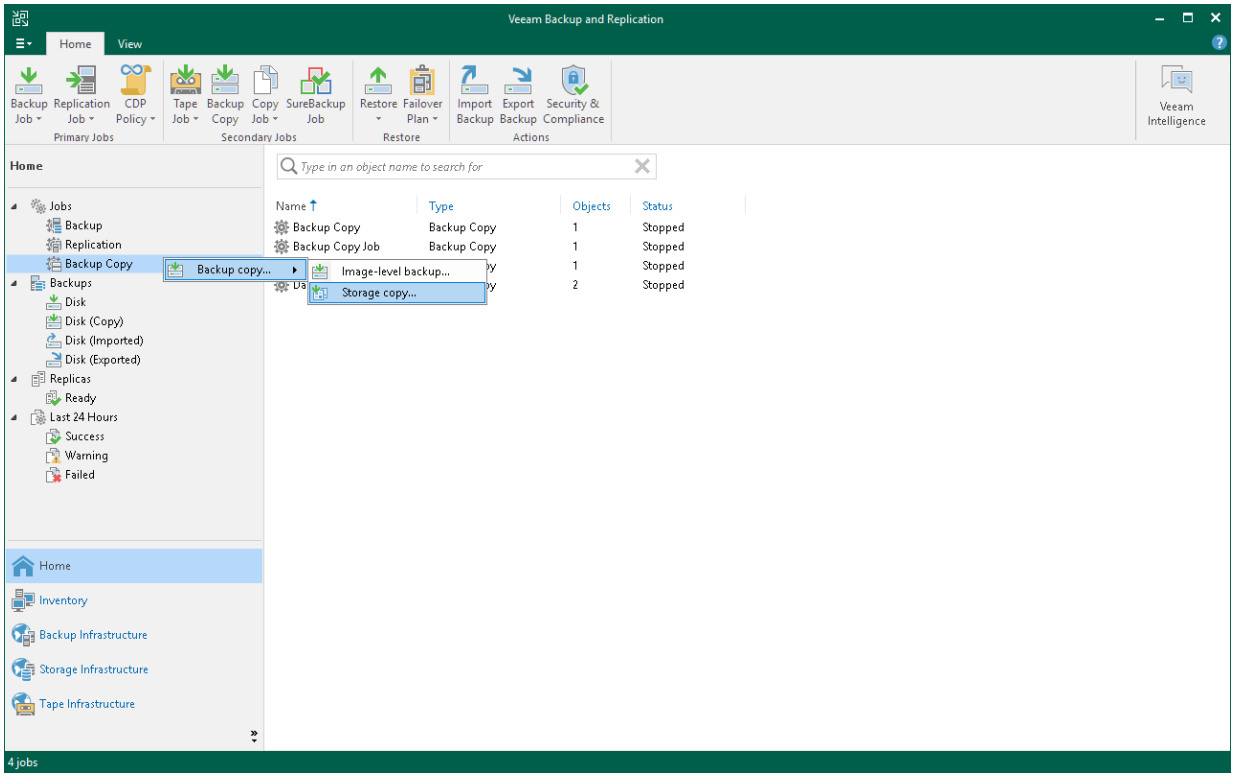

# Step 2. Specify Job Name and Description

At the **Job** step of the wizard, specify name and description for the backup copy job:

1. In the **Name** field, enter a name for the job.
2. In the **Description** field, enter a description for the job. The default description contains information about the user who created the job, date and time when the job was created.

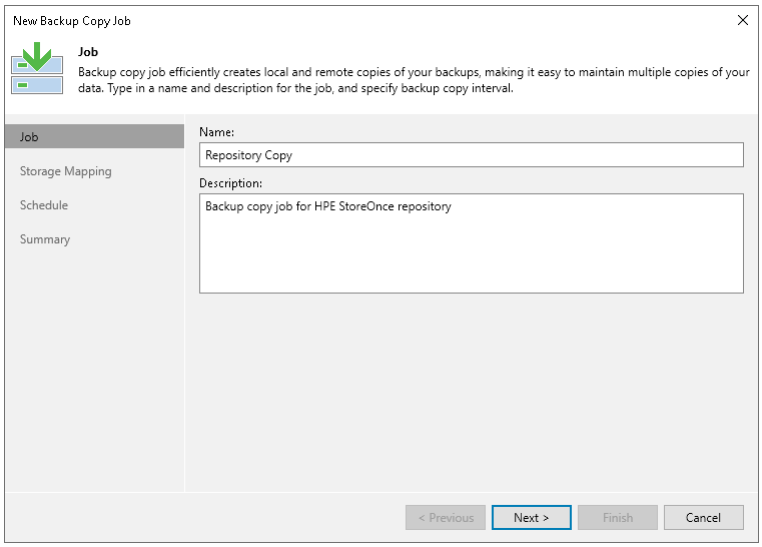

# Step 3. Select Source and Target Repositories

At the **Storage Mapping** step of the wizard, select a source repository from which you want to copy backups and a target repository where you want to store the copies.

1. Click **Add** to open the **Add Repository** window.
2. From the **Source backup repository** list, select a backup repository from which you want to copy backup files. The unsupported backup repositories are not shown in the list.
3. From the **Target backup repository** list, select a backup repository where you want to store the copies. The unsupported backup repositories are not shown in the list.

> **IMPORTANT**
>
> Veeam Backup & Replication does not copy all types of backups. For the list of supported backup types, see Creating Backup Copy Jobs for HPE StoreOnce Repositories.

## Requirements and Limitations for Source and Target Repositories

When you create backup copy jobs, check the following prerequisites and limitations for the source and target repositories:

- The source and target backup repositories must be HPE StoreOnce backup repositories or scale-out backup repositories that consist of HPE StoreOnce repositories only.
- [For scale-out backup repositories] The file placement policy must be Data locality. For more information, see Data locality.
- Within one backup copy job, you can use each repository as a source only once.
- You must not create the same pairs of source and target repositories, even in different backup copy jobs.

# Requirements for Data Flow

When you create backup copy jobs, check that you do not create loops in data flow across all backup copy jobs. This means that data copied from one repository must not be copied to it again. The following image shows backup copy jobs configured correctly (without loops) and incorrectly (with loops).

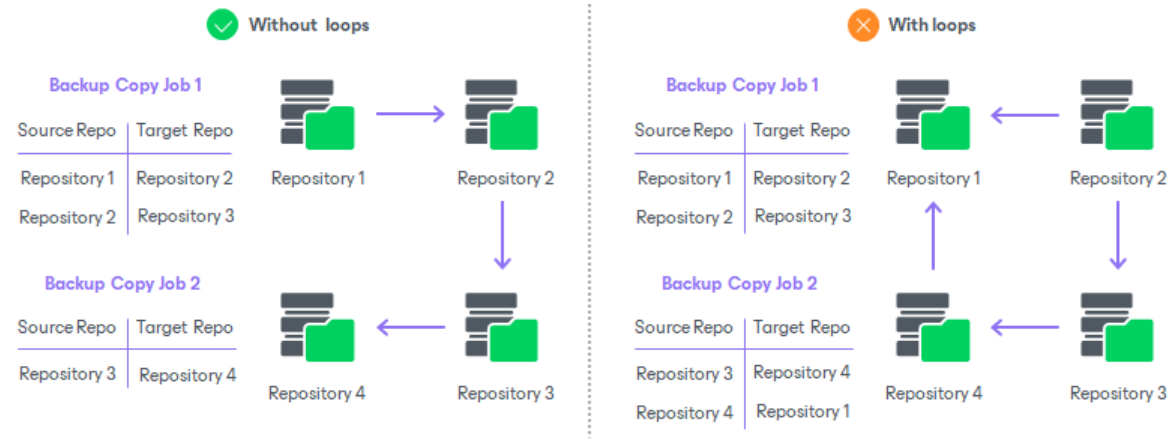


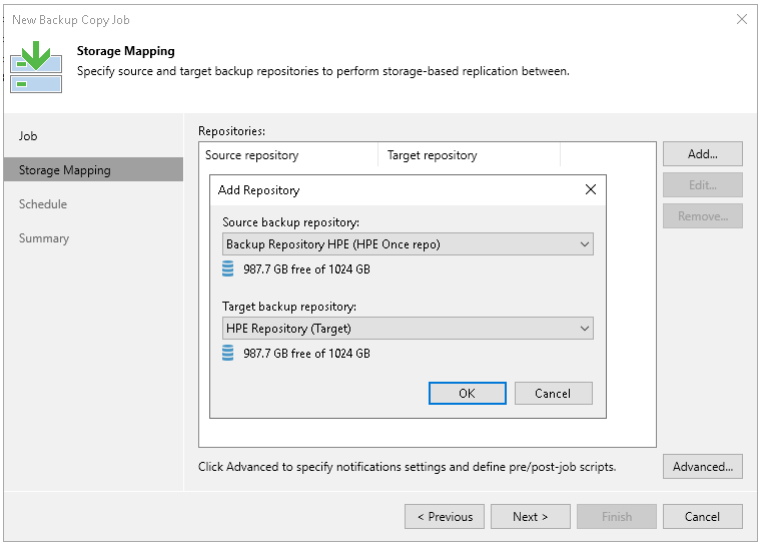

# Step 4. Specify Advanced Settings

At the **Storage Mapping** step of the wizard, you can specify the following settings for the backup copy job:

- Maintenance Settings
- Notification Settings
- Script Settings

> **TIP**
>
> After you specify necessary settings for the backup copy job, you can save them as default settings. To do this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new backup copy job, Veeam Backup & Replication will automatically apply the default settings to the new job.

## Maintenance Settings

In the maintenance settings, you can configure whether to perform a health check and after which period delete from the target HPE StoreOnce repository files deleted from the source repository. Note that the health check may lower the performance of the target repository. For details, see Health Check for Backup Files.

1. At the **Storage Mapping** step of the wizard, click **Advanced**.
2. To disable the health check, clear the **Perform backup files health check** check box. By default the health check is enabled.
3. In the **Keep secondary copies for additional** field, specify after which period of time Veeam Backup & Replication will delete files from the target repository once they have been deleted from the source repository.

   The backup copy job waits the specified number of days since the backup file creation on the source repository and after deletes the backup files from the target repository. If the specified number of days has already passed at the moment of deletion, the backup copy job deletes the backup file immediately.

   For example, if the option is set to 30 days and the backup file on the source repository is deleted after 15 days, Veeam Backup & Replication will keep the backup file on the target repository for an additional 15 days.

4. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

> **IMPORTANT**
>
> Veeam Backup & Replication does not perform the health check for encrypted and compressed backup files.

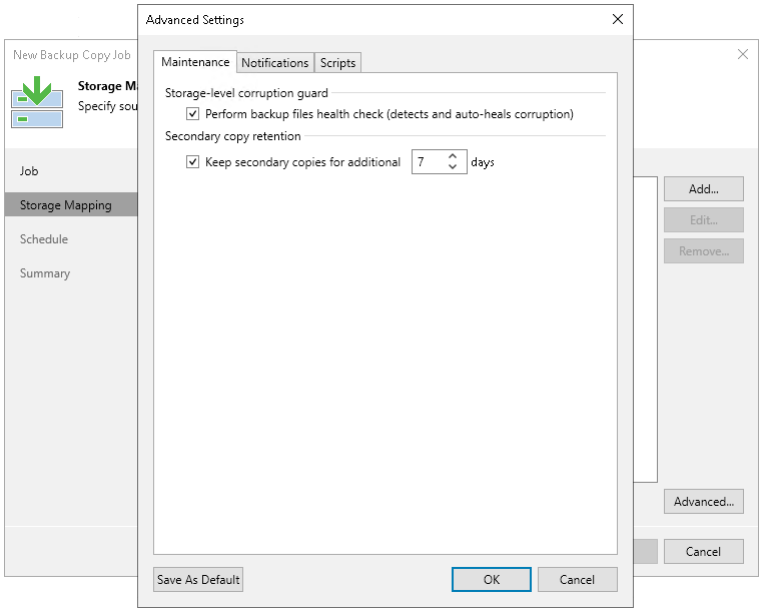


# Notification Settings

To specify notification settings for the backup copy job:

1. At the **Storage Mapping** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when data from each source repository is copied.

   SNMP traps will be sent if you specify global SNMP settings in Veeam Backup & Replication and configure software on the recipient machine to receive SNMP traps. For more information, see Specifying SNMP Settings.
4. Select the **Send email notifications to the following recipients** check box if you want to receive email notifications when copying data of each source repository finishes with *Success*, *Warning* or *Failed* status.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication. For more information, see Configuring Global Email Notification Settings.
5. You can choose whether to use global notification settings or specify custom notification settings.
   - To receive typical notifications for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server.
   - To configure custom notifications, select **Use custom notification settings specified below**. You can specify the following notification settings:
     i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%Time%* (completion time), *%JobName%*, *%JobResult%*, *%VmCount%* (number of machines in the job) and *%Issues%* (number of machines in the job that have been processed with the *Warning* or *Failed* status).

     ii. Select the **Notify on success**, **Notify on warning** and **Notify on error** check boxes to receive email notifications if data processing completes successfully, fails or completes with a warning.

6. If you want to save this set of settings as the default one, click **Save as default**. When you create a new job, the saved settings will be offered as the default. This also applies to all users added to the backup server.

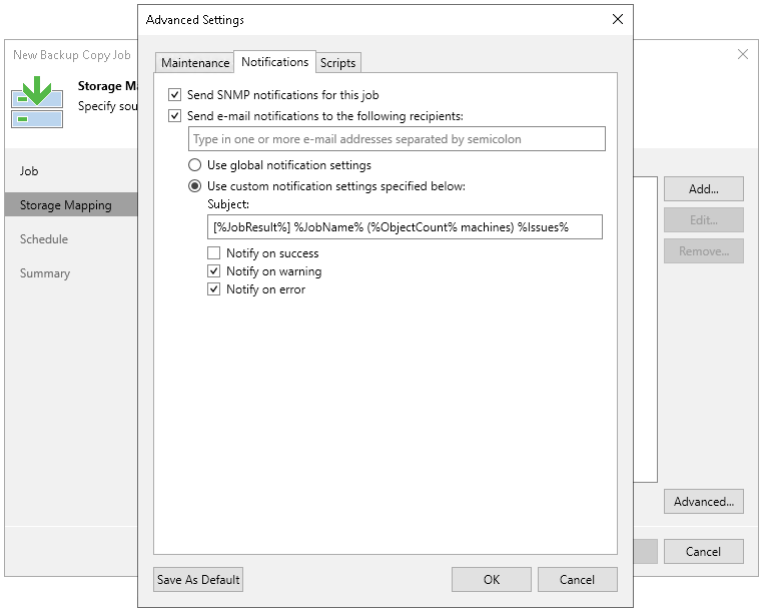


# Script Settings

To specify script settings for the backup copy job:

1. At the **Storage Mapping** step of the wizard, click **Advanced**.
2. Click the **Scripts** tab.
3. Select the **Run the following script before the job** and **Run the following script after the job** check boxes to execute custom scripts before and after copying data of each source repository finishes.

   Then click **Browse** and select executable files from a local folder on the backup server. The scripts are executed on the backup server after the transformation processes are completed on the target repository.
4. You can change how often the scripts must be executed:
   - To run scripts after a specific number of backup copy sessions, select **Run scripts every... backup session** option and specify the number of sessions.
   - To run scripts on specific days, select the **Run scripts on selected days only** option and click the **Days** button to specify week days.

     > **NOTE**
     >
     > If you select the **Run scripts on the selected days only** option, Veeam Backup & Replication executes scripts only once on each selected day — when the job runs for the first time. During subsequent job runs, scripts are not executed.

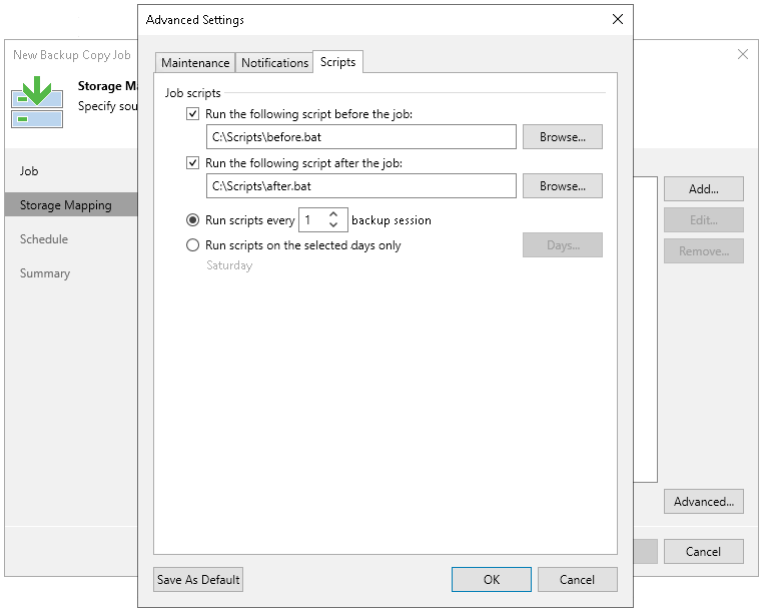
New Backup Copy Job
Storage M
Specify sou
Job
Storage Mapping
Schedule
Summary
Add...
Edit...
Remove...
Advanced...
Cancel
Advanced Settings
Maintenance
Notifications
Scripts
Job scripts
Run the following script before the job:
C:\Scripts\before.bat
Browse...
Run the following script after the job:
C:\Scripts\after.bat
Browse...
Run scripts every 1 backup session
Run scripts on the selected days only
Days...
Saturday
Save As Default
OK
Cancel

# Step 5. Define Backup Copy Window

At the **Schedule** step of the wizard, you can define a time span in which the backup copy job will transport data between source and target backup repositories. For more information, see Backup Copy Window.

To define a 'prohibited' period for the backup copy job:

1. Select the **During the following time periods only** option.
2. In the schedule box, select the desired time area.
3. Use the **Enabled** and **Disabled** options to mark the selected time segments as allowed or prohibited for the backup copy job.

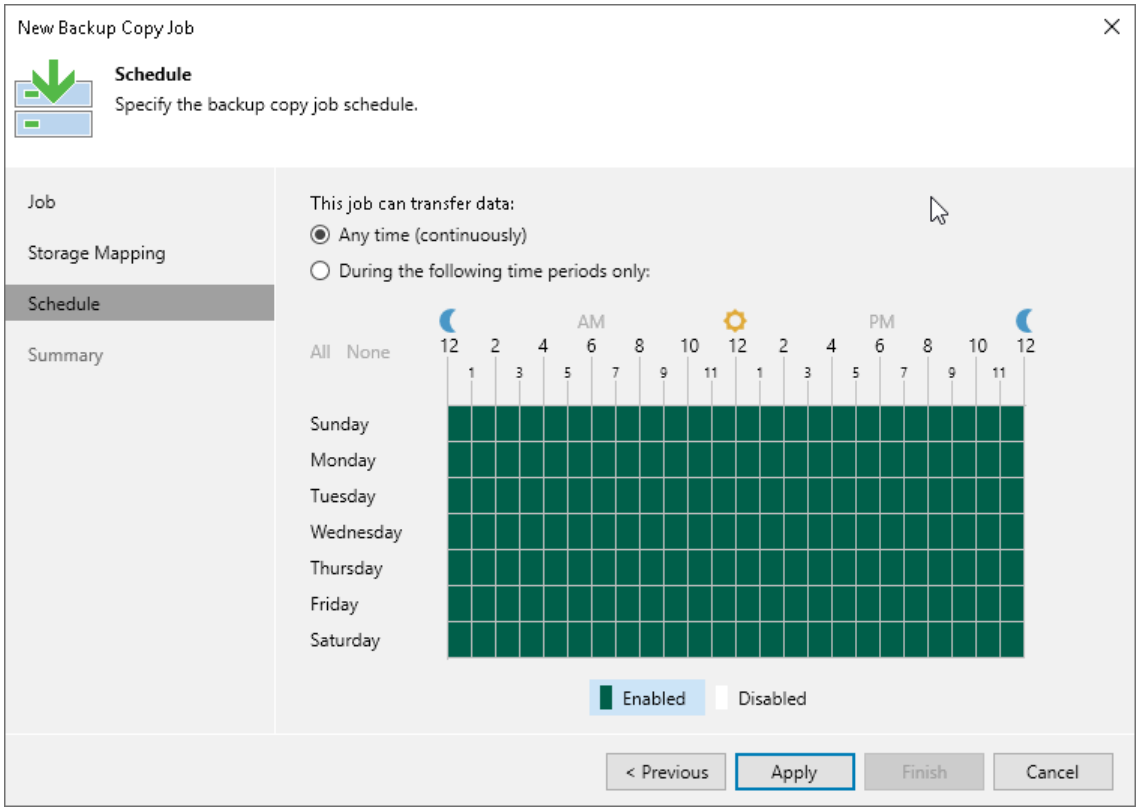

# Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of backup copy job configuration:

1. Review details of the backup copy job.
2. Select the **Enable the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

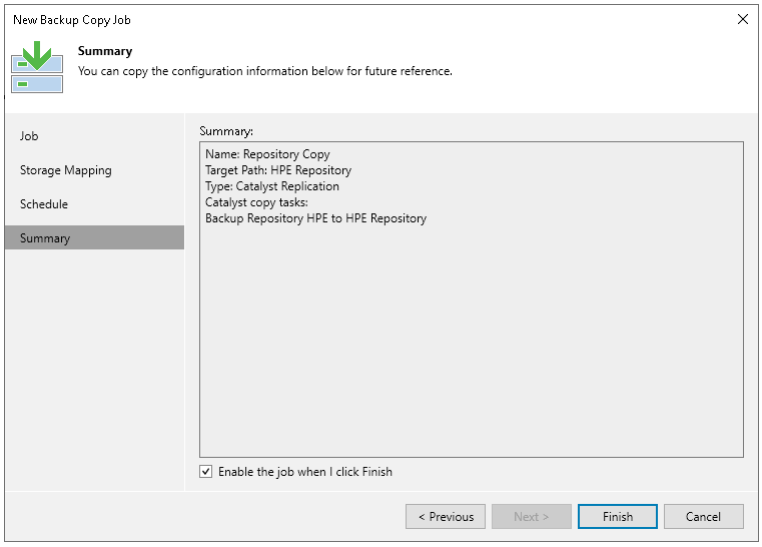

# Creating Backup Copy Jobs for Veeam Plug-ins

To create copies of Veeam Plug-in backups of Oracle, SAP HANA and Microsoft SQL Server databases, you must configure a backup copy job. For more details, see the following sections of the Veeam Plug-ins for Enterprise Applications Guide:

- Backup Copy for Oracle RMAN Backups
- Backup Copy for SAP on Oracle Backups
- Backup Copy for SAP HANA Backups
- Backup Copy for Microsoft SQL Server Backups

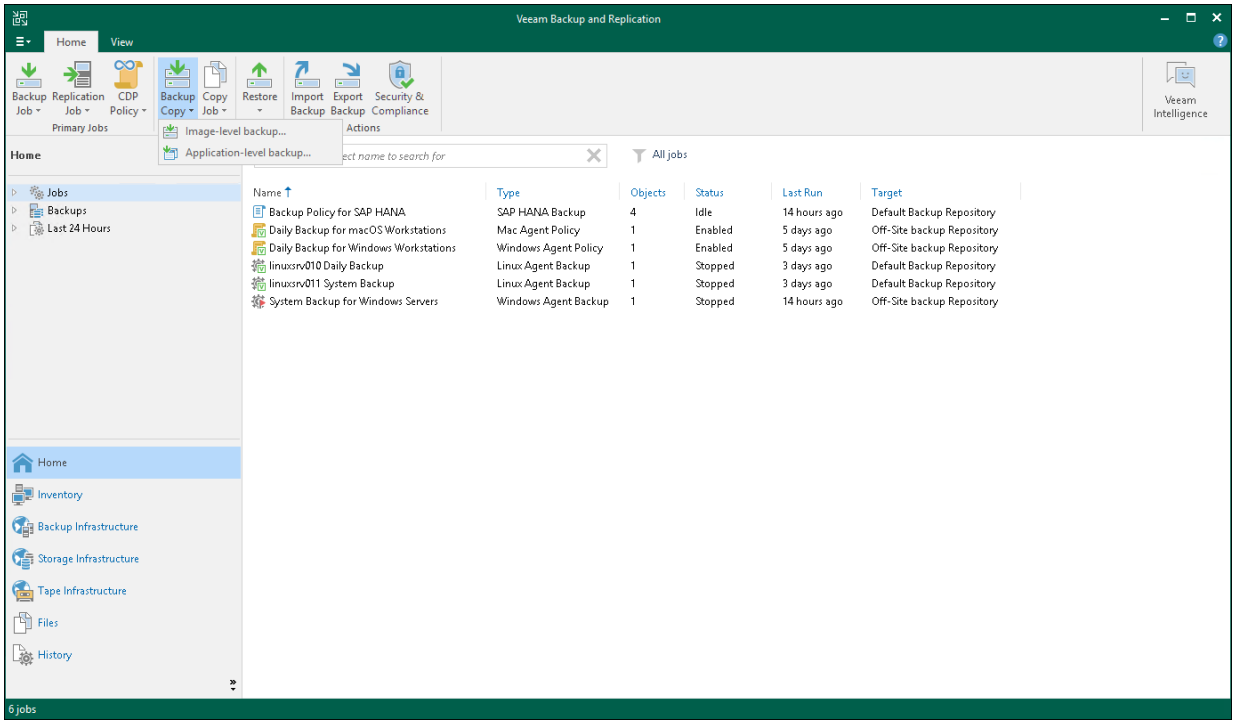

# Encrypting Backup Copy Jobs

The encryption process for a backup job includes the following steps:

1. Job encryption is configured in the advanced job settings. When you enable the encryption option, you need to enter a password or select a KMS server.
2. If you enter a password, Veeam Backup & Replication generates a secret key. If you select a KMS server, Veeam Backup & Replication asks the KMS server to generate an asymmetric KMS key. A secret key or a KMS key will be used to encrypt data encryption keys during the job. For more information, see How Data Encryption Works.
3. When you start a job, Veeam Backup & Replication encrypts backup files with data encryption keys. By default, this operation is performed on the source backup repository.
4. Encrypted data is passed to the target backup repository and stored to a resulting backup file.

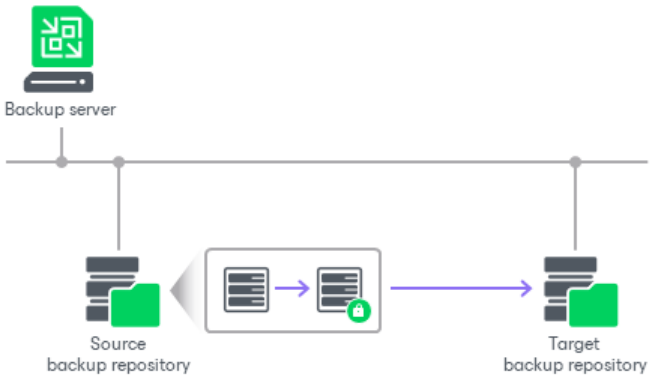


**NOTE**

If you plan to store backups in Veeam Data Cloud Vault, you must enable job encryption.

If you use WAN accelerators for backup copy jobs, the encryption process is performed on the target WAN accelerator.

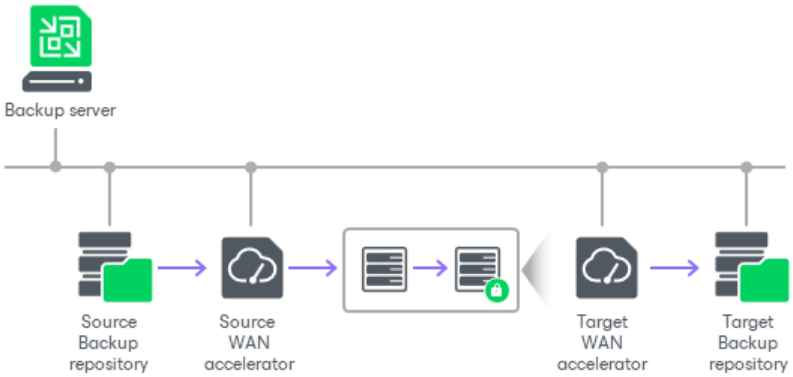

An encrypted backup copy job may use an encrypted backup file as a source. In this situation, Veeam Backup & Replication decrypts encrypted source backup file, encrypts it again and and transfers to the target backup repository. Even if encryption is disabled in the backup copy job, Veeam Backup & Replication will decrypt data blocks of the encrypted source backup files.

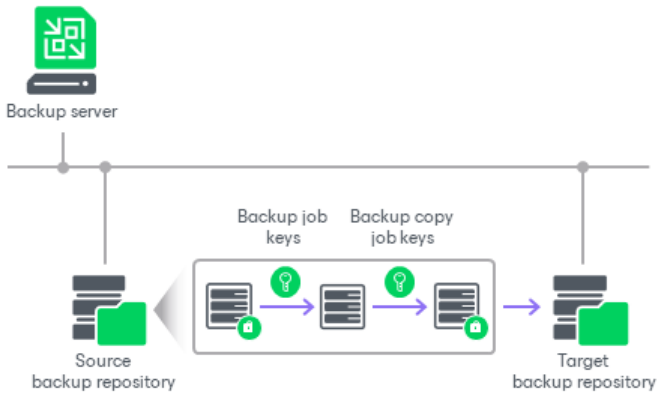


If you update encryption settings for an existing backup copy job, consider the following:

- If you enable encryption, during the next job session Veeam Backup & Replication will automatically create a full backup file. The created full backup file and subsequent incremental backup files in the backup chain will be encrypted with the specified password or KMS key.

  Note that Veeam Backup & Replication does not encrypt the previous backup chain created by this job. If you want to start a new chain so that the unencrypted previous chain can be separated from the encrypted new chain, follow this Veeam KB article.

- If you change the password or start using KMS keys for the already encrypted job, during the next job session Veeam Backup & Replication will create a new incremental backup file. The created backup file and subsequent backup files in the backup chain will be encrypted with the new password or KMS key.

- If you disable encryption, during the next job session Veeam Backup & Replication will automatically create a full backup file.

# Linking Backup Jobs to Backup Copy Jobs

You can link backup jobs to backup copy jobs. This option lets you create a secondary target for the backup job and store backups created by the backup job in the secondary backup repository.

When you link a backup job to the backup copy job, Veeam Backup & Replication automatically updates properties of the backup copy job and adds to it the backup job as a source of data. During every backup copy session, the backup copy job checks the source backup repository for new restore points. As soon as a backup job session is finished and a new restore point appears in the source backup repository, the backup copy job automatically copies this restore point to the target backup repository.

You can link a backup job to an existing backup copy job using the **Backup Job** wizard. To link jobs:

1. Open the backup job settings for editing. For more information, see Editing Job Settings.
2. Navigate to the **Storage** step.
3. Select the **Configure secondary destination for this job** check box.

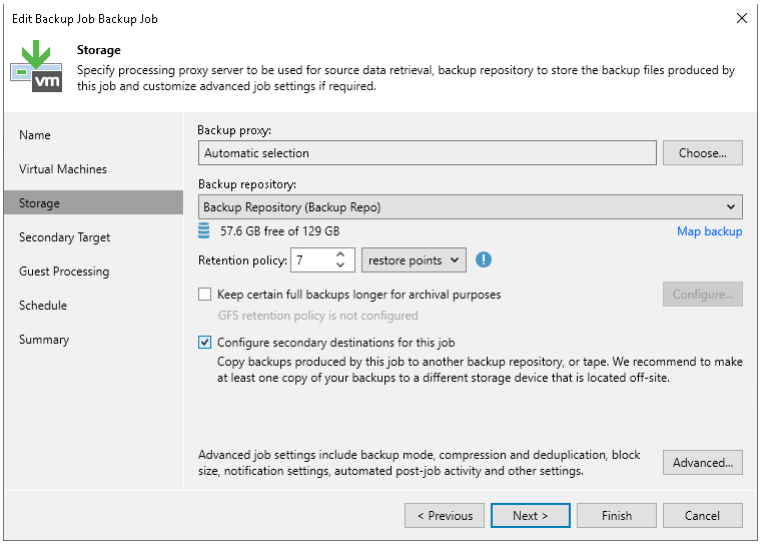


4. At the **Secondary Target** step of the wizard, click **Add** and choose a backup copy job to which the backup job must be linked. The backup copy job must be already configured on the backup server.

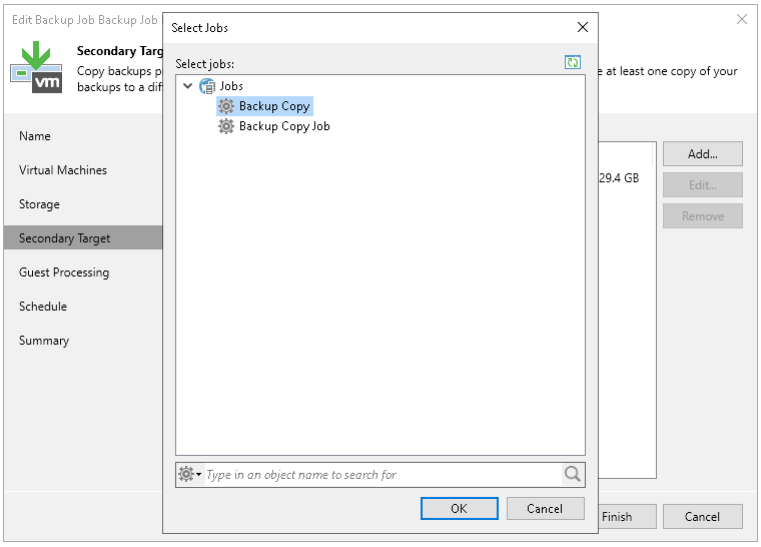

# Managing Backups

To view all backups created by backup copy jobs, open the **Home** view and select the **Backups** > **Disk (Copy)** node in the inventory pane. The list of available backups is displayed in the working area. You can view backup properties, remove unnecessary backups and remove missing restore points.

# Viewing Backup Properties

You can view summary information about backups created by backup copy jobs. The summary information provides the following data: available restore points, date of restore points creation, compression and deduplication ratios, data size and backup size.

> **NOTE**
>
> **Data Size** represents the size of a backup file before compression and deduplication. **Backup Size** represents the size of a backup file after compression and deduplication.

In the summary information, Veeam Backup & Replication displays data about restore points created by the short-term retention scheme and archive restore points created by the GFS retention scheme (if GFS retention is enabled). Archive restore points are marked with the following letters:

- R — full backups created with the short-term retention scheme or active full backups
- W — weekly backups
- M — monthly backups
- Y — yearly backups

In the summary information, you can also see the following icons:

| Icon | State |
|---|---|
| | Full restore point |
| | Incremental restore point |
| | Reverse incremental restore point |
| | Missing full restore point |
| | Missing incremental restore point |

To view summary information for a backup copy:

1. Open the **Home** view.
2. In the inventory pane, select **Backups** > **Disk (copy)**.

3. In the working area, right-click the backup copy and select **Properties**.

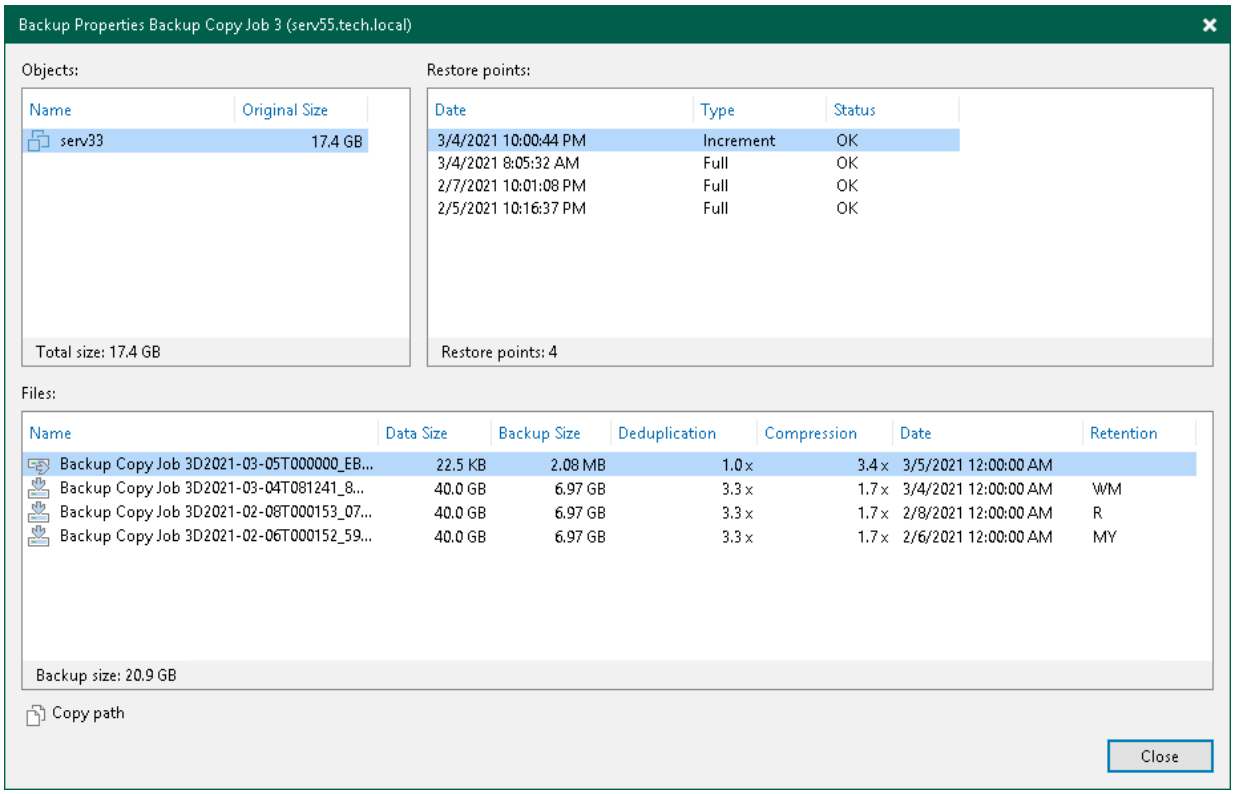

# Upgrading Backup Chain Formats

Veeam Backup & Replication supports the following ways to store backup files: per-machine backup with separate metadata files, per-machine backup with single metadata file and single-file backup. For more information, see Backup Chain Formats.

Changing the **Use per-machine backup files** option for existing backup repositories or moving backups to other repositories does not change the backup chain format. To change the format, follow the instructions in this section.

> **NOTE**
>
> Consider the following:
>
> - Before upgrading the backup chain format of backup copy backups, upgrade the backup chain format of source backups first.
> - [For immediate copy mode] If the backup chain is mixed, consists of single-file backups and per-machine backups with single metadata file, and the last part is per-machine backups with single metadata file, you cannot change the backup chain format. You need to wait until the backup chain becomes per-machine.
> - Before upgrading the backup chain format, Veeam Backup & Replication disables the job to which the backups belong. After the upgrade, the job stays disabled, you need to enable it manually.
> - You must disable log backup jobs manually before changing the format.
> - [For format upgrade using mapping] During the upgrade, Veeam Backup & Replication synthesizes full backups in the per-machine backup with separate metadata files format using the existing backup (the source for mapping). That is why you need enough space on the repository to store full backups. After the successful change operation, the backup used as the source for mapping is placed to the node with the (Orphaned) postfix. After the successful upgrade, you can delete this orphaned backup.
> - [For periodic copy mode] You can not upgrade the backup chain format of Windows agent failover cluster backups.
> - You can stop an upgrade session only for the backups for which the session has not started. If the session has started, Veeam Backup & Replication continues the upgrade to avoid backup corruption.
> - You cannot upgrade the backup chain format for backups stored on repositories with rotated drives.

## Immediate Copy Mode: Per-Machine Backup with Single Metadata File to Per-Machine with Separate Metadata Files

To upgrade the backup chain format from per-machine backup with single metadata file to per-machine backup with separate metadata files, use the upgrade functionality:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** node.
3. In the working area, select the necessary backup copy jobs. Note that the selected jobs must be targeted to the same repository.
4. Right-click one of the selected jobs and click **Upgrade backup chain format**. Alternatively, click **Upgrade Backup Chain Format** on the ribbon.

Veeam Backup & Replication will generate new metadata files for the existing backups. After the upgrade, the job will continue the backup chain and will create per-machine backups with separate metadata files.

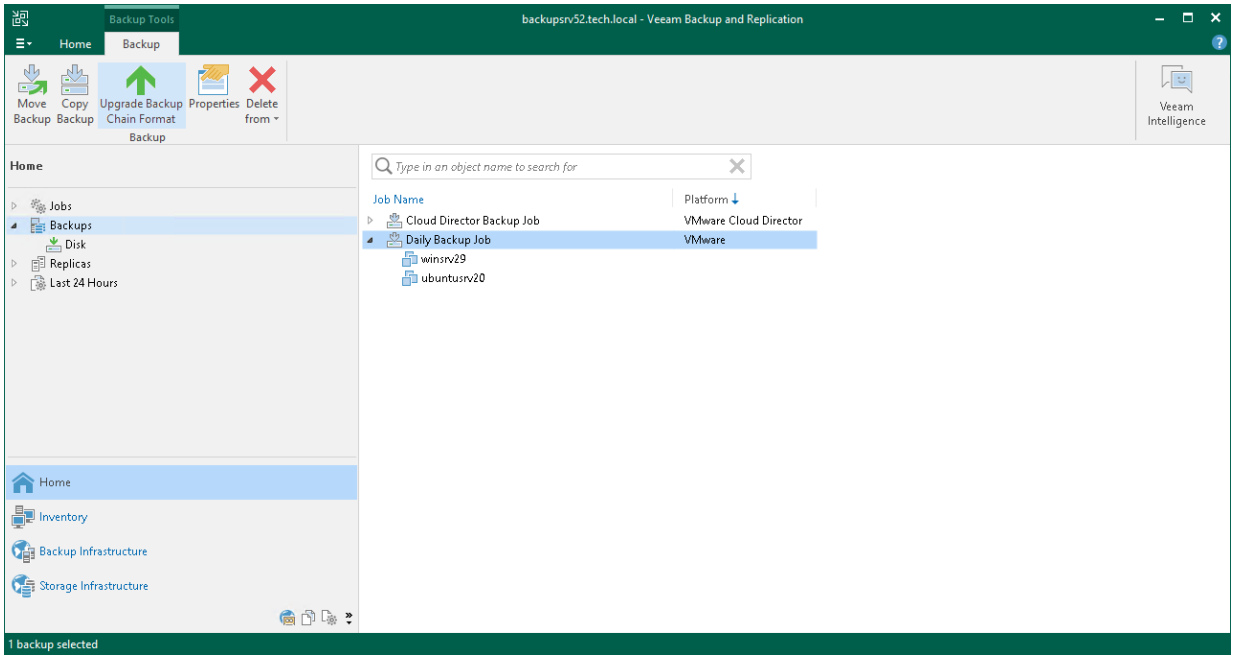


# Immediate Copy Mode: Single-File Backup to Per-Machine with Separate Metadata Files

To upgrade the backup chain format from single-file backup to per-machine backup with separate metadata files, do the following:

1. Detach backups from the backup copy job for whose backups you want to change the backup chain format. For more information, see Detaching Backups from Jobs.
2. Edit the backup copy job from which you have detached backups.
3. At the **Target** step of the wizard, map the backup copy job to the detached backups as described in section Map Backup File.

Alternatively, you can create a new backup copy job as described in subsection Periodic Copy Mode.

Veeam Backup & Replication will start a new backup chain in the required format and will retain the detached backups according to the background retention. For more information on when the background retention applies, its limitations and considerations, see Background Retention.

# Periodic Copy Mode

To upgrade the backup chain format for the periodic copy mode, do the following:

1. Create a new backup copy job.
2. At the **Objects** step of the wizard, select the same sources as in the backup copy job whose backups you want to upgrade. Note that you can also add other sources.
3. At the **Target** step of the wizard, select the backup repository where the backups whose format you want to change are stored.
4. At the **Target** step of the wizard, map the backup copy job to the backups as described in section Map Backup File.

Veeam Backup & Replication will start a new backup chain in the required format and will place the backups used for mapping to the orphaned node and retain them according to the background retention. For more information on when the background retention applies, its limitations and considerations, see Background Retention.

# Moving Backups

Veeam Backup & Replication allows you to move all backups of a backup copy job to another repository or to move specific workloads and their backups to another backup copy job.

## Moving Backups to Another Repository

For information on the move operation, how it works and its limitations, see Backup Move.

To move backups to another repository and target a job to this repository, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups** node.
3. In the working area, select the necessary job.
4. Right-click the job and select **Move backup**. Alternatively, click **Move Backup** on the ribbon.
5. In the **Move Backup to Another Location** window, select the repository to which you want to move backups.

   Veeam Backup & Replication will reconfigure and target the backup job to the selected repository.
6. Click **OK**.

Alternatively, you can change the repository in the job settings.

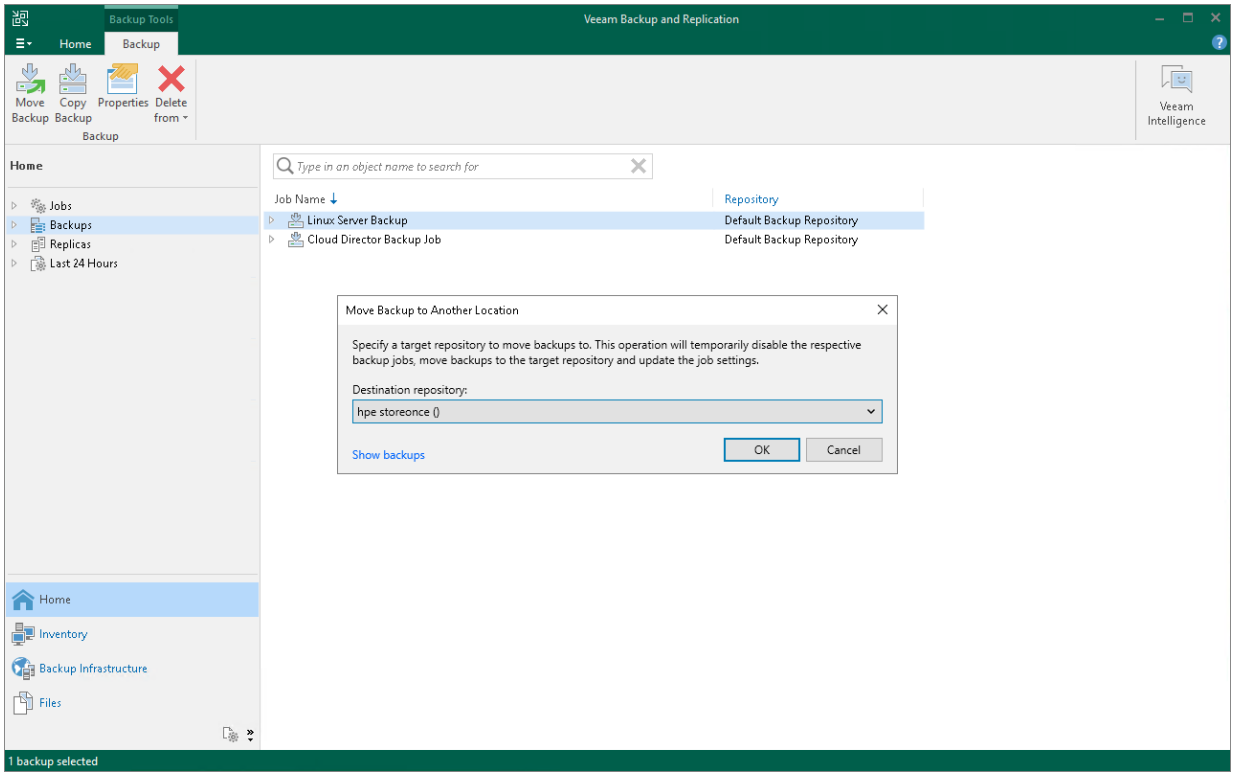


## Moving Backups to Another Job

For information on the move operation, how it works and its limitations, see Backup Move.

To move backups to another job:

7. Open the **Home** view.
8. In the inventory pane, select the **Backups** node.
9. In the working area, expand the necessary backup job and select workloads.

> **NOTE**
>
> You can move individual workloads and their backups only if backups are per-machine with separate metadata files for each workload.

4. Right-click one of the selected workloads and click **Move backup**. Alternatively, click **Move Backup** on the ribbon.
5. In the **Move Backup to Another Job** window, select the backup job to which you want to move backups.

   Veeam Backup & Replication will include workloads into the selected job and exclude workloads from the original job. Backups of the selected workloads will be moved to the repository to which the selected job is targeted.
6. Click **OK**.

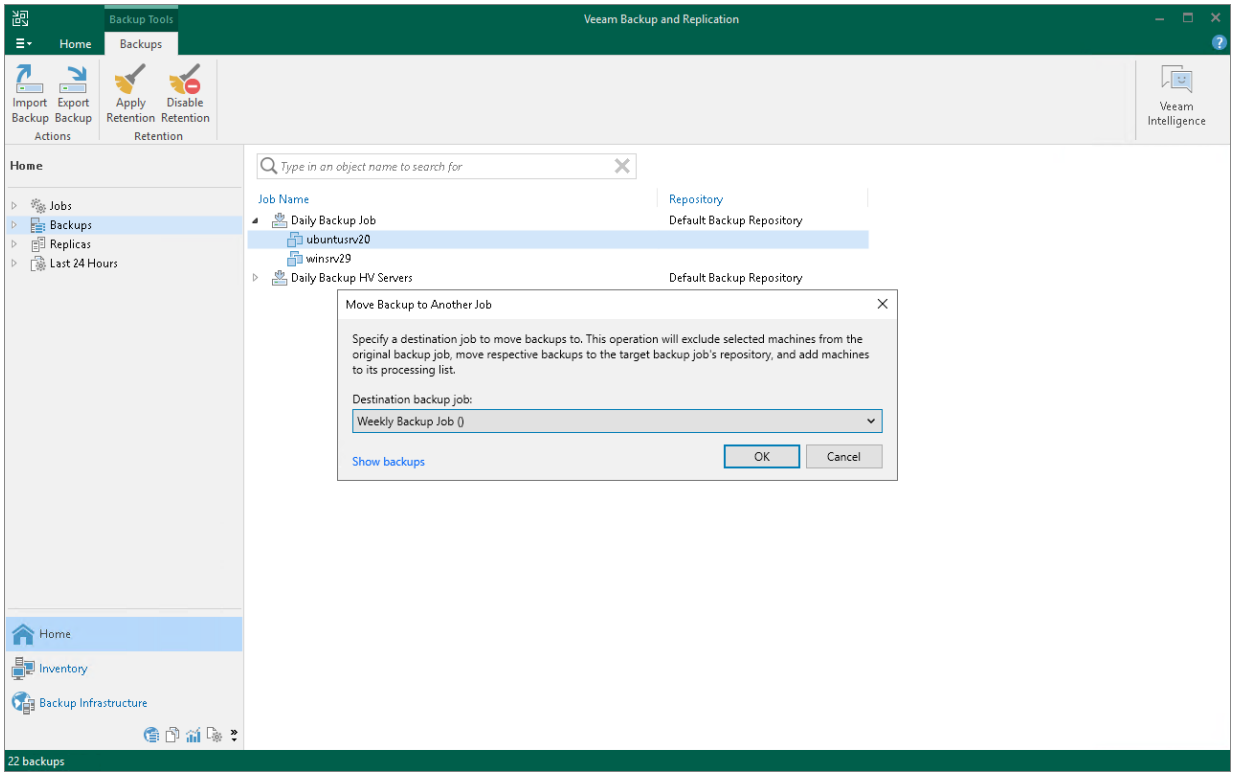


# Managing Failed Activities

If the move operation fails, Veeam Backup & Replication assigns the *User action required* status to it. In this case, you need to decide how to finish the operation: retry the move operation for failed workloads, cancel all changes, or move failed workloads but detach their backups. If you detach failed backups, the target job creates active full backups for failed workloads and continues backup chains for other workloads. The detached backups are shown in a node with the **(Orphaned)** postfix in the inventory pane.

> **NOTE**
>
> The original job will still be in the disabled state until you finalize the failed move operation.

To finalize the move operation:

10. Open the **Home** view.
3. In the inventory pane, select the **Last 24 Hours** node.
4. Right-click the failed move session and select the required action. Alternatively, select the required action on the ribbon.

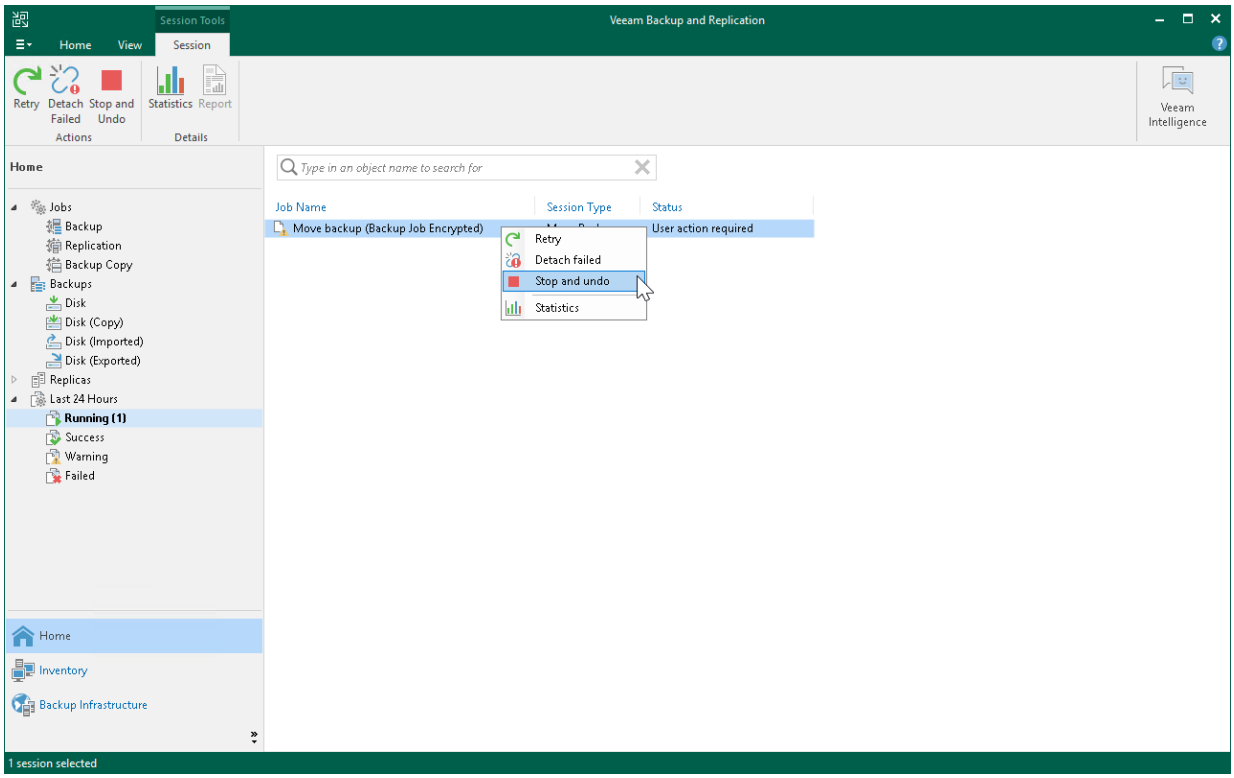

# Copying Backups

Copying backups can be helpful if you want to copy backups of a workload or backup job to another repository, local or shared folder. Veeam Backup & Replication copies the whole backup chain. If you want to convert a specific restore point into a single VBK file, use backup export. For more information, see Exporting Backups.

When Veeam Backup & Replication performs the copy operation, it disables the job, copies files to the target location and then enables the job. After the copy operation finishes, the copied backups are shown in a node with the **(Exported)** postfix in the inventory pane.

> **NOTE**
>
> This section is about one-time copy operation. If you want to copy backups on a schedule, create a backup copy job. For more information, see Backup Copy.

## Requirements and Limitations

Consider the following:

- The copy operation does not change the backup chain format (single-file backup, per-machine with single metadata file or per-machine with separate metadata files). If you copy backups between repositories with and without the **Use per-machine backup files** check box enabled, backups preserve their formats.
- If you copy backups from a scale-out backup repository and some backups are stored on extents in the Maintenance mode, such backups are not copied.
- Veeam Backup & Replication copies backups only from the performance tier of the scale-out backup repository. If you want to copy data from the capacity tier, you first must download it to the performance tier. For more information, see Downloading Data from Capacity Tier.
- You cannot copy backups between extents of a scale-out backup repository. To learn how to manage backups within the scale-out backup repository, see Scale-Out Backup Repositories.
- You cannot copy backups stored in Veeam Cloud Connect repositories. For more information on Veeam Cloud Connect repositories, see the Cloud Repository section in the Veeam Cloud Connect Guide.
- The Copy backup operations to, from and between cloud repositories are not supported. For more information on Veeam Cloud Connect's limitations, see the Veeam Cloud Connect Backup section in the Veeam Cloud Connect Guide.
- [For VMware Cloud Director] You can copy backups of a whole job or individual vApps. You cannot move backups of VMs.
- You cannot copy backups created by a backup job managed by Veeam Agent (backup policy).
- You cannot copy backups created by Veeam plug-ins for Enterprise applications, Veeam Cloud Plug-ins (Veeam Backup for AWS, Veeam Backup for Google Cloud, Veeam Backup for Microsoft Azure), and Veeam Kasten.
- Traffic throttling is not supported for backup copy operations.

## Copying Backups

To copy backups, do the following:

1. Open the **Home** view.

2. In the inventory pane, select the **Backups** node.
3. In the working area, select the necessary job or workload. Note that you can copy backups of an individual workload only if its backups are per-machine backups with separate metadata files.
4. Right-click the job and select **Copy backup**. Alternatively, click **Copy Backup** on the ribbon.
5. In the **Copy Backup to Another Location** window, choose where you want to copy backups — to a repository or to a local or shared folder.
6. If you want to delete the copied backups after a specific time period, select the **Delete this backup automatically in** check box and specify the time period.

   Backups that fall out of the specified retention policy will be removed automatically. If you do not specify the time period for deletion, copies will be stored until you remove them manually.

   > **TIP**
   >
   > You can customize retention period values in the drop-down list as described in this Veeam KB article.

7. Click **OK**.

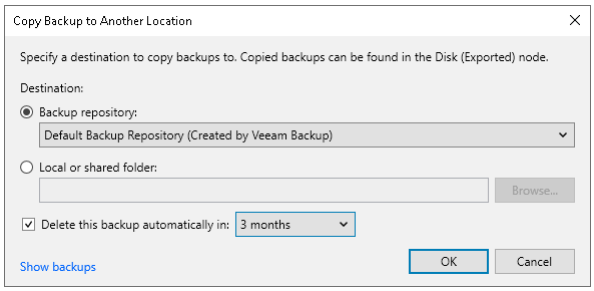


# Managing Failed Activities

If the copy operation fails, Veeam Backup & Replication assigns the *User action required* status to it. In this case, you need to decide how to finish the operation: retry the copy operation for failed backups, skip the failed backups or cancel all changes.

> **NOTE**
>
> The original job will still be in the disabled state until you finalize the failed copy operation.

To finalize the copy operation:

1. Open the **Home** view.
2. In the inventory pane, select the **Last 24 Hours** node.

3. Right-click the failed copy session and select the required action. Alternatively, select the required action on the ribbon.

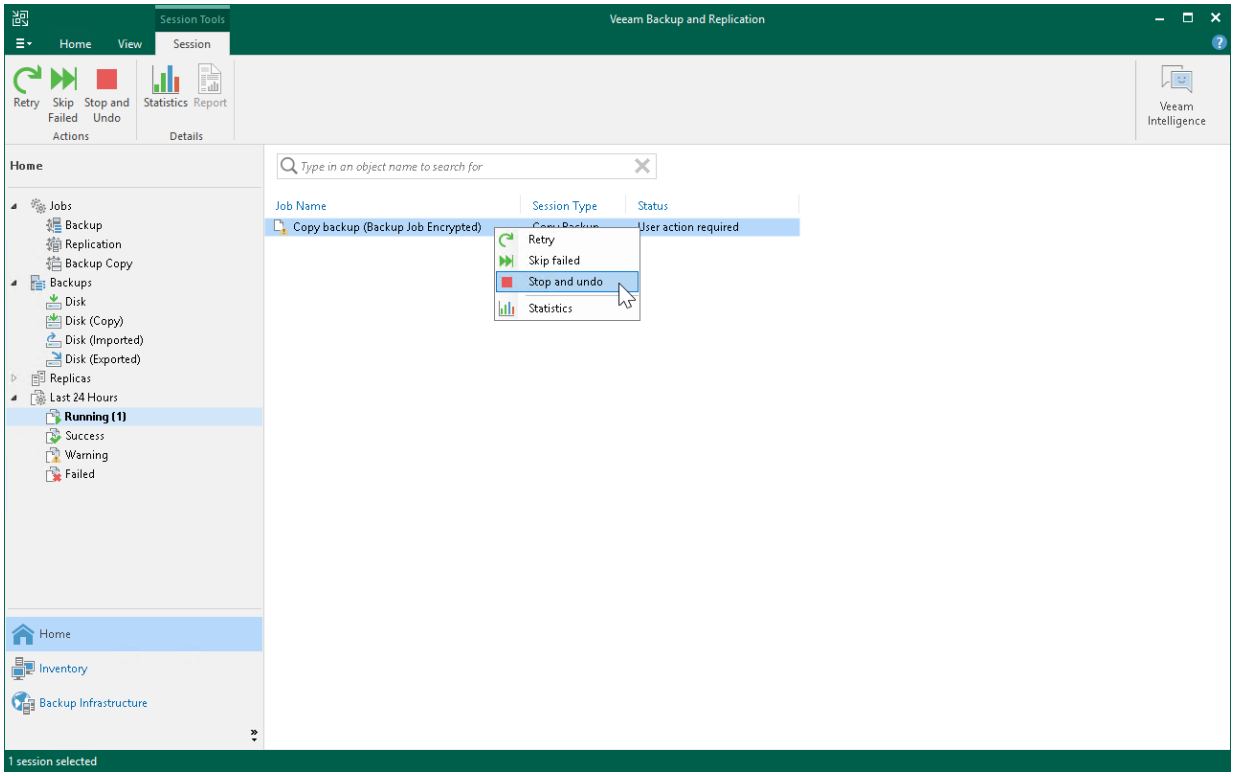

# Removing Backups

You can detach backups from backup copy jobs, permanently delete backups from the target backup repositories or remove records about backups from the Veeam Backup & Replication console and configuration database using the **Remove from configuration** operation.

## Detaching from Job

When you detach backup copies from a backup job, the job stops processing these backup files. During the next run, the job will start a new backup chain, that is, will create active full backup.

Veeam Backup & Replication detaches the whole backup chain including GFS backups. The detached backup files remain in the backup repository and also in the Veeam Backup & Replication console.
Veeam Backup & Replication shows the detached backups in the inventory pane in the node with the **(Orphaned)** postfix. These backups are retained according to the background retention process. For more information, see Background Retention.

To detach backups from a backup copy job:

1. Open the **Home** view.
2. In the inventory pane, select **Backups** > **Disk (Copy)**.
3. In the working area, right-click the backup copy and select **Detach from job**. Alternatively, click **Remove from** > **Job** on the ribbon.

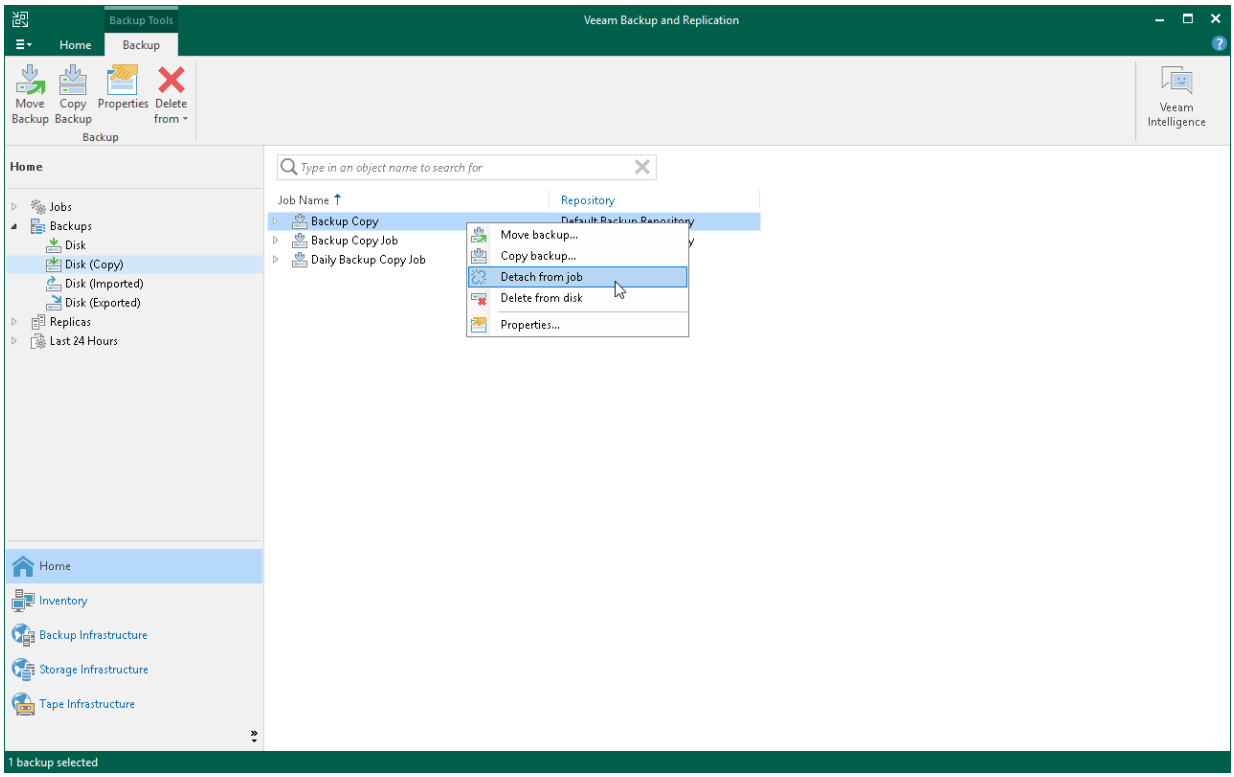

# Deleting from Disk

When you use the **Delete from disk** option, you delete records about backup copies from the Veeam Backup & Replication console and configuration database, and delete backup files from the target backup repository. This option can be used for the whole backup copy or for some workloads in the backup copy.

> **IMPORTANT**
>
> For scale-out backup repositories the **Delete from disk** operation will remove the backups not only from the performance tier but also from the capacity and archive tier. If you want to remove backups from the performance tier only, you should move those backups to the capacity tier instead. For details, see Manually Moving Backups to Capacity Tier.

To permanently remove backup copies from the target backup repository:

1. Open the **Home** view.
2. In the inventory pane, select **Backups** > **Disk (Copy)**.
3. In the working area, right-click the backup copy or a workload in the backup copy and select **Delete from disk**.
4. To remove all weekly, monthly, quarterly and yearly backups from disk, select the **Include archived full backups** check box and click **Yes**.

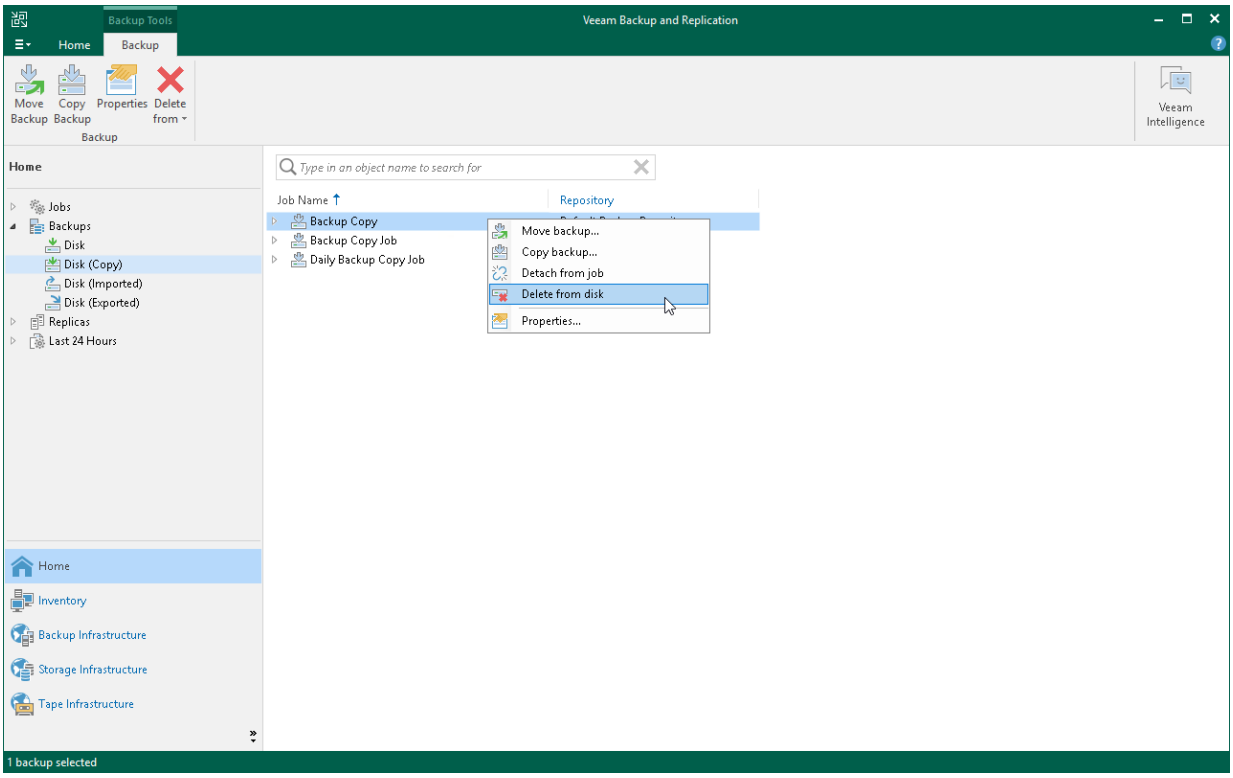

# Removing Backups from Configuration

> **IMPORTANT**
>
> Removing backups from configuration is designed for experienced users only. Consider using **Detach from job** or **Delete from disk** operations.
>
> Create encrypted configuration backup before removing backups from configuration.

When you remove a backup from the configuration, backup files remain in the backup repository. You can import the backup to Veeam Backup & Replication at any time later and restore data from it.

When you remove an encrypted backup from configuration, Veeam Backup & Replication removes encryption keys from the configuration database. If you import such backup on the same backup server or another backup server, you will have to specify the password or unlock the backup with Veeam Backup Enterprise Manager. For more information, see Importing Encrypted Backups.

To remove a backup from the configuration:

1. Open the **Home** view.
2. In the inventory pane, select **Backups** > **Disk (Copy)**.
3. In the working area, press the [Ctrl] key, right-click the backup that you want to remove and select **Remove from configuration**.

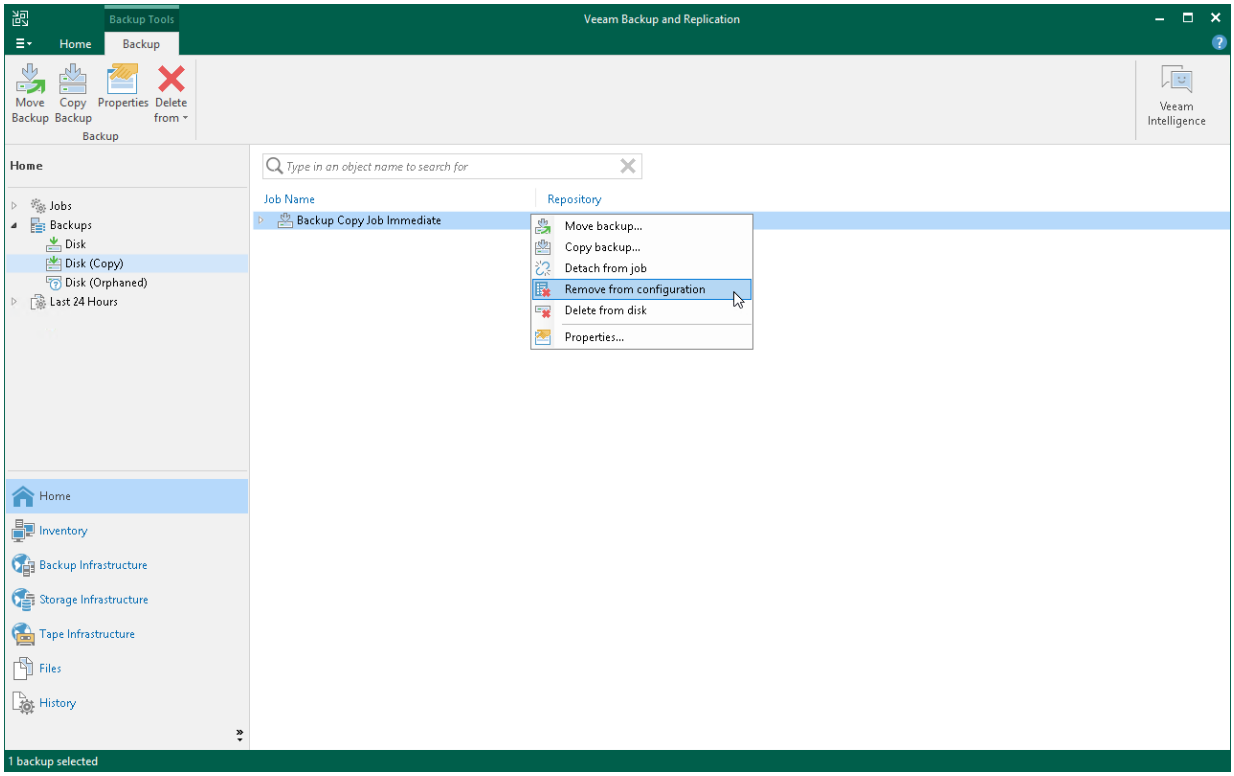

# Removing Missing Restore Points

In some cases, one or more restore points in the backup chain may be inaccessible. This can happen, for example, if the backup repository is put to the Maintenance mode (for scale-out backup repositories), the backup repository is not available or some backup file is missing in the backup chain. Backup chains that contain missing restore points get corrupted — you cannot perform backup copy or restore data from the missing restore point, and restore points that depend on the missing restore point.

You can perform the following with missing restore points:

- Forget — you can remove records about missing restore points from the configuration database. Veeam Backup & Replication will ignore the missing restore points and will not display them in the console. The backup files will remain on disk (if backup files are still available).
- Remove — you can remove records about missing restore points from the Veeam Backup & Replication console and configuration database, and delete backup files from disk (if backup files are still available).

> **NOTE**
>
> Consider the following:
>
> - The **Forget** and **Remove from disk** options are available only for restore points that are missing from the backup chain or that depend on missing ones.
> - You can manually update information about missing restore points. For this, disable a backup copy job and rescan the backup repository that is the target for the backup copy job. For more information, see Disabling and Deleting Backup Copy Jobs and Rescanning Backup Repositories.
>
>  Manual update can be required because Veeam Backup & Replication requires some time to update information in the configuration database for restore points that were removed from a backup chain or became inaccessible. That is why such restore points may not be displayed in the console as missing restore points.

To remove records about missing restore points from the configuration database:

1. Open the **Home** view.
2. In the inventory pane, select **Disk (copy)** under **Backups**.
3. In the working area, select the backup and click **Properties** on the ribbon or right-click the backup and select **Properties**.
4. In the **Backup Properties** window, right-click the missing restore point and select **Forget**.
    - To remove only the selected restore point and restore points that depend on it (that is, a part of the backup chain starting from this restore point), select **This and dependent backups**.

- To remove all missing restore points, select **All unavailable backups**.

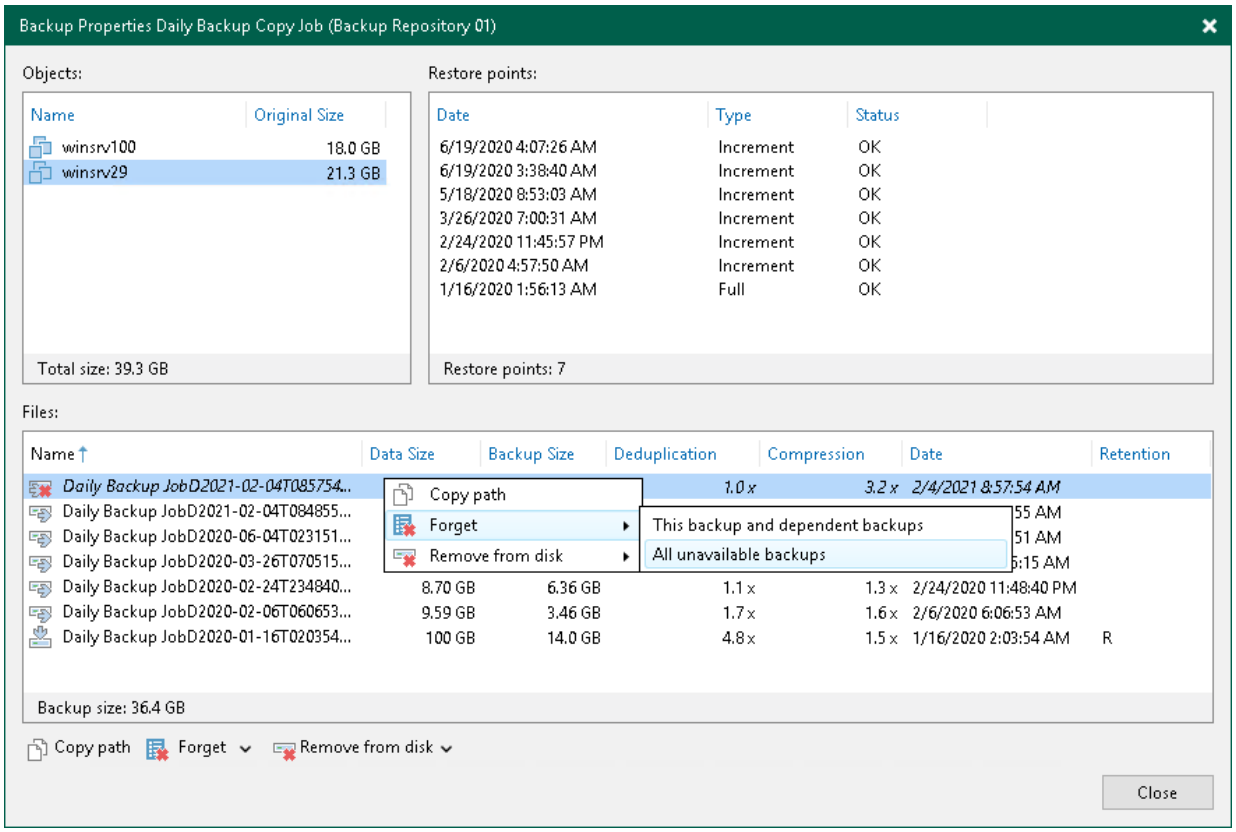


To remove missing restore points from the configuration database and disk:

1. Open the **Home** view.
2. In the inventory pane, click **Disk (copy)** under **Backups**.
3. In the working area, select the backup and click **Properties** on the ribbon or right-click the backup and select **Properties**.
4. In the **Backup Properties** window, right-click the missing restore point and select **Remove from disk**.
   - To remove only the selected restore point and restore points that depend on it (that is, a part of the backup chain starting from this restore point), select **This and dependent backups**.

- To remove all missing restore points, select **All unavailable backups**.

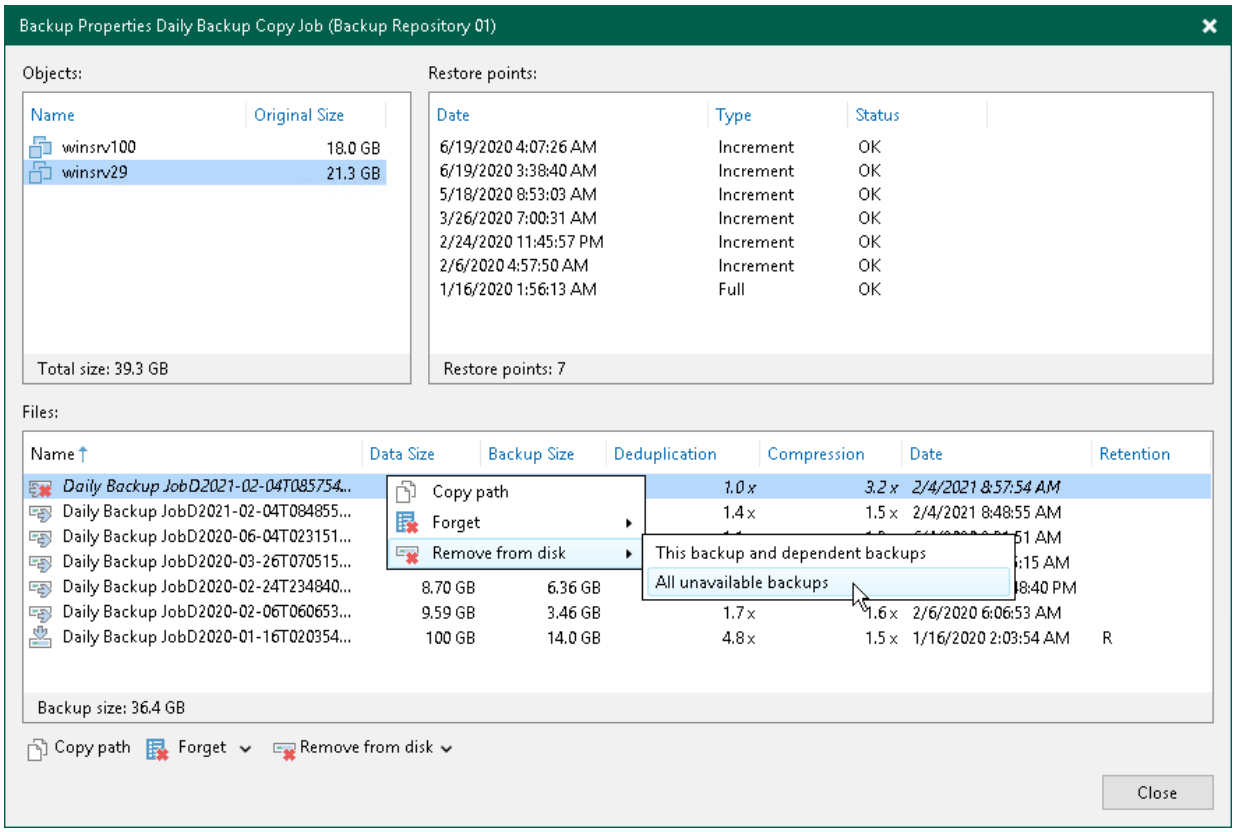

# Managing Backup Copy Jobs

To view all configured jobs, open the **Home** view and select the **Jobs** > **Backup Copy** node in the inventory pane. The list of available jobs is displayed in the working area. You can edit job properties, start and stop jobs, and delete unnecessary jobs.

# Editing Backup Copy Job Settings

You can edit backup copy job settings after you create it. For example, you may want to change scheduling settings for the job or add some machines to the job.

To edit job settings:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs** > **Backup Copy**.
3. In the working area, select the job and click **Edit** on the ribbon or right-click the job and select **Edit**.

You will follow the same steps as you have followed when creating the job and can change job settings as required.

## Related Topics

- Creating Backup Copy Jobs for VMs and Physical Machines
- Creating Backup Copy Jobs for HPE StoreOnce Repositories
- Creating Backup Copy Jobs for Veeam Plug-ins

# Disabling GFS Scheme

If you disable the **Keep certain full backups longer for archival purposes** option, and you already have archive full backups in the target backup repository, Veeam Backup & Replication will offer you to remove existing archive full backups.

- Click **Yes** to remove archive full backups from the target backup repository. Archive full backups will be removed during the next retention cycle (next backup copy session). The backup copy job will not create archive full backups.
- Click **No** to keep archive full backups in the target backup repository. Archive full backups will be displayed under the **Backups** > **Disk (Imported)** node in the Veeam Backup & Replication console. The backup copy job will not create archive full backups.

> **NOTE**
>
> If you disable the **Keep certain full backups longer for archival purposes** option and enable it again later, archive full backups that remained on disk will not be linked to the backup copy job. They will still be displayed under the **Backups** > **Disk (Imported)** node in the Veeam Backup & Replication console.

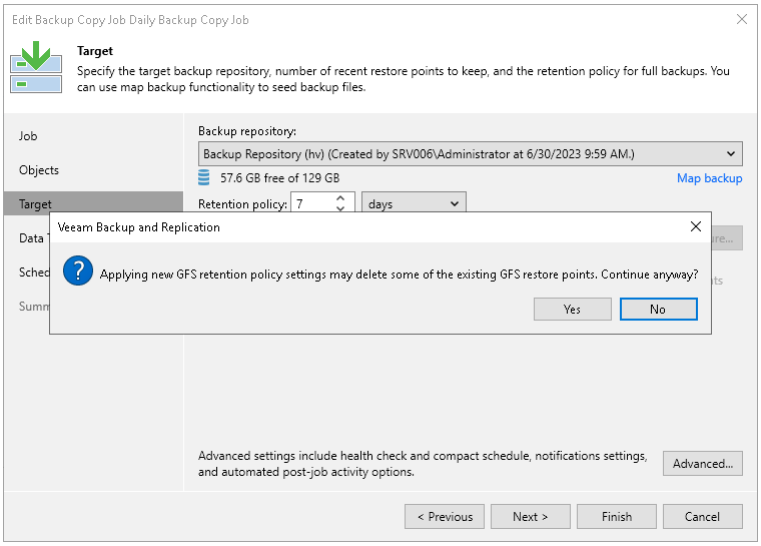
Edit Backup Copy Job Daily Backup Copy Job
Target
Specify the target backup repository, number of recent restore points to keep, and the retention policy for full backups. You can use map backup functionality to seed backup files.
Job
Objects
Target
Data
Sched
Summ
Backup repository:
Backup Repository (hv) (Created by SRV006\Administrator at 6/30/2023 9:59 AM.)
57.6 GB free of 129 GB
Map backup
Retention policy: 7
days
Veeam Backup and Replication
Applying new GFS retention policy settings may delete some of the existing GFS restore points. Continue anyway?
Yes
No
Advanced settings include health check and compact schedule, notifications settings, and automated post-job activity options.
Advanced...
< Previous
Next >
Finish
Cancel

# Starting and Stopping Backup Copy Jobs

Manual start can be helpful if the backup copy job was disabled for some time and restore points were not copied to the target repository. This procedure differs for the periodic and immediate copy modes. You can also manually stop the periodic backup copy job. For more information, see the following sections:

- Starting Jobs in Immediate Copy Mode
- Starting Jobs in Periodic Copy Mode
- Stopping Jobs in Periodic Copy Mode

> **NOTE**
>
> The immediate backup copy job can only be disabled and can not be stopped.

## Starting Jobs in Immediate Copy Mode

To start the backup copy job manually:

1. Open the **Home** view.
2. In the inventory pane, select the **Backup Copy** node.
3. In the working area, select the backup copy job and click **Sync now** on the ribbon or right-click the backup copy job and select **Sync now**.
4. In the opened window, do the following:
   - If you want to copy all restore points created by the source jobs but that were not copied since the last backup copy job session, click **All**.
   - If you want to copy only the latest restore point for each source job, click **Latest**.

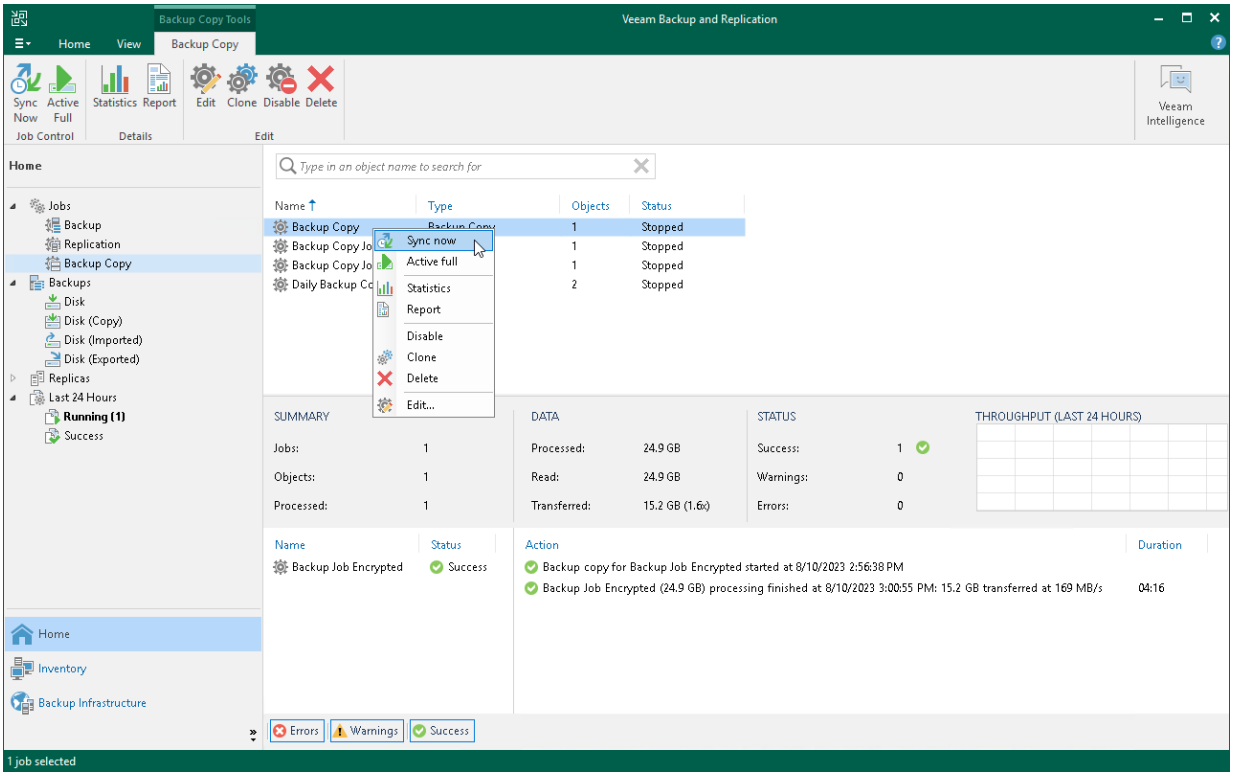

# Starting Jobs in Periodic Copy Mode

To start the backup copy job manually:

1. Open the **Home** view.
2. In the inventory pane, select the **Backup Copy** node under **Jobs**.
3. In the working area, select the backup copy job and click **Start** on the ribbon or right-click the backup copy job and select **Start**.

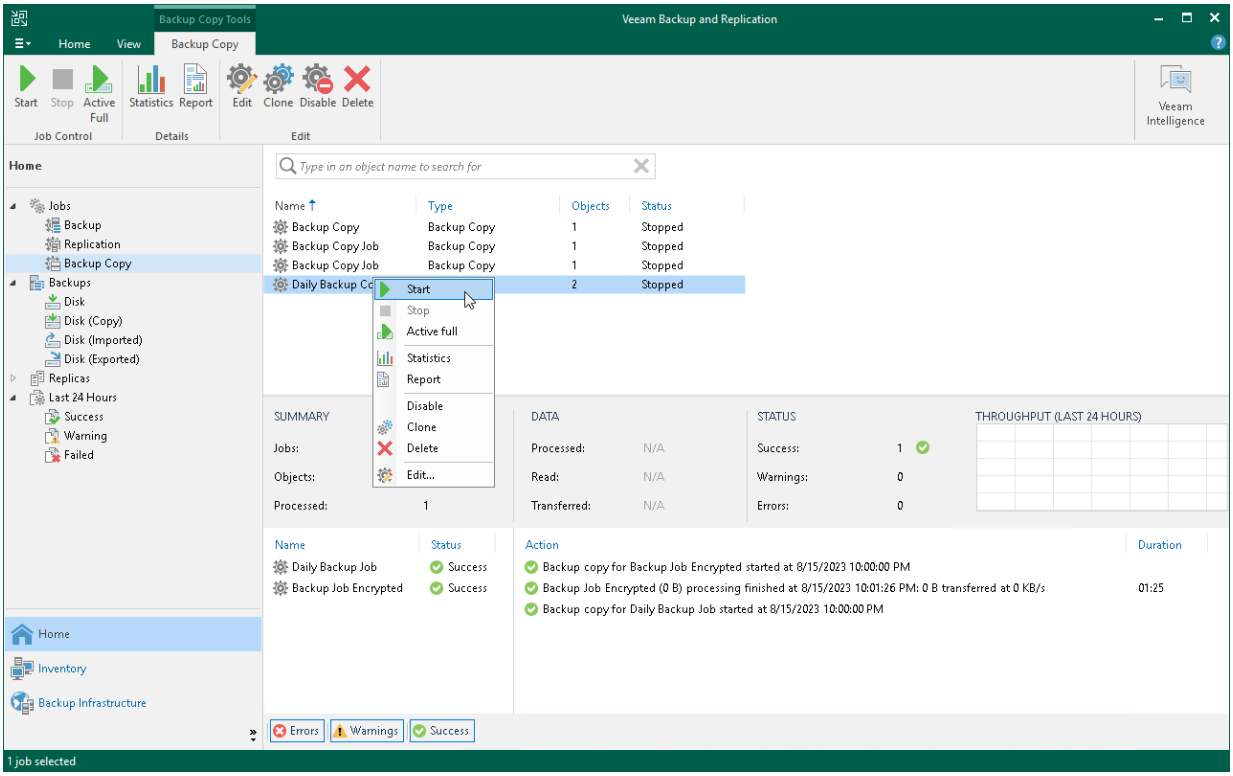


# Stopping Jobs in Periodic Copy Mode

To stop the backup copy job manually:

1. Open the **Home** view.
2. In the inventory pane, select the **Backup Copy** node under **Jobs**.

3. In the working area, select the backup copy job and click **Stop** on the ribbon or right-click the backup copy job and select **Stop**.

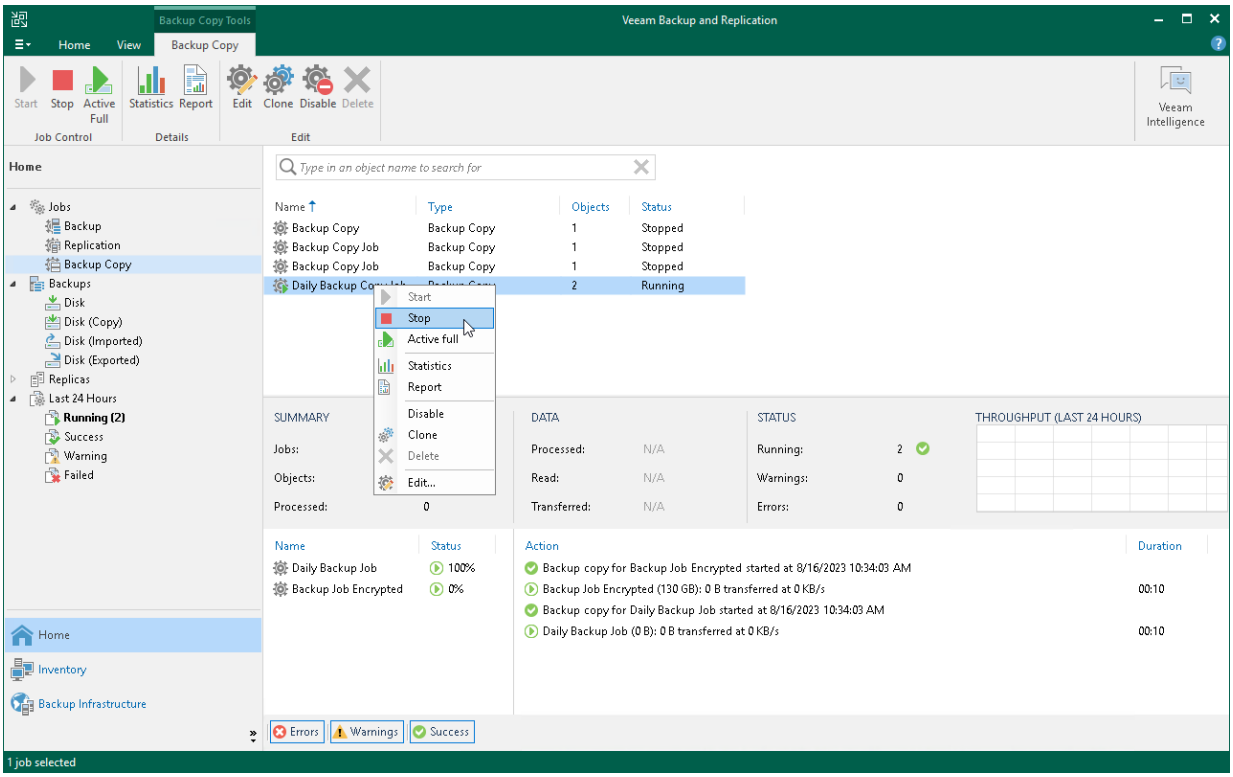

# Creating Active Full Backups

You can manually create an ad-hoc full backup — active full backup, and add it to the backup chain in the target backup repository. Active full backup can be helpful if you want to change backup copy job settings, for example, enable or disable encryption. Veeam Backup & Replication will apply new settings starting from this full backup.

To create an active full backup manually:

1. Open the **Home** view.
2. In the inventory pane, select the **Backup Copy** node under **Jobs**.
3. In the working area, select the backup copy job and click **Active full** on the ribbon or right-click the backup copy job and select **Active full**. Veeam Backup & Replication will start a new backup copy session, copy data from the source backup repository and save it in a full backup file in the target backup repository.

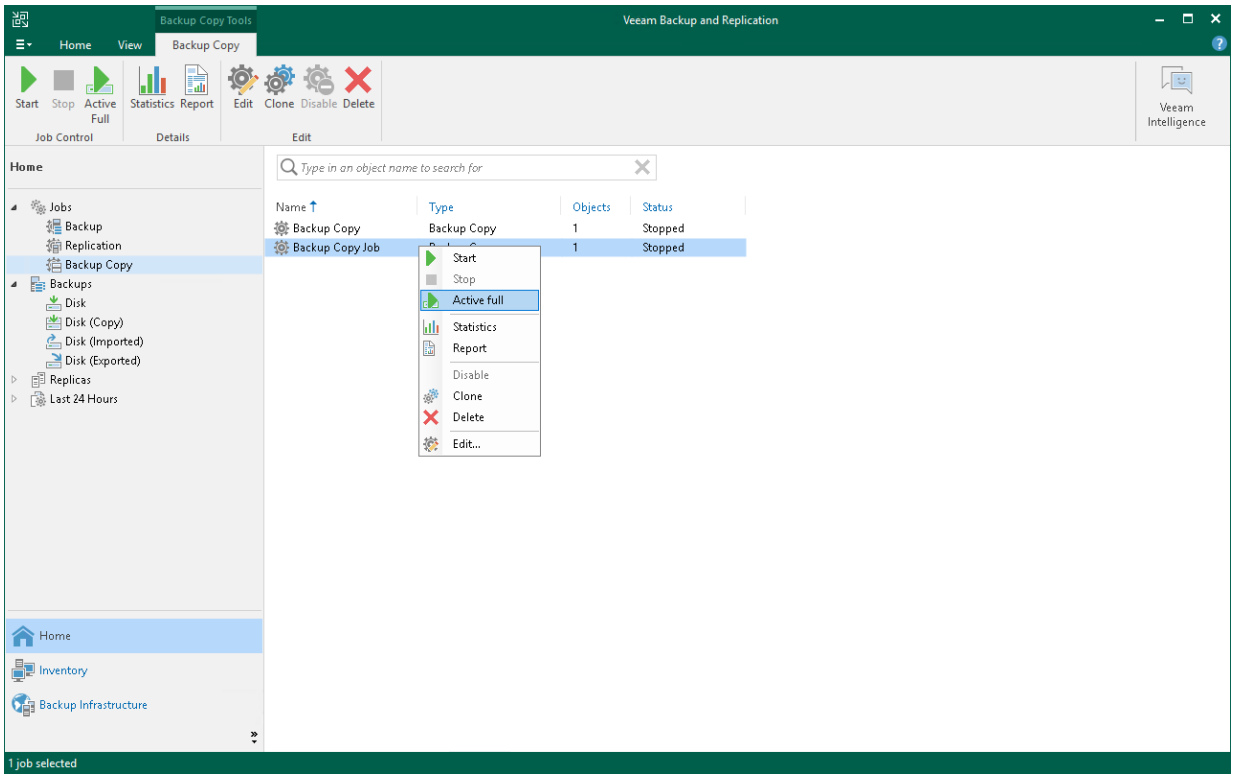

# Disabling and Deleting Backup Copy Jobs

You can temporarily disable backup copy jobs. The disabled job is paused for some period of time and is not run by the specified schedule. You can enable a disabled job at any time. You can also permanently delete a job from Veeam Backup & Replication and from the configuration database.

## Disabling Job

To disable a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs** > **Backup Copy**.
3. In the working area, select the job and click **Disable** on the ribbon or right-click the job and select **Disable**.

To enable a disabled job, select the job in the list and click **Disable** once again.

> **NOTE**
>
> If the health check is enabled in the job settings, it will continue to run according to its own schedule, regardless of the job status.

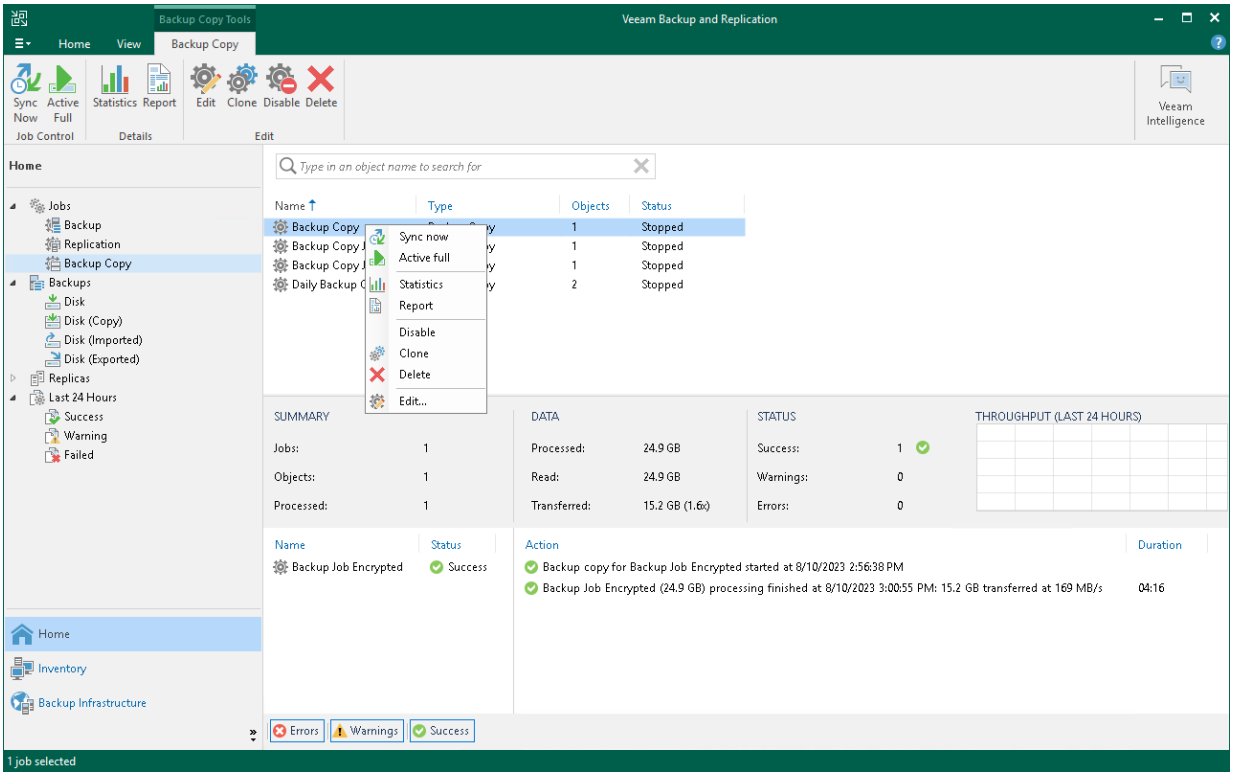


## Deleting Job

To delete a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs** > **Backup Copy**.
3. In the working area, select the job, click **Delete** on the ribbon or right-click the job and select **Delete**.

> **NOTE**
>
> If you want to delete an active backup copy job, you must first stop the synchronization process. To do this, disable the backup job. After the job is disabled, you can delete it.

After the job is deleted, the backups created by this job are displayed under the **Backups** > **Disk (Orphaned)** node. If the backup files created by this job were also stored in an object storage repository, they will also be displayed under the **Backups** > **Object Storage (Orphaned)** node.

# Cloning Backup Copy Job

You can create new backup copy jobs by means of job cloning. Job cloning allows you to create an exact copy of any job with the same job settings. Configuration information of the created job copy are written to the configuration database that stores information of the original job.

To create multiple jobs with similar settings, you can configure a set of jobs that will be used as 'job templates'. You can then clone these 'job templates' and edit settings of cloned jobs as required.

The name of the cloned job is formed by the following rule: *<job_name_clone1>*, where *job_name* is the name of the original job and *clone1* is a suffix added to the original job name. If you clone the same job again, the number in the name will be incremented, for example, *job_name_clone2*, *job_name_clone3* and so on.

When cloning job, Veeam Backup & Replication can change some job settings so that cloned jobs do not hinder original jobs.

- If the original job is scheduled to run automatically, Veeam Backup & Replication disables the cloned job. To enable the cloned job, select it in the job list and click **Disable** on the ribbon or right-click the job and select **Disable**.
- If the original job is configured to use a secondary target, the cloned job is created without the secondary target settings.

To clone a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup copy job and click **Clone** on the ribbon or right-click the job and select **Clone**.
4. After a job is cloned, you can edit all its settings, including the job name.

> **NOTE**
>
> Consider the following:
>
> - The job cloning functionality is available for all types of licenses.
> - [For Veeam Backup & Replication before 12.1 (build 12.1.0.2131)] The job cloning functionality is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.

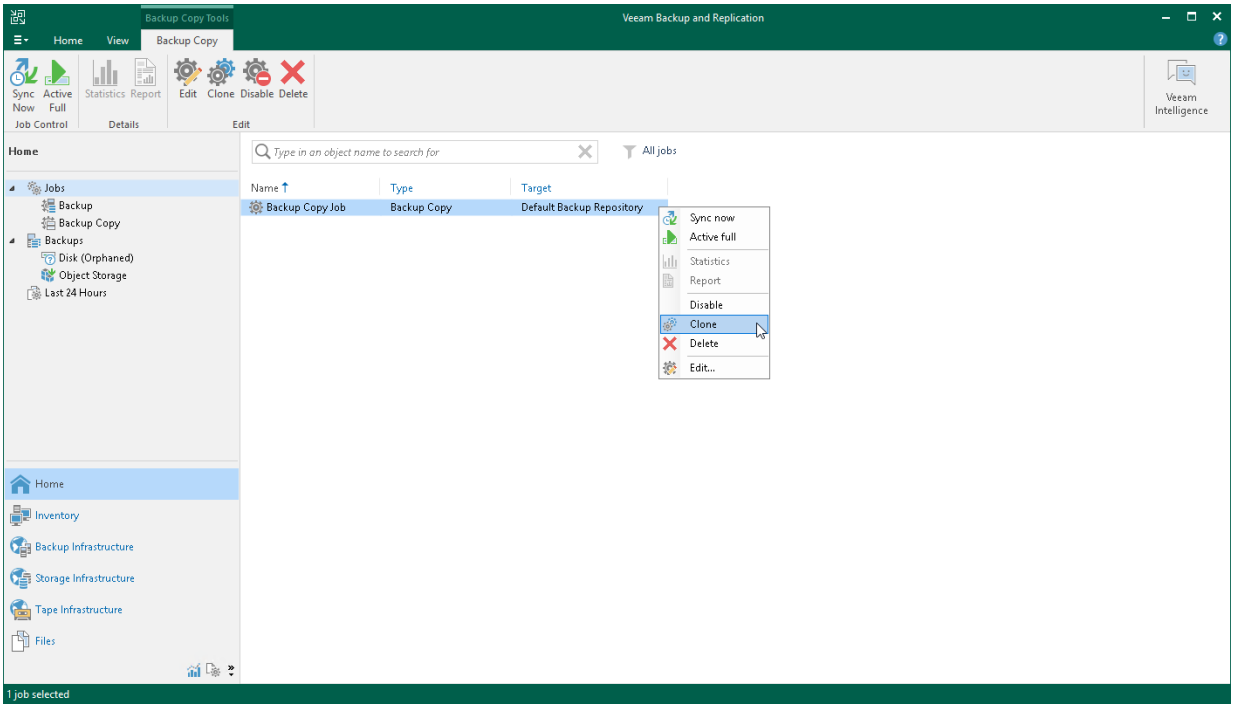
Backup Copy Tools
Veeam Backup and Replication
Home
View
Backup Copy
Sync Now
Active Full
Statistics
Report
Edit
Clone
Disable
Delete
Job Control
Details
Edit
Veeam Intelligence
Home
Type in an object name to search for
All jobs
Jobs
Backup
Backup Copy
Backups
Disk (Orphaned)
Object Storage
Last 24 Hours
Name
Type
Target
Backup Copy Job
Backup Copy
Default Backup Repository
Sync now
Active full
Statistics
Report
Disable
Clone
Delete
Edit...
Home
Inventory
Backup Infrastructure
Storage Infrastructure
Tape Infrastructure
Files
1 job selected

# Reporting

The process of performing reporting operations for backup copy jobs is the same as described in the Backup: Reporting section. However, the process differs for immediate backup copy jobs created in Veeam Backup & Replication 11 the following way:

- Viewing real-time statistics

  You can view real-time statistics for the whole job if you select the job in the working area of the **Jobs** node. The whole job report shows general information about the job itself and child jobs — tasks that copy backup jobs added as sources to the backup copy job. You can also view real-time statistics for an individual child job if you select the child job in the working area of the **Last 24 Hours** or **Running** node. This statistics shows detailed information about the selected child job including the processed VMs.

- Viewing job session results

  You select and view session results for a child job.

- Viewing job reports

  The job report shows results for the last job run and does not provide details on child jobs. If you want to get reports once a child job finishes, configure notifications. For more information, see Notification Settings.

- Viewing session reports

  You can view session reports only if you configured notifications for a job. In this case, you get reports once a child job finishes. For more information, see Notification Settings.

# VM Copy

With Veeam Backup & Replication, you can run a VM copy job to create an independent fully-functioning copy of a VM or VM container (host, cluster, folder, resource pool, VirtualApp, datastore or tag) on the selected storage. VM copying can be helpful if you want to move your datacenter, create a test lab and so on.

The produced copy of a VM is stored decompressed, in a native VMware vSphere format, so it can be started right away. Although VM copy is similar to replication in many respects, there are several important differences.

- VM copy is a single-use process (that is, every run of a VM copy job mirrors a VM in its latest state). Due to their nature, VM copy jobs do not support incremental runs.
- Veeam Backup & Replication does not create and maintain restore points for VM copies. If you schedule to run a VM copy job periodically, every new run will overwrite the existing copy.
- With the VM copy job, all VM disks are copied as thick, while replication allows you to preserve the format of disks or convert the disk format on the fly.
- There are no failover or failback possibilities for a VM copy.

VM copy jobs use the same infrastructure components as backup jobs (for details, see Backup Infrastructure for Backup). In addition to available scenarios, you can also copy VMs to a target folder on any server or host connected to the backup server.

# Copying VMs

With VM copy jobs you can create a fully-functional copy of a VM and store this copy in the backup repository or storage device. VM copying may be helpful if you want to move your datacenter to another location, archive a VM before decommissioning and so on.

To create a VM copy, you must configure a VM copy job. One job can be used to process one VM or more VMs.

You can configure a job and start it immediately or save the job to start it later. Jobs can be started manually or scheduled to run automatically at a specific time.

Before you create a VM copy job, check prerequisites. Then use the **New VM Copy Job** wizard to configure a VM copy job.

# Before You Begin

Before you create a VM copy job, check the following prerequisites:

- Backup infrastructure components that will take part in the VM copying process must be added to the backup infrastructure and properly configured. These include the source ESXi host and server or backup repository on which you plan to store the VM copy.
- The target storage device must have enough free space to store created VM copies. To receive alerts about low space on the storage device, configure global notification settings. For more information, see Specifying Other Notification Settings.
- If you plan to use pre-freeze and post-thaw scripts, you must create scripts before you configure the VM copy job.

Consider the following limitations:

- Due to Microsoft limitations, you cannot use Microsoft Entra ID (formerly Azure Active Directory) credentials to perform application-aware processing on VMs running Microsoft Windows 10 (or later).
- If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.

# Step 1. Launch VM Copy Job Wizard

To run the **VM Copy Job** wizard, do either of the following:

- On the **Home** tab, click **Copy Job** > **Virtual machine**.
- Open the **Home** view. In the inventory pane, right-click **Jobs** and select **VM Copy**.
- Open the **Inventory** view, in the working area select the VMs, click **Add to VM Copy** on the ribbon and select **New job** or right-click the VMs area and select **Add to VM copy job** > **New job**. In this case, the selected VMs will be automatically added to the VM copy job. You can add other VMs to the job when passing through the wizard steps.
- You can quickly add the VMs to an already existing job. To do this, open the **Inventory** view, in the working area select the VMs and click **Add to VM Copy** > *name of the job* on the ribbon or right-click the VMs and select **Add to VM copy job** > *name of the job*.

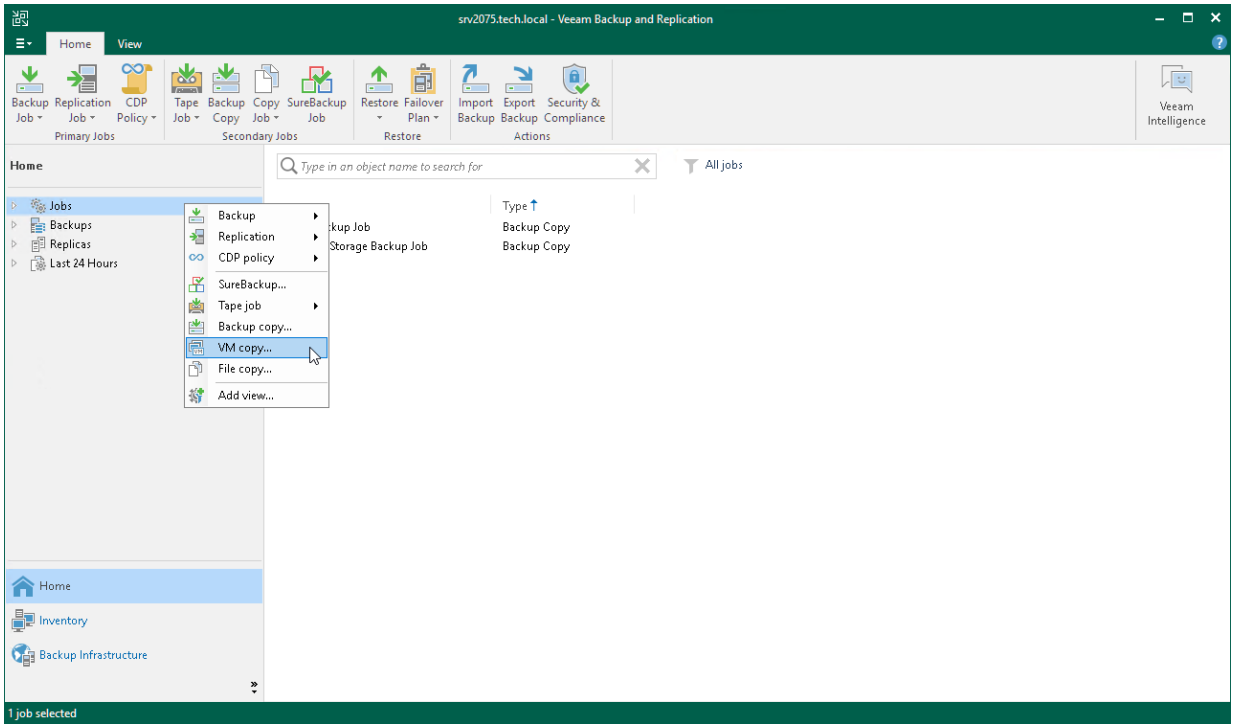

# Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a name and description for the VM copy job.

1. In the **Name** field, enter a name for the VM copy job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created a job, date and time when the job was created.

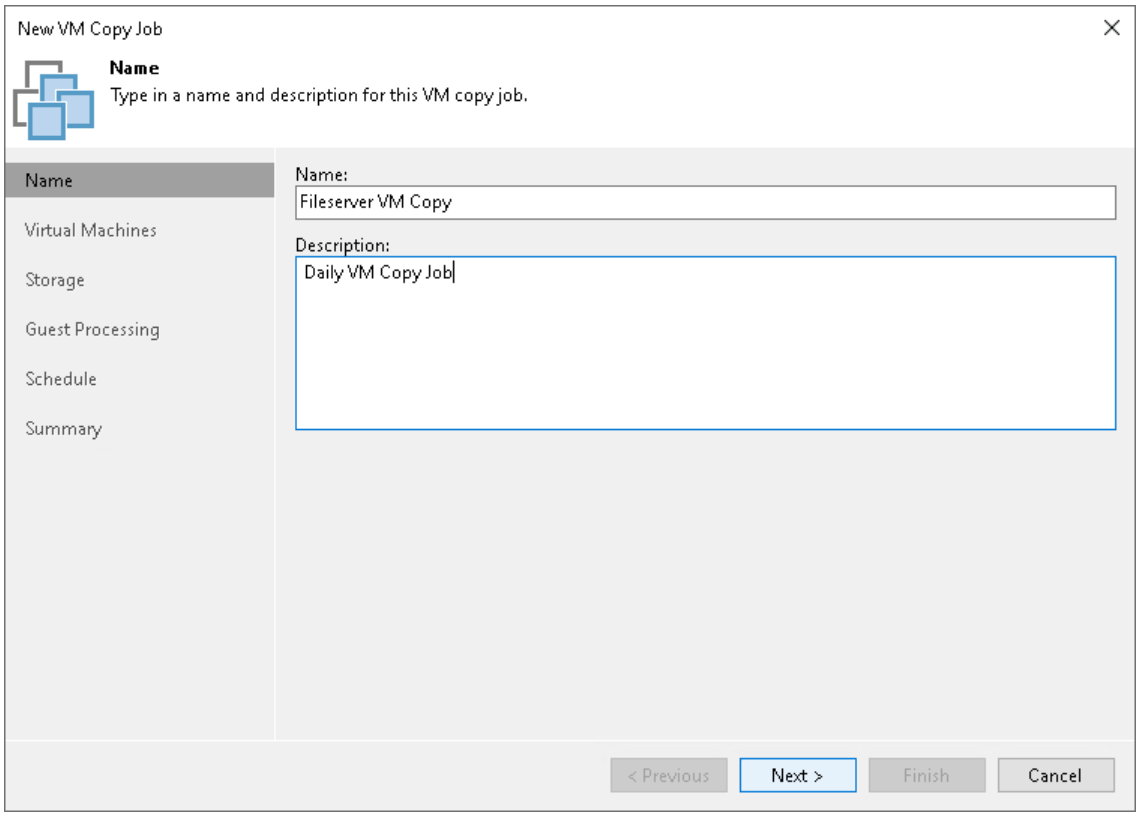

# Step 3. Select VMs to Copy

At the **Virtual Machines** step of the wizard, select VMs and VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags) that you want to copy.

Jobs with VM containers are dynamic in their nature. If a new VM is added to the container in the virtual infrastructure after the VM copy job is created, Veeam Backup & Replication will automatically update the job settings to include the added VM.

1. Click **Add**.
2. Use the toolbar at the top right corner of the window to switch between views: **Hosts and Clusters**, **VMs and Templates**, **Datastores and VMs** and **Tags**. Depending on the view you select, some objects may not be available. For example, if you select the **VMs and Templates** view, no resource pools, hosts or clusters will be displayed in the tree.
3. Select the object and click **Add**.

To quickly find the necessary object, you can use the search field at the bottom of the **Add Objects** window.

1. Click the button to the left of the search field and select the necessary type of object to search for: *Everything*, *Folder*, *Cluster*, *Host*, *Resource pool*, *VirtualApp* or *Virtual machine*.
2. Enter the object name or a part of it in the search field.
3. Click the **Start search** button on the right or press [Enter] on the keyboard.

The initial size of VMs and VM containers added to the VM copy job is displayed in the **Size** column in the list. The total size of objects is displayed in the **Total size** field. Use the **Recalculate** button to refresh the total size value after you add a new object to the job.

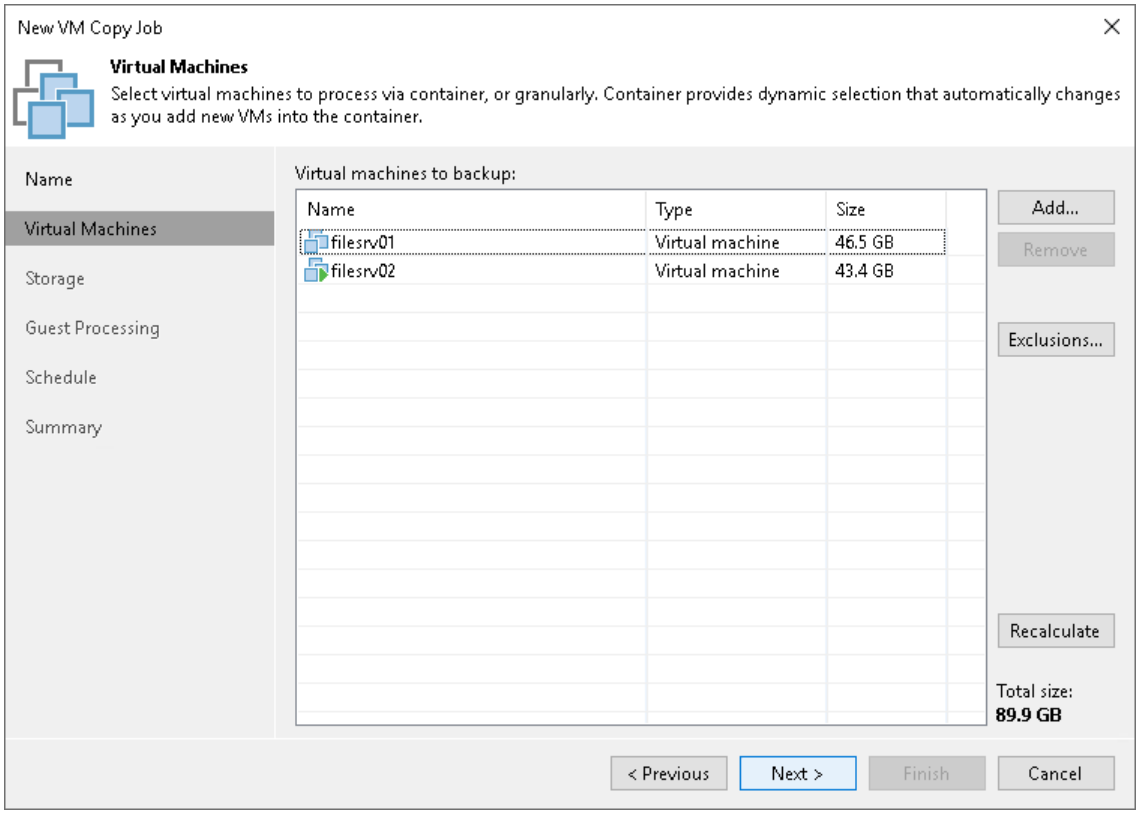

# Step 4. Exclude Objects from VM Copy Job

After you have added VMs and VM containers to the job, you can specify which objects you want to exclude from the VM copy. You can exclude the following types of objects:

- VMs from VM containers
- Specific VM disks
- VM templates

> **NOTE**
>
> Veeam Backup & Replication automatically excludes VM log files from VM copies to make copying process faster and reduce the size of the resulting file.

To exclude VMs from a VM container:

1. At the **Virtual Machines** step of the wizard, select a VM container added to the job and click **Exclusions**.
2. Click the **VMs** tab.
3. Click **Add**.
4. Use the toolbar at the top right corner of the window to switch between views: **Hosts and Clusters**, **VMs and Templates**, **Datastores and VMs** and **Tags**. Depending on the view you select, some objects may not be available. For example, if you select the **VMs and Templates** view, no resource pools, hosts or clusters will be displayed in the tree.
5. In the displayed tree, select the necessary object and click **Add**. Use the **Show full hierarchy** check box to display the hierarchy of all VMware Servers added to Veeam Backup & Replication.
6. Click **OK**.

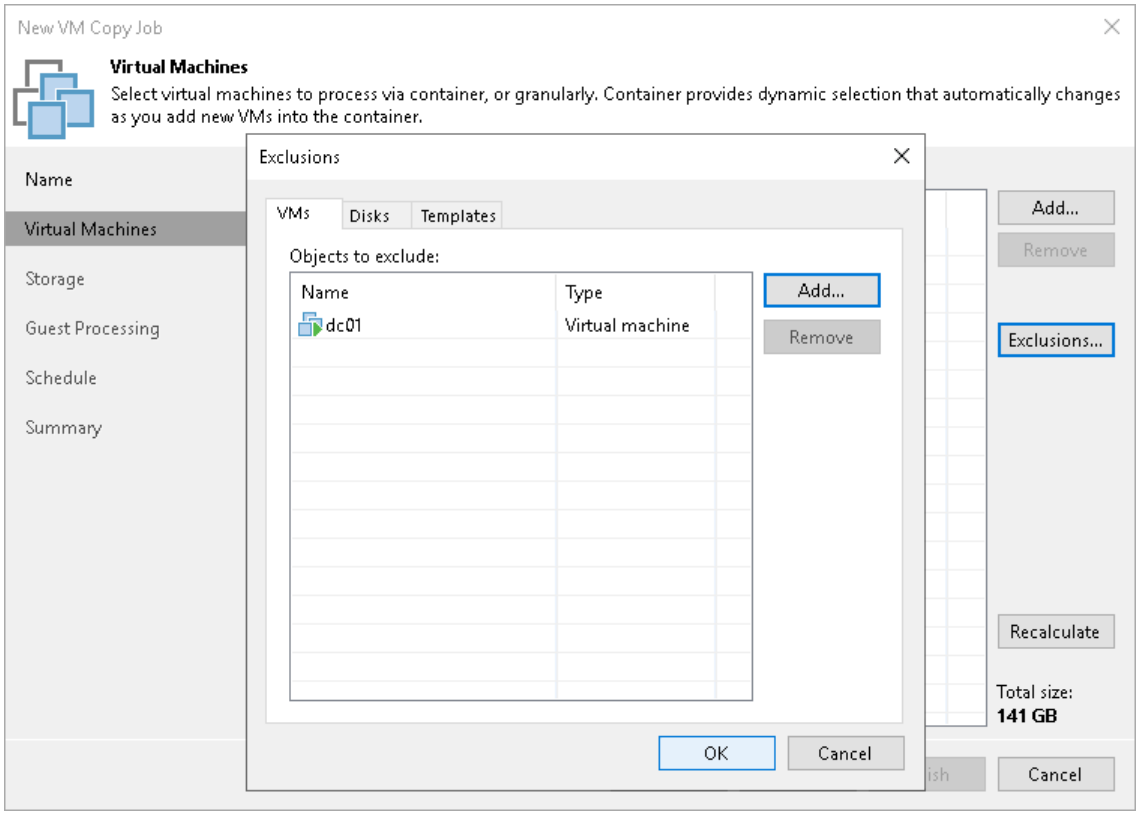

To exclude VM disks:

1. At the **Virtual Machines** step of the wizard, select a VM or VM container added to the job and click **Exclusions**.
2. Click the **Disks** tab.
3. Select the VM in the list and click **Edit**. If you want to exclude disks of a VM added as a part of the container, click **Add** to include the VM in the list as a standalone object.
4. Choose disks that you want to copy. You can choose to process all disks, 0:0 disks (typically, the system disks) or add to the list custom IDE, SCSI or SATA disks.
5. Select the **Remove excluded disks from VM configuration** check box. Veeam Backup & Replication will modify the VMX file of a copied VM to remove excluded disks from the VM configuration. If you use the VM copy to register the VM in a location where excluded disks are not accessible with the original paths, you will not have to manually edit the VM configuration file to be able to power on the VM.

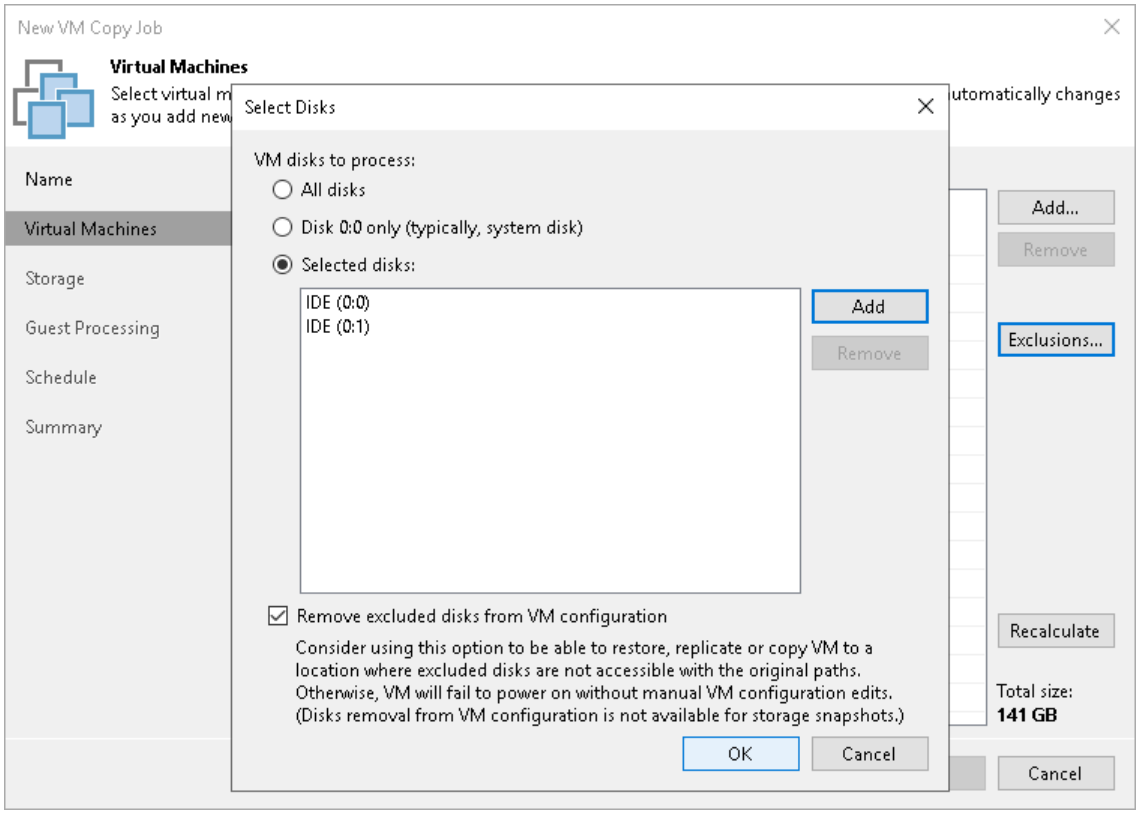


To exclude VM templates:

1. At the **Virtual Machines** step of the wizard, select a VM or VM container added to the job and click **Exclusions**.
2. Click the **Templates** tab.
3. Clear the **Backup VM templates** check box.

4. If you want to include VM templates into the full VM copy only, leave the **Backup VM templates** check box selected and select the **Exclude templates from incremental backup** check box.

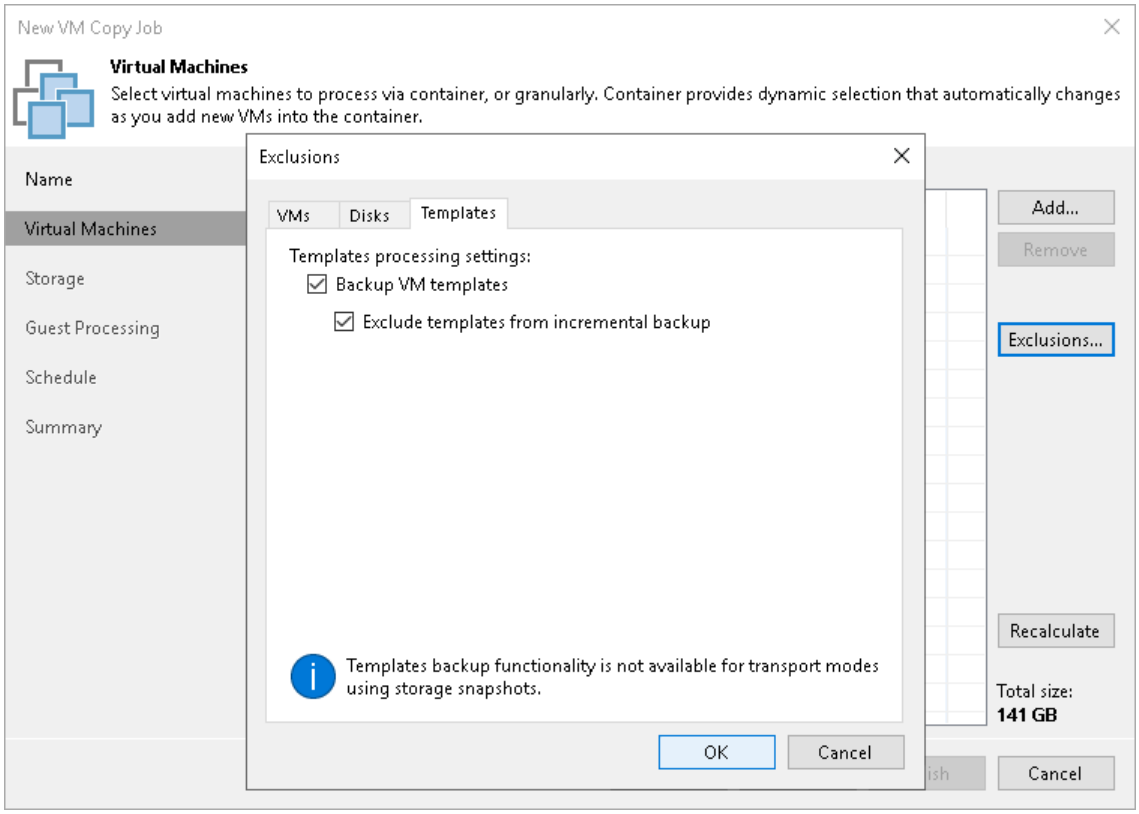

# Step 5. Specify Copy Destination

At the **Storage** step of the wizard, select which backup proxy must be used for VM data transporting and specify the destination for the VM copy.

1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that have access to the source datastore and automatically assign an optimal backup proxy for processing VM data.

     Veeam Backup & Replication assigns backup proxies to VMs included in the VM copy job one by one. Before processing a new VM in the VM list, Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use for data retrieval and the current workload on the backup proxies to select the most appropriate one for VM processing.
   - If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that the job must use. It is recommended that you select at least two backup proxies to ensure that the VM copy job starts if one of the proxies fails or loses its connectivity to the source datastore.
2. In the **Copy destination** section, select a location where the created VM copy must be stored.
   - Select a backup repository from the list if you want to create a VM copy in the backup repository configured in the backup infrastructure. When you select a backup repository, Veeam Backup & Replication automatically checks how much free space is available on it.

     > **IMPORTANT**
     >
     > You cannot use a scale-out backup repository as the copy destination.

- Select **Server** if you want to store the VM copy on a disk or storage device attached to the server. From the **Server** list, select a server added to the backup infrastructure. In the **Path to folder** field, specify a folder on the server where the created VM copy must be stored.

  Use the **Check** button to see how much free space is available in the target location.

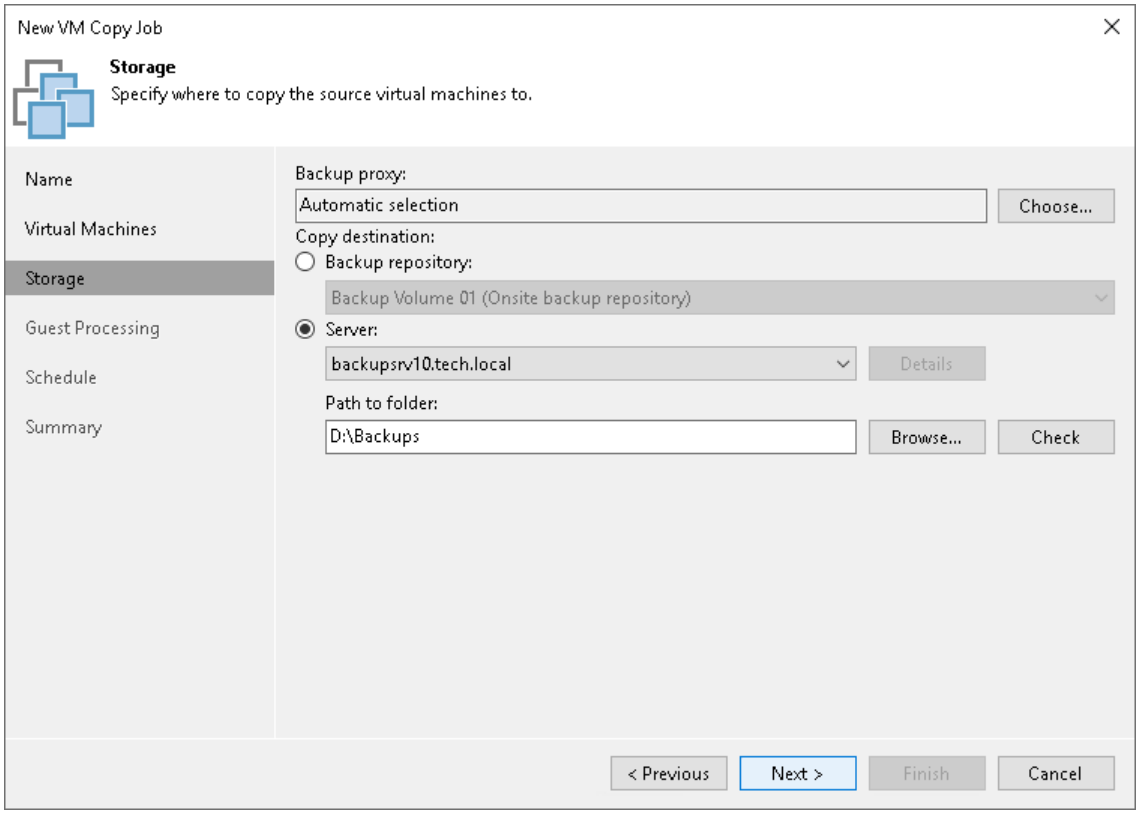

# Step 6. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, you can enable the following settings for VM guest OS processing:

- Application-aware processing
- Microsoft SQL Server transaction log settings
- Oracle archived log settings
- PostgreSQL settings
- Use of pre-freeze and post-thaw scripts

To coordinate guest processing activities, Veeam Backup & Replication deploys non-persistent runtime components or uses (if necessary, deploys) persistent agent components on the VM guest OS.

The non-persistent runtime components run only during guest processing and are stopped immediately after the processing is finished (depending on the selected option, during the VM copy job session or after the replication job completes).

You must specify a user account that will be used to connect to the VM guest OS and deploy the non-persistent runtime components or connect to (if necessary, deploy) persistent agent components:

1. From the **Guest OS credentials** list, select a user account that has enough permissions. For more information on the permissions and requirements for the user account, see Permissions for Guest Processing.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. For more information on adding credentials, see the Credentials Manager section.

2. By default, Veeam Backup & Replication uses the same credentials for all VMs in the job. If some VM requires a different user account, click **Credentials** and enter custom credentials for the VM.

> **IMPORTANT**
>
> Credentials for application-aware processing and guest OS file indexing for Microsoft Windows VMs must be specified in the following format:
>
> - For Active Directory accounts — *DOMAIN*|*Username*
> - For local accounts — *Username* or *HOST*|*Username*

3. If you have added Microsoft Windows VMs to the job, specify which guest interaction proxy Veeam Backup & Replication can use to deploy the non-persistent runtime components or connect to (if necessary, deploy) persistent agent components on the VM guest OS. On the right of the **Guest interaction proxy** field, click **Choose**.
   - Leave **Automatic selection** to let Veeam Backup & Replication automatically select the guest interaction proxy.
   - Select **Use the selected guest interaction proxy servers only** to explicitly define which servers will perform the guest interaction proxy role. The list of servers contains Microsoft Windows servers added to the backup infrastructure.

To check if Veeam Backup & Replication can communicate with VMs added to the job and deploy the non-persistent runtime components or connect to (if necessary, deploy) persistent agent components on their guest OSes, click **Test Now**. Veeam Backup & Replication will use the specified credentials to connect to all VMs in the list.

NOTE

The guest interaction proxy functionality is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.

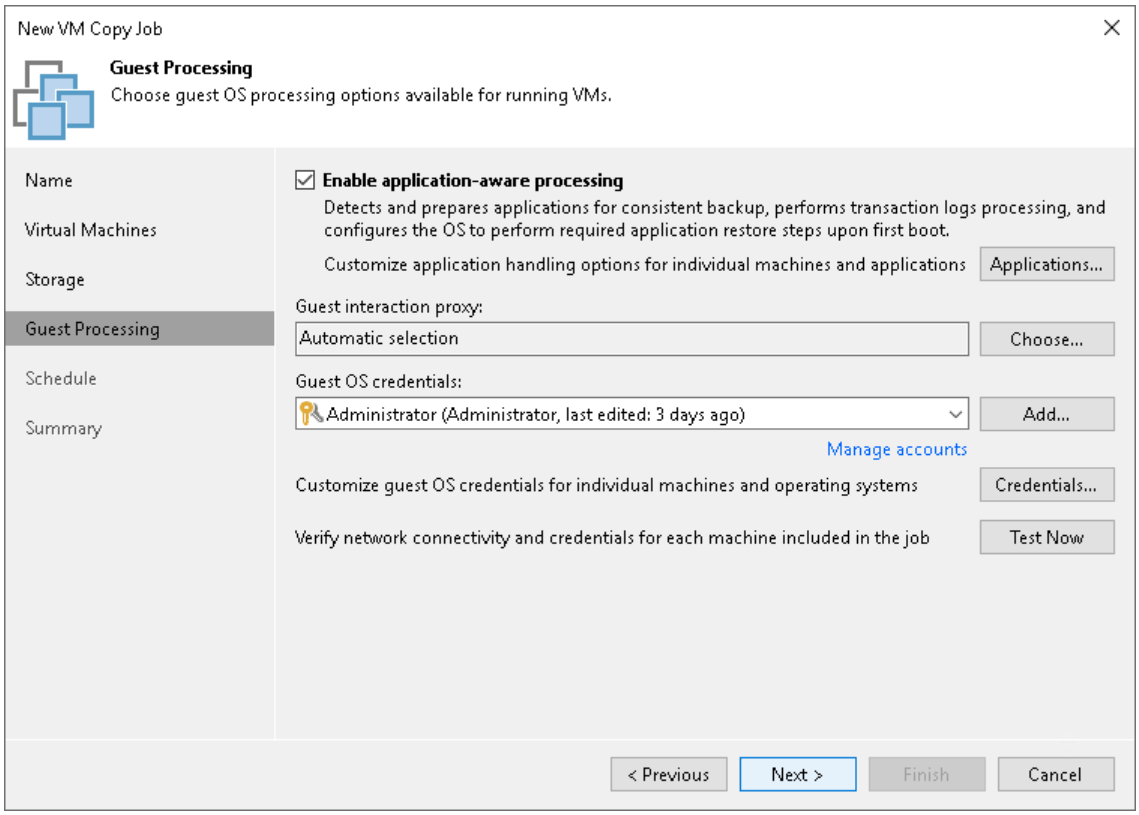


# Application-Aware Processing

If you add to the VM copy job VMs running VSS-aware applications, you can enable application-aware processing to create a transactionally consistent VM copy. The transactionally consistent VM copy guarantees proper recovery of applications on VMs without data loss.

To enable application-aware processing:

1. Select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the VM and click **Edit**.

   To define custom settings for a VM added as a part of the VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose a VM whose settings you want to customize. Then select the VM in the list and define the necessary settings.
4. On the **General** tab, in the **Applications** section specify the VSS behavior scenario:
   - Select **Require successful processing** if you want Veeam Backup & Replication to stop the VM copy process if any VSS errors occur.
   - Select **Try application processing, but ignore failures** if you want to continue the VM copy process even if VSS errors occur. This option is recommended to guarantee completion of the job. The created VM image will not be transactionally consistent but crash consistent.

- Select **Disable application processing** if you do not want to enable quiescence for the VM.

5. [For Microsoft Exchange, Microsoft SQL and Oracle VMs] In the **Transaction logs** section, specify if Veeam Backup & Replication must process transaction logs or copy-only VM copies must be created.

   a. Select **Process transaction logs with this job** if you want Veeam Backup & Replication to process transaction logs.

      [For Microsoft Exchange VMs] With this option selected, the non-persistent runtime components or persistent components running on the VM guest OS will wait for the VM copy job to complete successfully and then trigger truncation of transaction logs. If the VM copy job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

      [For Microsoft SQL Server VMs and Oracle VMs] You will have to specify settings for transaction log handling on the **SQL** and **Oracle** tabs of the **VM Processing Settings** window. For more information, see Microsoft SQL Server Transaction Log Settings and Oracle Archived Log Settings.

   b. Select **Perform copy only** if you use another backup tool to perform VM guest level backup or replication, and this tool maintains consistency of the database state. Veeam Backup & Replication will create a copy-only VM image for the selected VMs. The copy-only VM image preserves the chain of full/differential backup files and transaction logs on the VM. For more information, see Microsoft Docs.

6. In the **Persistent guest agent** section, specify if Veeam Backup & Replication must use persistent guest agents on each protected VM for application-aware processing.

   By default, Veeam Backup & Replication uses non-persistent runtime components.
   Veeam Backup & Replication deploys runtime components on each protected VM when the backup job starts, and removes the runtime components as soon as the backup job finishes.

Select the Use persistent guest agent check box to enable persistent agent components for guest processing. For more information, see Non-Persistent Runtime Components and Persistent Agent Components.

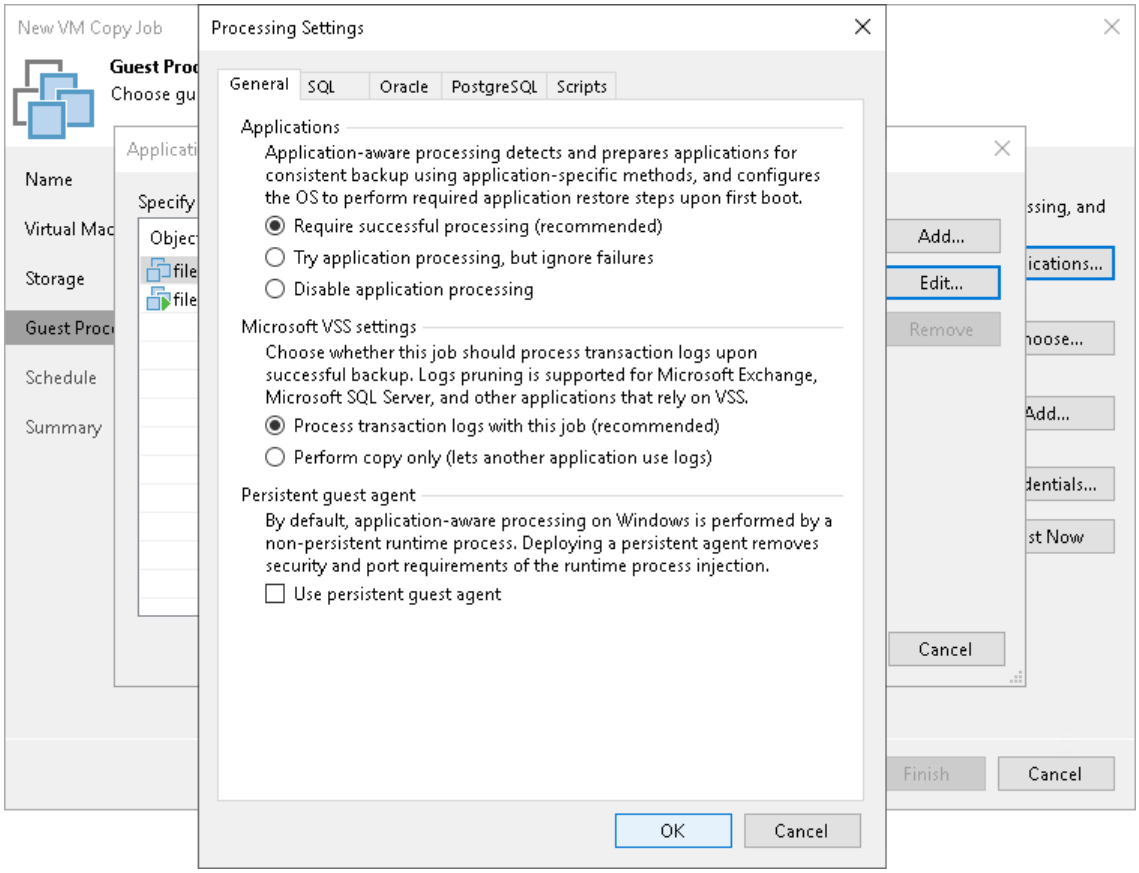


# Microsoft SQL Server Transaction Log Settings

If you copy a Microsoft SQL VM, you can specify how Veeam Backup & Replication must process transaction logs:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Microsoft SQL Server VM and click **Edit**.
4. In the **Transaction logs** section, select **Process transaction logs with this job**.
5. In the **VM Processing Settings** window, click the **SQL** tab.
6. Specify how transaction logs must be processed:
   - Select **Truncate logs** if you want Veeam Backup & Replication to trigger truncation of transaction logs only after the job completes successfully. In this case, the non-persistent runtime components or persistent components will wait for the job to complete and then trigger truncation of transaction logs. If the VM copy job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

- Select **Do not truncate logs** if you do not want Veeam Backup & Replication to truncate logs at all. This option is recommended if you are using another backup tool to perform VM guest-level backup or replication, and this tool maintains consistency of the database state. In such scenario, Veeam Backup & Replication will not trigger transaction log truncation. After you fail over to the necessary restore point of the VM copy, you will be able to apply transaction logs to get the database system to the necessary point in time between VM copy job sessions.

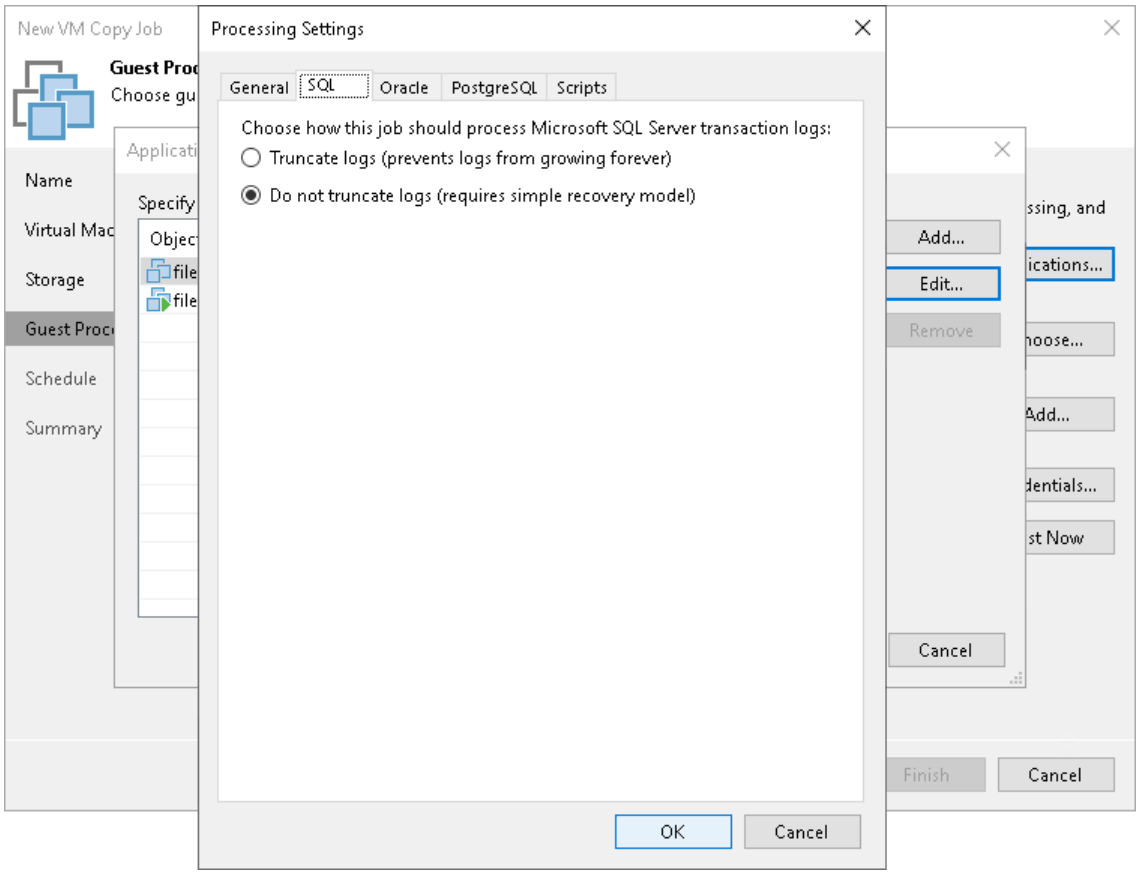


# Oracle Archived Log Settings

If you copy an Oracle VM, you can specify how Veeam Backup & Replication must process transaction logs:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Oracle VM and click **Edit**.
4. In the **Transaction logs** section, select **Process transaction logs with this job**.
5. In the **VM Processing Settings** window, click the **Oracle** tab.
6. In the **Specify Oracle account with SYSDBA privileges** section, specify a user account that Veeam Backup & Replication will use to connect to the Oracle database. The account must have SYSDBA rights on the Oracle database.

   You can select **Use guest credentials** in the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle database.

7. In the **Archived logs** section, specify if Veeam Backup & Replication must truncate transaction logs on the Oracle VM:

   - Select **Do not truncate archived logs** if you want Veeam Backup & Replication to preserve archived logs on the VM guest OS. When the VM copy job completes, the non-persistent runtime components or persistent components will not truncate transaction logs.

     It is recommended that you select this option for databases for which the ARCHIVELOG mode is turned off. If the ARCHIVELOG mode is turned on, transaction logs on the VM guest OS may grow large and consume all disk space. In this case, the database administrators must take care of transaction logs themselves.

   - Select **Truncate logs older than <N> hours** or **Truncate logs over <N> GB** if you want Veeam Backup & Replication to truncate archived logs that are older than <N> hours or larger than <N> GB. The non-persistent runtime components or persistent components running on the VM guest OS will wait for the VM copy job to complete successfully and then trigger transaction logs truncation using Oracle Call Interface (OCI). If the job does not manage to copy the Oracle VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

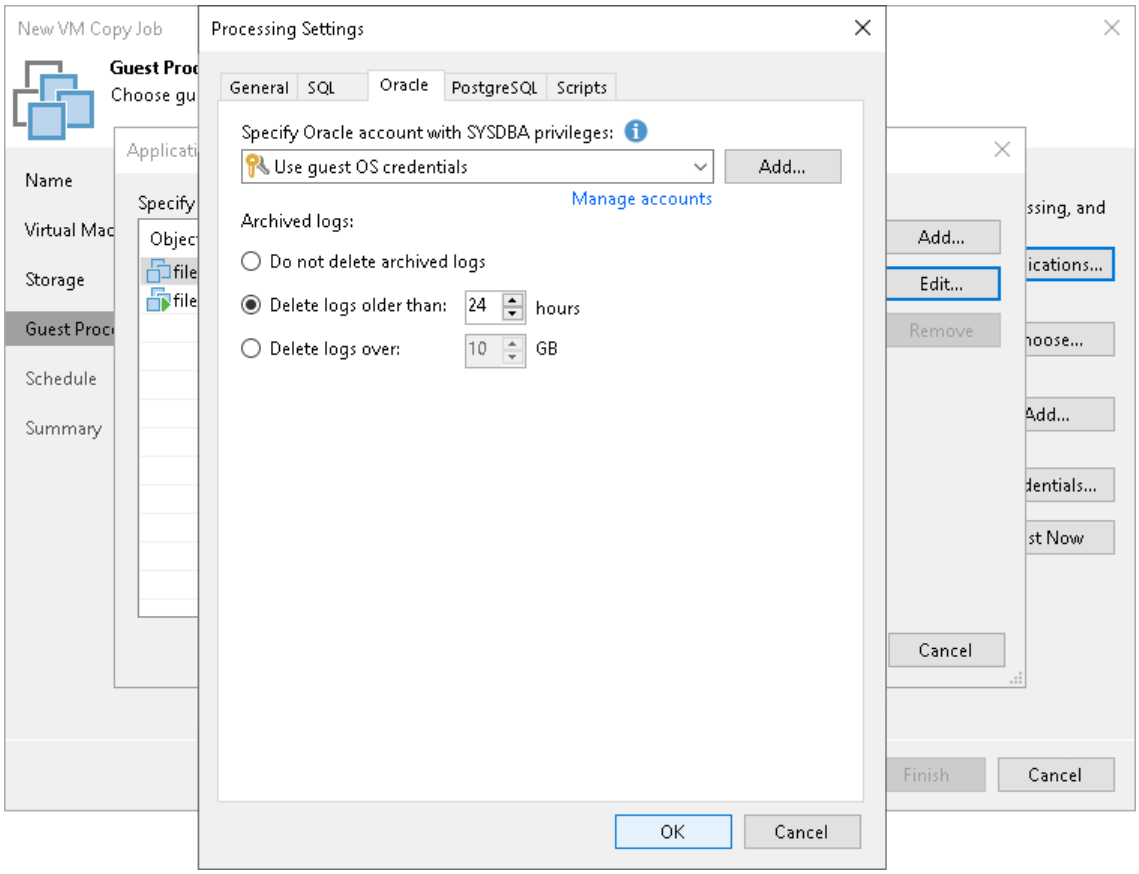


# PostgreSQL WAL Files Settings

To create a transactionally consistent backups of PostgreSQL VM, you must enable application-aware processing and define settings of WAL files processing.

# Enabling Application-Aware Processing

Before configuring WAL files processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the PostgreSQL VM and click **Edit**.

# Specifying WAL Files Settings

To define how Veeam Backup & Replication will process WAL files on this VM, do the following:

1. In the **Processing Settings** window, click the **PostgreSQL** tab.
2. From the **Specify PostgreSQL account with superuser privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the PostgreSQL instance. The account must have privileges described in section Permissions. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add credentials.

   By default, the **Use guest credentials** option is selected in the list. With this option selected, Veeam Backup & Replication will connect to the PostgreSQL instance under the account In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the PostgreSQL instance.

   Note that if you plan to select the **System user without password file (peer)** authentication method at the step 3 of this procedure, you can add a user account in the Credentials Manager without specifying the password for the account.
3. In the **Specified user is** section, specify how the user will authenticate against the PostgreSQL instance:
   - Select **Database user with password** if the account that you specified at the step 2 is a PostgreSQL account, and you entered the password for this account in the **Credentials Manager**.
   - Select **Database user with password file (.pgpass)** if the password for the account that you specified at the step 2 is defined in the `.pgpass` configuration file on the PostgreSQL VM. For more information about the password file, see PostgreSQL documentation.

- Select **System user without password file (peer)** if you want Veeam Backup & Replication to use the peer authentication method. In this case, Veeam Backup & Replication will apply the OS account as the PostgreSQL account.

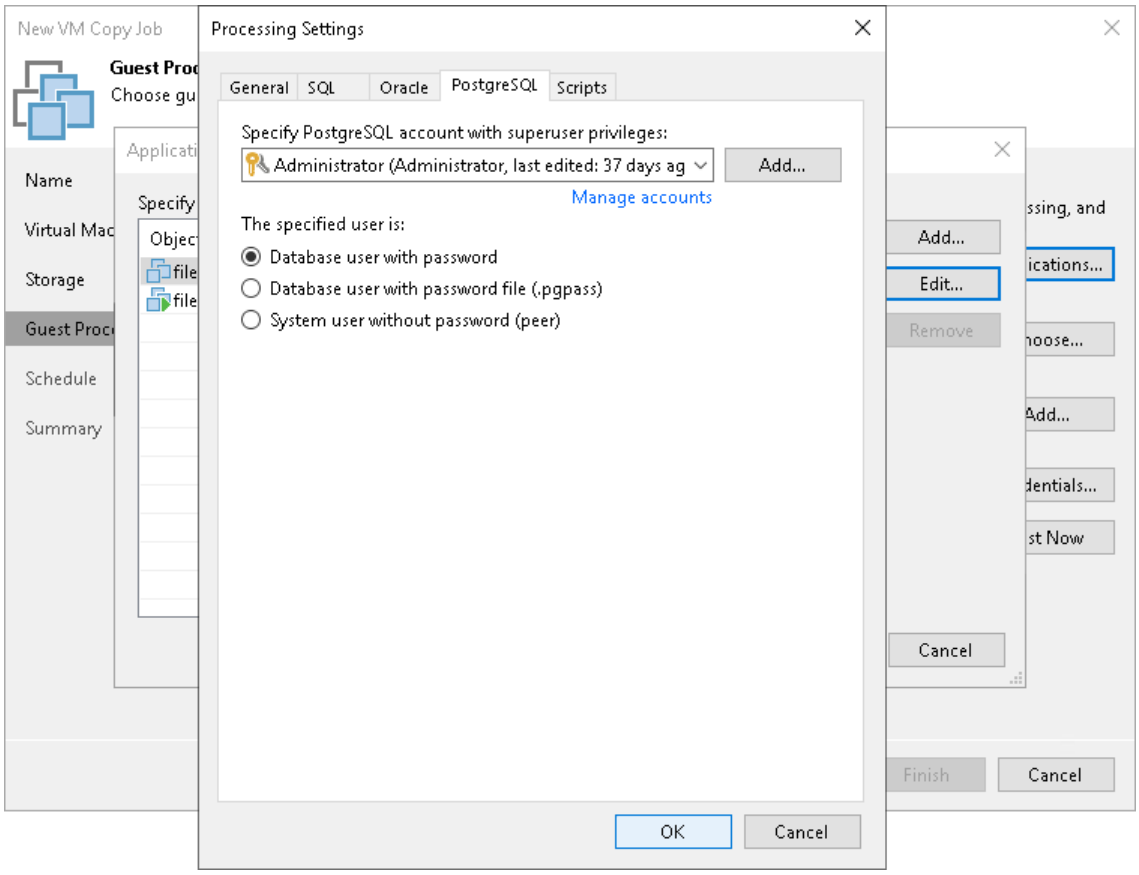


# Pre-Freeze and Post-Thaw Scripts

If you plan to copy VMs running applications that do not support VSS, you can instruct Veeam Backup & Replication to run custom pre-freeze and post-thaw scripts for these VMs. The pre-freeze script quiesces the VM file system and application data to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script brings the VM and applications to their initial state.

To specify pre-freeze and post-thaw scripts for the job:

1. At the **Guest Processing** step, click **Applications**.
2. In the displayed list, select the VM and click **Edit**.
3. Click the **Scripts** tab.
4. In the **Script processing mode** section, specify the scenario for scripts execution:
   - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the VM copy process if the script fails.
   - Select **Ignore script execution failures** if you want to continue the VM copy process even if script errors occur.
   - Select **Disable script execution** if you do not want to run scripts for the VM.

5. In the **Windows scripts** section, specify paths to pre-freeze and post-thaw scripts for Microsoft Windows VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

6. In the **Linux scripts** section, specify paths to pre-freeze and post-thaw scripts for Linux VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

   If you have added to the job a VM container with Microsoft Windows and Linux VMs, you can select to execute both Microsoft Windows and Linux scripts for the VM container. When the job starts, Veeam Backup & Replication will automatically determine what OS type is installed on the VM and apply required scripts to quiesce this VM.

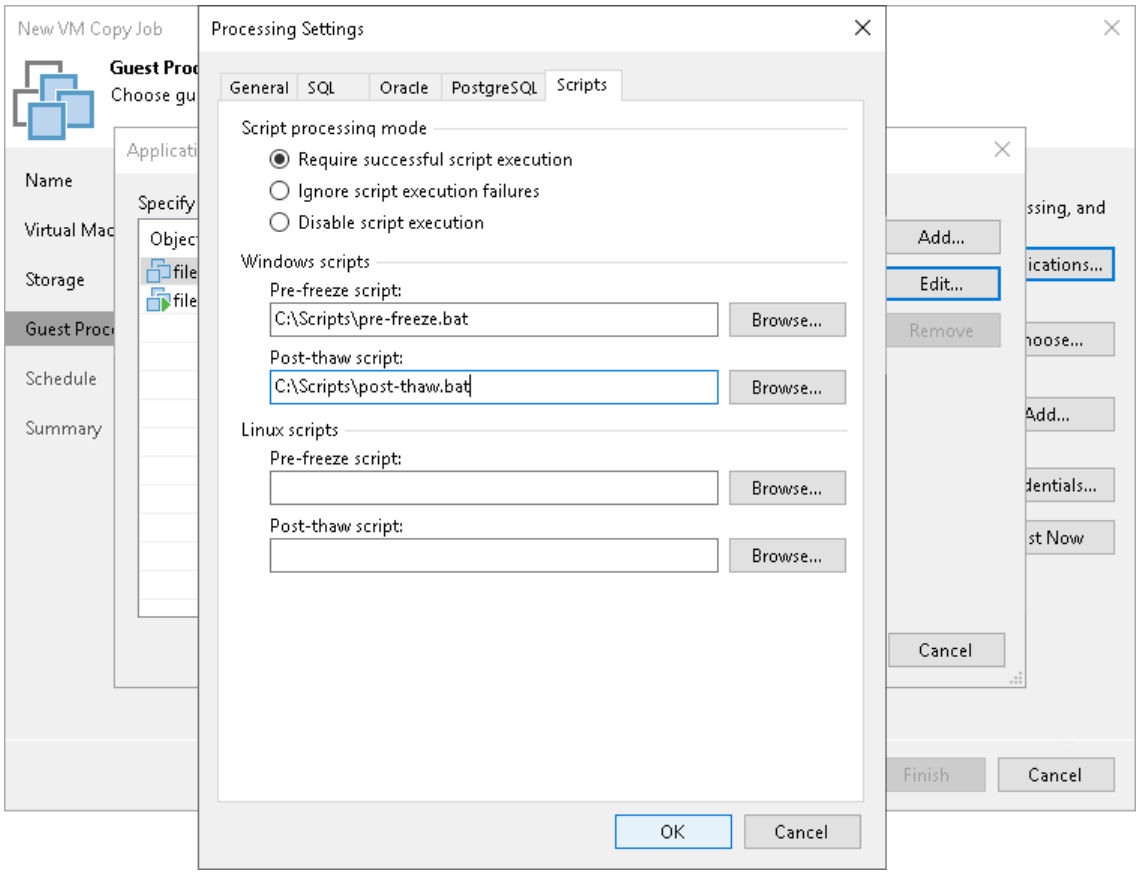

# Step 7. Define Job Schedule

At the **Schedule** step of the wizard, select to run the VM copy job manually or schedule the job to run on a regular basis.

To specify the job schedule:

1. Select the **Run the job automatically** check box. If this check box is not selected, you will have to start the job manually to perform VM replication.
2. Define scheduling settings for the job:
   - To run the job at specific time daily, on defined week days or with specific periodicity, select **Daily at this time**. Use the fields on the right to configure the necessary schedule.
   - To run the job once a month on specific days, select **Monthly at this time**. Use the fields on the right to configure the necessary schedule.

   > **NOTE**
   >
   > When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a set time interval, select **Periodically every**. In the field on the right, select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** and use the time table to define the permitted time window for the job. In the **Start time within an hour** field, specify the exact time when the job must start.

     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run the job by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.

   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.
   - To chain jobs, use the **After this job** field. In the common practice, jobs start one after another: when job *A* finishes, job *B* starts and so on. If you want to create a chain of jobs, you should define the time schedule for the first job in the chain. For the rest of the jobs in the chain, select the **After this job option** and choose the preceding job from the list.
3. In the **Automatic retry** section, define whether Veeam Backup & Replication must attempt to run the job again if the job fails for some reason. During a job retry, Veeam Backup & Replication processes failed VMs only. Enter the number of attempts to run the job and define time spans between them. If you select continuous schedule for the job, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job sessions.

4. In the **Backup window** section, determine a time interval within which the job must be completed. The backup window prevents the job from overlapping with production hours and ensures it does not provide unwanted overhead on your production environment. To set up a backup window for the job:

   a. Select the **Terminate job outside of the backup window** check box and click **Window**.

   b. In the **Time Periods** window, define the allowed hours and prohibited hours for VM copying. If the job exceeds the allowed window, it will be automatically terminated.

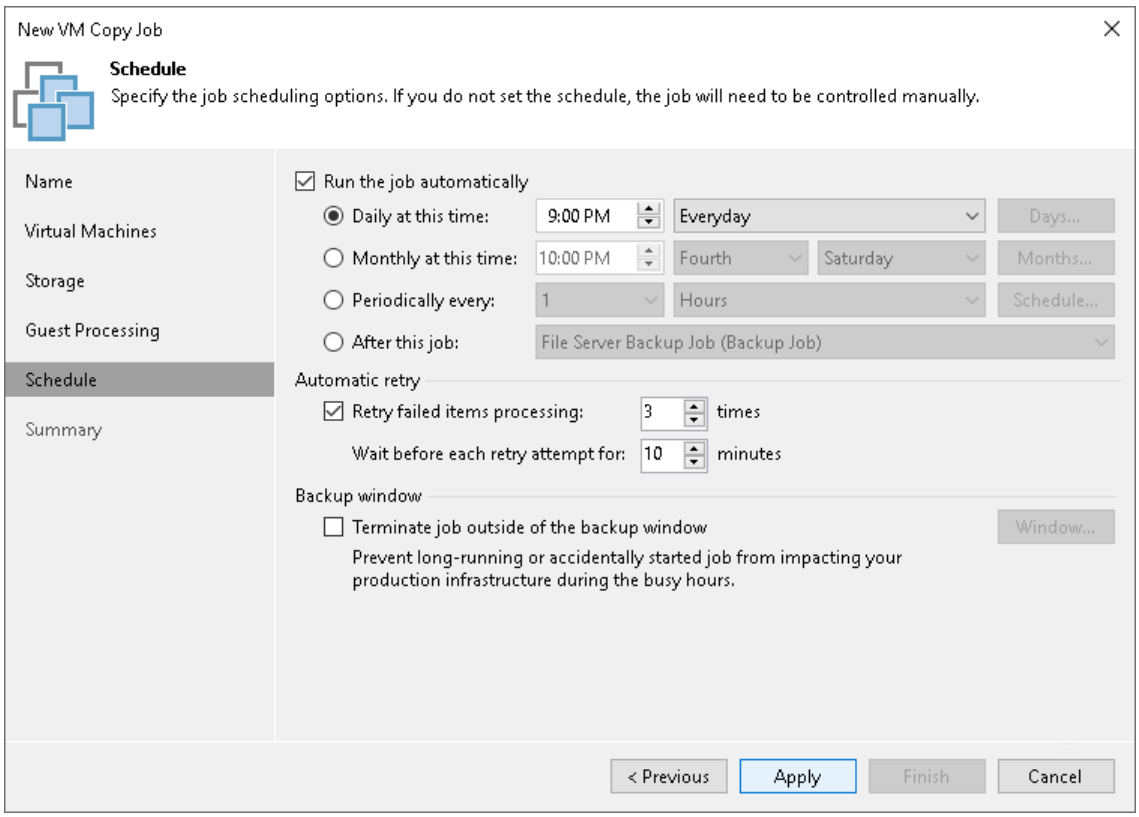

# Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of VM copy job configuration.

1. Review details of the VM copy job.
2. Select the **Run the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

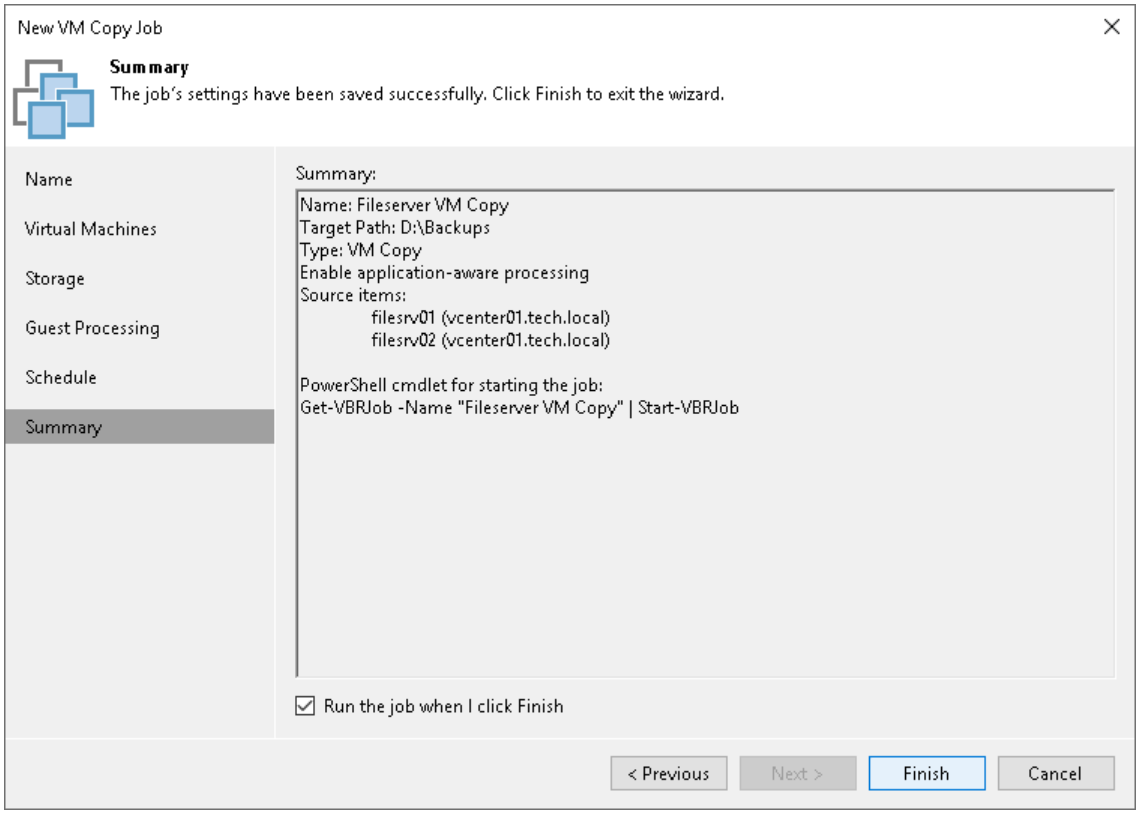

# File Copy

You can copy and move files and folders between servers and hosts added to the backup infrastructure. For file copying operations, Veeam Backup & Replication offers a Windows Explorer-like user interface familiar to a Microsoft Windows user. You can copy files manually or schedule file copy jobs to run automatically by the defined schedule.

The file copy functionality is not intended for creating backups of VM guest OS files. Use backup jobs to create VM image-level backups instead.

# Creating File Copy Jobs

To schedule a copying process for files and folders, you must configure a file copy job. You can run the file copy job immediately after its creation, schedule or save the job.

File copy jobs let you copy files between the following backup infrastructure objects:

- Virtualization hosts
- Microsoft Windows servers
- Linux servers
- ExaGrid storage appliances used as backup repositories
- Quantum DXi
- Fujitsu ETERNUS CS800
- Infinidat InfiniGuard

Before you configure a file copy job, check prerequisites. Then use the **New File Copy Job** wizard to create a job.

# Before You Begin

Before you configure a file copy job, check the following prerequisites:

Backup infrastructure components that will take part in the file copying process must be added to the backup infrastructure and properly configured. These include a source and target host or server between which files and folders will be copied.

Consider the following limitations:

- File copy is not supported for Unix systems, for example, Solaris, FreeBSD and AIX.
- Veeam Backup & Replication does not preserve the Access Control List (ACL) settings for copied guest OS folders. The ACL settings are preserved for files only.

> **TIP**
>
> You can restore the ACL settings for recovered guest OS files and folders using Guest OS File Restore.

# Step 1. Launch New File Copy Job Wizard

To launch the **New File Copy Job** wizard, do either of the following:

- On the **Home** tab, click **Copy Job > File**.
- Open the **Files** view, in the working area right-click the necessary files and select **Add to File Copy Job > New job**. Veeam Backup & Replication will start the **New File Copy Job** wizard and add selected files and folders to this job. You can add other files and folders to the job later on, when you pass through the wizard steps.

You can add files and folders to already existing jobs. To do this, open the **Files** view, in the working area right-click necessary objects and select **Add to File Copy Job > Name of the job**.

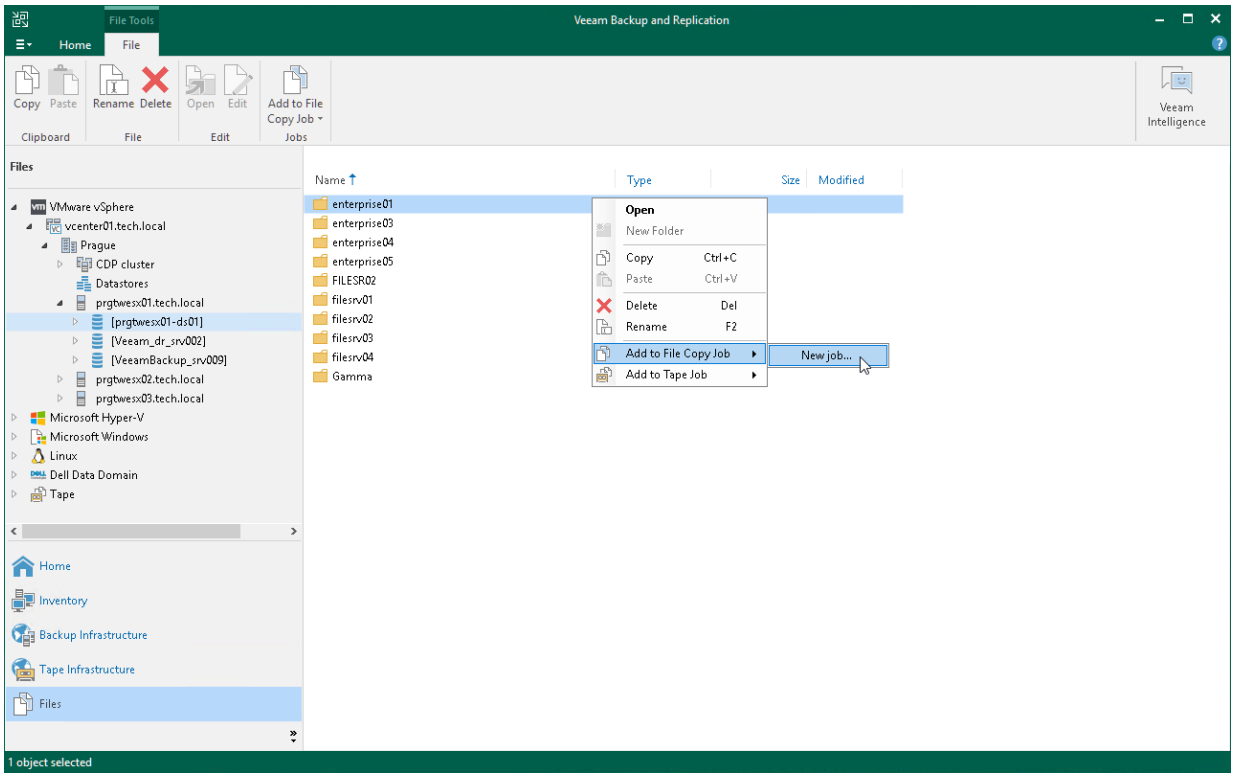

# Step 2. Specify Job Name and Description

At the **Name** step of the wizard, enter a name and description of the created job.

1. In the **Name** field, enter a name for the file copy job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created a job, date and time when the job was created.

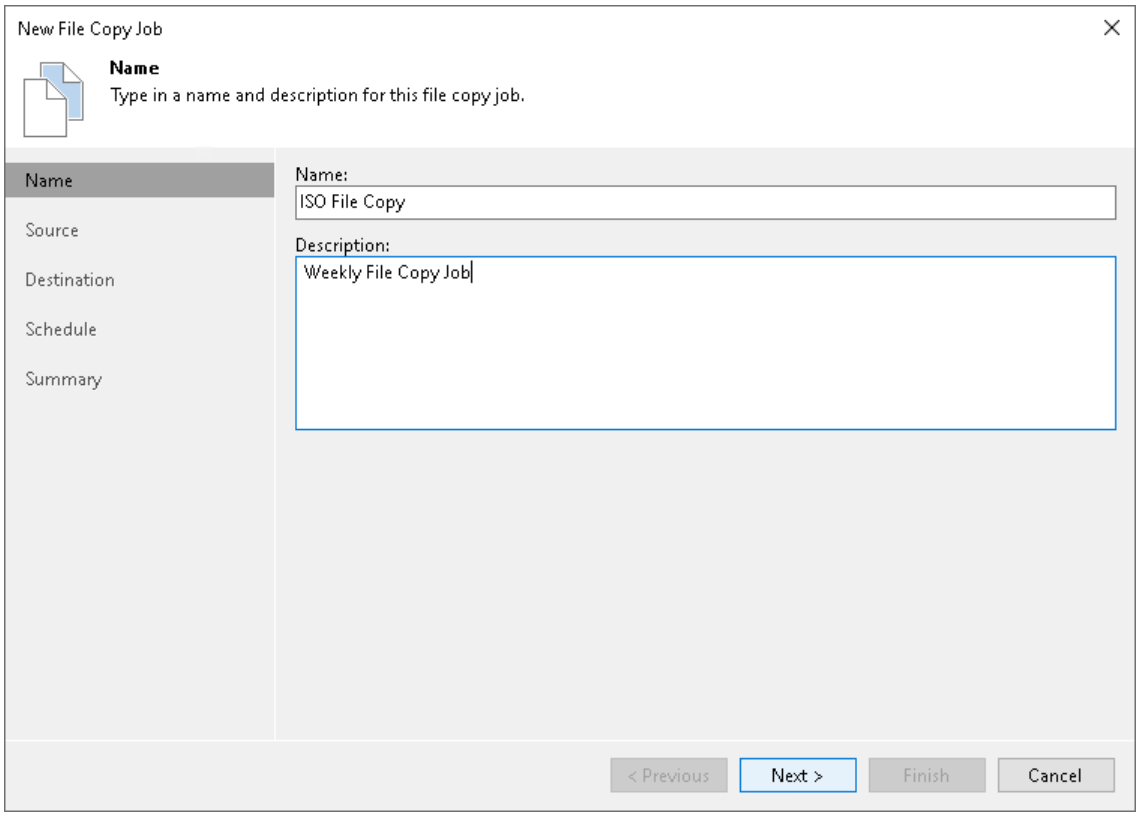

# Step 3. Select Files and Folders to Be Copied

At the **Source** step of the wizard, select files and folders that you want to copy.

To select files and folders that you want to copy:

1. Click **Add** to open the **Select File or Folder** window.
2. From the **Server** list, choose a host or server on which files or folders that you want to copy reside.
3. Select files or folders that must be copied. The selected items will be added to the list.

> **IMPORTANT**
>
> If the list contains files/folders with the same names and extensions, Veeam Backup & Replication copies only one instance of a file/folder. This limitation applies even if you add files/folders from different hosts or servers. To avoid this limitation, you can rename files/folders on the source or add parent folders to the list.

To remove a file or folder from the list, select it and click **Remove**.

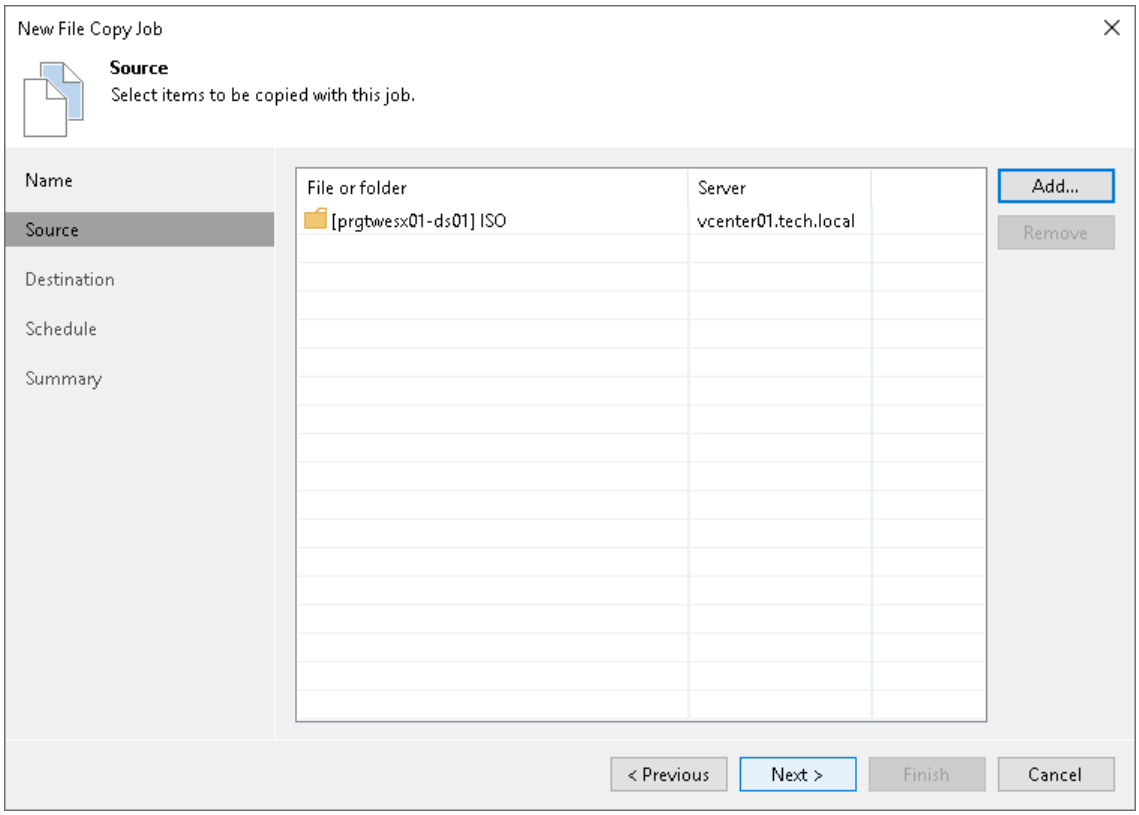

# Step 4. Select Destination for Copying

At the **Destination** step of the wizard, select a destination to which files or folders must be copied.

1. From the **Server** list, select the target host or server.
2. Click **Details** on the right of the **Server** field to view or edit server properties.
3. Click **Browse** next to the **Path to folder** field and select a folder where copied items must be stored. To create a dedicated folder for copied files or folders, use the **New Folder** button at the bottom of the **Select Folder** window.

**IMPORTANT**

If the target folder already contains files/folders with the same names and extensions as the files/folder that must be copied, Veeam Backup & Replication will replace files/folders in the target folder with new files/folders.

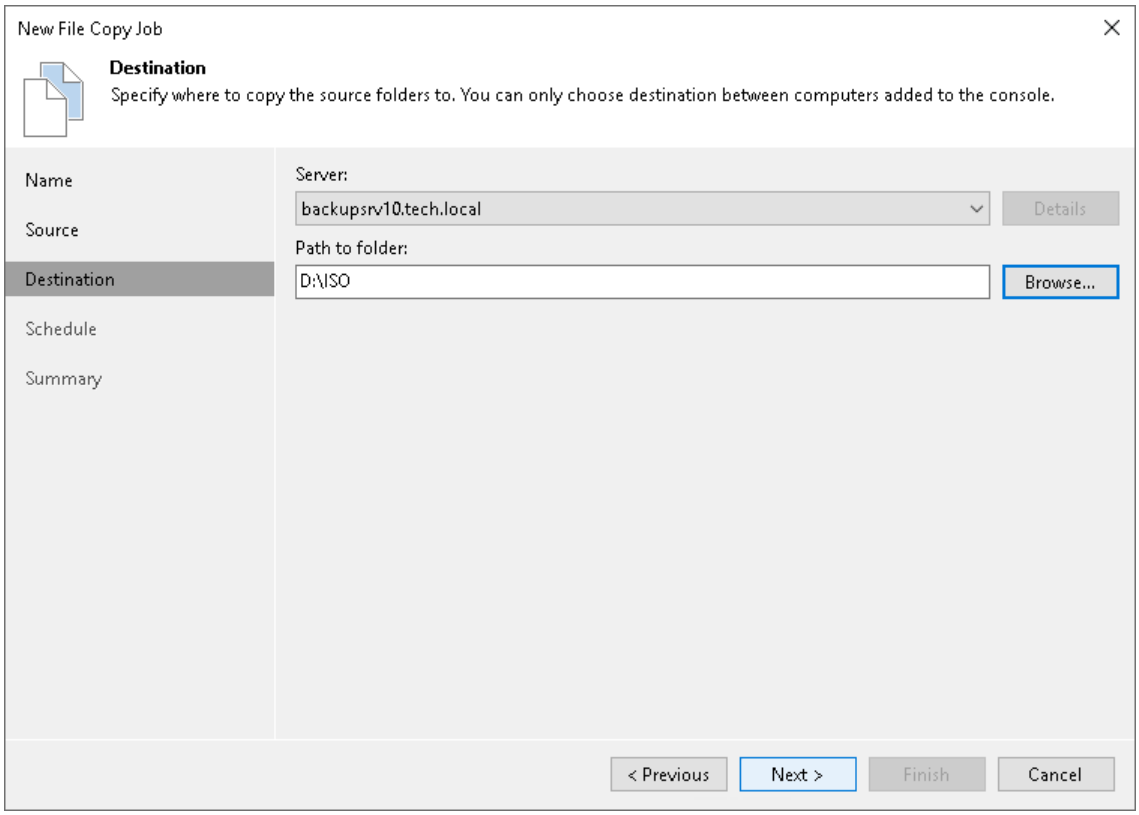

# Step 5. Define Job Schedule

At the **Schedule** step of the wizard, you can select to run the file copy job manually or schedule the job to run on a regular basis.

To specify the job schedule:

1. Select the **Run the job automatically** check box. If this check box is not selected, you will have to start the job manually to copy files or folders.
2. Define scheduling settings for the job:
   - To run the job at specific time daily, on defined week days or with specific periodicity, select the **Daily at this time** option. Use the fields on the right to configure the necessary schedule.
   - To run the job once a month on specific days, select the **Monthly at this time** option. Use the fields on the right to configure the necessary schedule.

     **NOTE**

     When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a set time interval, select the **Periodically every** option. In the field on the right, select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** and use the time table to define the permitted time window for the job. In the **Start time within an hour** field, specify the exact time when the job must start.

     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run the job by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.

   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.

- To chain jobs, use the **After this job** field. In the common practice, jobs start one after another: when job *A* finishes, job *B* starts and so on. If you want to create a chain of jobs, you should define the time schedule for the first job in the chain. For the rest of the jobs in the chain, at the **Schedule** step of the wizard, select the **After this job option** and choose the preceding job from the list.

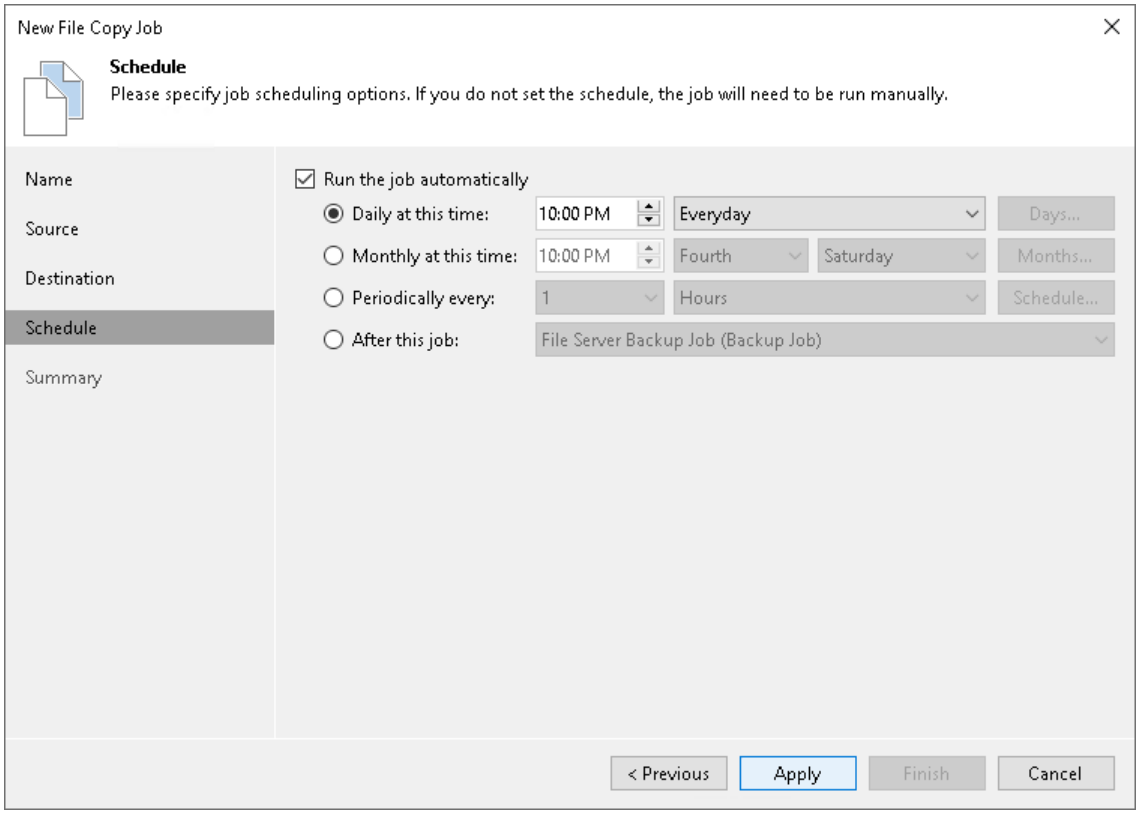

# Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of file copy job configuration.

1. Review details for the created file copy job.
2. Select the **Run the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

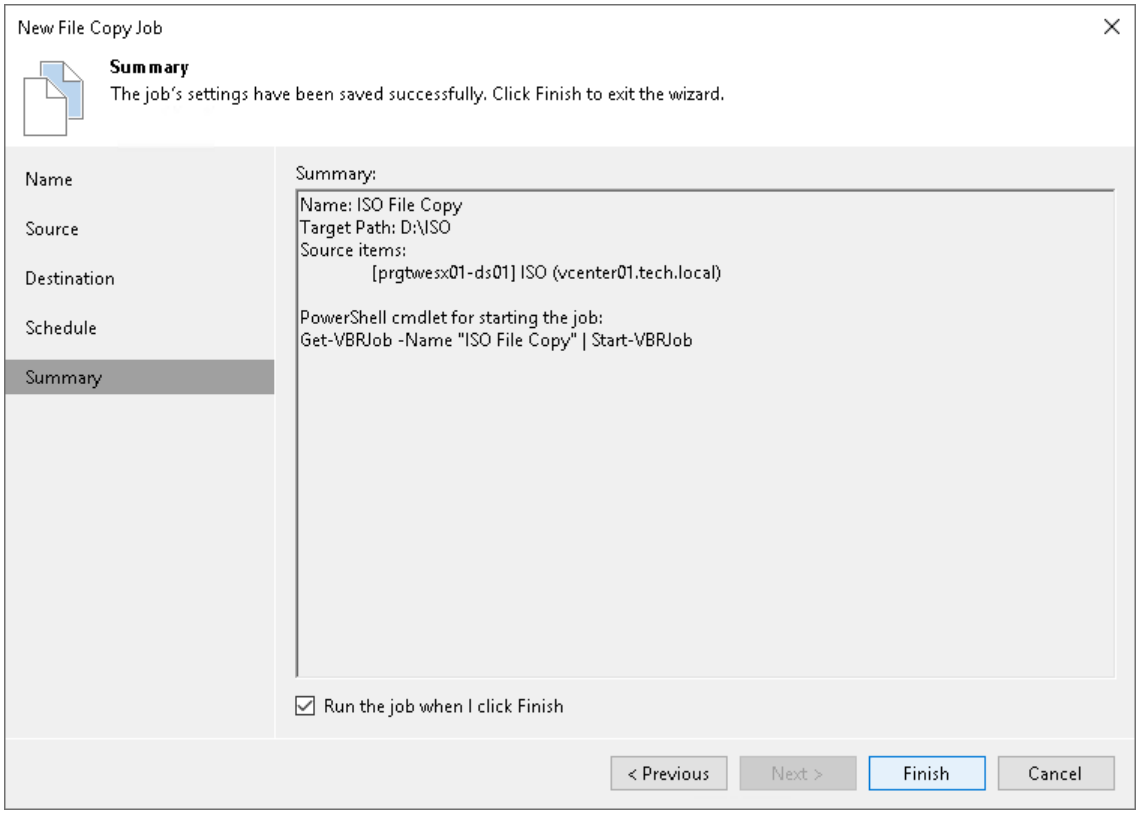

# Copying Files and Folders Manually

You can manually copy and move files and folders between servers and hosts added to the backup infrastructure.

Veeam Backup & Replication lets you copy files manually between the following backup infrastructure objects:

- Virtualization hosts
- Microsoft Windows servers
- Linux servers
- Deduplicating storage appliances used as backup repositories

> **IMPORTANT**
>
> You cannot copy backup files (VBK, VIB and VRB) to HPE StoreOnce storage appliances used as backup repositories. To copy such files, use backup copy jobs.

To copy files and folders:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the source server or host.
3. Right-click files and folders that you want to copy and select **Copy**.
4. In the inventory pane, expand the file tree of the target server or host.
5. Right-click a destination folder and select **Paste**.

You can also use a drag-n-drop operation to copy files and folders between the source and target hosts or servers.

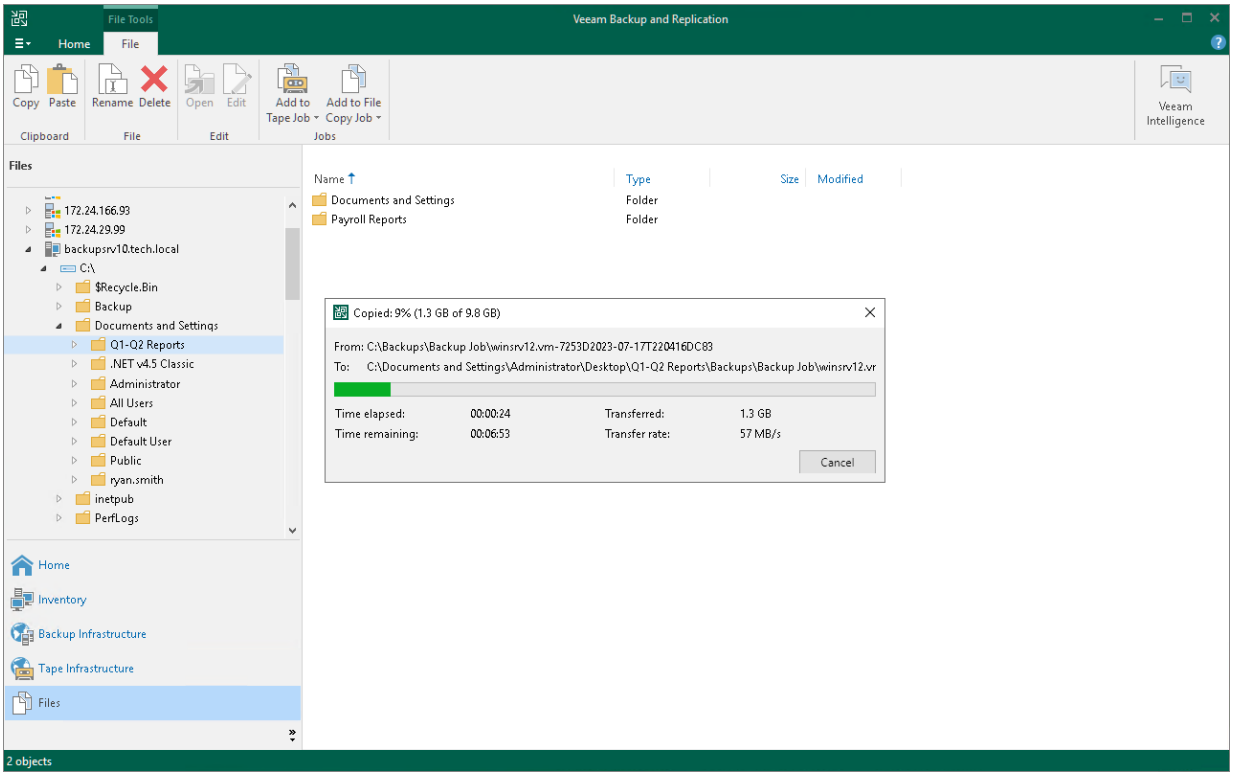

# Managing Folders

You can create, rename and delete folders in the **Files** view of Veeam Backup & Replication.

To create a folder:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the necessary server or host.
3. In the working area, right-click anywhere on the blank area and select **New Folder**.

To rename a folder:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the necessary server or host.
3. In the working area, select the folder and click **Rename** on the ribbon or right-click the folder and select **Rename**.
4. Enter a new name for the folder and press [Enter] on the keyboard.

To remove a folder:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the necessary server or host.
3. In the working area, select the folder and click **Delete** on the ribbon or right-click the folder and select **Delete**.

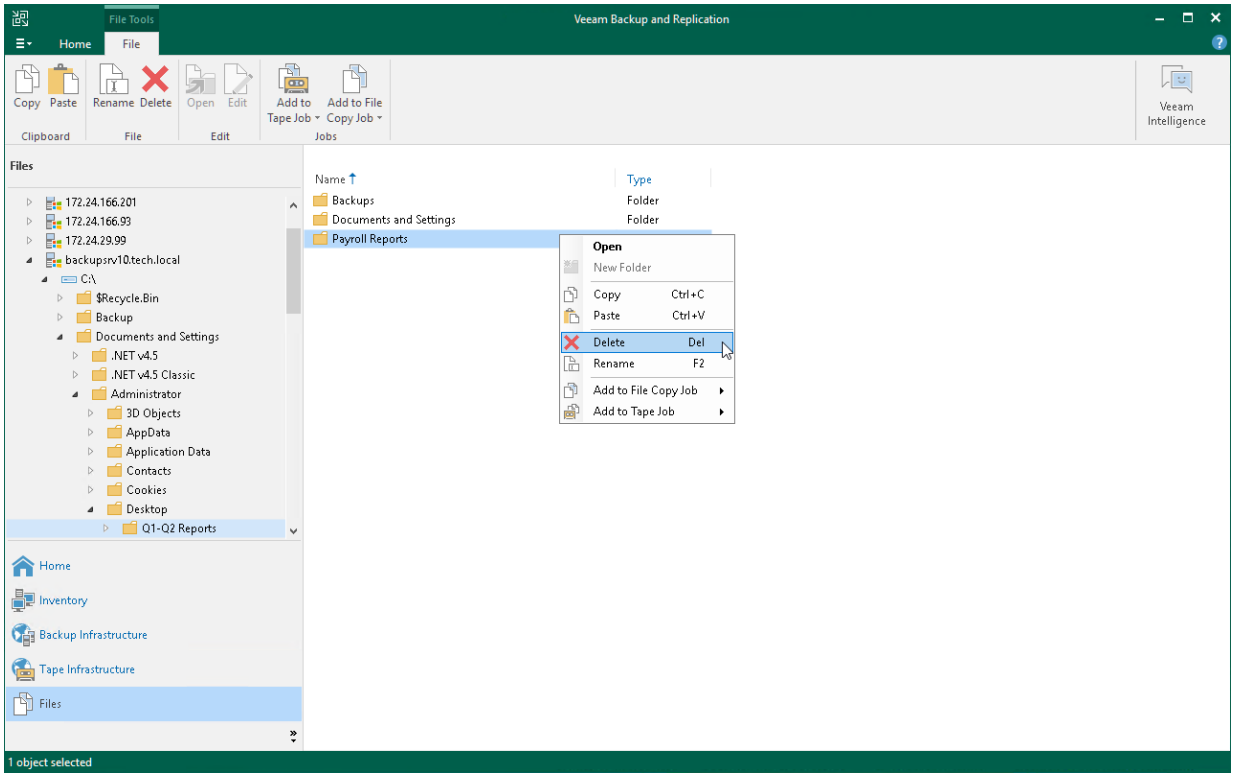

# Editing and Deleting Files

You can edit files and delete them in the **Files** view of Veeam Backup & Replication. For example, you may want to edit a configuration file of the VM (VMX) or need to delete from the storage files of unused VMs.

To edit a file:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the necessary server or host.
3. In the working area, select the file and click **Edit** on the ribbon or right-click the folder and select **Edit**.
4. Veeam Backup & Replication will open the selected file in the editor. Edit the file as required and click **Save** on the file editor toolbar.

To delete a file:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the necessary server or host.
3. In the working area, select the file and click **Delete** on the ribbon or right-click the folder and select **Delete**.

> **NOTE**
>
> To delete a folder on VSAN, you must remove a real folder, not a symbolic link to this folder. The real folder is named with GUID, for example, *c07a2953-8096-5b20-a11a-002590c5857c*, while the symbolic link contains the folder name, for example, *srv02_vm*. If you delete the folder symbolic link, the delete operation will fail, and the folder will not be removed.

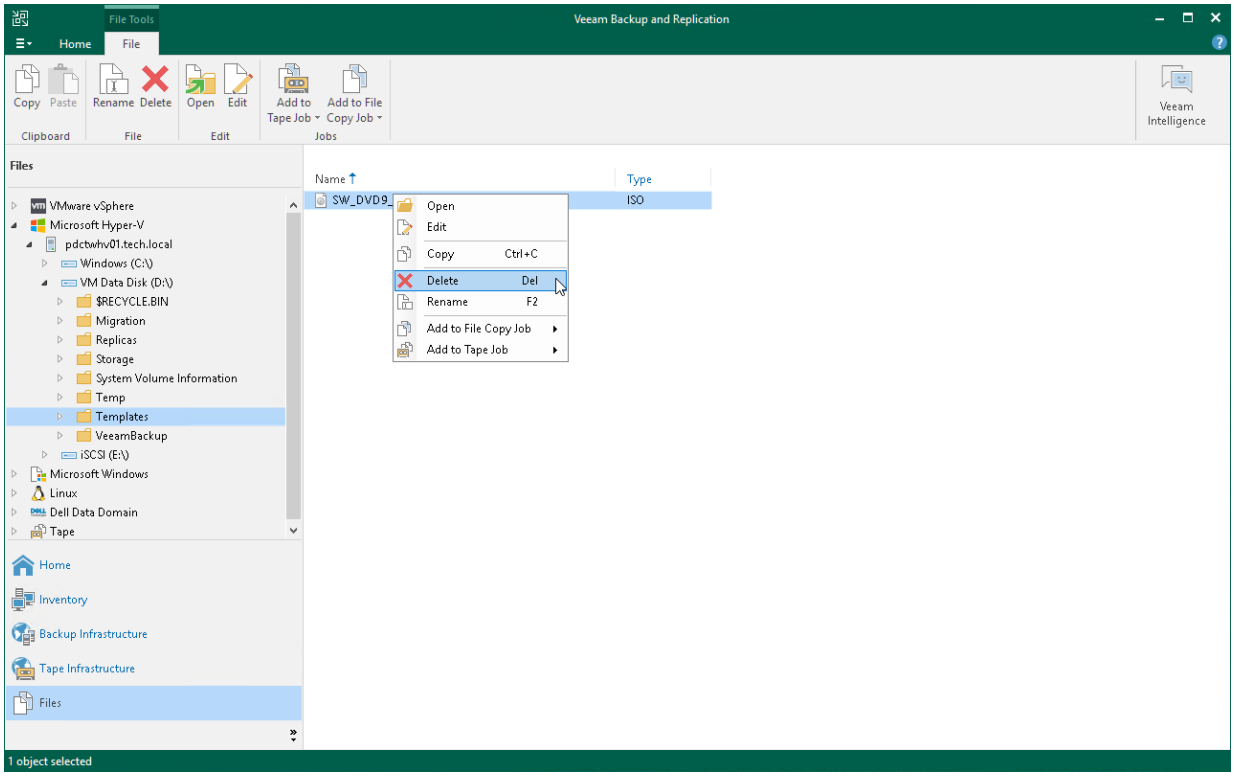

# Quick Migration

Quick Migration allows you to migrate VMs or virtual disks between ESXi hosts and datastores.
Veeam Backup & Replication supports migration of VMs or their disks in any state with minimum disruption to business operations. You can use Quick Migration as a self-contained capability or as a way to finalize the Instant Recovery and Instant Disk Recovery processes.

When you perform Quick Migration, Veeam Backup & Replication analyzes your virtual environment, its configuration, the state of VMs and selects the most appropriate VM relocation method:

- **vMotion and Storage vMotion**

  vMotion and Storage vMotion are native migration mechanisms of VMware vCenter.
  Veeam Backup & Replication uses these methods whenever it is possible.

- **Veeam Quick Migration**

  Veeam Quick Migration is the Veeam Backup & Replication proprietary technology.
  Veeam Backup & Replication uses this method when VMware vCenter methods cannot be used. For example, if your VMware vSphere license does not provide support for vMotion and Storage vMotion, or you need to migrate VMs from one standalone ESXi host to another.

  Veeam Quick Migration supports two modes of VM migration:

  - **SmartSwitch**

    With SmartSwitch, Veeam Backup & Replication suspends a VM, then moves the VM configuration file and copies changes made to the VM disk after snapshot creation to the target host. After the migration is completed, the VM is resumed on the target host.

  - **ColdMigration**

    With ColdMigration, Veeam Backup & Replication stops the VM, then copies changes made to the VM disk after snapshot creation to the new host. After, the VM is started on the target host.

## Veeam Quick Migration of VMs

Migration of a VM using the Veeam Quick Migration method includes the following steps:

1. Veeam Backup & Replication copies VM configuration file (.VMX) to the target host and registers the VM.
2. Veeam Backup & Replication triggers a VM snapshot creation and copies VM disk content to the new destination.
3. Veeam Backup & Replication uses different modes when moving the VM between hosts with compatible and non-compatible CPUs.
   - If you move a VM between two hosts with compatible CPUs, Veeam Backup & Replication uses the SmartSwitch mode.
   - If you move a VM between two hosts with non-compatible CPUs or VM RAM is more than 8 GB, Veeam Backup & Replication uses the ColdMigration mode.

## Veeam Quick Migration of Virtual Disks

Migration of a VM disk using the Veeam Quick Migration method includes the following steps:

1. A temporary VM is created on the target datastore.

2. Veeam Backup & Replication copies the disk from the original datastore to the temporary VM on the target datastore.
3. If the original VM is powered on, Veeam Backup & Replication suspends it.
4. If changes were made to the original disk during the copy process performed at step 3, Veeam Backup & Replication copies these changes to the disk of the temporary VM.
5. On the original VM, Veeam Backup & Replication replaces old path to the disk with the path to the disk of the temporary VM.
6. If the original VM was suspended during the migration, Veeam Backup & Replication powers the VM on.

**IMPORTANT**

Before the migration, VM disks must be recovered using Instant Disk Recovery.

# Veeam Quick Migration of First Class Disks (FCDs)

Migration of First Class Disks (FCDs) supports only the vMotion method and includes the following steps:

1. Veeam Backup & Replication gets information about the target datastore.
2. Veeam Backup & Replication connects to one of the ESXi hosts of a cluster to which the target datastore is mounted.
3. Veeam Backup & Replication migrates FCDs from the source datastore to the target datastore.
4. If redo logs are stored on a custom datastore, Veeam Backup & Replication will delete FCDs snapshots after migration.
5. Veeam Backup & Replication assigns storage policies to FCDs.

# Quick Migration Architecture

Quick Migration architecture in a VMware vSphere environment comprises the following components:

- Source host and target host with associated datastores
- One or two backup proxy servers

Similar to the backup architecture, Quick Migration uses two-service architecture: the source-side Veeam Data Mover interacts with the source host, and the target-side Veeam Data Mover interacts with the target host. To perform on-site migration, you can deploy one backup proxy for data processing and transfer. This backup proxy must have access to the source host and to the target host at the same time. In this scenario, the source-side Veeam Data Mover and the target-side Veeam Data Mover are started on the same backup proxy.

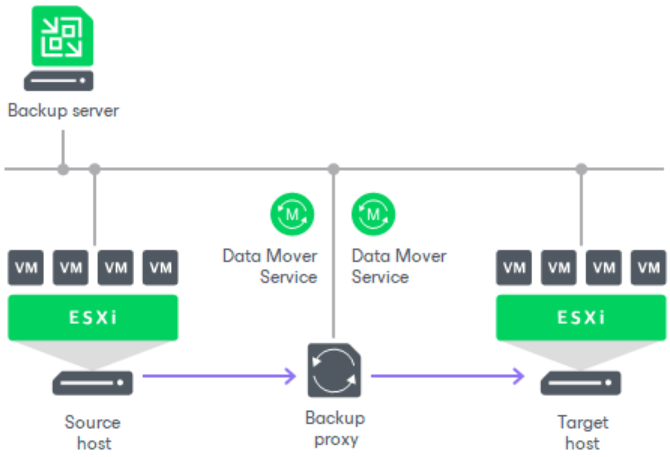


The common requirement for off-site migration is that one Veeam Data Mover runs in the production site (closer to the source host and datastore), and the other Veeam Data Mover runs in the remote target site (closer to the target host and datastore). During backup, Veeam Data Movers maintain a stable connection, which allows for uninterrupted operation over WAN or slow links.

For off-site migration, you need to deploy at least one local backup proxy in each site: a source backup proxy in the production site, and a target backup proxy in the remote target site.

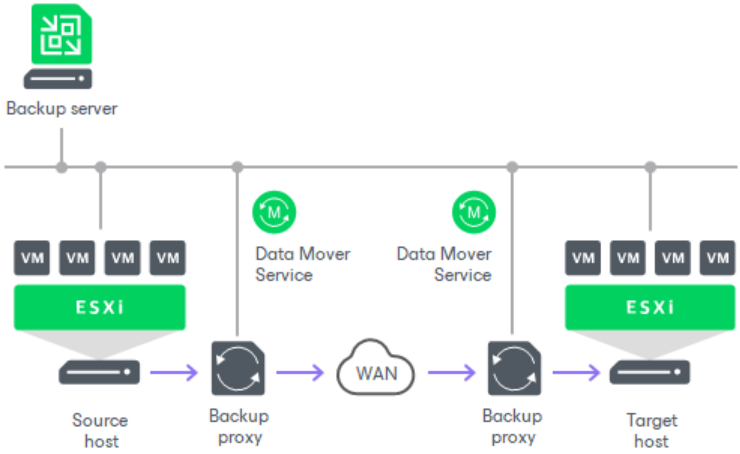

# Migrating VMs

You can relocate one or more VMs with Quick Migration. Quick migration can be used to move VMs from one ESXi host to another one. You can perform "hot" Quick Migration for running VMs or "cold" Quick Migration for VMs that are powered off.

Quick Migration is not job-driven: it cannot be saved as a job or scheduled to run later.
Veeam Backup & Replication will start relocating VMs immediately after you finish working with the **Quick Migration** wizard.

Before you start Quick Migration, check prerequisites. Then use the **Quick Migration** wizard to migrate VMs.

# Before You Begin

Before you perform Quick Migration, check the following prerequisites and limitations:

- Backup infrastructure components that will take part in Quick Migration must be added to the backup infrastructure and properly configured. These include the source and target ESXi hosts.
- The target datastore must have enough free space to store disks of the VMs that you plan to migrate. To receive alerts about low space on the target datastore, configure global notification settings. For more information, see Specifying Other Notification Settings.
- Quick Migration is not possible for VMs that use SCSI bus sharing. This is because Quick Migration requires snapshots, which VMware does not support for such VMs.
- If you perform Quick Migration from a vSAN datastore, or if you store redo logs on a vSAN datastore during Instant Recovery, expect significant delays during migration. That is because of the specific way vSAN organizes data storage. Veeam Backup & Replication cannot get the difference between the source and target VM. Veeam Backup & Replication needs to read the entire disks instead of delta disks.
- If you want to use VMware vSphere vMotion to relocate VMs between hosts and VMware vSphere Storage vMotion to relocate VM disks between datastores, make sure that you have a VMware vSphere license covering these features.
- If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.

## Encryption

Veeam Backup & Replication does not keep encryption settings if a VM is migrated with VMware vMotion. After the migration process is finished, you will need to enable encryption for the migrated VM manually.

## Integration with Instant Recovery

When you restore a VM using Instant Recovery, Veeam Backup & Replication starts the VM directly from a compressed and deduplicated backup file. To finalize recovery of a VM, you still need to move it to a new location. Moving the VM with VMware Storage vMotion or hot replication may require a lot of time and resources, or it may cause loss of valuable data.

Veeam Quick Migration was designed to complement Instant Recovery. Instead of pulling data from vPower NFS datastore, Quick Migration registers the VM on the target host, restores the VM contents from the backup file located in the backup repository and synchronizes the VM restored from backup with the running VM.

For more information, see Instant Recovery to VMware vSphere.

> **NOTE**
>
> Virtual appliance (HotAdd) transport mode cannot be used if the role of the backup proxy and mount server or backup repository where the backup file is stored are assigned to the same VM.

# Step 1. Launch Quick Migration Wizard

If you use Quick Migration as a way to finalize the Instant Recovery or Instant Disk Recovery processes, you must launch Quick Migration as described in sections Finalizing Instant Recovery to VMware vSphere and Finalizing Instant Disk Recovery.

To launch Quick Migration as a self-contained capability:

1. Open the **Inventory** view.
2. In the infrastructure tree, select a host or VM container (host, cluster, folder, resource pool, VirtualApp, datastore or tag) in which the VMs that you want to relocate reside.
3. In the working area, select the VM and click **Quick Migration** on the ribbon or right-click the VMs and select **Quick Migration**.

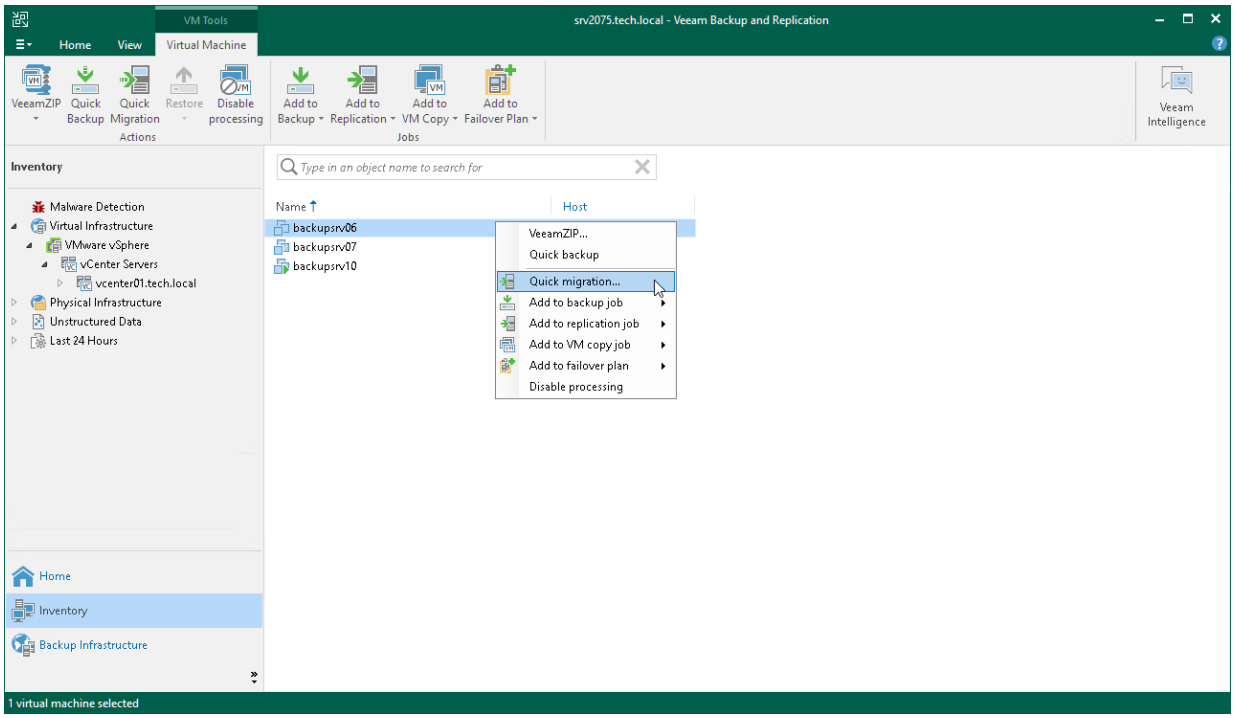

# Step 2. Select VMs to Relocate

At the **Virtual Machines** step of the wizard, select the VMs and VM containers that you want to relocate:

1. Click **Add**.
2. Use the toolbar at the top right corner of the window to switch between views: **Hosts and Clusters**, **VMs and Templates**, **Datastores and VMs** and **Tags**. Depending on the view you select, some objects may not be available. For example, if you select the **VMs and Templates** view, no resource pools, hosts or clusters will be displayed in the tree.
3. Select the necessary object and click **Add**.

To quickly find the necessary object, you can use the search field at the bottom of the **Add Objects** window.

1. Click the button to the left of the search field and select the necessary type of object to search for: *Everything*, *Folder*, *Cluster*, *Host*, *Resource pool*, *VirtualApp* or *Virtual machine*.
2. Enter the object name or a part of it in the search field.
3. Click the **Start search** button on the right or press [Enter] on the keyboard.

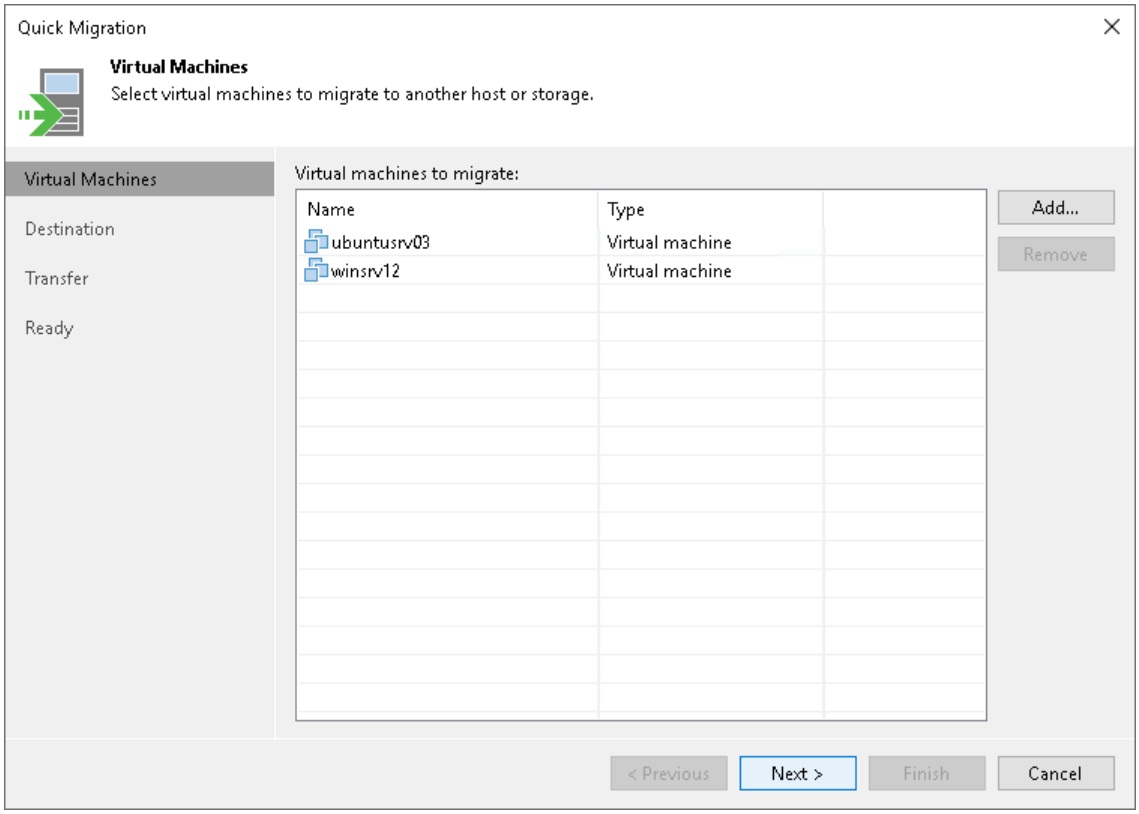

# Step 3. Specify VM Destination

At the **Destination** step of the wizard, select the destination to which the selected VMs must be relocated:

1. Click **Choose** next to the **Host or cluster** field and select an ESXi host or cluster where the relocated VM must be registered.
2. If all or majority of relocated VMs must belong to the same resource pool, click **Choose** next to the **Resource pool** field and select the target resource pool.

   If you want to place relocated VMs to different resource pools:

   a. Click the **Pick resource pool for selected VMs** link.

   b. In the **Choose Resource Pool** window, click **Add VM** on the right and select the VMs.

   c. Select the added VM in the **VM resource pool** list and click **Resource Pool** at the bottom of the window.

   d. From the list of available resource pools, select the target resource pool.
3. If all or majority of relocated VMs must be placed to the same folder, click **Choose** and select the folder.

   If you want to place relocated VMs to different folders:

   a. Click the **Pick VM folder for selected VMs** link.

   b. In the **Choose Folder** window, click **Add VM** on the right and select the VMs.

   c. Select the added VM in the **VM folder** list and click **VM Folder** at the bottom of the window.

   d. From the list of available folders, select the target folder.

   The **VM folder** section is disabled if you selected a standalone ESXi host as a target for VM relocation.
4. If all or majority of relocated VMs must be stored on the same datastore, click **Choose** and select the datastore. Veeam Backup & Replication displays only those datastores that are accessible by the selected ESXi host. If you have chosen to relocate VMs to a cluster, Veeam Backup & Replication will display only shared datastores.

   > **IMPORTANT**
   >
   > [For Instant Recovery finalization] If you migrate VMs to the same datastore cluster that is used as the destination for redirecting virtual disk updates (the Datastore step of the **Instant Recovery to VMware** wizard), you must enable the **Force Veeam transport usage** check box at the Transfer step of the **Quick Migration** wizard. Veeam Backup & Replication will use Veeam Quick Migration instead of Storage vMotion. This will help to prevent data loss due to a bug in VMware Storage vMotion. If you do not enable the **Force Veeam transport usage** check box, your relocated VM may be deleted.

   If you want to place relocated VMs to different datastores:

   a. Click the **Pick datastore for selected virtual disks** link.

   b. In the **Choose VM Files Location** window, click **Add VM** on the right and select the VM that must be placed on datastores.

   c. Select the added VM in the **Files location** list and click **Datastore** at the bottom of the window.

d. From the list of available datastores, select the target datastore.

You can also place the configuration file and individual disk files of a VM to different datastores:

a. Add a VM to the **Files location** list, expand the VM and select the required files.

b. Click **Datastore** at the bottom of the window and choose the destination for the files.

5. By default, Veeam Backup & Replication saves disks of relocated VMs in the thin format. If necessary, you can change the disk format. For example, if the original VM uses thick disks, you can change the disk format of the relocated VM to thin provisioned and save on disk space required to store VM data.

   Disk format change is available only for VMs using virtual hardware version 7 or later.

   To change VM disk format:

   a. Click the **Pick datastore for selected virtual disks** link.

   b. In the **Choose VM Files Location** window, click **Add VM** on the right and select the VM whose disk format you want to change.

   c. Select the added VM in the list and click **Disk type** at the bottom of the window.

   d. In the **Disk Type Settings** section, choose the format that will be used to restore VM disk files: same as the source disk, thin or thick.

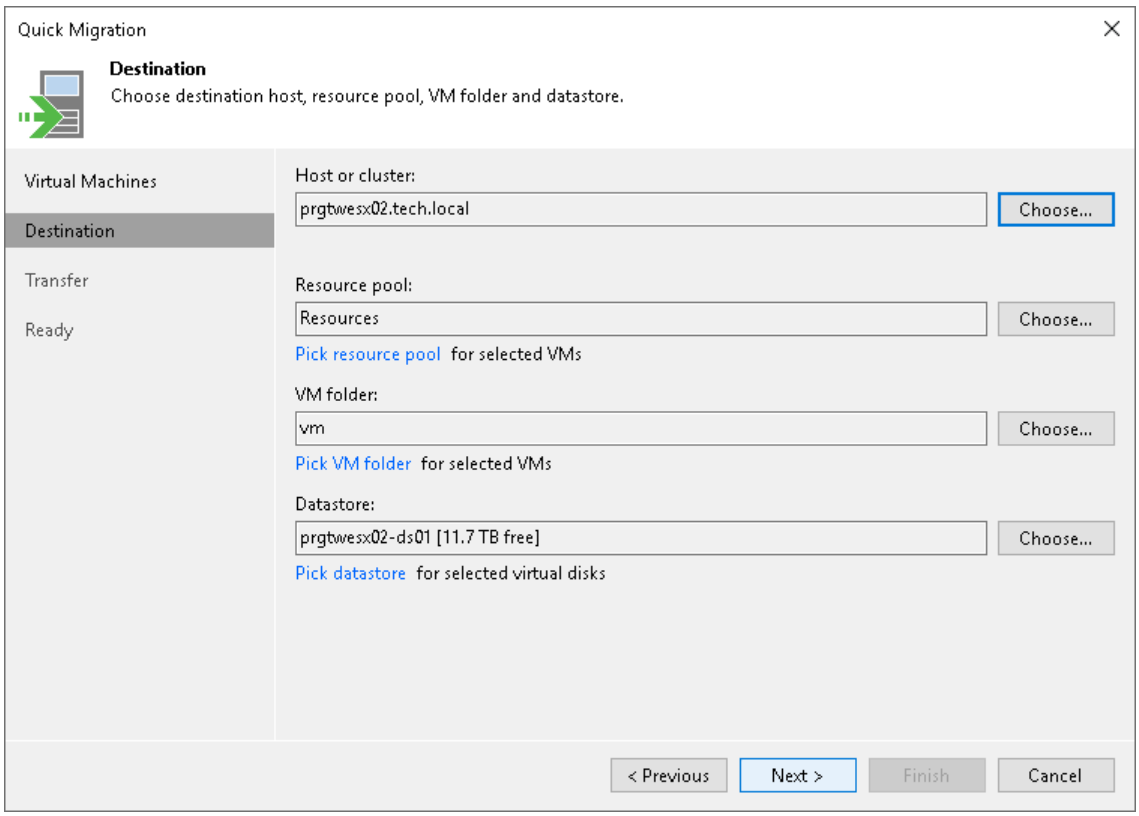

# Step 4. Select Infrastructure Components for Data Transfer

At the **Transfer** step of the wizard, assign infrastructure components to relocate the VMs:

1. In the **Data transfer** section, select backup proxies that must be used to transfer VM data from source to target.

   If you plan to migrate VMs within one site, the same backup proxy can act as the source backup proxy and target backup proxy. For off-site migration, you must deploy at least one backup proxy in each site to establish a stable connection across the sites for data transfer.

   Click **Choose** next to the **Source proxy** and **Target proxy** fields to select backup proxies for migration. In the **Backup Proxy** window, you can choose automatic proxy selection or assign proxies explicitly.

   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that have access to the source datastore and will automatically assign optimal proxy resources for processing VM data.

     During migration, VMs are processed one by one. Before processing a new VM in the VM list, Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use for data retrieval and the current workload on the backup proxies to select the most appropriate resource for VM processing.

   - If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that must be used to perform migration.

2. Select which migration mechanism to use: VMware vMotion or Veeam Quick Migration. Veeam Backup & Replication can use VMware vMotion only if your VMware license covers this functionality.

   - If you want to use VMware vMotion to relocate the VMs, leave the **Force Veeam transport usage** check box not selected. Veeam Backup & Replication will attempt to use the VMware vMotion mechanism to migrate the selected VMs. If VMware vMotion cannot be used for some reason (for example, if using it can cause data loss or if you do not have a VMware vSphere license for this functionality), Veeam Backup & Replication will fail over to its native migration mechanism.

   - If you do not want to use VMware vMotion, select the **Force Veeam transport usage** check box. Veeam Backup & Replication will use its native migration mechanism.

> **IMPORTANT**
>
> If you use a native Veeam mechanism to relocate a VM, Veeam Backup & Replication suspends the initial VM on the source ESXi host (SmartSwitch) or powers off the initial VM (cold switch) for a short period of time during Quick Migration. For more information, see Quick Migration.

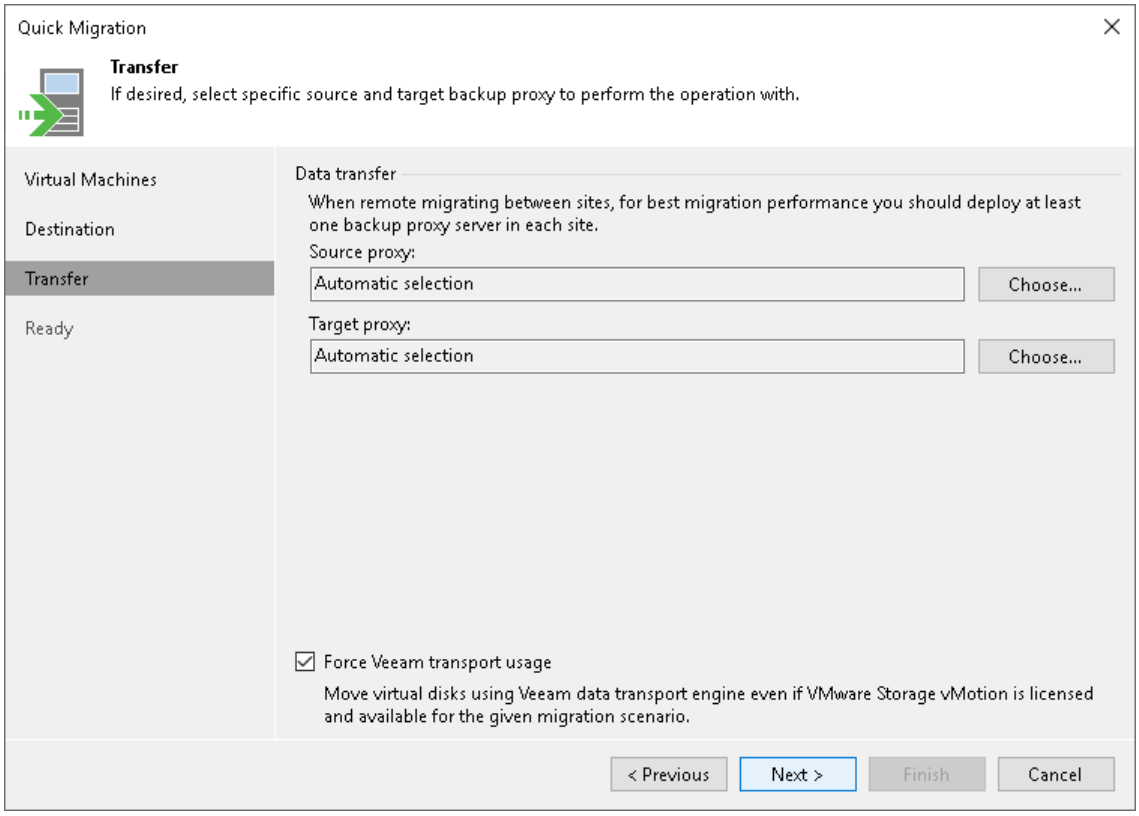
Quick Migration
Transfer
If desired, select specific source and target backup proxy to perform the operation with.
Virtual Machines
Destination
Transfer
Ready
Data transfer
When remote migrating between sites, for best migration performance you should deploy at least one backup proxy server in each site.
Source proxy:
Automatic selection
Choose...
Target proxy:
Automatic selection
Choose...
Force Veeam transport usage
Move virtual disks using Veeam data transport engine even if VMware Storage vMotion is licensed and available for the given migration scenario.
< Previous
Next >
Finish
Cancel

# Step 5. Finish Working with Wizard

At the **Ready** step of the wizard, your actions differ depending on the method you have chosen to use for migration.

If VMware vMotion is used for migration, review details and click **Finish**. In this case, all existing jobs to which the source VM is added will switch to the VM on the target host (target VM). The backup chains will be continued, thus, the next job sessions for the VM will be incremental.

If Veeam Quick Migration is used, you must also choose whether you want to delete files of the VM for which the migration was launched after Veeam Backup & Replication receives a heartbeat signal from the VM on the target host (target VM):

- If you want to delete the VM for which the migration was launched (source VM; for Instant Recovery, this is the VM created during Instant Recovery but before finalization), leave the **Delete source VM files upon successful migration** check box selected. All jobs to which the VM is added will switch to the target VM. The backup chains will be continued, thus, the next job sessions for the VM will be incremental.

  If the heartbeat signal is not received from the target VM, the source VM will not be deleted and the target VM will not be added to any jobs. To protect the target VM, you must add it to a backup job manually.

- If you want to leave the VM for which the migration was launched (source VM; for Instant Recovery, this is the VM created during Instant Recovery but before finalization), clear the selection of the **Delete source VM files upon successful migration** check box. In this case, the source VM will not be deleted. All jobs to which the source VM is added will still process this VM and continue existing backup chains. To protect the target VM, you must add it to a backup job manually.

During Veeam Quick Migration, Veeam Backup & Replication names the VMs according to the following patterns:

- **Source VM**: `source_vm_name` at the start of migration; `source_vm_name_migrated` at the end of migration.
- **Target VM**: `source_vm_name_GUID` at the start of migration; `source_vm_name` at the end of migration.

> **NOTE**
>
> The **Delete source VM files upon successful migration** option applies if the following conditions are met:
>
> - VMware vMotion is not used.
> - VMs are powered on before migration.
> - VMware Tools are installed on VMs.

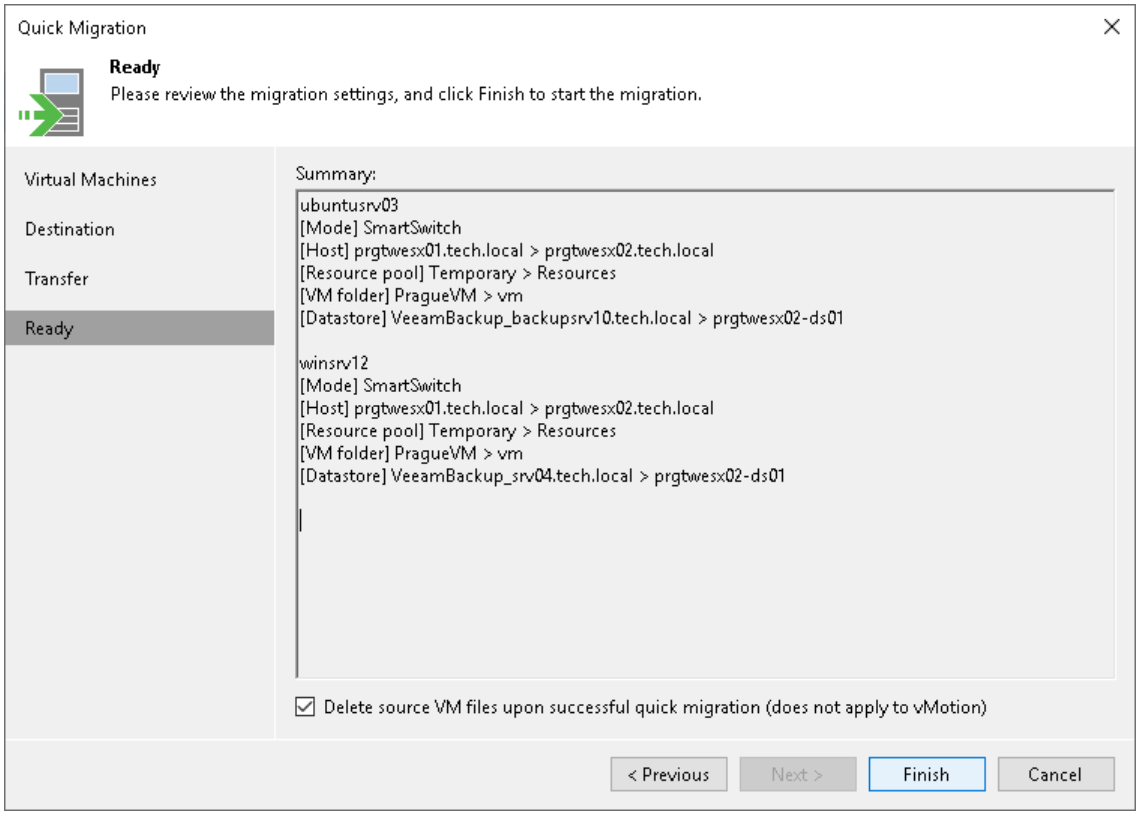
Quick Migration
Ready
Please review the migration settings, and click Finish to start the migration.
Virtual Machines
Destination
Transfer
Ready
Summary:
ubuntusrv03
[Mode] SmartSwitch
[Host] prgtwesx01.tech.local > prgtwesx02.tech.local
[Resource pool] Temporary > Resources
[VM folder] PragueVM > vm
[Datastore] VeeamBackup_backupsrv10.tech.local > prgtwesx02-ds01
winsrv12
[Mode] SmartSwitch
[Host] prgtwesx01.tech.local > prgtwesx02.tech.local
[Resource pool] Temporary > Resources
[VM folder] PragueVM > vm
[Datastore] VeeamBackup_srv04.tech.local > prgtwesx02-ds01
Delete source VM files upon successful quick migration (does not apply to vMotion)
< Previous
Next >
Finish
Cancel

# Migrating First Class Disks (FCDs)

You can relocate First Class Disks (FCDs) between datastores using Quick Migration. Migration of FCDs supports only the vMotion method. Depending on whether FCDs that you want to migrate are attached to a VM or not, Veeam Backup & Replication defines the migration scenario:

- If FCDs are not attached, Veeam Backup & Replication will migrate only the FCDs that you have selected.
- If FCDs are attached to a VM, Veeam Backup & Replication will also look for other FCDs attached to this VM and will perform migration of these FCDs as well.

Quick Migration is not a job-driven process: it cannot be saved as a job or scheduled to run later. To initiate Quick Migration of FCDs, you must complete FCD Quick Migration wizard and launch the **FCD Quick Migration** wizard.

# Before You Begin

Before you perform FCD Quick Migration, check the following prerequisites and limitations.

- All FCDs that you plan to migrate, must be located on a vPower NFS datastore that is mounted to the target cluster.
- If redo logs are stored on a custom datastore, Veeam Backup & Replication checks that these redo logs are available on this datastore.
- The target datastore to which you want to migrate FCDs must meet the following requirements:
  - The datastore must be added to the infrastructure of the vCenter which you specified to perform FCD Quick Migration.
  - The datastore must be mounted to all ESXi hosts of a cluster which you have selected to perform FCD Quick Migration.
  - The datastore must have enough free space for all migrated FDCs.
  - You can not migrate FCDs to the vPower NFS datastore.
  - You must use a datastore other than the one selected for storing redo logs.
- You can not migrate FCDs to the same datastore where they have already been registered.

# Step 1. Launch FCD Quick Migration Wizard

To launch the **FCD Quick Migration** wizard

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click the VM to which disks were recovered and select **Migrate to production**. Veeam Backup & Replication will launch the FCD Quick Migration wizard.

After you finish working with the wizard, Veeam Backup & Replication migrates the disks with all changes made after the disk recovery and before its migration.

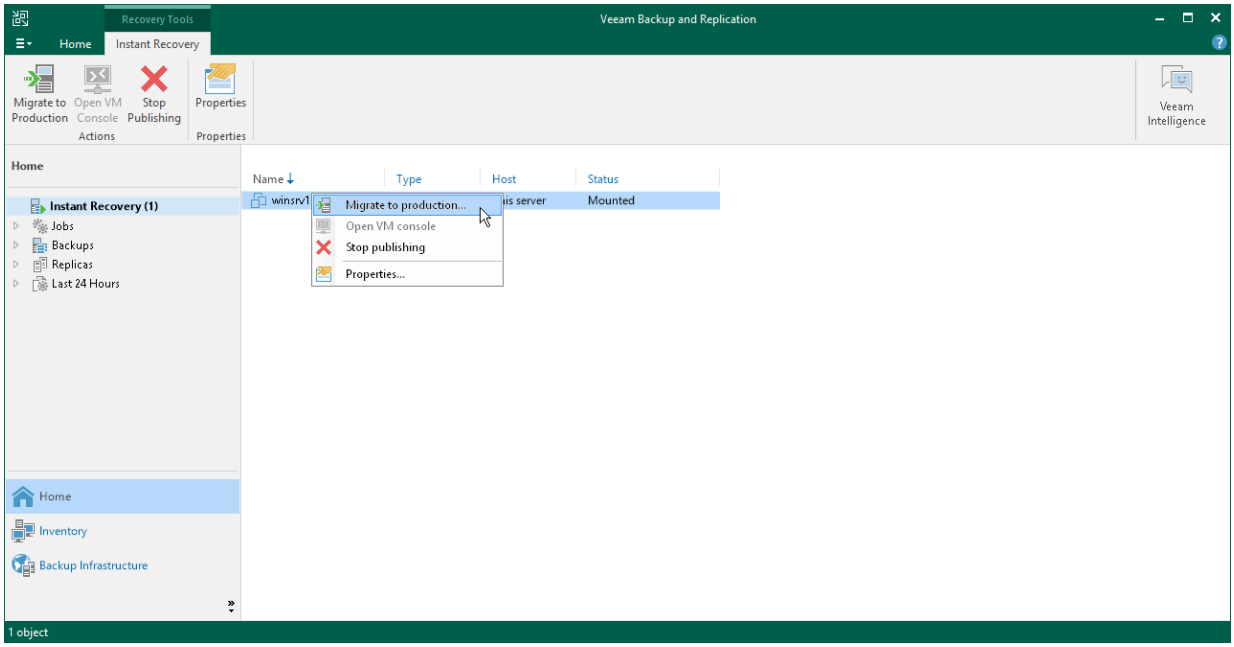

# Step 2. Select FCDs to Migrate

At the **Disks** step of the wizard, specify the FCDs that you want to migrate. To add FCDs, click **Add** and select necessary FCDs from the virtual environment. If you want to exclude specific FDCs from migration, select the necessary FCD and click **Remove**.

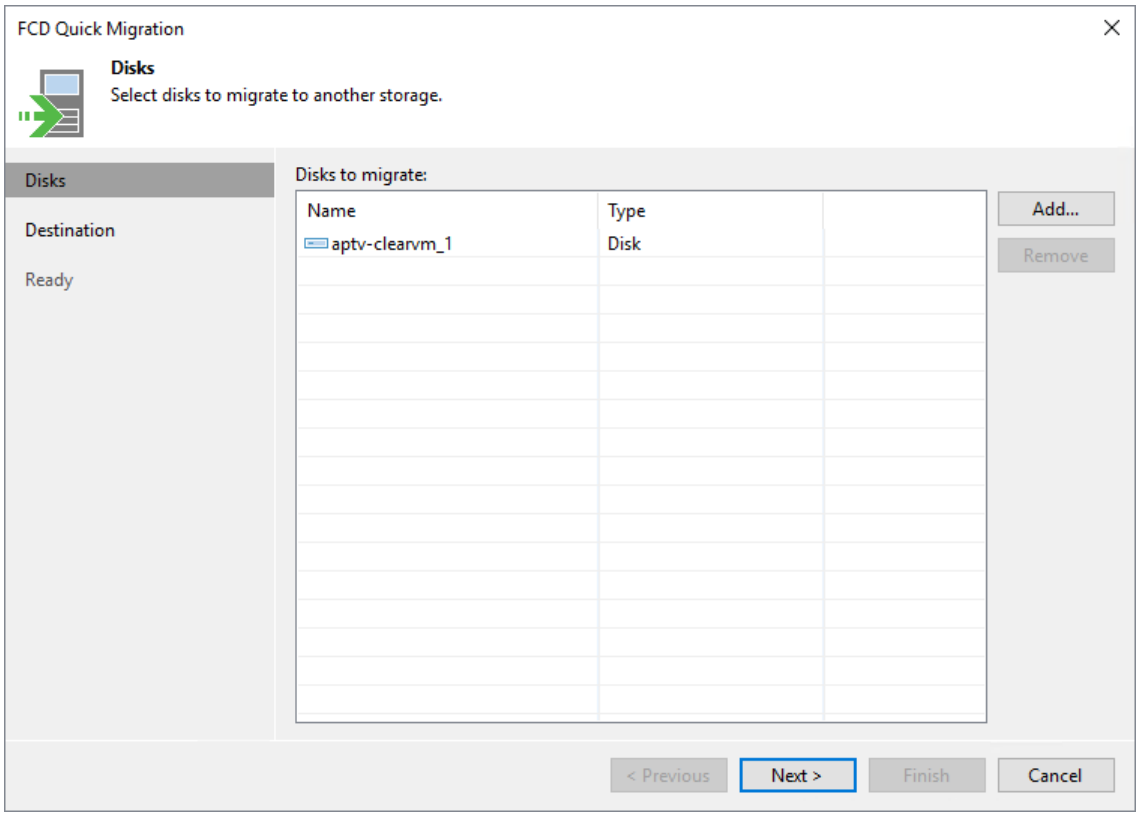

# Step 3. Specify FCD Destination

At the **Destination** step of the wizard, specify the datastore to which you want to migrate FCDs. You can migrate different FCDs to various datastores.

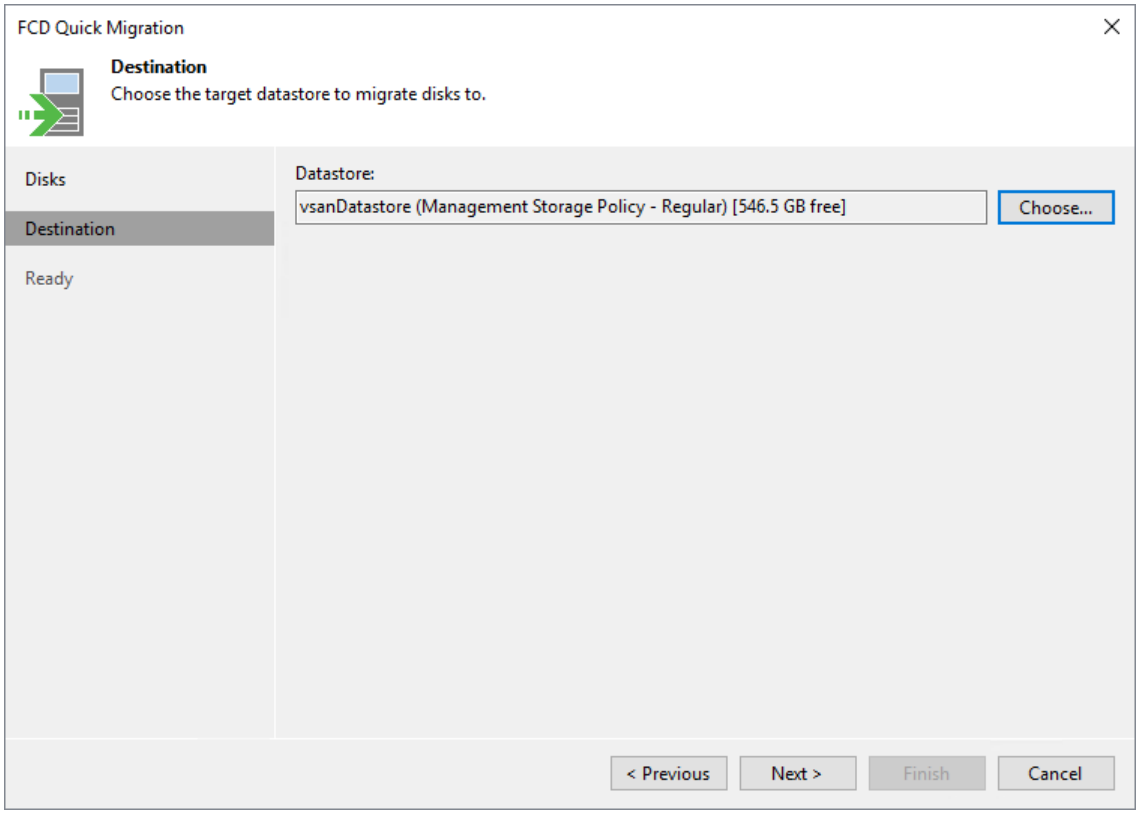

# Step 4. Finish Working with Wizard

At the **Ready** step of the wizard, review details on FCDs Quick Migration and click **Finish** to exit the wizard.

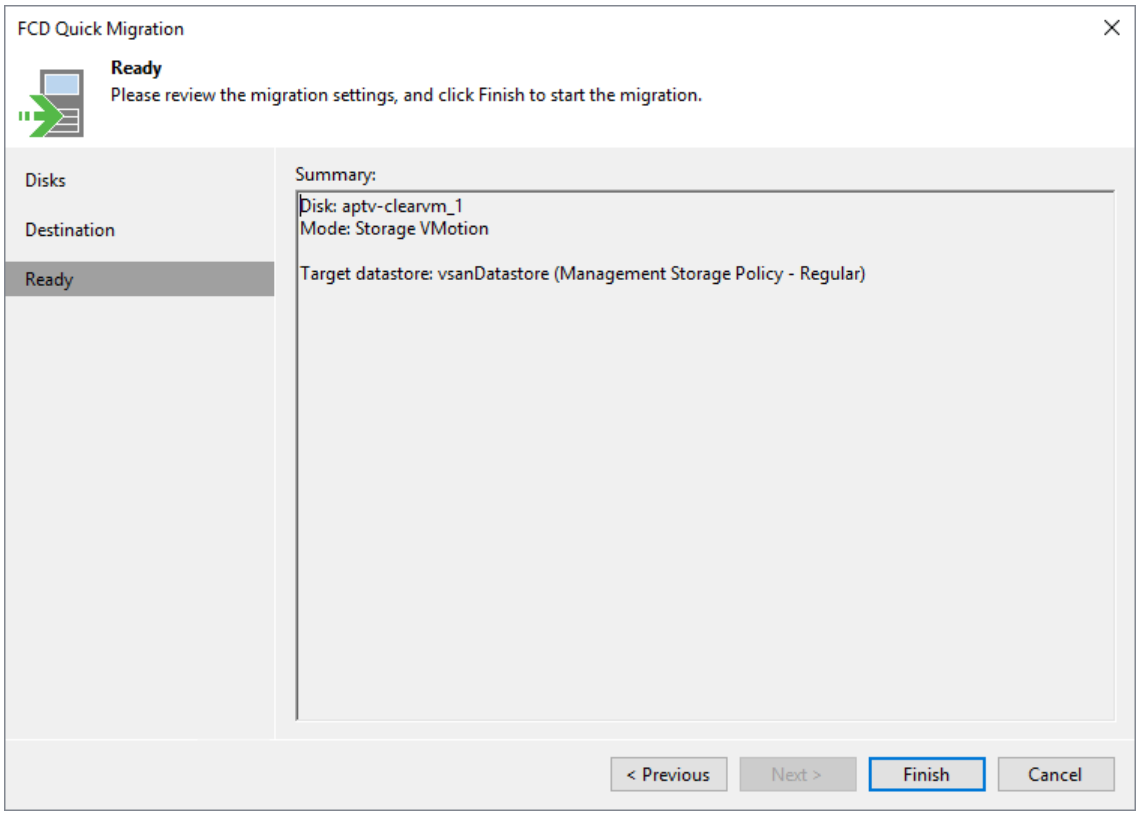

# Recovery Verification

Veeam Backup & Replication offers two technologies to verify recoverability of VM backups and replicas:

- SureBackup
- SureReplica

**IMPORTANT**

The recovery verification functionality is included in the Veeam Universal License.

When using a legacy socket-based license, Enterprise or higher edition is required. If you use the Standard edition, you can manually verify VM backups with Instant Recovery. For more information, see Manual Recovery Verification.

# SureBackup

SureBackup is the Veeam technology that allows you to test machines backups and check if you can recover data from them. You can verify any restore point of a backed-up machine.

# How SureBackup Works

For SureBackup, Veeam Backup & Replication uses a regular image-based backup. SureBackup job can operate in two different recovery verification modes:

- **Full recoverability testing**. Veeam Backup & Replication runs machines in an isolated environment directly from backup and performs tests against live applications.This mode ensures recoverability of your production workloads in a disaster recovery event.
- **Backup verification and content scan only**. Veeam Backup & Replication performs backup integrity check and its content analysis to detect traces of malware or any other unwanted or sensitive data. These tests do not require setting up a virtual lab or an application group.

## Full Recoverability Testing

During recovery verification, Veeam Backup & Replication performs the following actions:

1. If the SureBackup job is configured to perform malware scan, Veeam Backup & Replication scans backups of the machines from the application group with antivirus software and the specified YARA rule.
2. Veeam Backup & Replication publishes the machines from the application group and the machine under verification in the isolated environment — virtual lab. Machines are started directly from compressed and deduplicated backup files that reside in the backup repository. To achieve this, Veeam Backup & Replication utilizes the Veeam vPower NFS Service.
3. Veeam Backup & Replication performs a number of tests against the machines in the application group and machine under verification: heartbeat test, ping test and application test.
4. If the SureBackup job is configured to validate backup files, Veeam Backup & Replication performs a cyclic redundancy check for the backup file from which the machine under verification is started and, optionally, for backup files from which the machines in the application group are started. The backup file validation is performed after all verification tests are complete.
5. When the recovery verification process is over, Veeam Backup & Replication unpublishes machines and creates a report on their state. The report is sent to the backup administrator by email.

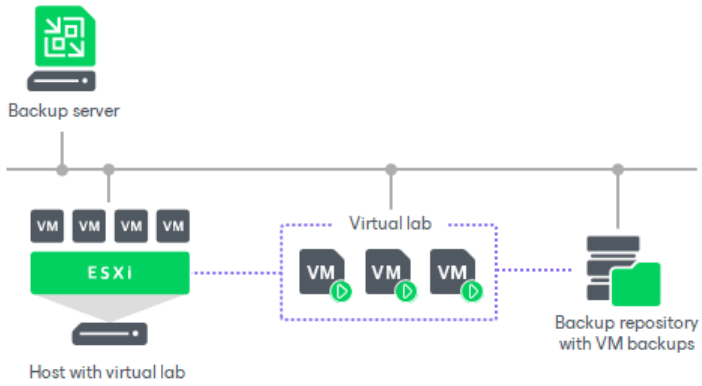


**NOTE**

If the machine has multiple disk partitions with different OS installed, SureBackup will only verify the partition with the highest boot priority.

During verification, a backed-up machine image remains in read-only state. All changes that take place when the machine is running are written to redo log files that are stored on the datastore selected in the virtual lab settings. When the recovery process is complete, the redo logs are removed.

To perform recovery verification, you need to create the following objects:

1. Application group. During recovery verification, the verified machine may need to be started with a group of machines on which it is dependent. The application group enables full functionality of applications running inside the machine and lets you run these applications just like in the production environment.
2. Virtual lab. The virtual lab is the isolated virtual environment in which the verified machine and machines from the application group are started and tested.
3. SureBackup job. The SureBackup job is a task to perform recovery verification. You can run the SureBackup job manually or schedule it to run automatically by schedule.

# Backup verification and content scan only

During the backup verification and content scan, Veeam Backup & Replication performs the following actions:

1. If the SureBackup job is configured to perform malware scan, Veeam Backup & Replication scans backups of the machines from the linked job with antivirus software and the specified YARA rule.
2. If the SureBackup job is configured to validate backup files, Veeam Backup & Replication performs a cyclic redundancy check for the backup file from which the machine under verification is started. The backup file validation is performed after all verification tests are complete.
3. When the recovery verification process is over, Veeam Backup & Replication creates a report on the machines' state. The report is sent to the backup administrator by email.

During verification, a backed-up machine image remains in read-only state. All changes that take place when the machine is running are written to the differencing disk (AVHD/AVHDX file), created for the recovered machine. When the recovery verification process is complete, the changes are discarded.

# Backup Recovery Verification Tests

To verify machines with a SureBackup job in **Full recoverability testing mode**, you can instruct Veeam Backup & Replication to run the following backup recovery verification tests:

- Predefined tests
- Custom verification scripts

When running a SureBackup job in **Backup verification and content scan only** mode, you can only instruct Veeam Backup & Replication to perform backup integrity check.

## Predefined Tests

Veeam Backup & Replication can verify machines with the following predefined tests:

- **Heartbeat test**. When the machine starts, Veeam Backup & Replication performs a heartbeat test. It waits for a heartbeat signal from VMware Tools installed inside the machine to determine that the machine guest OS is running. If the signal comes regularly at specific time intervals, the test is passed.
- **Ping test**. Veeam Backup & Replication sends ping requests to the machine from the backup server and checks if the machine can respond to them. If the machine responds to ping requests, the test is passed.
- **Application test**. Veeam Backup & Replication waits for applications inside the machine to start and runs a script against these applications. Veeam Backup & Replication uses two types of predefined scripts:
  - For DNS servers, domain controllers, Global Catalog servers, mail servers and web servers, Veeam Backup & Replication uses a script that probes an application-specific port. For example, to verify a domain controller, Veeam Backup & Replication probes port 389 for a response. If the response is received, the test is passed.
  - For Microsoft SQL Server, Veeam Backup & Replication uses a script that attempts to connect to instances and databases on the Microsoft SQL Server. For more information, see Microsoft SQL Server Checker Script.

> **NOTE**
>
> To run the heartbeat and ping tests, you must have VMware Tools installed on the machine. If VMware Tools are not installed, these tests will be skipped.

You can run verification tests for machines added to the application group or processed with a linked SureBackup job. You can specify and customize settings for verification tests in the application group or SureBackup job settings.

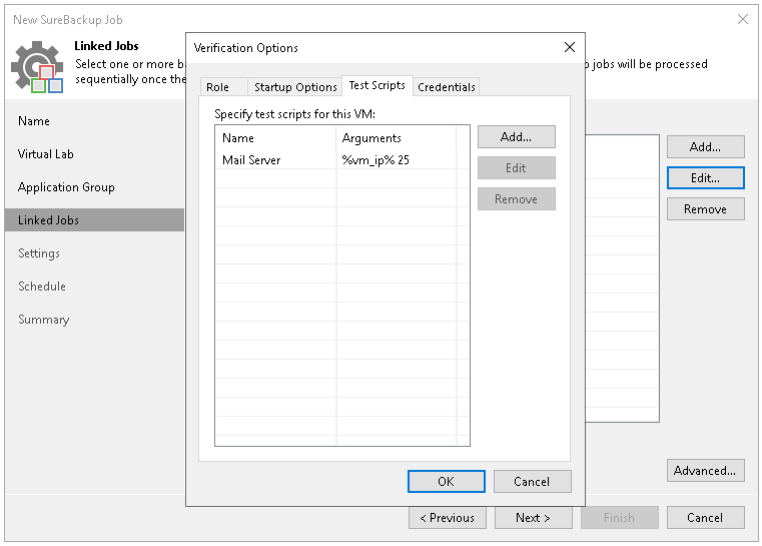


# Custom Verification Scripts

Veeam Backup & Replication can verify VMs with the following custom verification scripts:

- Microsoft SQL Server Checker script
- Backup file validation

**IMPORTANT**

Consider the following:

- Do not pass sensitive information using script arguments in a user interface.
- Windows Script Host must be enabled to run the Microsoft SQL Server Checker script.
- If you use the Microsoft SQL Server authentication mode, you may need to specify credentials of the account to connect to the machine on which Microsoft SQL Server is installed. To do this, use the Credentials tab in the application group or SureBackup job settings.
- You can not use domain credentials for the script execution if you have Veeam Backup & Replication deployed in Microsoft Azure and it is managed by Microsoft Intune. Use local credentials for Microsoft SQL Server instead.
- During the major version upgrade procedure (for example, from version 11 to version 12), Veeam Backup & Replication may replace the SQL checker script with the default one. If you have manually modified the SQL checker script before upgrading to a new version, reapply those modifications after the update.

# Microsoft SQL Server Checker Script

If you need to verify a virtualized Microsoft SQL Server, you can instruct Veeam Backup & Replication to run the Microsoft SQL Server Checker script against it during the SureBackup job. The script sequentially performs the following operations:

1. Connects to Microsoft SQL Server instances.
2. Enumerates databases on these instances.

3. Employs the USE SQL statement to connect to databases and check their availability.

The script is located on the backup server in the Veeam Backup & Replication product folder. The path by default: `C:\Program Files\Veeam\Backup and Replication\Backup\Veeam.Backup.SqlChecker.vbs`.

## Credentials for Script Execution

To execute the script, Veeam Backup & Replication connects to Microsoft SQL Server. By default, Veeam Backup & Replication uses the account under which the Veeam Backup Service is running. If you need to run the script under another account, you can specify credentials for this account. The script supports the following authentication methods:

- Microsoft Windows authentication mode. To use it, you must specify credentials for the account on the **Credentials** tab in the application group or SureBackup job settings.
- SQL Server authentication mode. To use it, you must pass credentials of the account to the script. You can pass credentials the following ways:
  - By running the Microsoft SQL Server Checker Script from the PowerShell. To do it, use the following command in the PowerShell console:

    ```
    cscript C:\Program Files\Veeam\Backup and Replication\Backup\Veeam.Back
    up.SqlChecker.vbs <sql server[\instance]> <username> <password>
    ```

    You can find the results of a script execution in the log file by the following path: `%programdata%\Veeam\Backup\<name of the job>\<VM name>_SQLChecker.log`. If necessary, you can change the log file location. To do this, specify a new path in the PowerShell command:

    ```
    cscript C:\Program Files\Veeam\Backup and Replication\Backup\Veeam.Back
    up.SqlChecker.vbs [C:\Logs] <sql server[\instance]> <username> <passwor
    d>
    ```

  - By modifying credentials in the `Veeam.Backup.SqlChecker.vbs` file. The script is located on the backup server in the Veeam Backup & Replication product folder. The default path: `C:\Program Files\Veeam\Backup and Replication\Backup\Veeam.Backup.SqlChecker.vbs`.

### Database Exclusion

By default, Veeam Backup & Replication verifies all databases on all instances of Microsoft SQL Server. However, you can exclude specific databases from verification. For example, vCenter Server database. To exclude an instance or a database, you must open the script in the text editor and edit the **Settings** section in the following way:

- To exclude a specific databases, uncomment the `gDBsToExclude.Push "dbname"` line in the script and specify the names of a database that you want to exclude. To exclude several databases, specify a separate line for each database.

  ```
  gDBsToExclude.Push "dbname1"
  gDBsToExclude.Push "dbname2"
  ```

- To exclude a specific instance, uncomment the `gInstancesToExclude.Push "instancename"` line in the script and specify the name of an instance that you want to exclude. To exclude several instances, specify a separate line for each instance.

```
gInstancesToExclude.Push "instancename1"
gInstancesToExclude.Push "instancename2"
```

- To exclude the default instance, uncomment the `gInstancesToExclude.Push "MSSQLSERVER"` line.

> **IMPORTANT**
>
> Instance and database names are case sensitive.

### Logging

To define whether the script has completed successfully or not, Veeam Backup & Replication publishes the following return codes in the SureBackup job session statistics:

- 0 — test is passed successfully.
- 1 — you use a wrong syntax for the script command.
- 2 — Veeam Backup & Replication is unable to connect to Microsoft SQL Server.
- 3 — all instances are excluded from the check.
- 4 — error occurred while Veeam Backup & Replication was getting the list of databases.
- 5 — unknown error.
- 6 — one or more databases are not accessible.

Results of script execution are written to the log file located by the following path: `%programdata%\Veeam\Backup\<name of the job>\<VM name>_SQLChecker.log`. If necessary, you can change the log file location. To do this, pass a new path to the log file in the `%log_path%` argument in the application group or SureBackup job settings.

# Backup File Validation

In addition to recovery verification tests, Veeam Backup & Replication allows you to perform backup file validation. For backup file validation, Veeam Backup & Replication performs a CRC check for backup files of machines verified by the SureBackup job. You can also validate backup files for machines from the application group with this test.

To validate the backup file, Veeam Backup & Replication uses the checksum algorithm. When Veeam Backup & Replication creates a backup file for a machine, it calculates a checksum for every data block in the backup file and stores this data in the backup file, together with machine data. During the backup file validation test, Veeam Backup & Replication decompresses the backup file, recalculates checksums for data blocks in the decompressed backup file and compares them with initial checksum values. If the results match, the test is passed.

The backup file validation test is started after recovery verification tests. As soon as Veeam Backup & Replication completes all "live" verification for all machines in the SureBackup job, it unpublishes machines and starts the backup file validation test.

The result of the backup file validation test impacts the state of the SureBackup job session. If the verification tests are completed successfully but the backup validation is not passed, Veeam Backup & Replication marks the SureBackup job session with the *Failed* status.

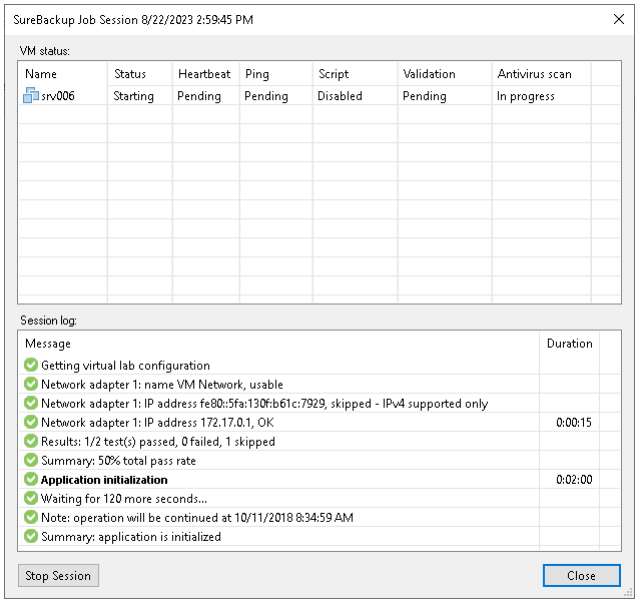

# Application Group

In most cases, a machine works not alone but in cooperation with other services and components. To verify such machine, you first need to start all services and components on which this machine is dependent. To this aim, Veeam Backup & Replication uses the application group.

The application group creates the "surroundings" for the verified machine. The application group contains one or several machines on which the verified machine is dependent. These machines run applications and services that must be started to enable fully functional work of the verified machine. Typically, the application group contains at least a domain controller, DNS server and DHCP server.

When you set up an application group, you specify a role of every machine and its boot priority. Additionally, you can specify what tests must be performed to verify machines in the application group.

When a SureBackup job is launched, Veeam Backup & Replication first starts in the virtual lab machines from the application group in the required order and performs necessary tests against them. This way, Veeam Backup & Replication creates the necessary environment for the verified machine. Only after all machines from the application group are started and tested, Veeam Backup & Replication starts the verified machine in the virtual lab.

For example, if you want to verify a Microsoft Exchange Server, you need to test its functionality in cooperation with the domain controller and DNS server. Subsequently, you must add to the application group a virtualized domain controller and DNS server. When Veeam Backup & Replication runs a SureBackup job, it will first start and verify the domain controller and DNS server in the virtual lab to make verification of the Microsoft Exchange Server possible.

# Creating Application Groups

Before you create an application group, check prerequisites. Then use the **New Application Group** wizard to create an application group.

## Before You Begin

Before you create an application group, consider the following:

- A valid license for Enterprise edition of Veeam Backup & Replication must be installed on the backup server.
- All applications and services on which verified machines are dependent must be virtualized in your environment.
- If you plan to scan machines data for malware, check requirements and limitations.
- If you plan to verify machines with a ping test, the firewall on tested machine must allow ping requests.
- If you plan to verify machines with a heartbeat test, VMware Tools must be installed in tested machines.
- If backup file that is used to add the machine to an application group cannot be found, Veeam Backup & Replication will search for other backup files containing this machine.
- VM replicas must be in the *Ready* state. If a VM replica is in the *Failover* or *Failback* state, you will not be able to add it to the application group.
- You cannot add to application groups VMs from backups of VMware Cloud Director VMs, backups created with backup copy jobs and backups stored in Cloud Connect backup repositories.
- [For storage snapshots] The storage system must be added to the backup infrastructure.

# Step 1. Launch New Application Group Wizard

To launch the **New Application Group** wizard, do one of the following:

- Open the **Backup Infrastructure** view, in the inventory pane select **SureBackup**. In the working area, click **Add Application Group**.
- Open the **Backup Infrastructure** view, in the inventory pane select **Application Groups** under **SureBackup** and click **Add App Group**.
- Open the **Backup Infrastructure** view, in the inventory pane right-click **Application Groups** under **SureBackup** and select **Add App Group**.

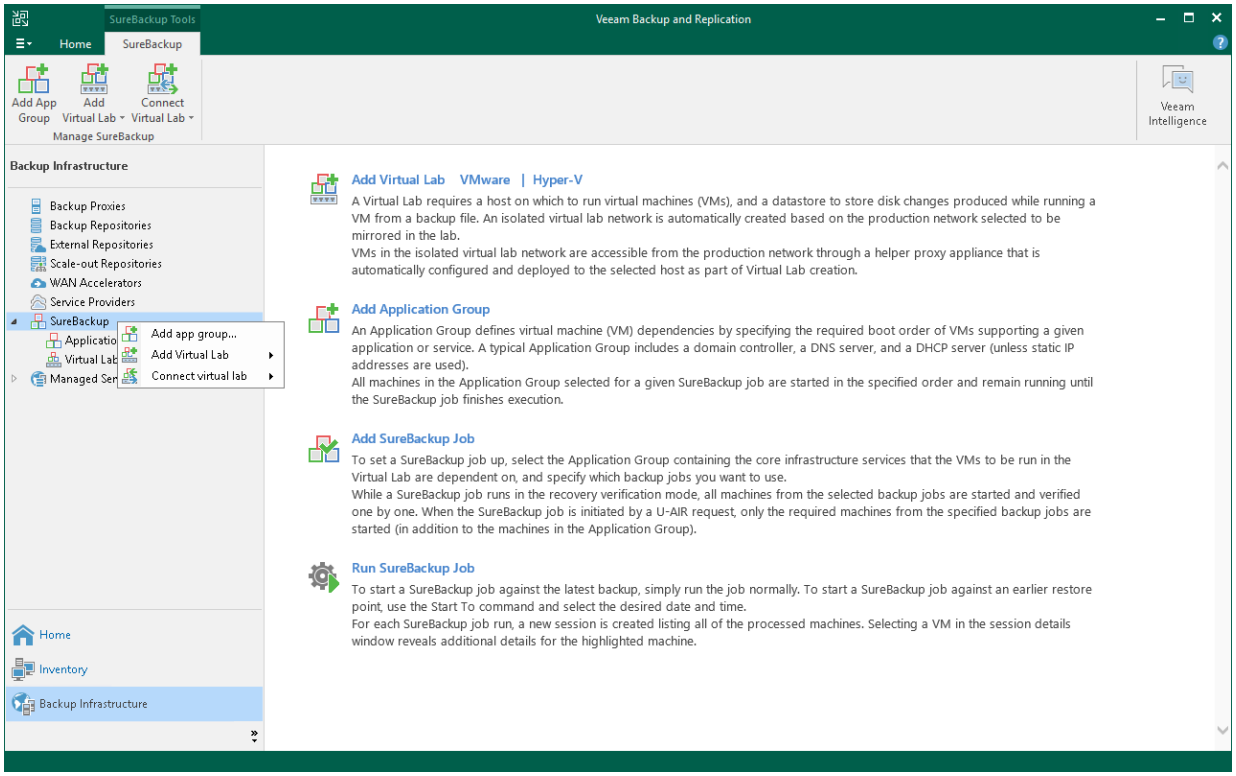

# Step 2. Specify Application Group Name and Description

At the **Name** step of the wizard, specify a name and description for the application group.

1. In the **Name** field, enter a name for the application group.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the group, date and time when the group was created.

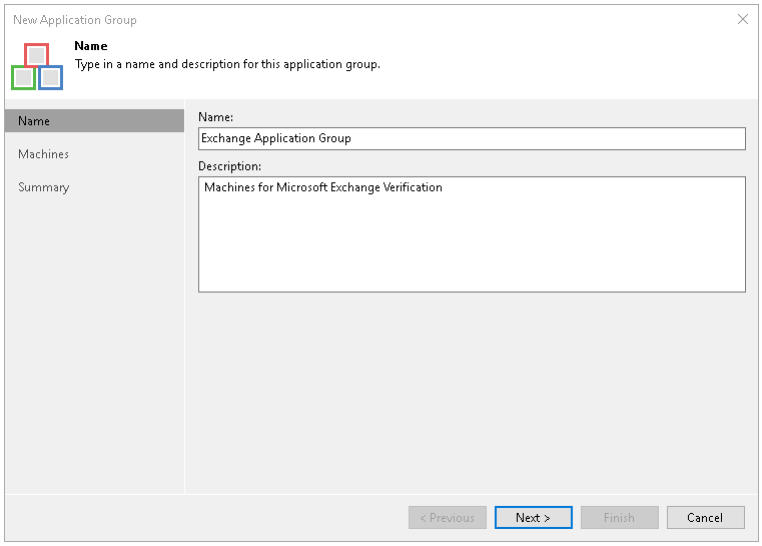

# Step 3. Add Machines to Application Group

At the **Machines** step of the wizard, add machines to the created application group. You can add machines from different sources:

- Machine backups
- VM replicas
- Storage snapshots

You can add machines from backups, storage snapshots and VM replicas to the same application groups. Keep in mind the following limitations:

- Machines must have at least one valid restore point or must reside on a storage snapshot.
- You cannot add the same machine twice. For example, if you add a machine from the storage snapshot, you will not be able to add the same machine from the backup.

To add machines to the application group:

1. Click **Add** and select **From backups**, **From replicas** or **From storage snapshots**.
2. In the displayed window, expand the job or storage snapshot, select the machine and click **Add**.
3. Machines in the list are specified in the order of their boot priority. To move a machine up and down in the list, select it and click **Move Up** or **Move Down**.

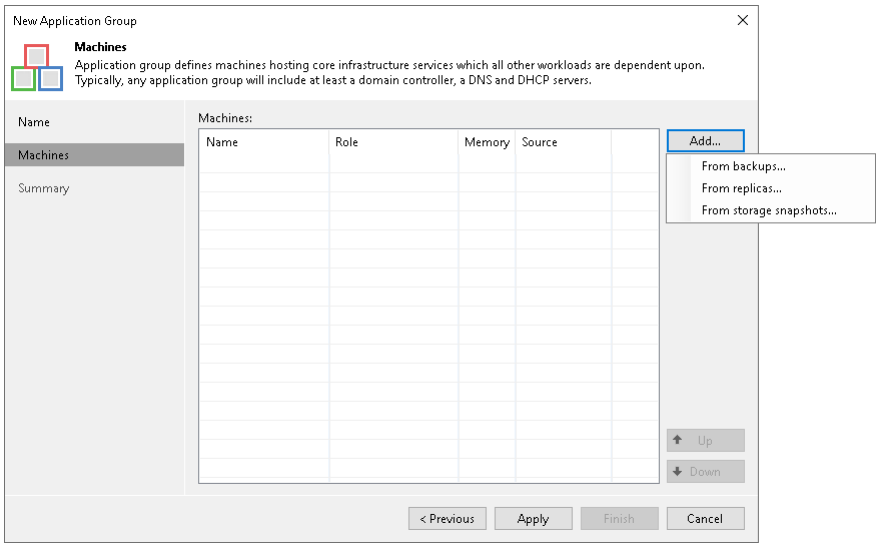

## Step 4. Specify Recovery Verification Options and Tests

You must specify verification options for every VM in the application group:

- Select a role that the VM performs.
- Configure startup settings.
- Select tests that must be performed for the VM.
- Specify credentials for running the verification script.

To specify recovery verification options:

1. At the **Machines** step of the wizard, select the machine in the list.
2. Click **Edit** on the right.
3. Use the **Verification Options** window to specify verification options.

# Role Settings

On the **Role** tab, select a role that the VM performs. Veeam Backup & Replication offers the following predefined roles for VMs:

- DNS Server
- Domain Controller (Authoritative Restore)

  In the Authoritative Restore mode, Veeam Backup & Replication starts a domain controller in the virtual lab and marks it as being authoritative to its replication partners. When other domain controllers (replication partners) are started in the virtual lab, they replicate data from the domain controller started in the Authoritative Restore mode.
- Domain Controller (Non-Authoritative Restore)

  In the Non-Authoritative Restore mode, Veeam Backup & Replication restores a domain controller in the virtual lab and marks it as being non-authoritative to its replication partners. Non-authoritative domain controllers then replicate data from a domain controller started in the Authoritative Restore mode.
- Global Catalog
- Mail Server
- SQL Server
- Veeam Backup for Microsoft 365 (machine on which Veeam Backup for Microsoft 365 is installed)
- Web Server

VM roles are described in XML files stored in the *%ProgramFiles%|Veeam|Backup and Replication|Backup|SbRoles* folder on the backup server. You can create your own roles. To do this, you must create a new XML file and specify role and test scripts settings in it. For more information, see Creating XML files with VM Roles Description.

After you select a role for the VM, Veeam Backup & Replication will automatically configure startup options and assign predefined test scripts for the chosen role. You can use these settings or specify custom settings on the **Startup Options** and **Test Scripts** tabs.

To verify VMs that perform roles other than those specified in the list, you will have to manually configure startup options and specify test scripts that must be run for these VMs.

**IMPORTANT**

If you want to add several domain controllers to the application group, you must assign the **Domain Controller (Authoritative Restore)** role to the first domain controller started in the virtual lab. Other domain controllers must have the **Domain Controller (Non-Authoritative Restore)** role.

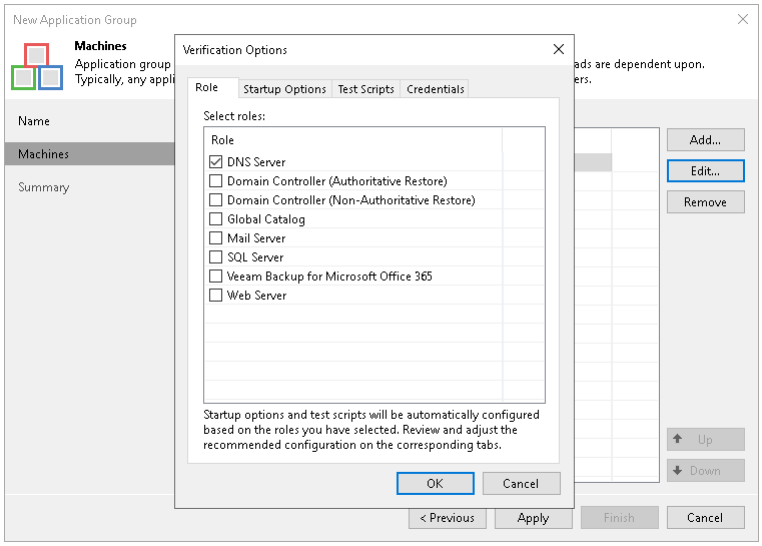


# Startup Settings

To configure VM startup settings:

1. In the **Verification Options** window, click the **Startup Options** tab.
2. In the **Memory** section, specify the amount of memory that you want to pre-allocate to the VM when this VM starts. The amount of pre-allocated memory is defined in percent. The percentage rate is calculated based on the system memory level available for the production VM. For example, if 1024 MB of RAM is allocated to the VM in the production environment and you specify 80% as a memory rate, 820 MB of RAM will be allocated to the verified VM on startup.

   Veeam Backup & Replication does not allow you to change machine CPU manually, it does this automatically. If the VM has more CPU than the host can provide, Veeam Backup & Replication scales down the CPU of the VM.
3. In the **Startup time** section, specify the **Maximum allowed boot time** value and the **Application initialization timeout** value.

   **Application initialization timeout**. Veeam Backup & Replication waits for the applications installed in the VM, for example, Microsoft SQL Server, to start. Depending on the software installed in a VM, the application initialization process may require more time than specified in the job settings. If applications installed in a VM are not initialized within the specified period of time, test scripts can be completed with errors. If such an error situation occurs, you need to increase the **Application initialization timeout** value and start the job once again.

   **Maximum allowed boot time**. Be careful when specifying the **Maximum allowed boot time** value. Typically, a VM started by the SureBackup job requires more time to boot than a VM started in the production environment. If VM is not initialized within the specified interval of time, the recovery verification process fails with the timeout error. If such error occurs, you need to increase the **Maximum allowed boot time** value and run the job again.

4. In the **Boot verification** section, specify when the VM must be considered to have been booted successfully:
   - **VM heartbeat is present**. If you enable this option, Veeam Backup & Replication will perform a heartbeat test for the verified VM.
   - **VM responds to ping on any network interface**. If you enable this option, Veeam Backup & Replication will perform a ping test for the verified VM.
   - **Automatically disable Windows Firewall**. If you select this option, Veeam Backup & Replication will disable Windows Firewall for the verified VM.

   If you enable both heartbeat test and ping test options, Veeam Backup & Replication will require that both tests are completed successfully.

> **NOTE**
>
> Veeam Backup & Replication performs a heartbeat test only if a VM has VMware Tools installed. If VMware Tools are not installed, the VM will be started but the test will not be performed. VMs without VMware Tools can still be used as auxiliary VMs: they can be started to enable proper work of other machines. In this case, you do not need to select any role for such machines.

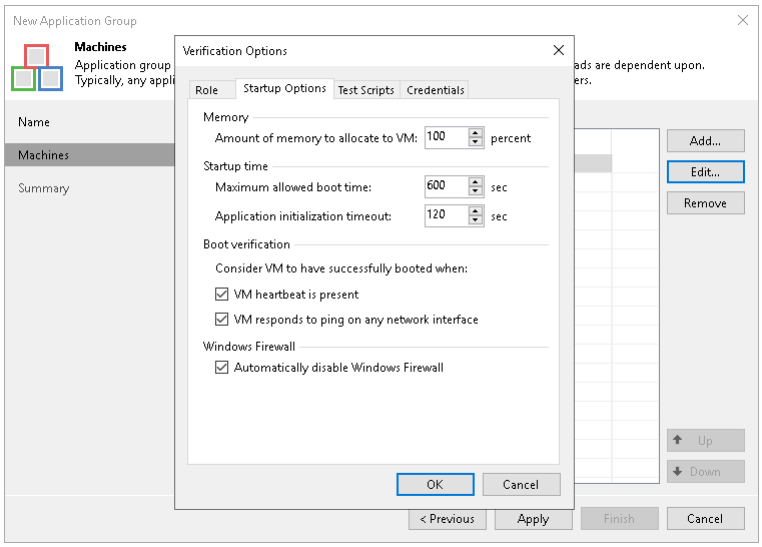


# Test Script Settings

When you select a VM role, Veeam Backup & Replication automatically assigns a predefined script which verifies applications inside this VM. The script is launched by the account that controls the Veeam Backup Service. If you want to verify a VM that has some other role not listed on the **Role** tab:

1. In the **Verification Options** window, click the **Test Scripts** tab.
2. Click **Add**.
3. In the **Test Scripts** window, select **Use the following test script**.
4. In the **Name** field, specify a name for the script.
5. In the **Path** field, define a path to an executable script file that must be run to verify the VM. You can do one of the following:
   - If you have your own custom script, define a path to it in the **Path** field.

- If you do not have a custom script, you can use a standard utility by Veeam, *Veeam.Backup.ConnectionTester.exe*, that probes application communication ports. The utility is located in the installation folder of Veeam Backup & Replication: *%ProgramFiles%|Veeam|Backup and Replication|Backup|Veeam.Backup.ConnectionTester.exe*. Specify this path in the **Path** field.

6. In the **Arguments** field, specify an IP address of the verified VM and the port that you want to probe (if the selected test probes the port). You can use the `%vm_ip%` variable to define the VM IP address or the `%vm_fqdn%` variable to define the VM fully qualified domain name.

   For Microsoft SQL Server, you can also specify a path to the log file in the `%log_path%` argument. For more information, see Backup Recovery Verification Tests.

7. Click **OK** to add the configured test.

To edit test settings, select the test in the list and click **Edit**. To delete a test, select the test in the list and click **Remove**.

**IMPORTANT**

Do not pass sensitive information using script arguments in a user interface.

**NOTE**

If a VM performs several roles and runs a number of applications, you can add several verification scripts to verify work of these applications. It is recommended that you specify the maximum startup timeout value and allocate the greatest amount of memory for such VMs.

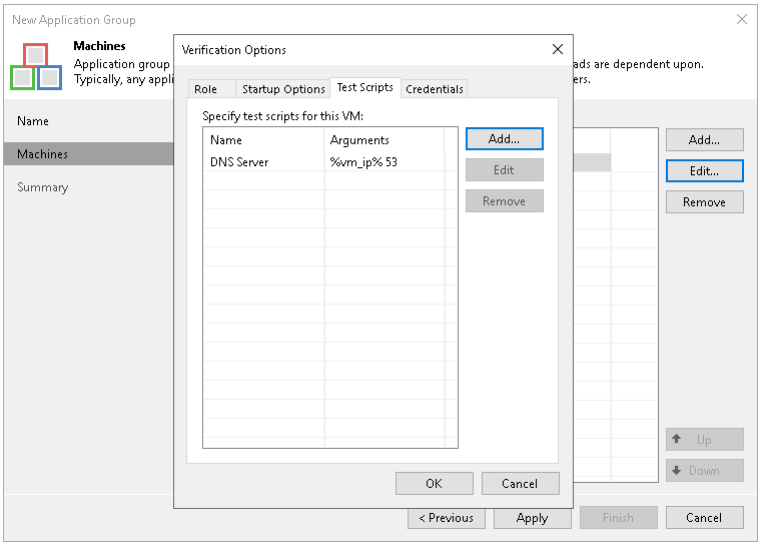


# Credentials Settings

By default, to run the verification script Veeam Backup & Replication uses the account under which the Veeam Backup Service is running. If you need to run the script under some other account, you can specify credentials for this account in the application group settings.

1. Click the **Credentials** tab.

2. From the **Credentials** list, select credentials for the account under which you want to run the script.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add the credentials. For more information, see Credentials Manager.

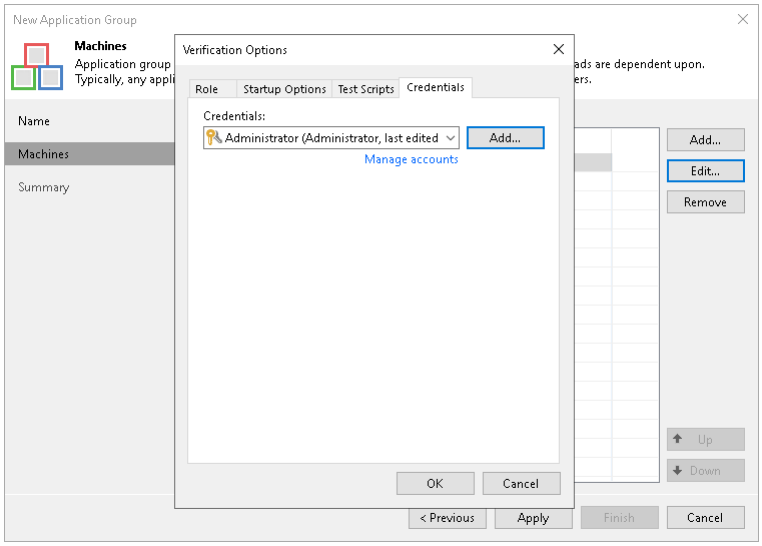

# Step 5. Review Application Group Settings and Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of application group configuration.

1. Review details of the application group.
2. Click **Finish** to save the application group settings and close the wizard.

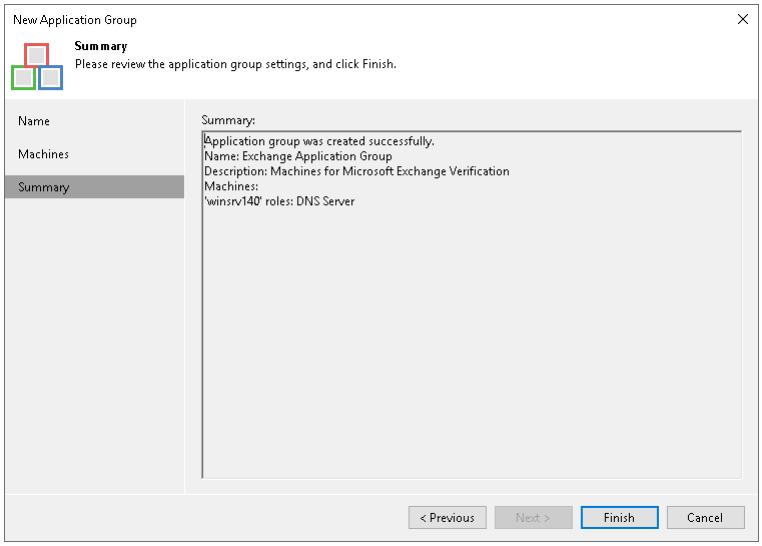

# Virtual Lab

The virtual lab is an isolated virtual environment in which Veeam Backup & Replication verifies VMs. In the virtual lab, Veeam Backup & Replication starts VMs from the application group and the verified VM. The virtual lab is used not only for the SureBackup verification procedure, but also for U-AIR, On-Demand Sandbox and staged restore.

The virtual lab itself does not require that you provision extra resources for it. However, VMs running in the virtual lab consume CPU and memory resources of the ESXi host where the virtual lab is deployed. All VM changes that take place during recovery verification are written to redo log files. By default, Veeam Backup & Replication stores redo logs on the datastore selected in the virtual lab settings and removes redo logs after the recovery process is complete.

The virtual lab is fully fenced off from the production environment. The network configuration of the virtual lab mirrors the network configuration of the production environment. For example, if verified VMs and VMs from the application group are located in two logical networks in the production environment, the virtual lab will also have two networks. The networks in the virtual lab will be mapped to necessary production networks.

VMs in isolated networks have the same IP addresses as in the production network. This lets VMs in the virtual lab function just as if they function in the production environment.

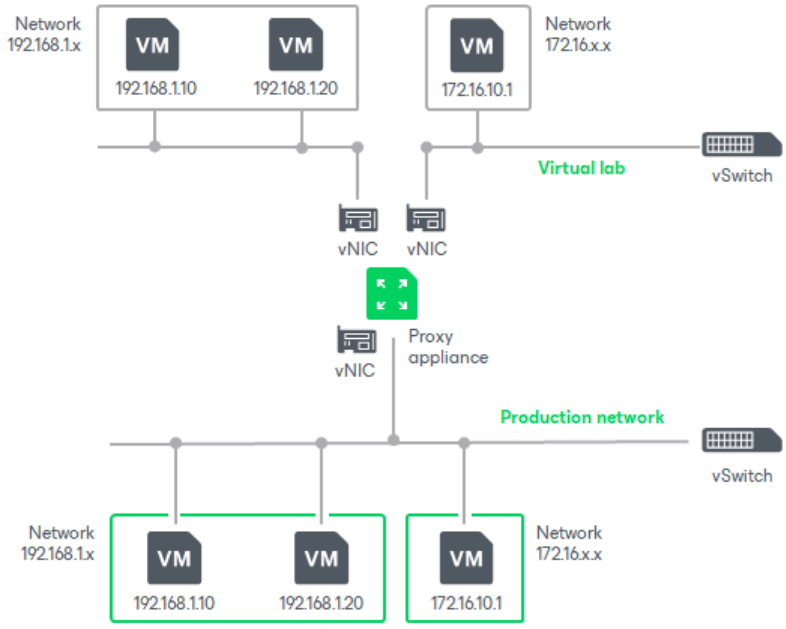


# Proxy Appliance

To enable communication between the production environment and isolated networks in the virtual lab, Veeam Backup & Replication uses a proxy appliance. The proxy appliance is an auxiliary Linux-based VM that is deployed on the ESXi host where the virtual lab is created. The proxy appliance VM is assigned an IP address from the production network and placed to the dedicated virtual lab folder and resource pool on the ESXi host.

The proxy appliance is connected to the production network and to the isolated network, so that it has visibility of the production environment and virtual lab. In essence, the proxy appliance acts as a gateway between the two networks — it routes requests from the production environment to VMs in the virtual lab.

To connect to isolated networks, the proxy appliance uses network adapters. Veeam Backup & Replication adds to the proxy appliance one network adapter per every isolated network. For example, if there are two networks in the virtual lab, Veeam Backup & Replication will add two network adapters to the proxy appliance. The network adapter gets an IP address from the isolated network. Typically, this IP address is the same as the IP address of the default gateway in the production network.

The proxy appliance is an optional component for recovery verification. Technically, you can create a virtual lab without the proxy appliance. However, in this case, you will not be able to perform automatic recovery verification of VMs. VMs will be started from backups in the virtual lab; you will have to access them using the VM console and perform necessary tests manually.

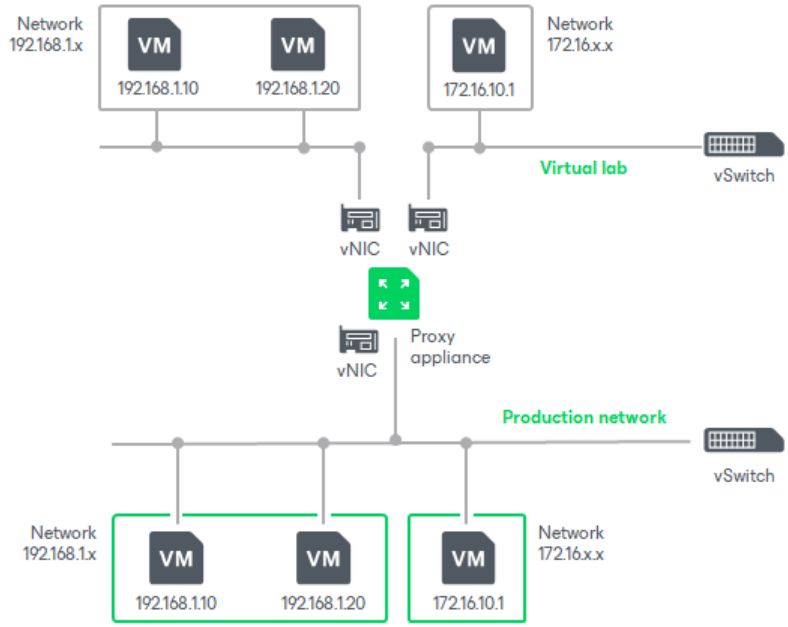


# IP Masquerading

To let the traffic into the virtual lab, Veeam Backup & Replication uses masquerade IP addressing.

Every VM in the virtual lab has a masquerade IP address, along with the IP address from the production network. The masquerade IP address resembles the IP address in the production network. For example, if the IP address of a VM is 172.16.1.13, the masquerade IP address may be 172.18.1.13.

The masquerade IP address is an "entry point" to the VM in the virtual lab from the production environment. When you want to access a specific VM in the virtual lab, Veeam Backup & Replication addresses it by its masquerade IP address.

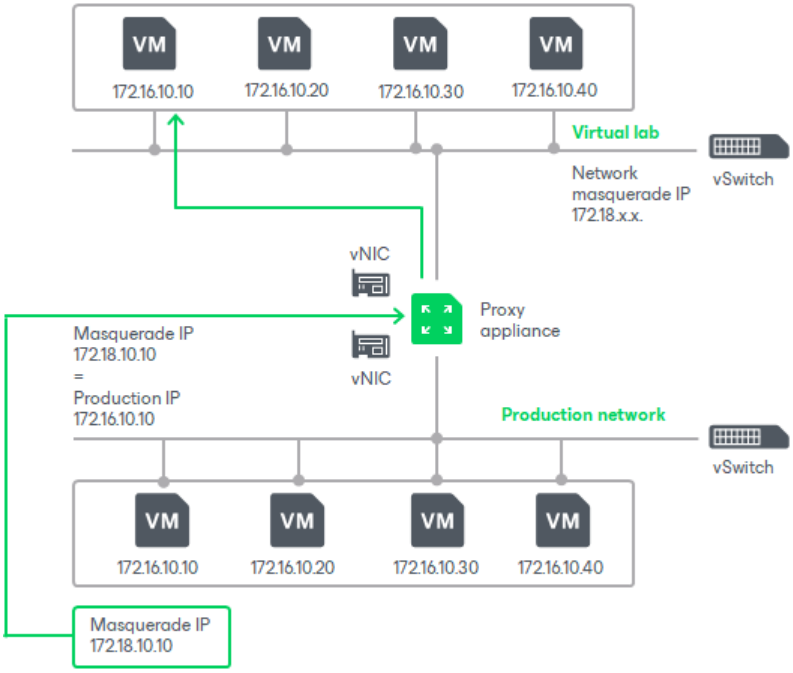


The rules that route requests to VMs in the virtual lab are specified in the routing table on the server from which you want to access VMs in the virtual lab. The routing table can be updated on the following servers:

- **Backup server**. Veeam Backup & Replication automatically creates the necessary static routes in the routing table on the backup server at the moment you launch a SureBackup job and Veeam Backup & Replication starts the virtual lab.
- **Client machine**. If you want to provide your users with access to VMs in the virtual lab, you need to manually update routing tables on their machines and add new static routes. For more information, see Static IP Mapping.

The added static route destines the masquerade network traffic to the proxy appliance. The proxy appliance acts as a NAT device: it resolves the masquerade IP address, replaces it with "real" IP address of the VM from the production network and then directs the request to the necessary VM in the virtual lab. The static route is non-persistent: when you power off the virtual lab, the route is removed from the routing table on the backup server or client machine.

For example, when trying to access a VM with IP address 172.16.10.10 in the isolated network, Veeam Backup & Replication sends a request to the masquerade IP address 172.18.10.10. According to the routing rule added to the IP routing table, all requests are first sent to the next hop — the proxy appliance. The proxy appliance performs address translation, substitutes the masquerade IP address with the IP address in the isolated network, and forwards the request to the necessary VM in the isolated network — in this example, to 172.16.10.10.

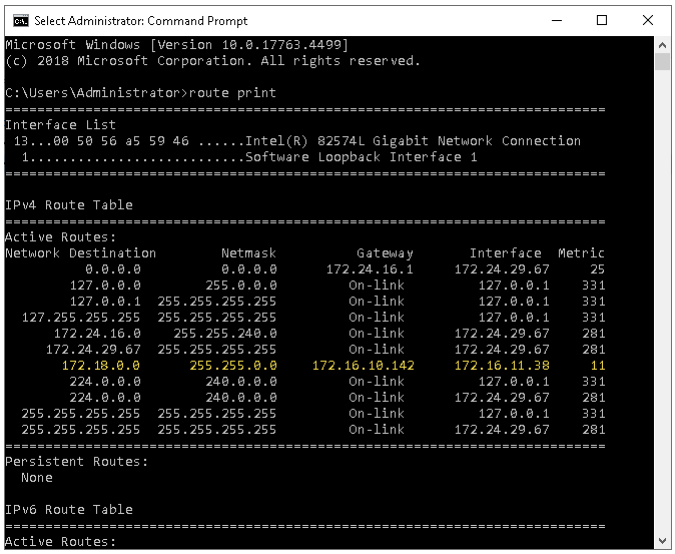


# Static IP Mapping

Sometimes it is necessary to provide many clients with access to a restored VM, which is especially the case for user-directed application item-level recovery. For example, you may want to provide users with access to the Microsoft Exchange Server started in the virtual lab using web-based access (like Outlook Web Access). Technically, you may update the routing table on every client machine; however, this will demand a lot of administrative effort.

For such situations, Veeam Backup & Replication lets you get access to a VM in the virtual lab directly from the production environment. To access a VM in the virtual lab, you must reserve a static IP address in the pool of production IP addresses and map this IP address to the IP address of a VM in the virtual lab.

The static IP address is assigned to the proxy appliance network adapter connected to the production network. IP traffic directed to the specified static IP address is routed by the proxy appliance to the VM in the isolated network.

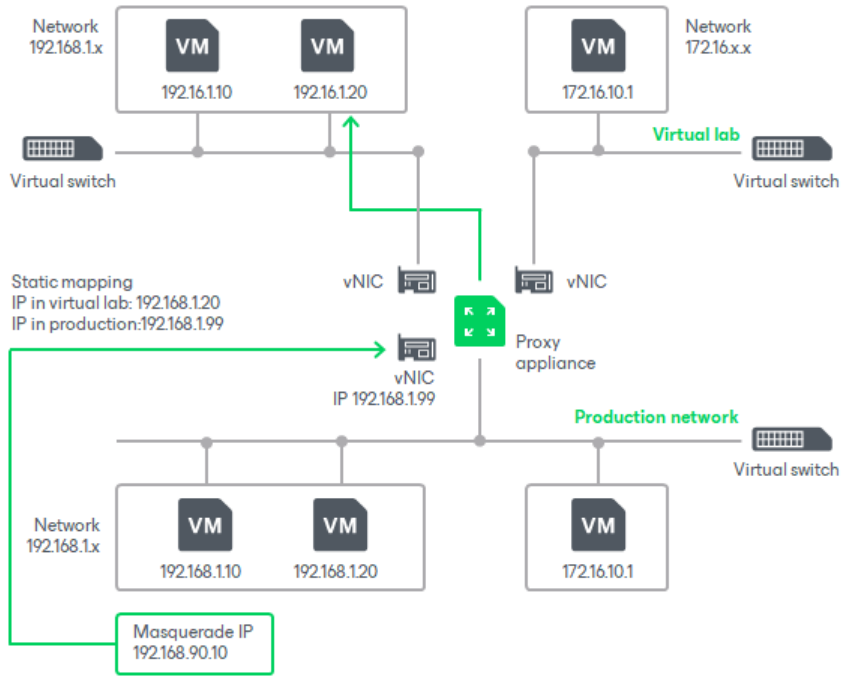


For example, for a VM with IP address 192.168.1.20 in the isolated network, you can reserve IP address 192.168.1.99 (a free IP address from the production network). As a result, you will be able to use IP address 192.168.1.99 to access the VM in the virtual lab from the production side.

You can also register an alias record in the production DNS server for the reserved IP address. For example, you can register backup.exchange.local as an alias for the IP address 192.168.1.99, and users will be able to access Microsoft Exchange Server by this alias.

# Virtual Lab Configuration

For SureBackup recovery verification, Veeam Backup & Replication offers two types of the virtual lab configuration:

- Basic single-host virtual lab
- Advanced single-host virtual lab

> **NOTE**
>
> You can also verify VM backups in Advanced Multi-Host virtual labs with DVS. This scenario can be helpful if you want to test VM backups and VM replicas in the same virtual lab or want to add verified VM backups and replicas to the same SureBackup job.
> For more information, see Advanced Multi-Host Virtual Labs.

## Basic Single-Host Virtual Labs

The basic single-host virtual lab can be used if all VMs that you want to verify, VMs from the application group and the backup server are connected to the same network.

For the basic single-host virtual lab, Veeam Backup & Replication creates one virtual network that is mapped to the necessary production network. Veeam Backup & Replication automatically adds a number of new objects on the ESXi host where the virtual lab is created:

- A resource pool
- A VM folder
- A standard vSwitch

The vSwitch is only used by the VMs started in the virtual lab. There is no routing outside the virtual lab to other networks.

Veeam Backup & Replication automatically configures all settings for the basic single-host virtual lab. The proxy appliance is also created and configured automatically on the ESXi host where the virtual lab is created.

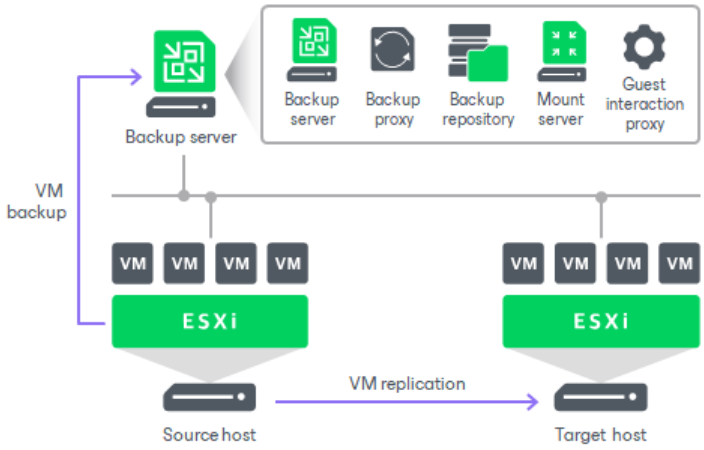


## Advanced Single-Host Virtual Labs

The advanced single-host virtual lab can be used if VMs that you want to verify and VMs from the application group are connected to different networks.

In the advanced single-host virtual lab, Veeam Backup & Replication creates several virtual networks for the virtual lab. The number of virtual networks corresponds to the number of production networks to which verified VMs are connected. Networks in the virtual lab are mapped to production networks.

Veeam Backup & Replication automatically adds a number of new VMware objects on the ESXi host where the virtual lab is created:

- A resource pool
- A VM folder
- A standard vSwitch

The vSwitch is only used by the VMs started in the virtual lab. There is no routing outside the virtual lab to other networks.

When you create an advanced single-host virtual lab, Veeam Backup & Replication configures basic settings for networks that are created in the virtual lab. You need to review these settings and manually adjust them.

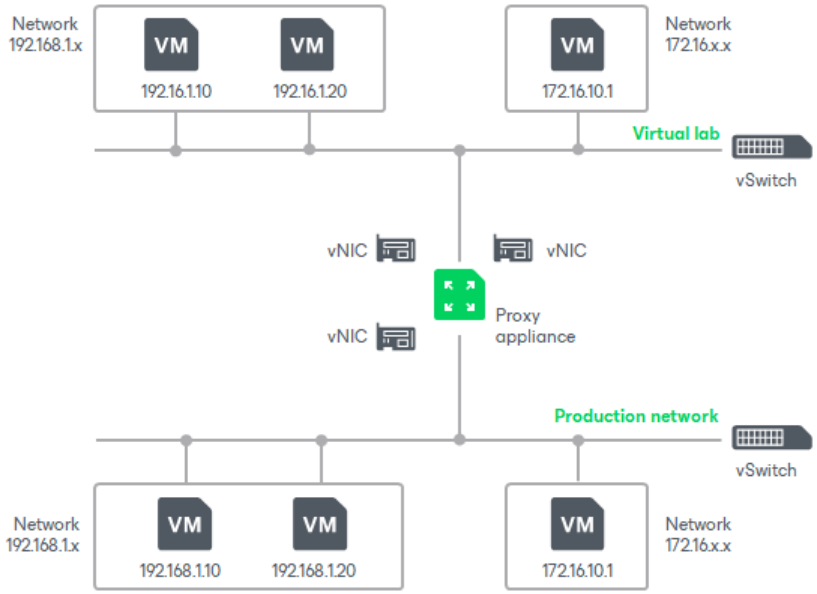


# Creating Virtual Lab

Before you create a new virtual lab, check prerequisites. Then use the **New Virtual Lab** wizard to create a virtual lab.

## Before You Begin

Before you create a virtual lab, check the following prerequisites:

- A valid license for Enterprise edition of Veeam Backup & Replication must be installed on the backup server.
- The ESXi host on which you plan to deploy a virtual lab must have a VMkernel interface. Otherwise the vPower NFS datastore will not be mounted on the ESXi host. For more information, see Veeam vPower NFS Service.
- If you plan to use the advanced multi-host networking mode for VM replicas verification, you must configure a DVS beforehand. For more information, see Advanced Multi-Host Virtual Labs.

> **NOTE**
>
> Starting from Veeam Backup & Replication 12, you can verify Microsoft Hyper-V backups and Veeam Agent backups using VMware vSphere virtual lab. To learn about SureBackup limitations for Veeam Agent backups, see Veeam Agent Management Guide.

# Step 1. Launch New Virtual Lab Wizard

To launch the **New Virtual Lab** wizard, do one of the following:

- Open the **Backup Infrastructure** view, in the inventory pane select **SureBackup**. In the working area, click **Add Virtual Lab** > **VMware**.
- Open the **Backup Infrastructure** view, in the inventory pane select **Virtual Labs** node under **SureBackup** and click **Add Virtual Lab** > **VMware vSphere**.
- Open the **Backup Infrastructure** view, in the inventory pane right-click **Virtual Labs** under **SureBackup** and select **Add Virtual Lab** > **VMware vSphere**.

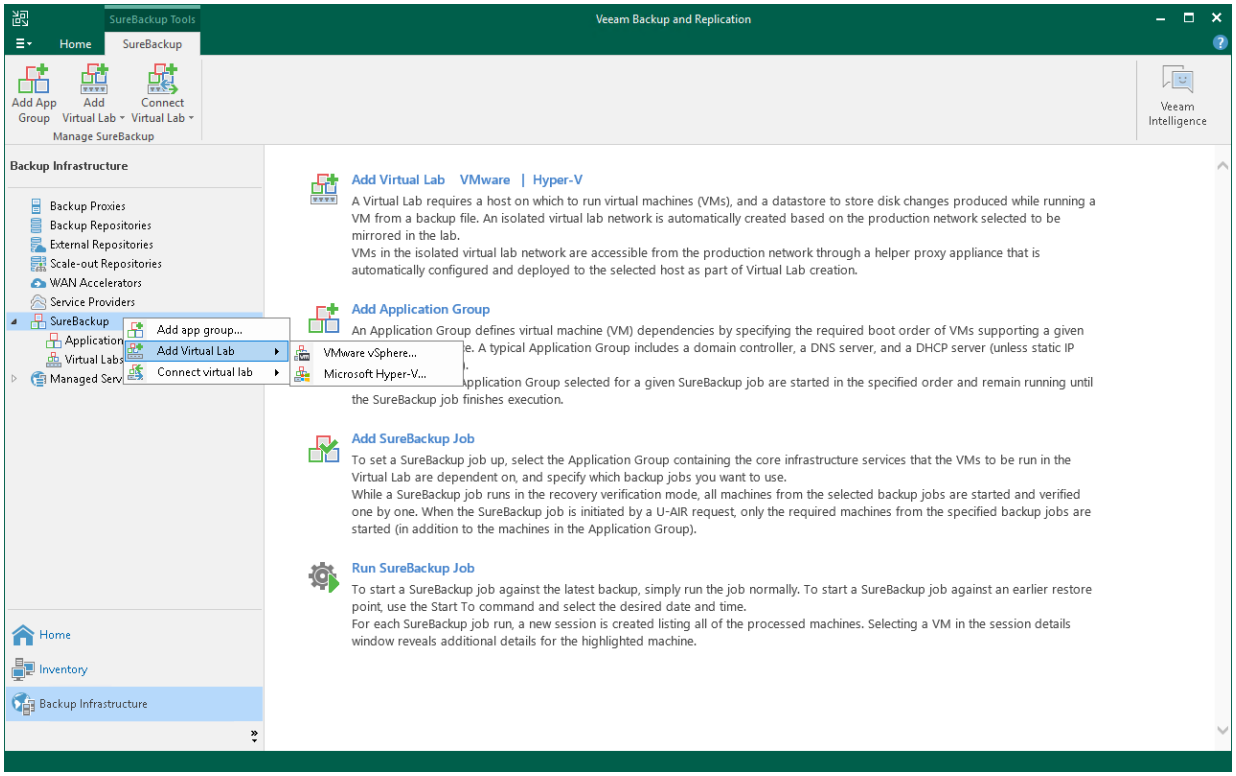

# Step 2. Specify Virtual Lab Name and Description

At the **Name** step of the wizard, specify a name and description for the virtual lab.

1. In the **Name** field, enter a name for the virtual lab.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the virtual lab, date and time when the lab was created.

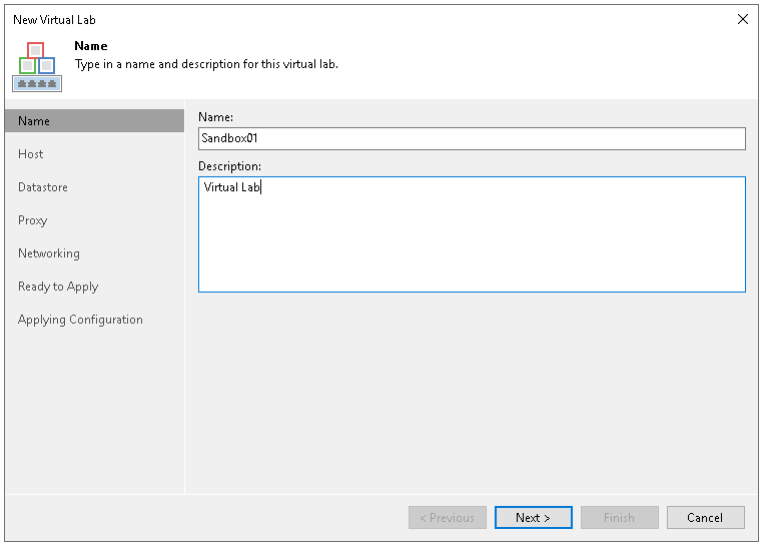

## Step 3. Select Host

At the **Host** step of the wizard, select an ESXi host on which the virtual lab must be created.

To select an ESXi host:

1. Click **Choose**.
2. Select an ESXi host on which the new virtual lab must be created. You can select a standalone ESXi host or an ESXi host being part of a cluster or vCenter Server hierarchy.
3. For every new virtual lab, Veeam Backup & Replication creates a dedicated folder and resource pool on the ESXi host. By default, the folder and pool have the same name as the virtual lab. To change the name of the folder and resource pool, click **Configure**. In the **Destination Options** window, enter the necessary names.

> **IMPORTANT**
>
> You cannot create resource pools in clusters with disabled DRS. If the target host is a part of such a cluster, the **Create a designated resource pool** option will be disabled in the **Destination Options** window.
> You cannot create folders on standalone ESXi hosts or ESXi hosts that are managed by the vCenter Severs but are added to Veeam Backup & Replication as standalone hosts. To overcome this situation, add the vCenter Server to Veeam Backup & Replication.

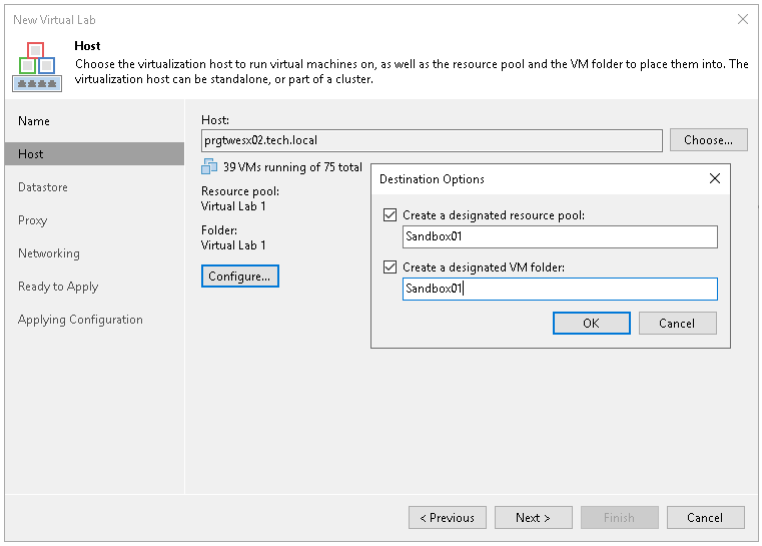


# Selecting an ESXi Host for VM Replicas Verification

When you select an ESXi host for the virtual lab where VM replicas will be verified, mind the location of verified VM replicas and VM replicas added to the application group:

- If verified VM replicas and VM replicas from the application group are located on the same ESXi host, you must select the ESXi host on which these VM replicas are registered. Verified VM replicas and VMs from the application group will be started on the selected ESXi host. If the application group contains VMs added from VM backups or storage snapshots, these VMs will also be started on the selected ESXi host.

  For this type of virtual lab configuration, you need to choose one of single-host networking modes: Basic single-host or Advanced single-host. For more information, see Selecting a Networking Mode.

- If verified VM replicas and VM replicas from the application group are located on different ESXi hosts, you can select any ESXi host in your virtual environment. Veeam Backup & Replication will create the virtual lab on the selected ESXi host. Verified VM replicas and VM replicas from the application group will be started on ESXi hosts where they are registered and connected to the virtual lab with the help of VMware DVS technology.

  The ESXi host on which the virtual lab is created must meet the following requirements:

  - The ESXi host must be located in the same datacenter where VM replicas are registered.
  - The ESXi host must have enough CPU and RAM resources. If the application group contains VMs that are started from backups or storage snapshots, these VMs will be started on the same ESXi host where the virtual lab is located, which will require a lot of resources.
  - For this type of virtual lab configuration, you must use the Advanced multi-host networking mode. For more information, see Selecting a Networking Mode.

# Step 4. Select Datastore

At the **Datastore** step of the wizard, you can select where redo logs for verified VMs must be stored. Redo logs are auxiliary files used to keep changes that take place when VMs run in the virtual lab. By default, redo logs are stored on the vPower NFS server. However, you can store redo logs on any datastore in the virtual environment. Redirecting redo logs improves verification performance. As soon as a recovery verification job completes, Veeam Backup & Replication deletes redo logs.

To redirect redo logs:

1. Select the **Redirect write cache** check box.

   If you perform staged restore, the **Redirect write cache** option allows you to select a datastore where VM delta files will be stored. Delta files are auxiliary files that keep changes made to a VM during script execution. For more information, see Staged Restore.

2. Click **Choose** and select a datastore from the list.

> **IMPORTANT**
>
> If disks of verified VMs are greater than 2 TB, you must not place redo logs on a VSAN datastore. Otherwise, Veeam Backup & Replication will fail to create snapshots for verified VMs. For more information, see VMware Docs.

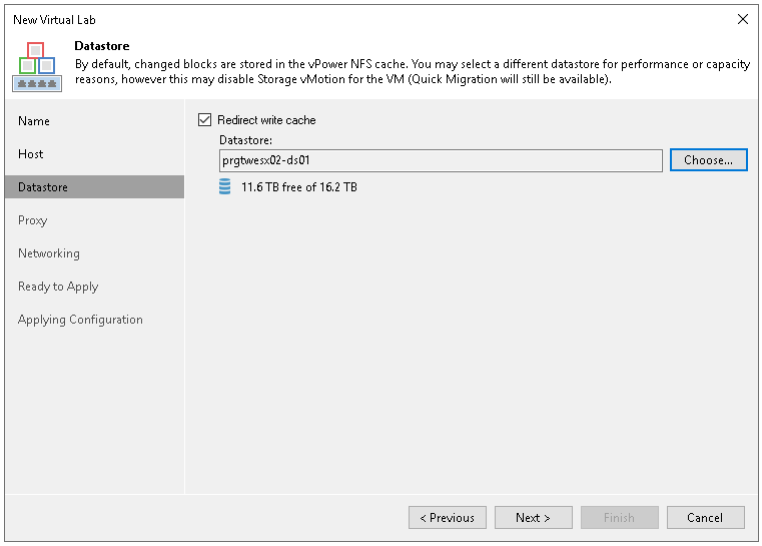

## Step 5. Set Up Proxy Appliance

At the **Proxy** step of the wizard, configure proxy appliance settings.

1. Select the **Use proxy appliance in this virtual lab** check box to enable automatic recovery verification of VMs. The proxy appliance acts as a gateway that provides access from the backup server to VMs in the virtual lab. If you do not select this check box, during recovery verification Veeam Backup & Replication will only start VMs in the virtual lab and perform the heartbeat test for VMs. You will have to manually test VMs or perform manual item-level restore over the VM console.
2. By default, the proxy appliance is placed on a datastore with the maximum amount of free space. The default name of the proxy appliance is the virtual lab name that you have specified at the **Name** step of the wizard. To change a name or a datastore for the proxy appliance, click **Edit** and specify a new name or choose a different datastore.
3. Click **Configure** and select a production network in which the proxy appliance will be created. Select IPv4 or IPv6 or use both. Specify an IP address for the proxy appliance in the production network and settings of the DNS server to be used. You can choose to automatically obtain an IP address for the backup proxy and DNS server settings or set them manually.

**IMPORTANT**

Consider the following:

- If you assign to the proxy appliance an IP address from the same network where the backup server is located, Veeam Backup & Replication will automatically add a new route to the routing table on the backup server. If you assign to the proxy appliance an IP address from a different network, you will have to manually add a new route to the routing table on the router in the production network. If you do not add a new route, tests and application scripts will fail and you will not be able to access VMs in isolated networks.
- When Veeam Backup & Replication starts a virtual lab, it verifies if the proxy appliance is available by sending a ping request to it. If the route is not added to the routing table, the SureBackup job will fail.
- You cannot edit the production network after you create the virtual lab.

4. By default, VMs in the virtual lab work in the isolated environment and do not have access to internet. If you want to let VMs in the virtual lab access the internet, select the **Allow proxy appliance to act as internet proxy for virtual machines in this** lab check box. In the **Port** field, specify a port for HTTP traffic. By default, port 8080 is used. In the **Production proxy** field, you can optionally specify an IP address or a fully qualified domain name of an internet-facing proxy server that VMs must use to access internet.
5. On every VM that you plan to start in the virtual lab, adjust proxy settings in the internet options. To do this, on the VM open **Internet Options** > **Connections** > **LAN Settings** > **Proxy server** and specify an IP address of the proxy appliance on the isolated network and port number.

**NOTE**

When you allow the proxy appliance to act as an internet proxy, you enable the HTTP(S) internet access for VMs in the virtual lab. The proxy appliance does not proxy other protocols (such as ICMP protocol used for ping tests) for VMs in the virtual lab.

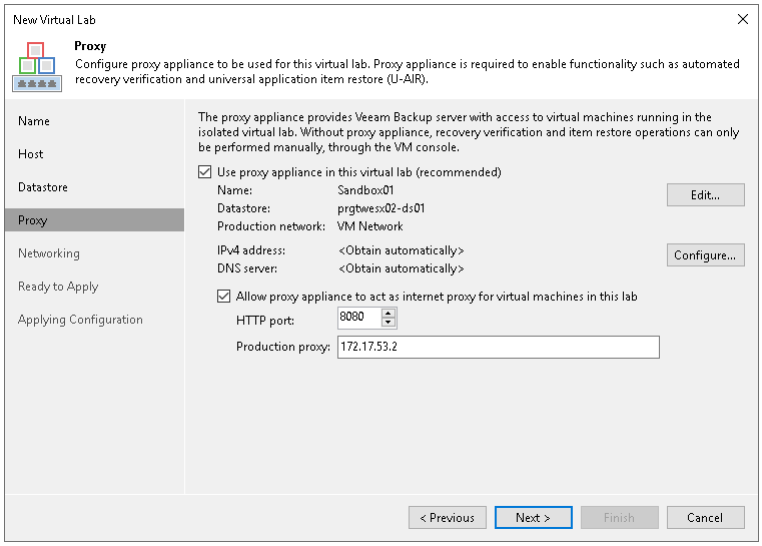
New Virtual Lab
Proxy
Configure proxy appliance to be used for this virtual lab. Proxy appliance is required to enable functionality such as automated recovery verification and universal application item restore (U-AIR).
Name
Host
Datastore
Proxy
Networking
Ready to Apply
Applying Configuration
The proxy appliance provides Veeam Backup server with access to virtual machines running in the isolated virtual lab. Without proxy appliance, recovery verification and item restore operations can only be performed manually, through the VM console.
Use proxy appliance in this virtual lab (recommended)
Name: Sandbox01
Datastore: prgtwesx02-ds01
Production network: VM Network
Edit...
IPv4 address: <Obtain automatically>
DNS server: <Obtain automatically>
Configure...
Allow proxy appliance to act as internet proxy for virtual machines in this lab
HTTP port: 8080
Production proxy: 172.17.53.2
< Previous
Next >
Finish
Cancel

## Step 6. Select Networking Mode

At the **Networking** step of the wizard, select the type of network settings configuration. The virtual lab configuration depends on objects that you plan to verify in the virtual lab:

- Backups
- Replicas
- VMs from storage snapshots

# Selecting Networking Mode for Verifying Backups

Veeam Backup & Replication offers two networking modes for the virtual lab in which VMs from backups can be verified:

- **Basic single-host**. This networking mode is recommended if all VMs that you plan to verify, VMs from the application group and the backup server are located in the same production network. In this case, Veeam Backup & Replication will automatically define all networking settings for the virtual lab.
- **Advanced single-host**. This networking mode is recommended if VMs that you plan to verify and VMs from the application group are located in different networks. In this case, you will have to manually define settings for isolated networks in the virtual lab.

If you select the **Advanced single-host** option, the **New Virtual Lab** wizard will include additional steps for customizing network settings.

> **NOTE**
>
> You can also verify VM backups in Advanced Multi-Host virtual labs with DVS. This scenario can be helpful if you want to test VM backups and replicas in the same virtual lab or want to add verified VM backups and replicas to the same SureBackup job.
> For more information, see Advanced Multi-Host Virtual Labs.

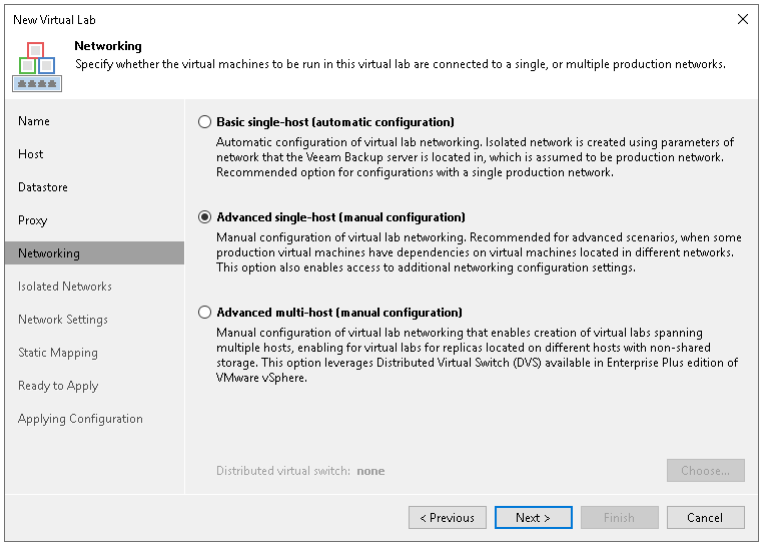

# Selecting Networking Mode for Verifying Replicas

Veeam Backup & Replication offers three networking modes for the virtual lab in which VM replicas are verified:

- **Basic single-host**. This type of networking is recommended if VM replicas that you plan to verify are located on the same ESXi host and are connected to the same production network. The backup server must also be located in this network. In this case, Veeam Backup & Replication will automatically define all networking settings for the virtual lab.
- **Advanced single-host**. This type of networking is recommended if VM replicas that you plan to verify are located on the same ESXi host but connected to different networks. In this case, you will have to manually define settings for isolated networks in the virtual lab.
- **Advanced multi-host**. This type of networking is recommended if VM replicas that you plan to verify are located on the different ESXi hosts. For multi-host configuration of the virtual lab, Veeam Backup & Replication uses VMware DVS technology.

  If you select the **Advanced multi-host** option, click **Choose** and select the necessary DVS in your virtual environment. Note that Veeam Backup & Replication does not configure DVS automatically: you must configure it beforehand.

You can use VMware NSX-T 3.0 or later as isolated or production networks in **Advanced multi-host** networking mode. In this case, you must configure DVS and NSX-T networks beforehand. NSX-T networks must be configured using VDS, as N-VDS is not supported.

If the **Advanced single-host** or **Advanced multi-host** option is selected, the **New Virtual Lab** wizard will include additional steps for customizing network settings.

> **IMPORTANT**
>
> For every isolated network in the virtual lab, Veeam Backup & Replication adds a new port group to the DVS. If you use a production DVS, you must isolate port groups created by Veeam Backup & Replication from the production environment. For more information, see Isolated Networks on DVS.

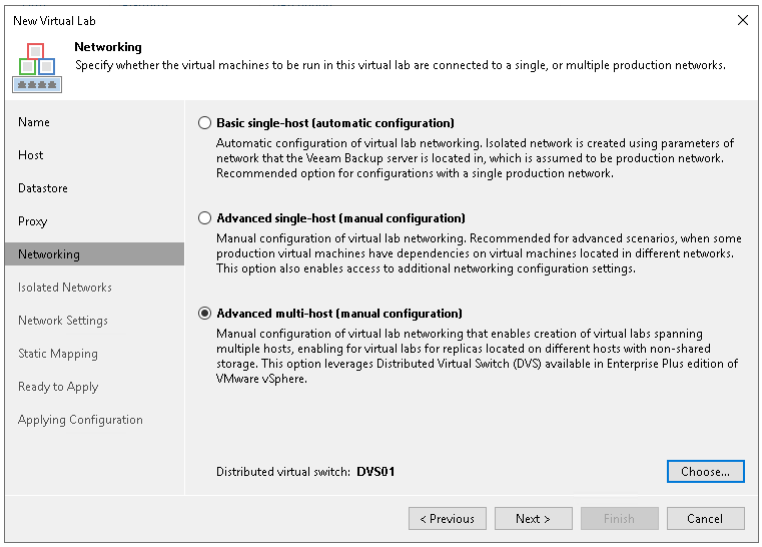


# Selecting Network Mode for Verifying VMs on Storage Snapshots

For verifying VMs from storage snapshots, you can select any networking mode.

# Step 7. Create Isolated Networks

The **Isolated Networks** step of the wizard is available if you have selected the advanced networking option at the Networking step of the wizard.

At the **Isolated Networks** step of the wizard, you must configure isolated networks to which verified VMs and VMs from the application group will be connected and map these networks to production networks where original VMs are located.

To add a network:

1. Click **Add**.
2. From the **Production network** list, select a production network in which a VM from the application group or verified VM resides.
3. In the **Isolated network** field, specify a name for the isolated network that must be mapped to the selected production network. You can specify a new name to create a new isolated network or select the existing isolated network from the drop-down list.
4. In the **VLAN ID** field, enter an ID for the created network. In the advanced multi-host virtual lab, VLAN IDs help ensure that the created network is isolated from the production environment. Alternatively, you can manually connect the DVS that you plan to use to the isolated network. For more information, see Advanced Multi-Host Virtual Labs.

**IMPORTANT**

Consider the following:

- Isolated network names cannot contain the following symbols: ~` !@#$%^&*+=;'><|?*:". Otherwise, Veeam Backup & Replication will return an error. For more information, see this Veeam KB article.
- You can map several production networks to the same isolated network. The production networks that you plan to map must have the same network masks and pools of IP addresses.
- You cannot map one production network to several isolated networks.
- You can assign only up to 9 isolated networks per virtual lab.
- Do not set the VLAN ID to match the production VLAN unless it is required. Running tests on the production VLAN can disrupt live network traffic and cause other potential issues.

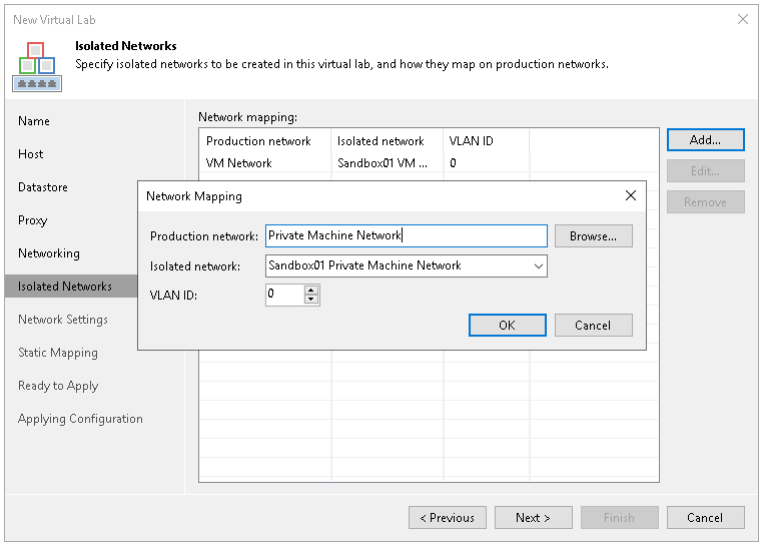

# Step 8. Specify Network Settings

The **Network Settings** step of the wizard is available if you have selected the advanced networking option at the Networking step of the wizard.

At the **Network Settings** step of the wizard, you must specify settings for every created isolated network and define how production networks map to isolated networks in the virtual lab.

Communication between production networks and isolated networks is carried out through vNIC adapters on the proxy appliance. A new vNIC adapter must be added for every isolated network. If you are planning to use IPv4 and IPv6 in your virtual lab, you must create a vNIC for each protocol.

To add a vNIC adapter for an isolated network:

1. At the **Network Settings** step of the wizard, click **Add**.
2. If you have both IPv4 and IPv6 enabled, select the required IP address.
3. Select a network to which the vNIC adapter must be connected. Specify an IP address that the proxy appliance must have in the isolated network and a subnet mask for this isolated network. For IPv6, specify a prefix length instead of a subnet mask. Typically, the IP address set for the proxy appliance coincides with the IP address of the gateway in the production network.
4. After you specify the IP address, Veeam Backup & Replication will automatically configure a masquerade IP address for accessing VMs in the virtual lab from the production network. You can change the masquerade network IP address if necessary.
5. Select the **Enable DHCP service on this interface** check box if you need to dynamically assign IP addresses for machines. The assigned IP addresses belong to the current isolated network.
6. Click the **DNS Servers** button and specify settings of a virtualized DNS server if necessary. Click **OK**.
7. To enable communication between isolated networks, select the **Route network traffic between vNICs** check box. Make sure that the IP address of the proxy appliance in the isolated network matches the IP address of the gateway in the production network.

**IMPORTANT**

Consider the following:

- You cannot assign more than one vNIC to a single isolated network for each protocol.
- Network addresses specified for different vNIC adapters must belong to different networks. For example, if the first network adapter has the 192.168.0.1 IP address and the network mask is 255.255.255.0, and the second one — 192.168.0.2 and the network mask is 255.255.255.0, such configuration will not work. In this situation, you need to assign to the second adapter an IP address from a different network, for example, 172.16.0.1.
- If you assign more than 1 network adapter (or IP address) to one production VM, Veeam Backup & Replication will apply the predefined tests only to one of these IP addresses. For more information on the predefined tests, see Predefined tests.

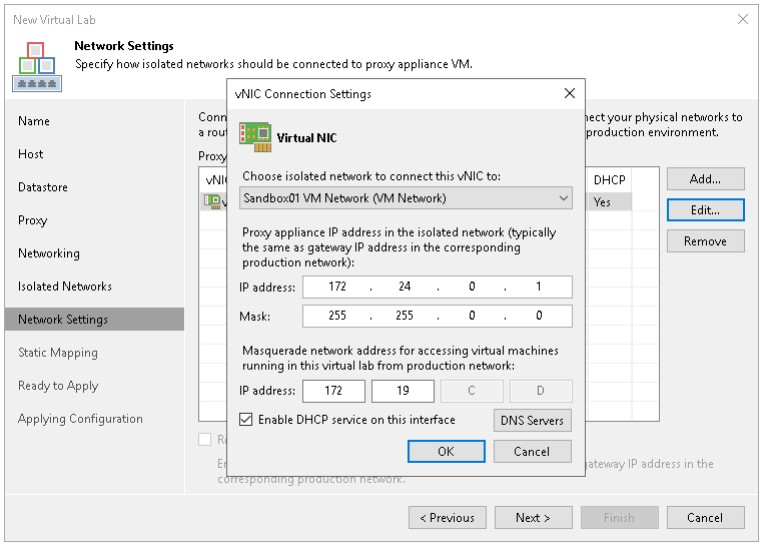
New Virtual Lab
Network Settings
Specify how isolated networks should be connected to proxy appliance VM.
Name
Host
Datastore
Proxy
Networking
Isolated Networks
Network Settings
Static Mapping
Ready to Apply
Applying Configuration
vNIC Connection Settings
Virtual NIC
Choose isolated network to connect this vNIC to:
Sandbox01 VM Network (VM Network)
Proxy appliance IP address in the isolated network (typically the same as gateway IP address in the corresponding production network):
IP address: 172 . 24 . 0 . 1
Mask: 255 . 255 . 0 . 0
Masquerade network address for accessing virtual machines running in this virtual lab from production network:
IP address: 172 19 C D
Enable DHCP service on this interface
DNS Servers
OK
Cancel
DHCP
Yes
Add...
Edit...
Remove
< Previous
Next >
Finish
Cancel

# Step 9. Specify Static IP Mapping Rules

The **Static Mapping** step of the wizard is available if you have selected the advanced networking option at the Networking step of the wizard.

At the **Static Mapping** step of the wizard, you can specify static IP address mapping rules to make VMs in the virtual lab accessible from any machine in the production network.

To add a new rule:

1. Select the **Define static IP address mapping** check box.
2. Click **Add**.
3. If you have both IPv4 and IPv6 enabled, select the required address mapping.
4. In the **IP Address Mapping** window, specify settings of a new rule:
   a. From the **Production network** drop-down list, select a production network in which a VM from the application group or verified VM resides.
   b. In the **Isolated IP** field, specify the IP address of the VM in the production network.
   c. In the **Access IP** field, specify the IP address in the production network that you want to use to access the VM in the virtual lab. You must use an IP address that is not allocated to any machine yet.

> **NOTE**
>
> It is recommended that you assign an access IP from the same subnet where the proxy appliance resides. In the opposite case, you will have to configure routing rules for the access IP manually.

For example, a DNS server that you plan to start in the virtual lab has IP address 172.17.53.15 in the production network. To set static mapping for the DNS server:

1. In the **Isolated IP** field, you must define its production IP address — 172.17.53.15.
2. In the **Access IP** field, you must define any free IP address from the production network, for example, 172.17.53.162.

After a virtual lab is created and VMs are started in the virtual lab, you will be able to access the DNS server in the virtual lab from the production environment by IP address 172.17.53.162.

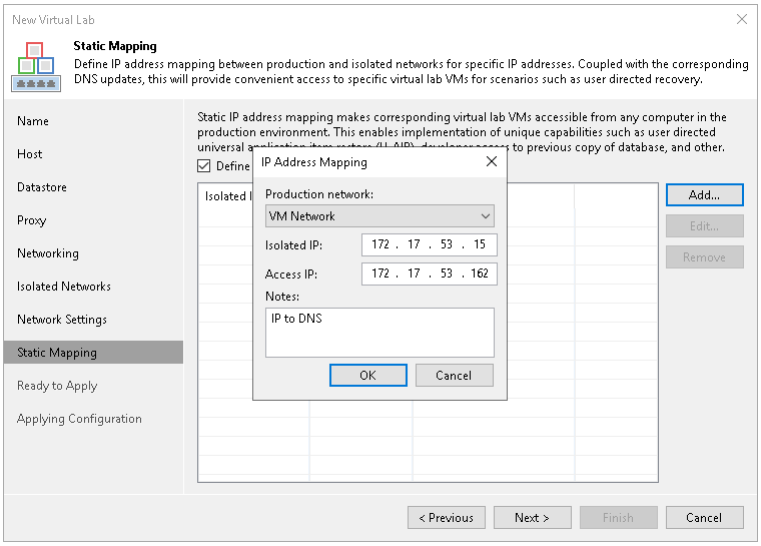

## Step 10. Apply Parameters

At the **Ready to Apply** step of the wizard, complete the procedure of virtual lab configuration.

1. Review details of the virtual lab.
2. Click **Apply** to create the virtual lab.
3. At the last step of the wizard, click **Finish** to exit the wizard.

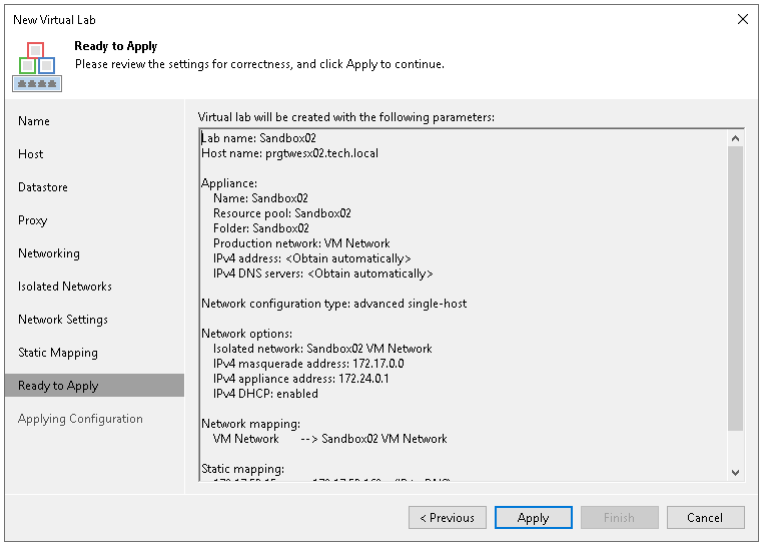


# Connecting to Existing Virtual Lab

You can connect an existing virtual lab and use this virtual lab for recovery verification. For example, you can connect to a virtual lab that has been created on another backup server.

To connect to a virtual lab:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Virtual Labs** under **SureBackup** and click **Connect Virtual Lab** > **VMware** on the ribbon or right-click **Virtual Labs** and select **Connect Virtual Lab** > **VMware**.

3. Select the virtual lab and click **Connect**. To quickly find a virtual lab, use the search field at the bottom of the **Select Virtual Lab** window: enter a virtual lab name or a part of it in the field at the bottom and press [Enter] on the keyboard.

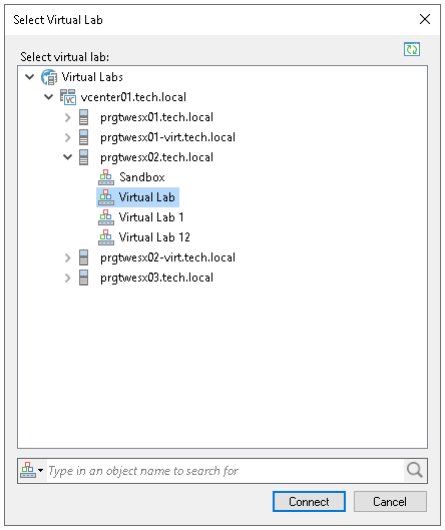


# Editing and Deleting Virtual Labs

You can edit settings of a virtual lab or delete the virtual lab.

Always use Veeam Backup & Replication to modify or delete a virtual lab. If you edit virtual lab settings or delete any of its components from outside, for example, in vSphere Client, the lab will be corrupted and its component such as the created vSwitch, resource pool will remain in the virtual infrastructure.

# Editing Virtual Labs

To edit settings of a virtual lab:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Virtual Labs** under **SureBackup**.
3. In the working area, select a virtual lab and click **Edit Virtual Lab** on the ribbon or right-click the virtual lab and select **Properties**.
4. Edit virtual lab settings as required.

# Removing Virtual Labs

To remove a virtual lab:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Virtual Labs** under **SureBackup**.
3. In the working area, select a virtual lab and click **Remove Virtual Lab** on the ribbon or right-click the virtual lab and select **Delete**.

4. If you want to remove virtual lab object from the virtual infrastructure, in the displayed window select the **Remove virtual lab objects from host** check box. If you do not enable this option, Veeam Backup & Replication will disconnect the virtual lab from the backup server. You will be able to connect to this virtual lab later.

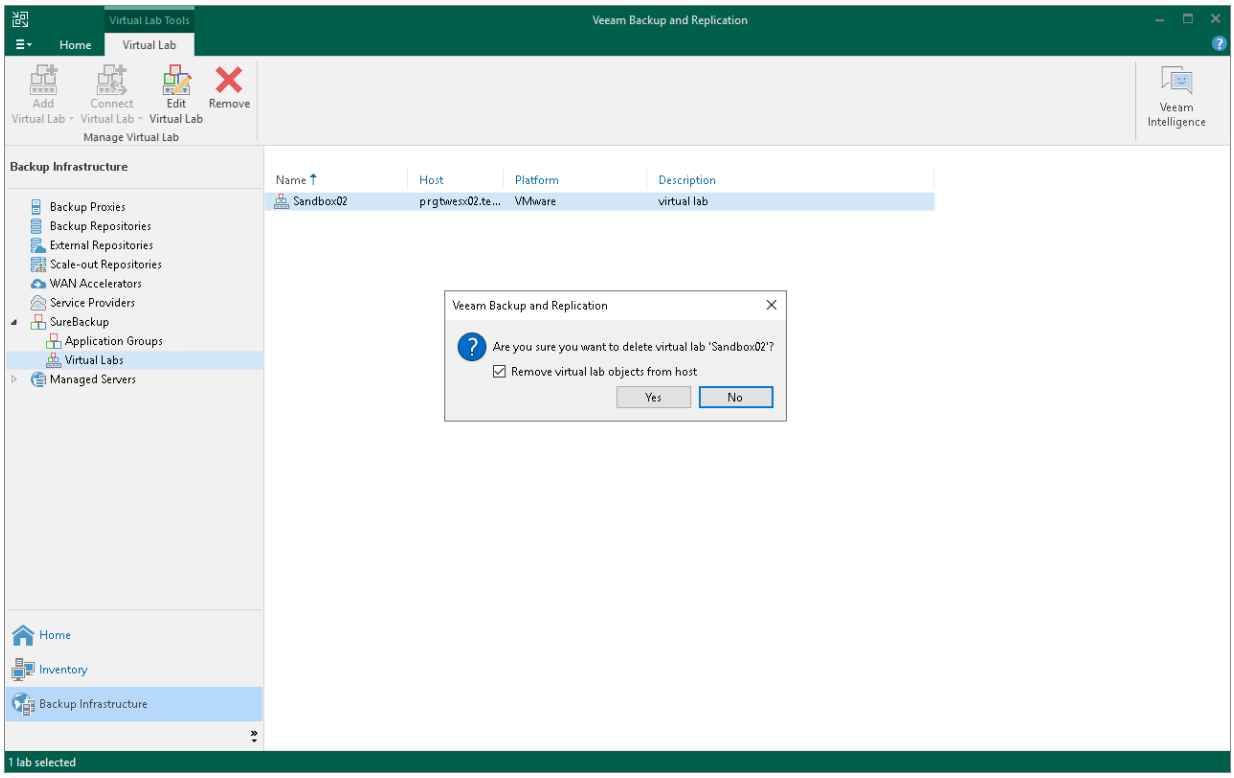

# SureBackup Job

A SureBackup job is a task for recovery verification. The SureBackup job can run in two different modes: **Full recoverability testing** and **Backup verification and content scan only**.

In the **Full recoverability testing** mode, the SureBackup job aggregates all settings and policies of the recovery verification task such as application group and virtual lab to be used, machine backups that must be verified in the virtual lab, and so on. It runs machines in an isolated environment directly from backup and performs tests against live applications. This mode ensures recoverability of your production workloads in a disaster recovery event.

In the **Backup verification and content scan only** mode, the SureBackup job performs only backup integrity check and its content analysis to detect traces of malware or any other unwanted or sensitive data. These tests do no require setting up a virtual lab or an application group.

You can run the SureBackup job manually or schedule it to run automatically.

When a SureBackup job runs in the **Full recoverability testing** mode, Veeam Backup & Replication first creates an environment for recovery verification:

1. Veeam Backup & Replication starts the virtual lab.
2. In the virtual lab, Veeam Backup & Replication starts machines from the application group in the required order. Machines from the application group remain running until verified machines (machines from the linked job) are booted from backups and tested.

   If Veeam Backup & Replication does not find a valid restore point for any of the machines from the application group, the SureBackup job fails.
3. When the virtual lab is ready, Veeam Backup & Replication starts verified machines to the necessary restore point and, depending on the job settings, verifies them one by one or creates several streams and verifies a number of machines simultaneously.

   If Veeam Backup & Replication does not find a valid restore point for any of the verified machines, verification of this machine fails, but the job continues to run.

By default, you can start and test up to three machines at the same time. You can also increase the number of machines to be started and tested simultaneously. Keep in mind that if these machines are resource demanding, performance of the SureBackup job as well as performance of the ESXi host on which the virtual lab resides may decrease.

Once the verification process is complete, machines from the application group are powered off. Optionally, you can leave machines from the application group running to perform manual verification or enable user-directed application item-level recovery.

In some cases, the SureBackup job schedule may overlap the schedule of the backup job linked to it. The backup file may be locked by the backup job and the SureBackup job will be unable to verify such backup. In this situation, Veeam Backup & Replication will not start the SureBackup job until the backup job is over.

To overcome the situation of job overlapping, you may chain the backup and SureBackup jobs or define the timeout period for the SureBackup job. For more information, see Specifying Job Schedule.

# SureBackup Job Processing

SureBackup job processing depends on the verification mode, **Full recoverability testing** mode or **Backup verification and content scan only** mode.

# Full recoverability testing

The recovery verification process includes the following steps:

1. **Getting virtual lab configuration**. Veeam Backup & Replication gets information about configuration of the virtual lab where verified VMs must be started.
2. **Starting virtual lab routing engine**. Veeam Backup & Replication starts a proxy appliance used as a gateway to provide access to the virtual lab.
3. **Performing malware scan**. If the recovery verification process includes malware scan, Veeam Backup & Replication scans VM data with antivirus software.

   After the malware scan is complete, Veeam Backup & Replication registers the VM on the selected ESXi host, powers it on, and runs recovery verification tests for this VM.

   Veeam Backup & Replication verifies VMs sequentially — one after another. For example, when the malware scan and recovery verification tests for VM *A* complete, Veeam Backup & Replication verifies VM *B*, and so on.
4. **Publishing**. Veeam Backup & Replication creates a vPower NFS datastore with a VM backup and registers it on the selected ESXi server. Veeam Backup & Replication does not deploy the whole VM from the backup file, it deploys VM configuration files only. Virtual disks are deployed per force and per required data blocks.
5. **Reconfiguring**. Veeam Backup & Replication updates configuration files for VMs that must be started in the isolated network.
6. **Registering**. Veeam Backup & Replication registers the verified VM on the selected ESXi host.
7. **Configuring DC**. If a verified VM has the Domain Controller or Global Catalog role, the VM is reconfigured.
8. **Powering on**. Veeam Backup & Replication powers on the verified VM in the isolated network.
9. **Performing heartbeat test**. Veeam Backup & Replication checks whether the VMware Tools heartbeat signal (green or yellow) is coming from the VM or not. If the VM has no VMware Tools, the test is not performed, and a notification is written to the session details.
10. **Running ping tests**. Veeam Backup & Replication checks if the VM responds to the ping requests or not. If the VM has no NICs and mapped networks for them and has no VMware Tools installed, the ping test is not performed, and a notification is written to the session details.
11. **Application initialization**. Veeam Backup & Replication waits for the applications installed in the VM, for example, Microsoft SQL Server, to start. The application initialization period is defined in settings of the SureBackup job and by default equals to 120 sec. Depending on the software installed in a VM, the application initialization process may require more time than specified in the job settings. If applications installed in a VM are not initialized within the specified period of time, test scripts can be completed with errors. If such an error situation occurs, you need to increase the **Application initialization timeout** value and start the job once again.
12. **Running test scripts**. Veeam Backup & Replication runs scripts to test whether the application installed in the VM is working correctly or not. If the VM has no VMware Tools installed or there are no NICs and mapped networks for them, Veeam Backup & Replication skips tests that use the *%vm_ip%* and *%vm_fqdn%* variables as the IP address and FQDN of the VM cannot be determined.

    Test results are written to the job session details. To define whether the script is completed successfully or not, Veeam Backup & Replication uses return codes. If the return code is equal to 0, the script is considered to complete successfully. Other values in the return code mean that the script failed.
13. **Powering off**. After all tests were performed, Veeam Backup & Replication powers off the verified VM.

14. **Unregistering**. Veeam Backup & Replication unregisters the verified VM on the selected ESXi host.
15. **Clearing redo logs**. Veeam Backup & Replication deletes redo logs from the datastore in the production environment. Redo logs store changes made to the VM while it is running from the backup file.
16. **Unpublishing**. Veeam Backup & Replication unpublishes the content of the backup file on the ESXi host.
17. **Running backup validation test**. After a VM was verified, powered off and unpublished, Veeam Backup & Replication runs a CRC check to verify the VM backup at the file level and make sure that this file is not corrupted.
18. **Stopping virtual lab engine**. Veeam Backup & Replication powers off the proxy appliance in the virtual lab.
19. **Deleting network routes**. Veeam Backup & Replication deletes added network routes from the routing table on the backup server.

## Stabilization Algorithm

Stabilization algorithm is the same for both recovery verification modes. To be able to perform tests for a verified VM without errors, Veeam Backup & Replication needs to know that the VM is ready for testing. To determine this, Veeam Backup & Replication waits for the VM to reach a stabilization point: that is, waits for the VM to boot completely and report it is ready for tests. After the stabilization point has been established, Veeam Backup & Replication can start performing heartbeat tests, ping tests and running test scripts against the VM.

Veeam Backup & Replication establishes the stabilization point with the help of VMware parameters that it gets from the VM. Depending on the VM configuration, it uses one of three algorithms to do that:

- **Stabilization by IP**. This algorithm is used if the VM has VMware Tools installed, there are NICs and mapped networks for these NICs. In this case, Veeam Backup & Replication waits for an IP address of the VM for mapped networks that is sent by VMware Tools running in the VM. The sent IP address must be valid and must not change for a specific period of time. For more information, see Backup Recovery Verification Tests.
- **Stabilization by heartbeat**. This algorithm is used if the VM has VMware Tools installed but there are no vNICs and mapped networks for them. In this case, Veeam Backup & Replication waits for a green or yellow heartbeat signal from the VM. The signal is sent by VMware Tools running in the VM.
- **Stabilization by Maximum allowed boot time**. This algorithm is used if the VM has neither VMware Tools installed, nor NICs and mapped networks for them. In this case, Veeam Backup & Replication waits for the time specified in the **Maximum allowed boot time** field, which is considered to be a stabilization period for the VM. Once this time interval is exceeded, Veeam Backup & Replication considers that the VM runs successfully and it is ready for testing.

When the stabilization point has been established, Veeam Backup & Replication runs ping, heartbeat tests and performs test scripts against the verified VM.

The stabilization process cannot exceed the time specified in the **Maximum allowed boot time** field. For this reason, you should be careful when specifying this value. Typically, a VM started by a SureBackup job requires more time to boot than a VM started in the production environment. If the stabilization point cannot be determined within the **Maximum allowed boot time**, the recovery verification process is finished with the timeout error. When such an error occurs, you need to increase the **Maximum allowed boot time** value and start the job again.

# Backup verification and content scan only

The recovery verification process includes the following steps:

1. **Performing malware scan**. If the recovery verification process includes malware scan and YARA rule scan, Veeam Backup & Replication scans VM data with antivirus software and the YARA rule.

   Veeam Backup & Replication verifies VMs sequentially — one after another. For example, when the malware scan and recovery verification tests for VM *A* complete, Veeam Backup & Replication verifies VM *B*, and so on.

2. **Running backup integrity check**. Veeam Backup & Replication runs a CRC check to verify the VM backup at the file level and make sure that this file is not corrupted.

# Creating SureBackup Job

To create a new SureBackup job, use the **New SureBackup Job** wizard.

## Before You Begin

Before you create and start a SureBackup job, check the following prerequisites:

- A valid Veeam Universal License of Veeam Backup & Replication must be installed on the backup server. When using a legacy socket-based license, Enterprise or higher edition is required.
- All applications and services on which verified machines are dependent must be virtualized in your environment.
- To perform full recoverability testing, you must create or connect a virtual lab. For more information, see sections Creating Virtual Lab and Connecting to Existing Virtual Lab.
- If you plan to scan machine data for malware, check requirements and limitations.
- If you plan to verify machines with a ping test, the firewall on tested machines must allow ping requests.
- If you plan to verify machines with a heartbeat test, VMware Tools must be installed in tested machines.
- [For storage snapshots] The storage system must be added to the backup infrastructure.

Consider the following limitations:

- Verified VM replicas must be in the *Ready* state. If a VM replica is in the *Failover* or *Failback* state, you will not be able to verify it with the SureBackup job.
- You cannot link the following machines backups to the SureBackup job: backups of VMware Cloud Director VMs, backups created by backup copy jobs and backups stored in cloud backup repositories.
- The source backup or replication job has a higher priority than the SureBackup job. If the source backup or replication job starts when the SureBackup job is running, and this job is about to modify the restore point from which the VM is started, Veeam Backup & Replication automatically powers off VMs in the virtual lab and completes the SureBackup job.
- The Microsoft SQL Server Checker script runs on the backup server side. For this reason, Named Pipes or TCP/IP connections must be enabled for the Microsoft SQL Server running in the virtual lab. For more information, see Microsoft Docs.

# Step 1. Launch New SureBackup Job Wizard

To launch the **New SureBackup Job** wizard, do either of the following:

- Open the **Backup Infrastructure** view, in the inventory pane select **SureBackup**. In the working area, click **Add SureBackup Job**.
- Open the **Home** view. On the **Home** tab, click **SureBackup Job** on the ribbon.
- Open the **Home** view. In the inventory pane, right-click **SureBackup** under **Jobs** and select **SureBackup**.

  You can use this method if you already have at least one SureBackup job. If there are no SureBackup jobs, the **SureBackup** node will not be displayed in the inventory pane. In this case, you can right-click **Jobs** in the inventory pane and select **SureBackup**.

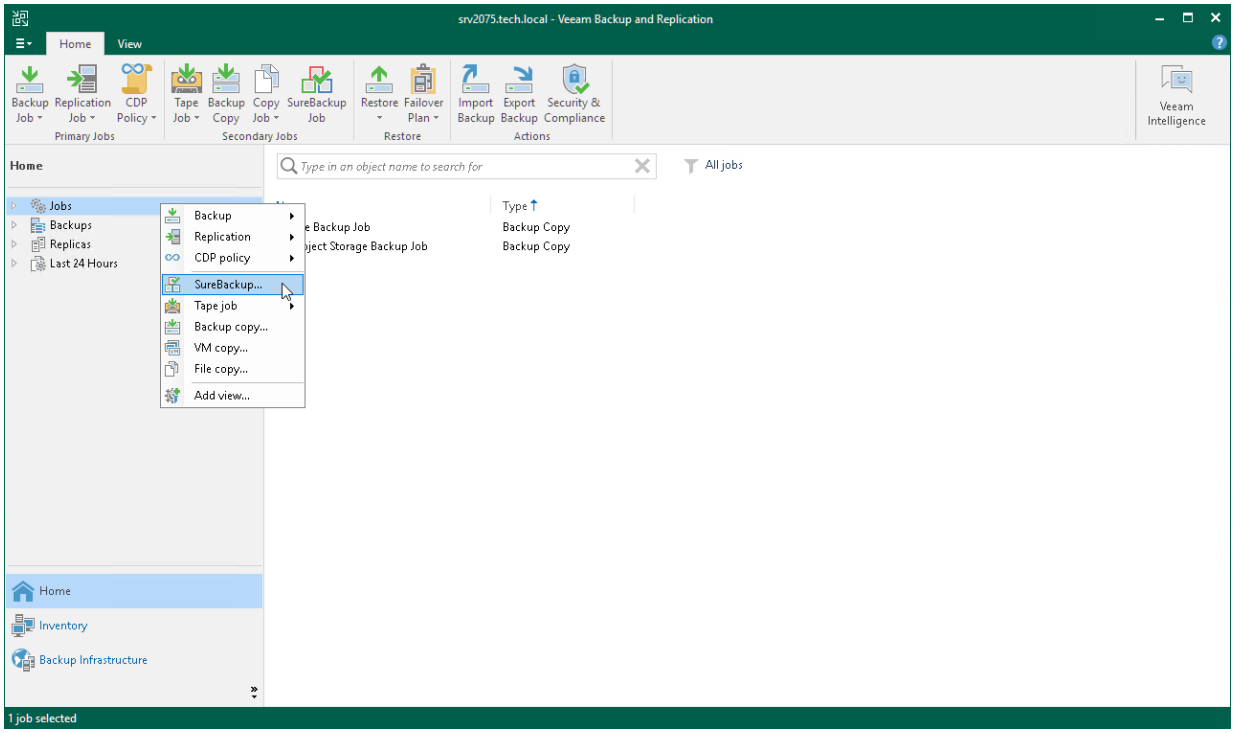

## Step 2. Specify Job Name, Description and Verification Mode

At the **Name** step of the wizard, specify a name, description and verification mode for the SureBackup job.

1. In the **Name** field, enter a name for the SureBackup job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the job, date and time when the job was created.
3. Select the verification mode, **Full recoverability testing** or **Backup verification and content scan only**.

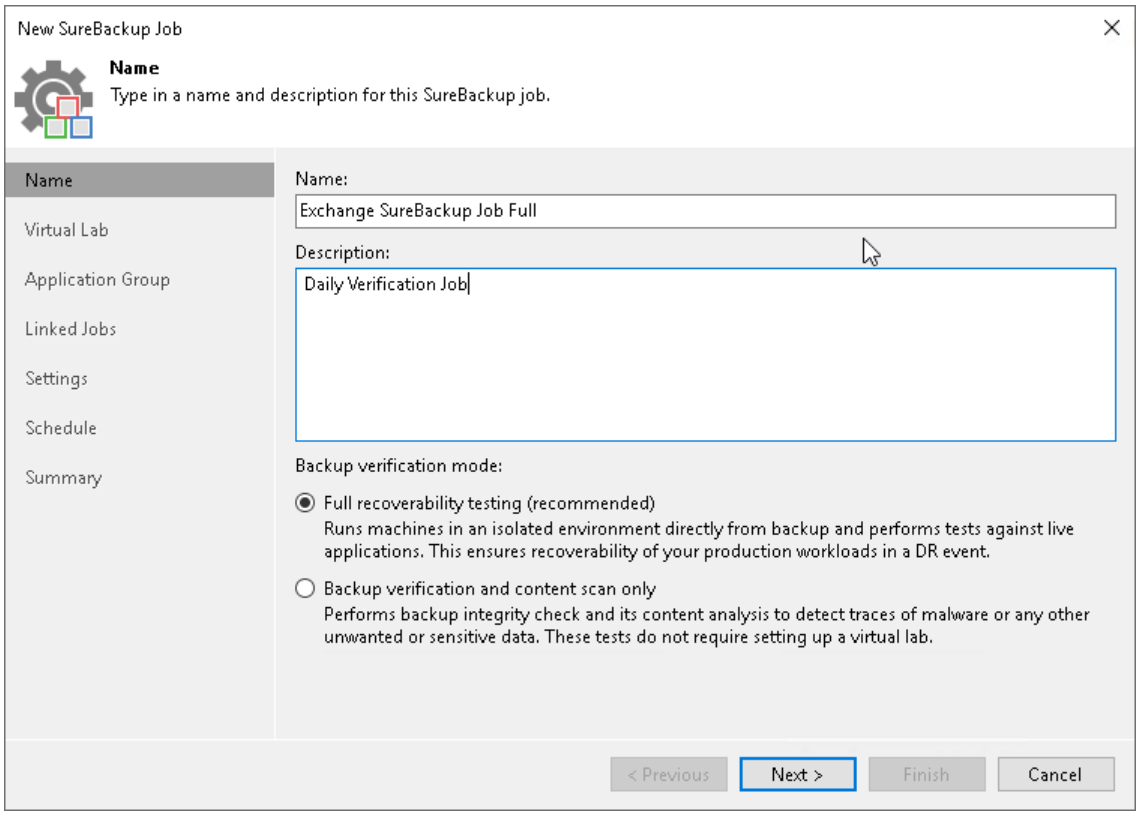

# Step 3. Select Virtual Lab

Selecting virtual lab is available only for SureBackup job working in **Full recoverability testing** mode.

At the **Virtual Lab** step of the wizard, select a virtual lab that you want to use for recovery verification.

If you do not have a virtual lab, click **Add New** and complete the New Virtual Lab wizard to create a new virtual lab.

1. From the **Virtual Lab** list, select a virtual lab. The list contains all virtual labs that are created or connected to the backup server.
2. In the **Virtual lab info** section, review information about the selected virtual lab.

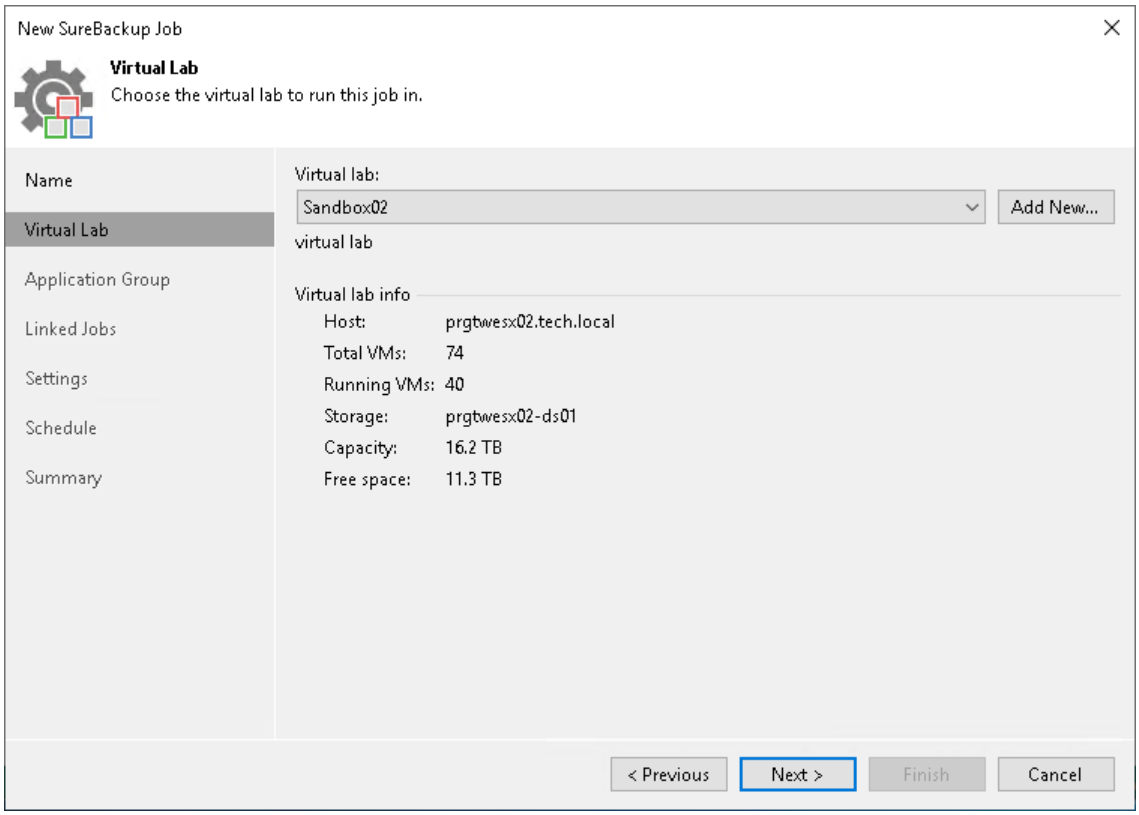

# Step 4. Select Application Group

Selecting application group is available only for SureBackup job working in **Full recoverability testing** mode.

At the **Application Group** step of the wizard, select an application group that you want to use for recovery verification.

You can select an application group or skip this step. If the application group is not selected, you must link at least one backup or replication job to the SureBackup job at the Linked Jobs step of the wizard. In this case, when the SureBackup job starts, Veeam Backup & Replication will only run machines from the linked job in the virtual lab and verify these machines.

To select an application group:

1. From the **Application group** list, select an application group. The list contains all application groups that are created on the backup server.
2. In the **Application group info** list, refer to the **Source Status** column to make sure that backups and replicas for VMs in the application group are created.
3. To leave machines from the application group running after the SureBackup job finishes, select the **Keep the application group running after the job completes** check box. With this option enabled, the lab will not be powered off when the SureBackup job completes, and you will be able to perform application item-level restore (U-AIR) and manually test machines started in the virtual lab.

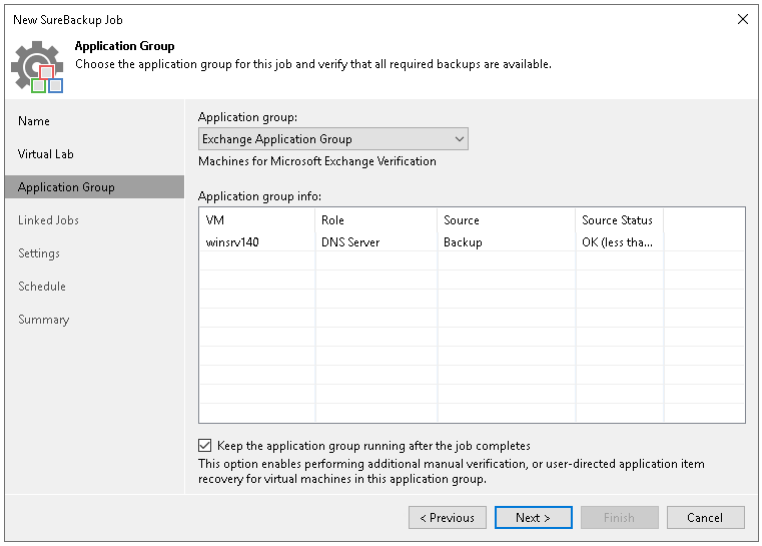

## Step 5. Link Backup or Replication Job

At the **Linked Jobs** step of the wizard, select backup or replication jobs with machines that you want to verify with the SureBackup job.

You can link a backup or replication job to the SureBackup job or skip this step. If you do not link a backup or replication job, in **Full recoverability testing** mode Veeam Backup & Replication will only start machines from the application group in the virtual lab and verify them. You have an option not to link a backup or replication job to the SureBackup job only if you have selected an application group at the Application Group step of the wizard.

# Linking Backup or Replication Job

To link a backup or replication job to the SureBackup job:

1. [For full recoverability testing mode] Select the **Test backups of the following backup jobs** check box.
2. Click **Add**.
3. In the **Select Jobs** window, select backup and replication jobs.
4. In the **Process simultaneously up to ... machines** field, specify the maximum number of machines (from a linked job) that can be processed at the same time. For example, if you choose to process three machines at the same time, Veeam Backup & Replication will create three streams — one stream per every verified machine. All three machines will be tested together. But if you choose to process one machine, then only one machine will be tested, and the next machine will be processed in the available stream. In **Full recoverability testing** mode, after all machines are verified, machines from the application group will be powered off or will be left running (if the **Keep the application group running after the job completes** option has been enabled at the Application Group step of the wizard).
5. Select **Process only randomly selected ... machines during each run** check box and specify the maximum number of machines you want to randomly test.

To exclude objects from the SureBackup job, perform the following steps:

1. Click **Exclusions**.
2. In the **Add Objects** window, select objects that you want to exclude.

   You can select the **Show full hierarchy** check box to display the hierarchy of all hosts added to Veeam Backup & Replication.
3. Click **Add**.
4. Click **OK**.

To remove a backup or replication job from the list, select it and click **Remove**.

> **NOTE**
>
> You can only exclude machines that are added to the backup or the replication job. Removed or deleted machines can no longer be excluded from the SureBackup job. Even if previously excluded, these machines will still be tested in the subsequent SureBackup job runs.

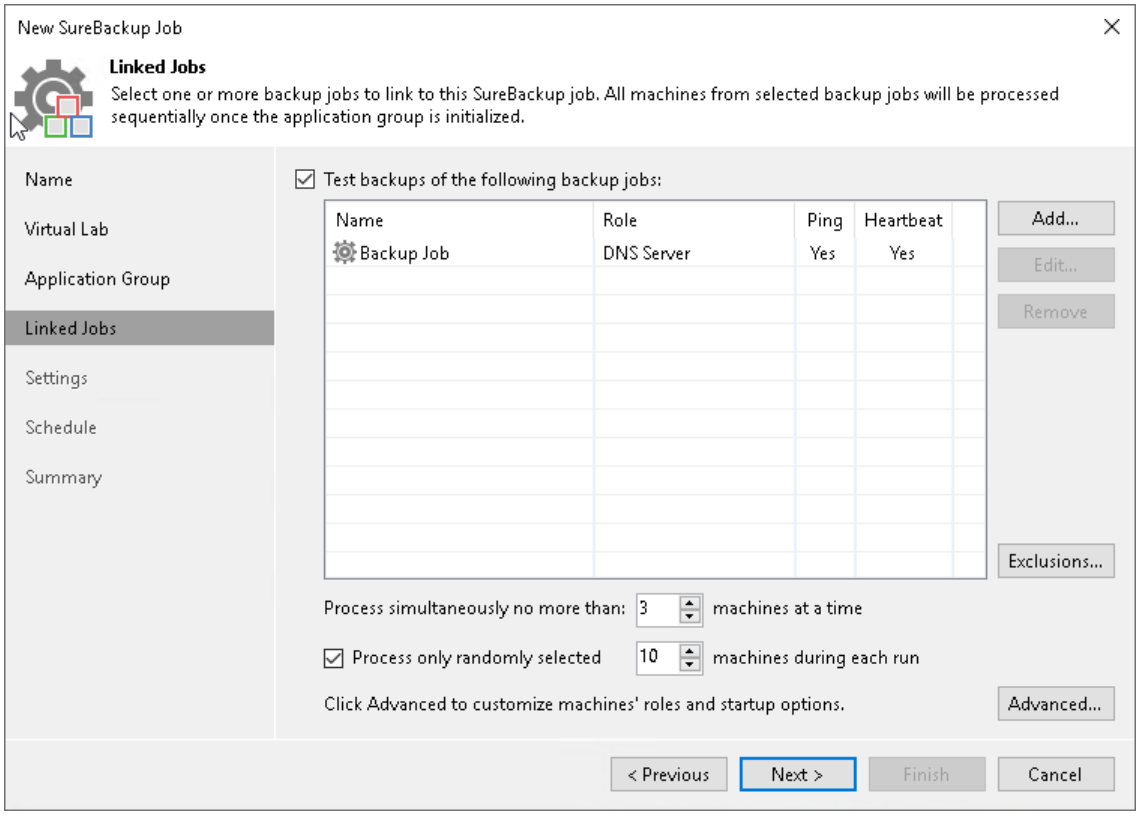
New SureBackup Job
Linked Jobs
Select one or more backup jobs to link to this SureBackup job. All machines from selected backup jobs will be processed sequentially once the application group is initialized.
Name
Virtual Lab
Application Group
Linked Jobs
Settings
Schedule
Summary
Test backups of the following backup jobs:
Name
Role
Ping
Heartbeat
Backup Job
DNS Server
Yes
Yes
Add...
Edit...
Remove
Exclusions...
Process simultaneously no more than: 3 machines at a time
Process only randomly selected 10 machines during each run
Click Advanced to customize machines' roles and startup options.
Advanced...
< Previous
Next >
Finish
Cancel

## Step 6. Specify Recovery Verification Options and Tests

Specifying recovery verification options and tests is available only for SureBackup job working in **Full recoverability testing** mode.

You must specify verification options for every machine from the jobs linked to the SureBackup job:

- Select a role that a machine performs.
- Configure startup settings.
- Select tests that must be performed for the machine.
- Specify credentials for running the verification script.

If all machines in the linked job perform the same role, you can specify startup options and test settings for the whole job in bulk. If machines have different roles, you can granularly specify startup options and test settings for every machine in the job.

- To specify startup options and select tests for the whole job, select a job in the list and click **Edit** on the right.
- To specify startup options and select tests for every machine in the job separately, select a job in the list and click **Advanced** on the right. Click **Add** and in the **Add Objects** window select a machine. Select the added machine in the list, click **Edit** and specify settings.

  If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.

> **IMPORTANT**
>
> If you specify startup options and tests individually for every machine, Veeam Backup & Replication will apply these options and tests only. Options and tests specified at the level of the SureBackup job will be ignored for this machine.

# Role Settings

On the **Role** tab, select the role that the machine performs. Veeam Backup & Replication offers the following predefined roles for machines:

- DNS Server
- Domain Controller (Authoritative Restore)

  In the Authoritative Restore mode, Veeam Backup & Replication starts a domain controller in the virtual lab and marks it as being authoritative to its replication partners. When other domain controllers (replication partners) are started in the virtual lab, they replicate data from the domain controller started in the Authoritative Restore mode.
- Domain Controller (Non-Authoritative Restore)

  In the Non-Authoritative Restore mode, Veeam Backup & Replication restores a domain controller in the virtual lab and marks it as being non-authoritative to its replication partners. Non-authoritative domain controllers then replicate data from a domain controller started in the Authoritative Restore mode.
- Global Catalog
- Mail Server
- SQL Server

- Veeam Backup for Microsoft 365 (machine on which Veeam Backup for Microsoft 365 is installed)
- Web Server

Machine roles are described in XML files stored in the *%ProgramFiles%|Veeam|Backup and Replication|Backup|SbRoles* folder. You can add your own roles. To do this, you need to create a new XML file and specify role and test scripts settings in it. For more information, see Creating XML files with Machine Roles Description.

After you select the necessary role, Veeam Backup & Replication will automatically configure startup options and assign predefined test scripts for the chosen role. You can use these settings or specify custom settings on the **Startup Options** and **Test Scripts** tabs.

To verify machines that perform roles other than those specified in the list, you will have to manually configure startup options and specify test scripts that must be run for these machines.

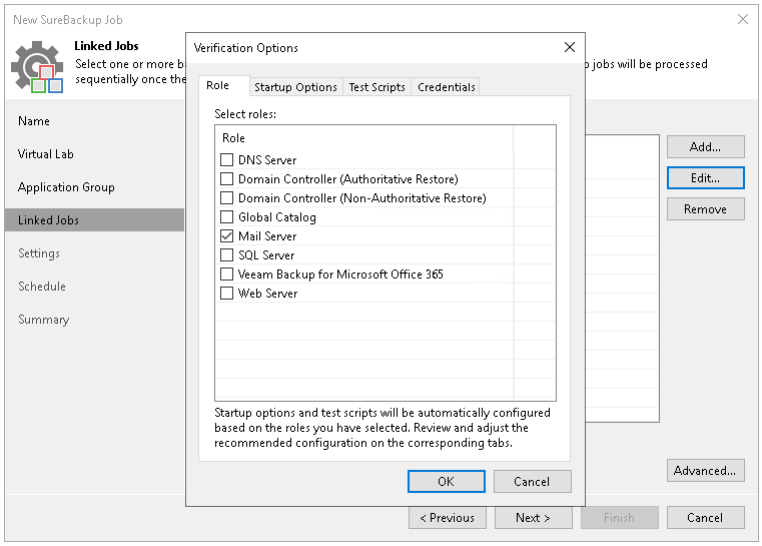


# VM Startup Settings

To configure VM startup settings:

1. In the **Verification Options** window, click the **Startup Options** tab.
2. In the **Memory** section, specify the amount of memory that you want to pre-allocate to the VM when this VM starts. The amount of pre-allocated memory is defined in percent. The percentage rate is calculated based on the system memory level available for the production VM. For example, if 1024 MB of RAM is allocated to the VM in the production environment and you specify 80% as a memory rate, 820 MB of RAM will be allocated to the verified VM on startup.

   Veeam Backup & Replication does not allow you to change VM CPU manually, it does this automatically. If the VM has more CPU than the host can provide, Veeam Backup & Replication scales down the CPU of the VM.
3. In the **Startup time** section, specify the allowed boot time for the VM and timeout to initialize applications on the VM.

   Be careful when specifying the **Maximum allowed boot time** value. Typically, a VM started by a SureBackup job requires more time to boot than a VM started in the production environment. If a VM fails to be initialized within the specified interval of time, the recovery verification process fails with the timeout error. If such error occurs, you need to increase the **Maximum allowed boot time** value and run the SureBackup job again.

4. In the **Boot verification** section, specify when the VM is considered to be successfully booted:
   - **VM heartbeat is present**. If you enable this option, Veeam Backup & Replication will perform a heartbeat test for the verified VM.
   - **VM responds to ping on any network interface**. If you enable this option, Veeam Backup & Replication will perform a ping test for the verified VM.
   - **Automatically disable Windows Firewall**. If you select this option, Veeam Backup & Replication will disable Windows Firewall for the the verified VM.

   If you enable both options, Veeam Backup & Replication will require that both tests are completed successfully: heartbeat test and ping test.

**IMPORTANT**

Do not pass sensitive information using script arguments in a user interface.

**NOTE**

Veeam Backup & Replication performs a heartbeat test only if a machine has VMware Tools installed. If VMware Tools are not installed, the machine will be started but the test will not be performed.

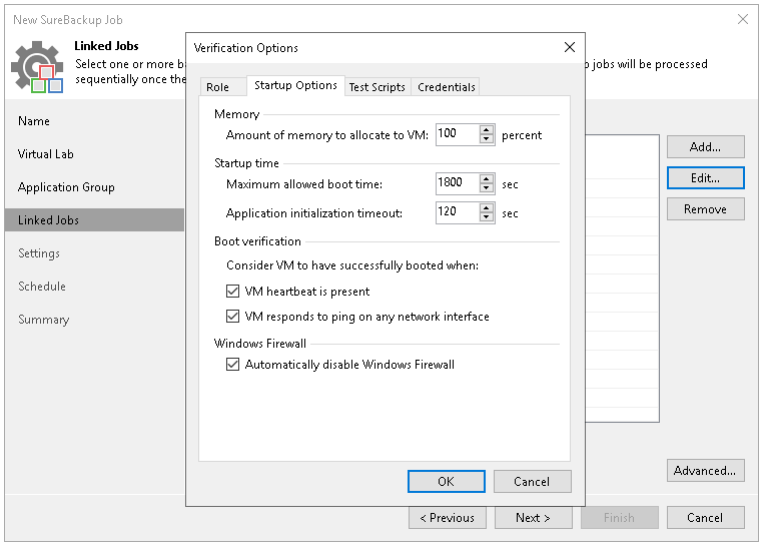


# Test Script Settings

When you select a machine role, Veeam Backup & Replication automatically assigns a predefined script that must be run to verify applications inside the machine. If you want to verify a machine that has some other role not listed on the **Role** tab, do the following:

1. In the **Verification Options** window, click the **Test Scripts** tab.
2. Click **Add**.
3. In the **Test Scripts** window, select **Use the following test script**.
4. In the **Name** field, specify a name for the script.
5. In the **Path** field, define a path to an executable script file that must be run to verify the machine. You can do one of the following:
   - If you have your own custom script, define a path to it in the **Path** field.

- If you do not have a custom script, you can use Veeam standard utility, `Veeam.Backup.ConnectionTester.exe`, that probes application communication ports. The utility is located in the installation folder of Veeam Backup & Replication: `%ProgramFiles%\Veeam\Backup and Replication\Backup\Veeam.Backup.ConnectionTester.exe`. Specify this path in the **Path** field.

6. In the **Arguments** field, specify an IP address of the verified machine and the port that you want to probe (if the selected test probes the port). You can use the *%vm_ip%* variable to define the machine IP address or the *%vm_fqdn%* variable to define the machine FQDN.

   For Microsoft SQL Server, you can also specify a path to the log file in the `%log_path%` argument. For more information, see Backup Recovery Verification Tests.

7. Click **OK** to add the configured test.

To edit test settings, select the test in the list and click **Edit**. To delete a test, select it in the list and click **Remove**.

If a VM performs several roles and runs a number of applications, you can add several verification scripts to verify work of these applications. It is recommended that you specify the maximum startup timeout value and allocate the greatest amount of memory for such VMs.

> **NOTE**
>
> Exit codes for SureBackup custom PowerShell scripts will return 0 or 1 depending on whether the script executed without exceptions. To utilize different error codes, set the `$?` variable value within your script. For more information, see Microsoft Docs.

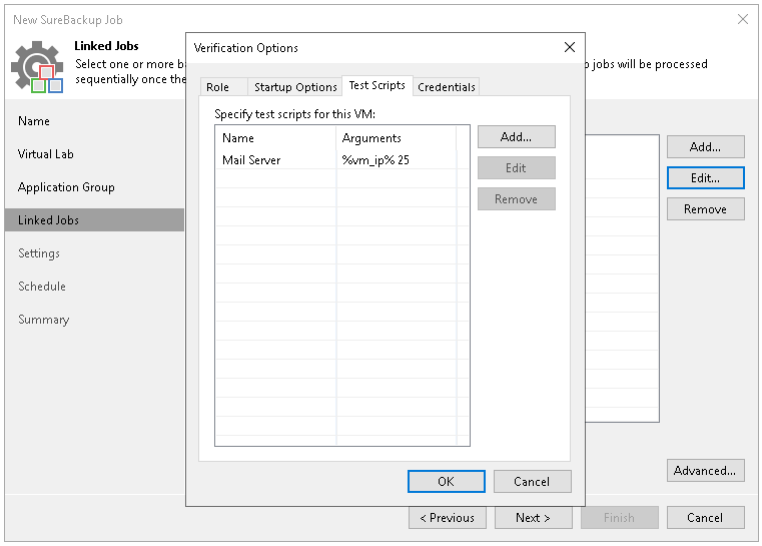


# Credentials Settings

In the **Credentials** tab, specify credentials to authenticate in the machine where you need to run the script.

1. Click the **Credentials** tab.

2. From the **Credentials** list, select credentials for the account under which you want to run the script.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add the credentials. For more information, see Credentials Manager.

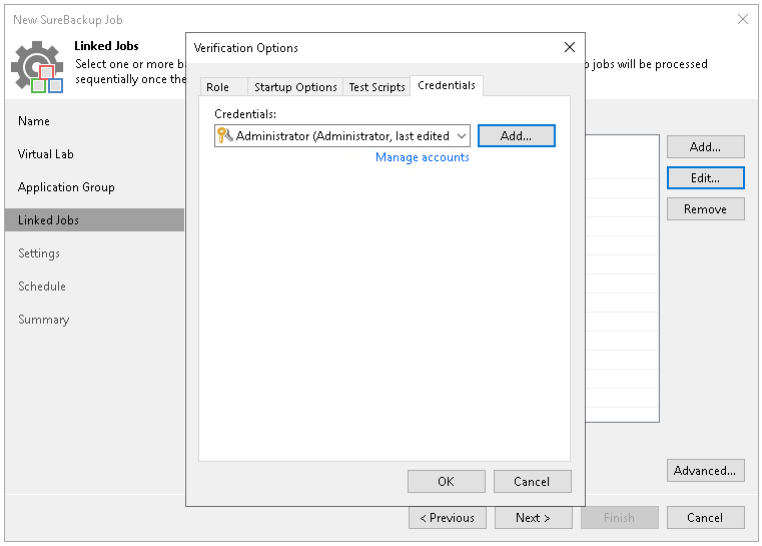

## Step 7. Specify Additional Job Settings

On the **Settings** step of the wizard, specify additional settings for the SureBackup job.

# Specifying Additional Job Settings

1. If you want Veeam Backup & Replication to scan VM data with antivirus software, select the **Scan backup content with Veeam Threat Hunter** or **Scan backup content with your existing software** check box. Name of the check box depends on the settings specified in Signature Detection. For more information, see Secure Restore.
   - If you want the antivirus software to continue scanning VM data after the first malware is found, select the **Continue scanning remaining files after the first occurrence** check box. For information on how to view results of the malware scan, see Viewing Recovery Verification Job Statistics.
   - [For full recoverability testing mode] If you do not want to scan VMs from the application group, select the **Skip application group machines from malware scan** check box. In this case, the antivirus will only scan VMs from linked jobs.

   Veeam Backup & Replication scans VM data with antivirus before running verification tests. Consider that the SureBackup job may take considerable time to complete if you are verifying backups of large sized VMs.
2. If you want Veeam Backup & Replication to scan VM data with YARA rule, select the **Scan backup content with the following YARA rule** check box and select the YARA rule from the drop-down list. For YARA rules to appear, they should be placed in the folder by the following path: `%Program Files%\Veeam\Backup and Replication\Backup\YaraRules`. Veeam Backup & Replication accepts only .yar and .yara extension. Veeam Backup & Replication can not scan VM replicas with YARA rule.

3. [For backups only] If you want to validate the backup file with a CRC check and make sure that the file is not corrupted, select the **Perform backup integrity check** check box. If you selected **Full recoverability testing** mode at the **Name** step of the wizard, you can optionally exclude VMs being a part of the application group from this test. To do this, select the **Skip integrity check for application group machines' backup** check box. For more information, see Backup Recovery Verification Tests.

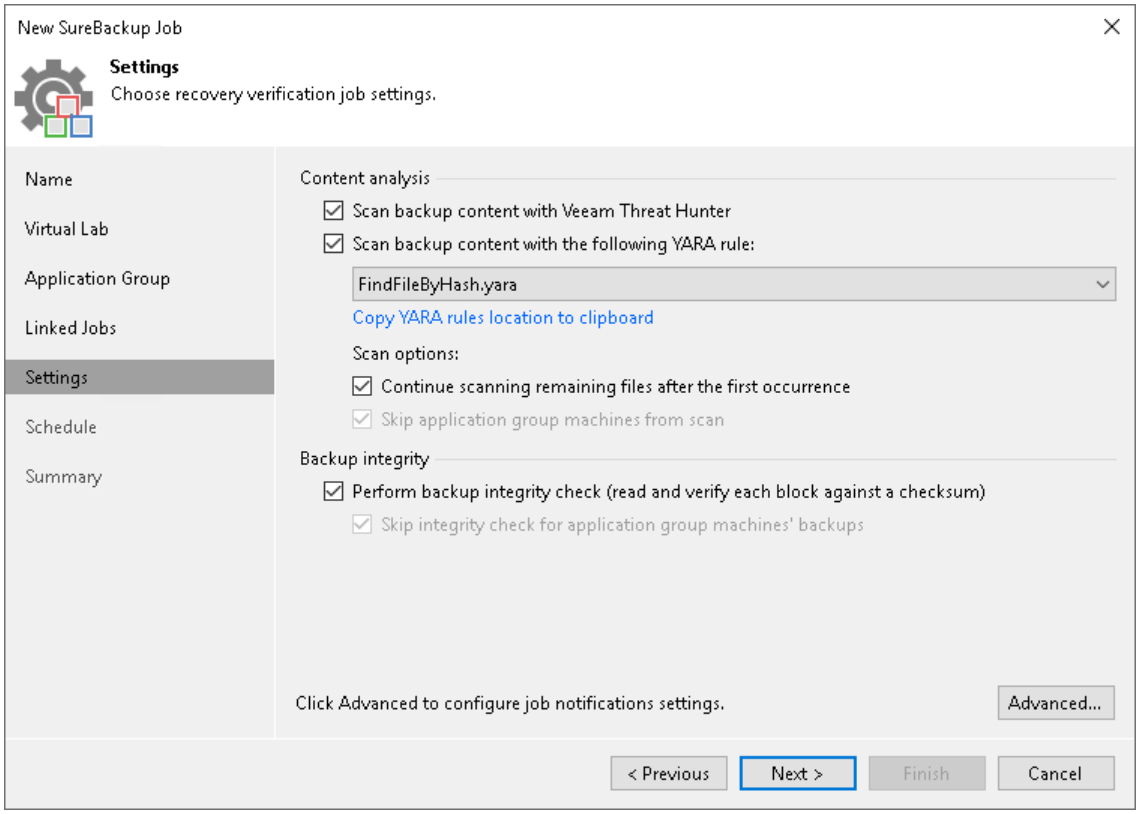


> **NOTE**
>
> [For full recoverability testing mode only] If you enable the **Keep the application group running after the job completes** option at the **Application Group** step of the wizard, the **Skip integrity check for application group machines' backup** option will be automatically enabled.

To configure notifications settings, click **Advanced** and specify notification settings:

1. If you want to receive SNMP notifications, select the **Send SNMP notifications for this job** check box.

   SNMP notifications will be sent only if you configure global SNMP settings in Veeam Backup & Replication and on recipient's computer. For more information, see Specifying SNMP Settings.

2. If you want to receive notifications by email, select the **Send email notifications to the following recipients** check box. In the field under the check box, specify the recipient email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent only if you configure global email notification settings in Veeam Backup & Replication. For more information, see Specifying Email Notification Settings.

3. If you want to receive notifications only if the job finishes with specific states, select the required check boxes. For example, **Notify on warning** and **Notify on error**. If you want to receive all notifications, select all the check boxes.

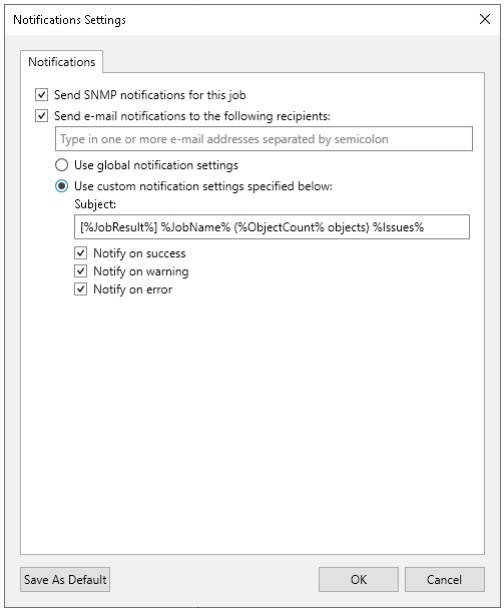

## Step 8. Specify Job Schedule

At the **Schedule** step of the wizard, select to manually run the SureBackup job or schedule the job at specific time, for example, after the backup or replication job completes. Keep in mind that SureBackup job will be stopped during synthetic operations (synthetic full backup, backup files merge and transformation) against the source backup chain, health check and replication session even if SureBackup job is scheduled to run.

1. To define a job schedule, select the **Run the job automatically** check box. If this check box is not selected, you will have to manually start the job to perform recovery verification.
2. Select the required schedule option:
   - **Daily at this time**. You can specify the time and select the days option from the drop-down list on which the SureBackup job will run: **Everyday**, **On weekdays** or **On these days**. If you select **On these days** option, click **Days** to specify them.
   - **Monthly at this time**. You can specify the time and select the day options from the drop-down lists on which the SureBackup job will run. Click **Months** to specify months on which the SureBackup job will run.
   - **Periodically every**. You can select the time options from the drop-down list or click **Schedule** to select the desired time area. Use the **Permitted** and the **Denied** options to mark the selected time segments. Use the **Start time within an hour** option to specify minutes. To run the job continuously, select **Continuously** from the drop-down list. A new job session will start as soon as the previous job session finishes.
   - **After this job**. If you choose this option, select the job from the drop-down list after which the SureBackup job will run. Typically, a SureBackup job should run after the linked backup or replication job completes. In this case, the SureBackup job will verify the VM backup or VM replica created by the source backup or replication job.

     To create a chain of jobs, you must define the time schedule for the first job in the chain. For the rest of the jobs in the chain, at the **Schedule** step of the wizard, select the **After this job** option and choose the preceding job from the list.

> **NOTE**
>
> When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

3. In some cases, the linked backup or replication job may not complete until the SureBackup job starts. If Veeam Backup & Replication finds out that the linked job is still running, the SureBackup job will fail to start. To overcome this situation, select the **If some linked backup jobs are still running, wait up to <N> minutes** check box and specify the necessary time period in the field on the right. If the linked job is still running, Veeam Backup & Replication will wait for the defined period of time and check the linked job after this period elapses.
   - If the linked job is finished within the specified period, the SureBackup job will start.

- If the linked job is still running, the SureBackup job will not start.

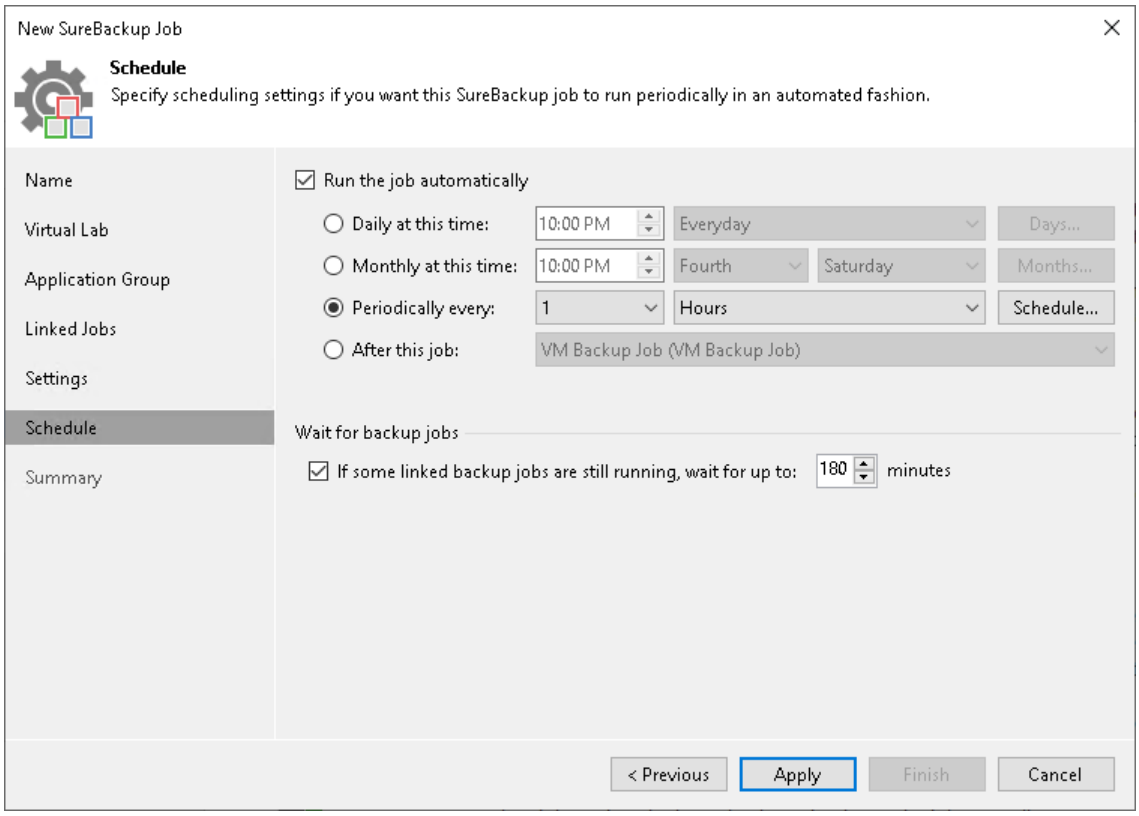

## Step 9. Review Job Summary and Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of SureBackup job configuration.

1. Review details of the SureBackup job.
2. If you want to start the job right after you finish working with the wizard, select the **Run the job when I click Finish** check box.
3. Click **Finish** to save the job settings and close the wizard.

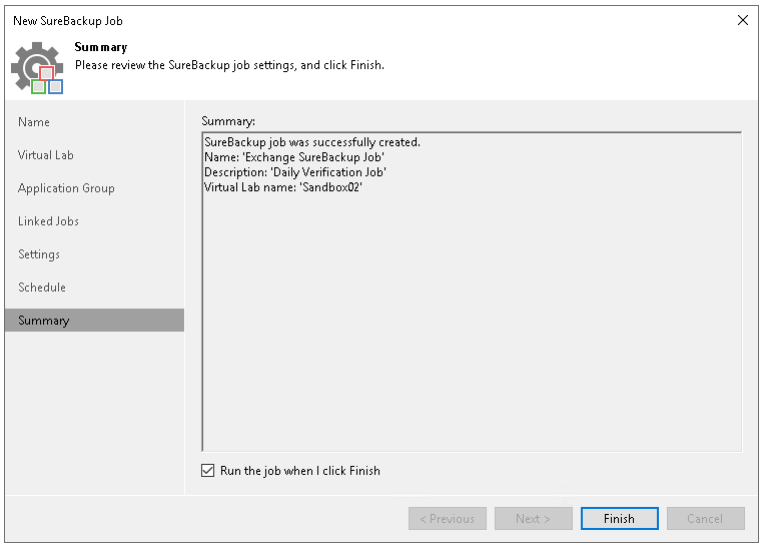


# Starting and Stopping SureBackup Job

You can instruct the SureBackup job to verify the latest restore point of a VM backup or VM replica or select a specific restore point to which the VM from the backup or VM replica must be started.

To start a VM from the latest restore point:

1. Open the **Home** view.
2. In the inventory pane, click **SureBackup** under **Jobs**.
3. In the working area, select the SureBackup job and click **Start** on the ribbon. You can also right-click the SureBackup job and select **Start**. Veeam Backup & Replication will start, verify and perform necessary tests for VMs from the latest restore point.

To start VMs from a specific point in time:

1. Open the **Home** view.
2. In the inventory pane, select **SureBackup** under **Jobs**.
3. In the working area, select the SureBackup job and click **Start to** on the ribbon. You can also right-click the SureBackup job and select **Start to**.
4. In the **Restore Point** window, select an approximate date of the restore point creation.
   Veeam Backup & Replication will pick the most recent restore point prior to this time and start, verify and perform necessary tests for VMs from this restore point.

5. If SureBackup job is configured to process VM replicas added to a CDP policy, specify the restore point type:
   - **Application-consistent**. SureBackup job will pick the most recent long-term application-consistent restore point prior to the time you specified.
   - **Crash-consistent**. SureBackup job will pick the most recent long-term crash-consistent restore point prior to the time you specified.

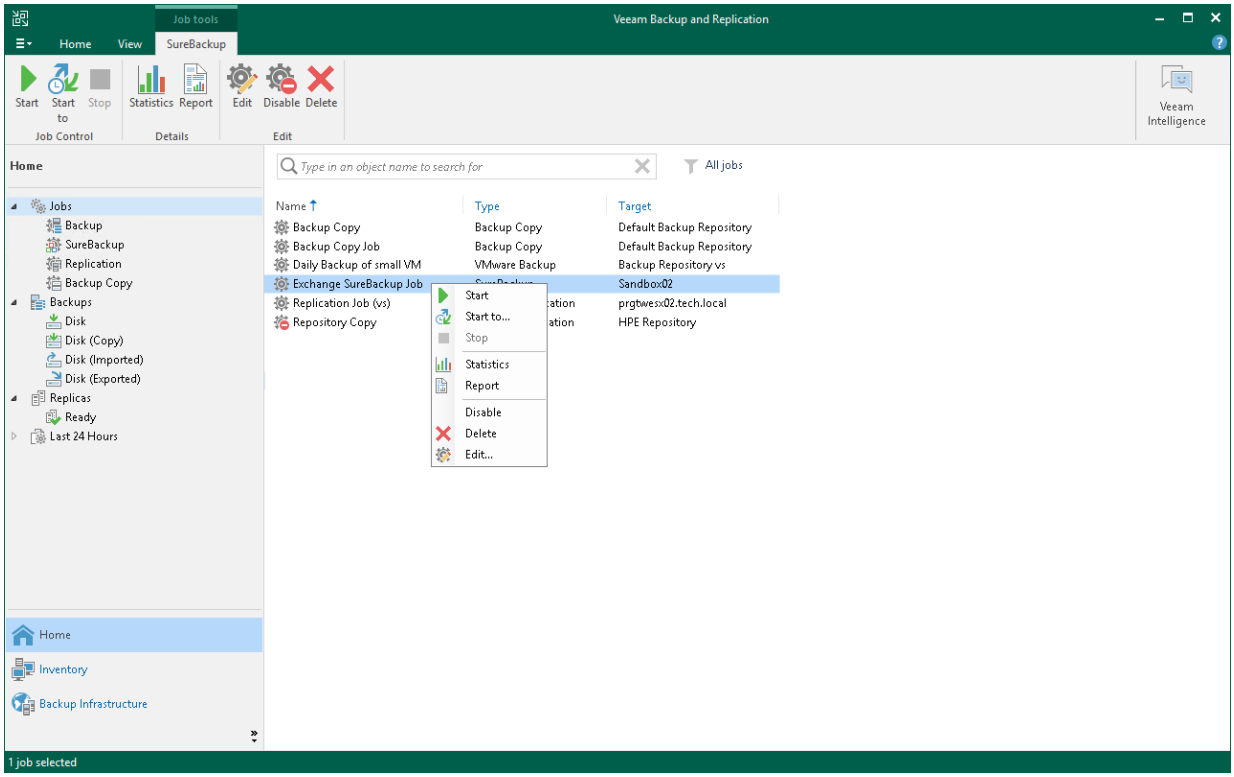


To stop a running SureBackup job session:

1. Open the **Home** view.
2. In the inventory pane, select **SureBackup** under **Jobs**.

3. In the working area, select the SureBackup job and click **Stop** on the ribbon. You can also right-click the SureBackup job and select **Stop**.

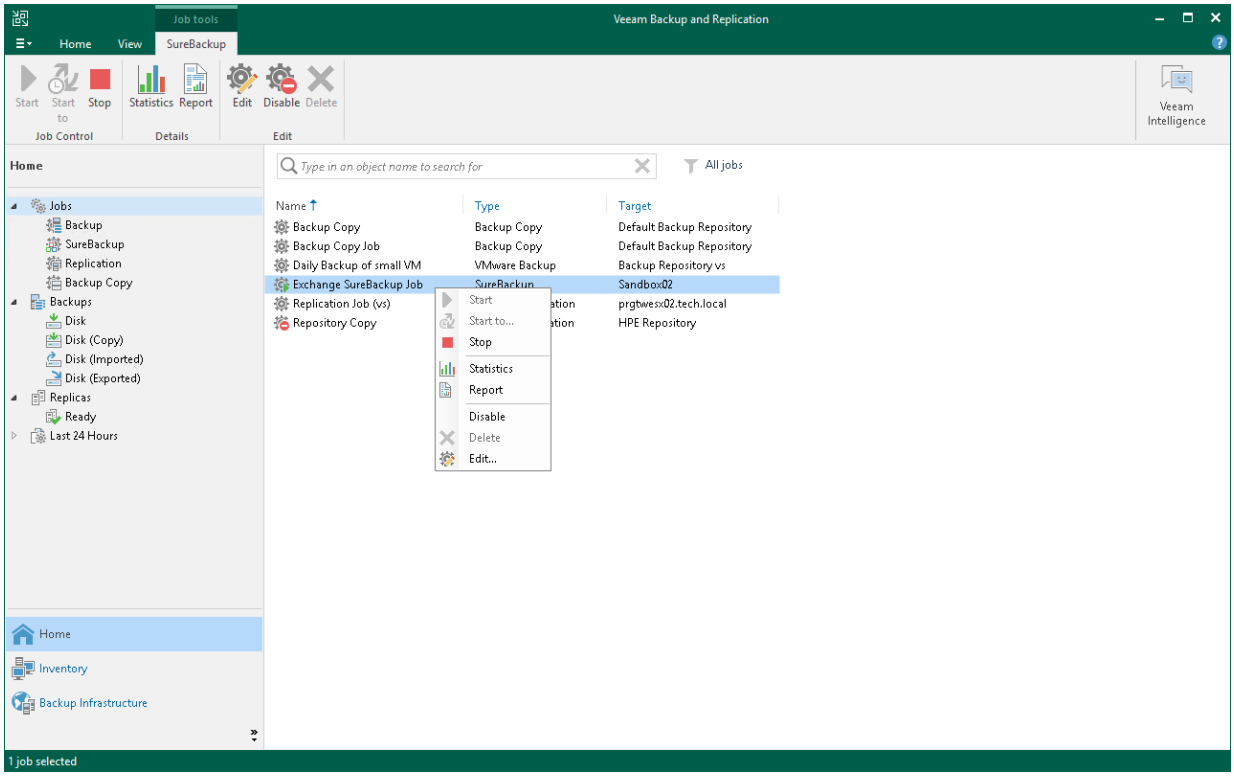


# Viewing Recovery Verification Job Statistics

You can monitor how tests for verified VMs are performed while a recovery verification job is running.

To see the status of VM tests:

1. Open the **Home** view.
2. In the inventory pane, select **SureBackup** under **Jobs**.
3. In the working area, right-click a recovery verification job and select **Statistics**. You can also double-click the job in the list.

The job session window displays statistics for all VMs that are started during the SureBackup job: VMs from the application group in the specified order and VMs from linked jobs. For your convenience, these VMs are marked with different icons.

After the verified VM is powered on, its name is displayed as a hyperlink. You can click the link to open the VM console to see what is happening inside the VM or perform manual testing.

If some VM fails to be verified automatically, you can start it manually when this VM is powered off. To start a VM, right-click the VM in the list and select Start. If the application group has already been powered off by that time, it will be started again. After that, you can open the VMware Remote Console (VMRC) and perform verification and testing manually.

> **IMPORTANT**
>
> VMware Remote Console is not included as part of Veeam Backup & Replication installation and must be installed separately. For details, see Install the VMware Remote Console Application.

If you enabled content analysis at the Settings step of the SureBackup job wizard, you can view the detailed logging of the scan process. To view logs, click the **Scan Log** button that will appear at the bottom of the job session window after the scan is complete.

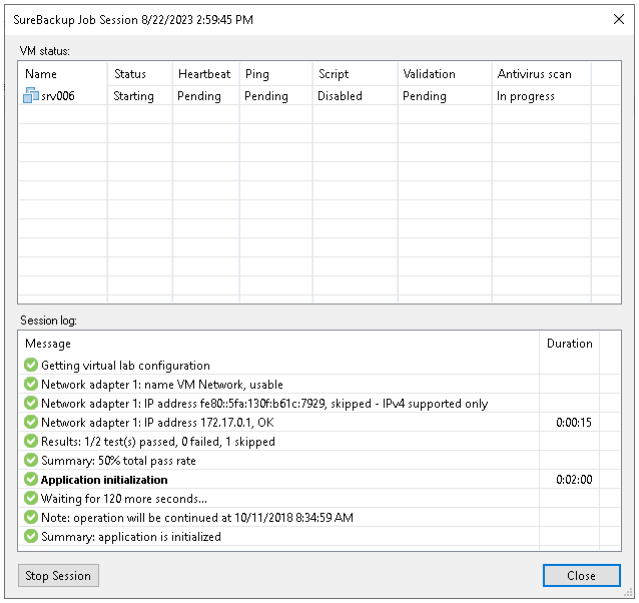


# Creating SureBackup Session Reports

You can generate HTML reports with statistics on the SureBackup job. A report contains detailed data on job sessions: job status, start and end time, details of the session performance, status of verified VMs and test results. You can generate a report for the whole SureBackup job or a specific job session/sessions.

The SureBackup job report contains data on all sessions initiated for a specific job. To generate a SureBackup job report:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the SureBackup job and click **Report** on the ribbon. You can also right-click the SureBackup job and select **Report**.

The session report contains data on a single job session. To generate a session report:

1. Open the **History** view.
2. In the inventory pane, select **Jobs**.

3. In the working area, select the session and click **Report** on the ribbon. You can also right-click the session and select **Report**.

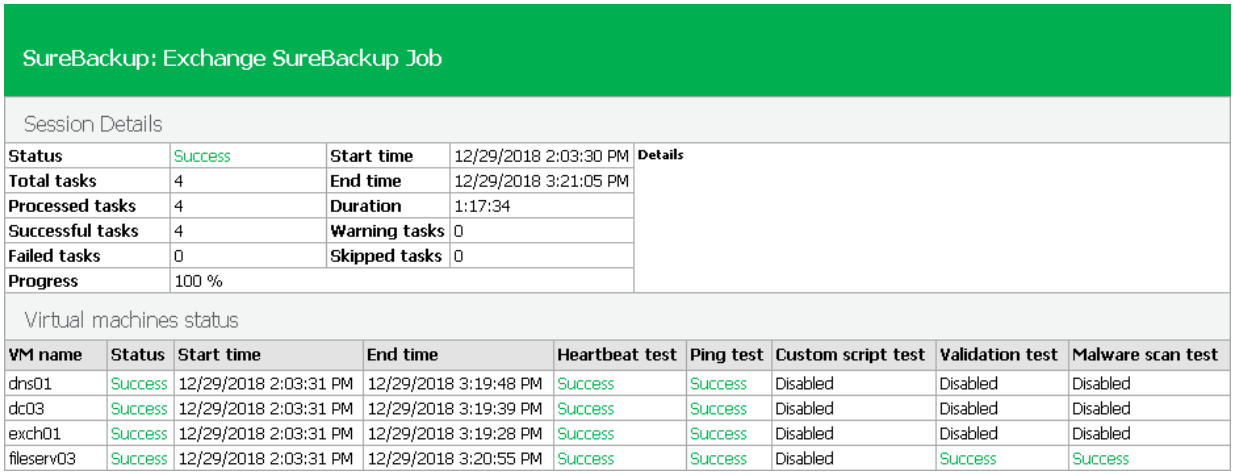

SureBackup: Exchange SureBackup Job

Session Details

| | | | | |
|---|---|---|---|---|
| **Status** | Success | **Start time** | 12/29/2018 2:03:30 PM | **Details** |
| **Total tasks** | 4 | **End time** | 12/29/2018 3:21:05 PM | |
| **Processed tasks** | 4 | **Duration** | 1:17:34 | |
| **Successful tasks** | 4 | **Warning tasks** | 0 | |
| **Failed tasks** | 0 | **Skipped tasks** | 0 | |
| **Progress** | 100 % | | | |

Virtual machines status

| VM name | Status | Start time | End time | Heartbeat test | Ping test | Custom script test | Validation test | Malware scan test |
|---|---|---|---|---|---|---|---|---|
| dns01 | Success | 12/29/2018 2:03:31 PM | 12/29/2018 3:19:48 PM | Success | Success | Disabled | Disabled | Disabled |
| dc03 | Success | 12/29/2018 2:03:31 PM | 12/29/2018 3:19:39 PM | Success | Success | Disabled | Disabled | Disabled |
| exch01 | Success | 12/29/2018 2:03:31 PM | 12/29/2018 3:19:28 PM | Success | Success | Disabled | Disabled | Disabled |
| fileserv03 | Success | 12/29/2018 2:03:31 PM | 12/29/2018 3:20:55 PM | Success | Success | Disabled | Success | Success |

# XML Files with Machine Roles Description

Machine roles that you can assign to verified machines and machines from the application group are described in XML files. These XML files are stored in the `%ProgramFiles%\Veeam\Backup and Replication\Backup\SbRoles` folder on the backup server.

To add a new role, you must create a new XML file and save it to the SbRoles subfolder on the backup server. Do not save the XML file on the machine where the Veeam Backup & Replication console is installed — this will not affect the list of roles in Veeam Backup & Replication.

XML files describing machine roles have the following structure:

```
<SbRoleOptions>
 <Role>
   <SbRole>
     <Id>4CDC7CC4-A906-4de2-979B-E5F74C44832F</Id>
     <Name>Web Server</Name>
   </SbRole>
 </Role>
 <Options>
   <SbVerificationOptions>
     <ActualMemoryPercent>100</ActualMemoryPercent>
     <MaxBootTimeoutSec>300</MaxBootTimeoutSec>
     <AppInitDelaySec>120</AppInitDelaySec>
     <TestScripts>
       <TestScripts>
         <TestScript>
           <Name>Web Server</Name>
           <Type>Predefined</Type>
           <TestScriptFilePath>Veeam.Backup.ConnectionTester.exe</TestScriptFi
lePath>
           <Arguments>%vm_ip% 80</Arguments>
         </TestScript>
       </TestScripts>
     </TestScripts>
     <HeartbeatEnabled>True</HeartbeatEnabled>
     <DisableWinFirewall>True</DisableWinFirewall>
     <PingEnabled>True</PingEnabled>
   </SbVerificationOptions>
 </Options>
</SbRoleOptions>
```

The XML file with the role description contains the following tags and parameters:

| Tag | Required/ Optional | Description |
|---|---|---|
| **<SbRoleOptions>** | Required | Encapsulates the machine role file. |
| **<Role>** | Required | Parent tag for a role assigned to a machine. *<SbRole>*, *<Id> and <Name>* are children of this tag. |
| **<SbRole>** | Required | Encapsulates basic information for a machine role: ID and name. |

| Tag | Required/ Optional | Description |
|---|---|---|
| **<Id>** | Required | Unique identifier of a machine role. |
| **<Name>** | Required | Name of a machine role. The machine role name is displayed in the roles list on the **Role** tab. |
| **<Options>** | Required | Parent tag for startup and test script options to be used for the defined role. *<SbVerificationOptions>*, *<ActualMemoryPercent>*, *<MaxBootTimeoutSec>*, *<AppInitDelaySec>*, *<TestScripts>*, *<Name>*, *<Type>*, *<TestScriptFilePath>*, *<Arguments>*, *<HeartbeatEnabled>*, *<PingEnabled>* are children of this tag. |
| **<SbVerificationOptions>** | Required | Encapsulates options data for a machine role. |
| **<ActualMemoryPercent>** | Optional | Percent of the original memory level that must be pre-allocated to a verified machine on the system boot. |
| **<MaxBootTimeoutSec>** | Optional | Maximum allowed time to boot a machine. |
| **<AppInitDelaySec>** | Optional | Duration of time for which Veeam Backup & Replication must wait after the machine is successfully booted in the virtual lab. After this time elapses, Veeam Backup & Replication will run test scripts. Time is specified in seconds. |
| **<TestScripts>** | Optional | Encapsulates test script data for a machine role. |
| **<Name>** | Optional | Name of a machine role. The machine role name is displayed on the **Test Scripts** tab. |
| **<Type>** | Optional | Type of the test script: *Predefined* or *Custom*. |
| **<TestScriptFilePath>** | Optional | Path to an executable file of the test script to be performed. The path can be absolute or relative. |
| **<Arguments>** | Optional | Arguments to be passed to the script. You can use the following variables:<br>• *%vm_ip%* — IP address of a verified machine.<br>or<br>%vm_fqdn% — a fully qualified domain name of a verified machine.<br>• *%log_path%* — path to a log file to which verification results are stored. |

| Tag | Required/ Optional | Description |
| --- | --- | --- |
| **<HeartbeatEnabled>** | Required | Must a heartbeat test be enabled for this machine role: *True* or *False*. |
| **<DisableWinFirewall>** | Required | Must a firewall be disabled for this machine role: *True* or *False*. |
| **<PingEnabled>** | Required | Must a ping test be enabled for this machine role: *True* or *False*. |

# Manual Recovery Verification

Beside automatic recovery verification, you can perform manual verification of machine backups. Manual verification can be performed with all editions of Veeam Backup & Replication.

## Boot Test

To perform a machine boot test, perform Instant Recovery for the verified machine. Power on the machine but do not connect the machine to the production network to avoid conflicts with the original machine.

## Application Test

To perform an application test:

1. Create an isolated network.
2. Use the **Instant Recovery** wizard to restore the verified machine. At the **Ready to Apply** step of the wizard, clear the **Connect VM to network** check box.
3. When the machine is started, connect it to the isolated network.

The same procedure must be performed for all machines that run applications on which the verified machine is dependent such as domain controller and DNS. All machines must be connected to the same isolated network and started in the correct order: for example, DNS > domain controller > verified machine.

# SureReplica

To guarantee recoverability of your data, Veeam Backup & Replication complements the SureBackup recovery verification technology with SureReplica.

SureReplica is in many respects similar to the SureBackup recovery verification. It lets you validate your disaster recovery environment without impacting the production infrastructure. You can automatically verify every created restore point of every VM replica and ensure that they are functioning as expected.

The SureReplica technology is not limited only to VM replica verification. Just like SureBackup, it provides the following capabilities:

- SureReplica: automated VM replicas verification, including storage snapshot replicas.
- On-Demand Sandbox: an isolated environment for testing VM replicas, training and troubleshooting.
- U-AIR: recovery of individual items from applications running on VM replicas.

# How SureReplica Works

SureReplica is Veeam's technology that lets you test a VM replica for recoverability. To ensure that the VM replica is functioning properly, Veeam Backup & Replication performs its "live" verification.
Veeam Backup & Replication automatically boots the VM replica from the necessary restore point in the isolated environment, performs tests against the VM replica, powers it off and creates a report on the VM replica state.

SureReplica supports regular VM replicas and VM replicas added to a CDP policy. SureReplica verification does not prevent CDP policy from running.

The SureReplica technology does not require the vPower engine. A VM replica is essentially an exact copy of a VM with a set of restore points. The VM replica data is stored in the raw decompressed format native to VMware. Therefore, to start a VM replica in the virtual lab, you do not need to present its data through the vPower NFS datastore to the ESXi host. Veeam Backup & Replication re-configures the VM replica settings for recovery verification, connects the VM replica to the isolated virtual lab and powers it on.

As there is no need to publish the VM from the backup file, the SureReplica processing is typically faster than SureBackup. Subsequently, the U-AIR and On-Demand Sandbox operations are faster, too.

During VM replica verification, Veeam Backup & Replication performs the following actions:

1. Veeam Backup & Replication triggers a VMware snapshot for a VM replica. The snapshot protects the VM replica from changes while it is running. All changes made to the VM replica are written to the delta file.

   During verification of a VM replica added to a CDP policy, Veeam Backup & Replication does not trigger a VMware snapshot of this VM replica.

2. If antivirus software scan is enabled, Veeam Backup & Replication performs antivirus scan and use backup server as the mount server during this operation.

3. Veeam Backup & Replication starts the VM replica in the virtual lab.

4. Veeam Backup & Replication performs tests against the verified VM replica.

5. When the verification process is over, Veeam Backup & Replication removes the delta file of the VM replica snapshot, powers off the VM replica and creates a report on its state. The report is sent to the backup administrator by email.

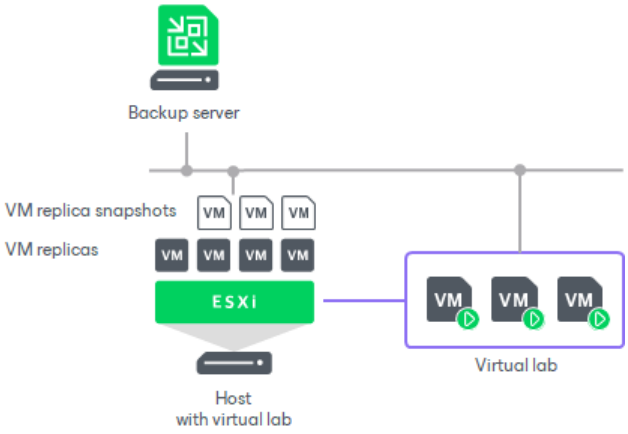

> **NOTE**
>
> Veeam Backup & Replication verifies only VM replicas in the *Ready* state. If a VM replica is in the *Failover* or *Failback* state, the verification process fails.
> When Veeam Backup & Replication verifies the VM replica, it puts the VM replica to the *SureBackup* state. You cannot perform failback and failover operations for a VM replica in the *SureBackup* state until recovery verification or the U-AIR process is over and the VM replica returns to the *Ready* state.

To perform VM replica verification, you need to create the following objects:

1. Application group. During recovery verification, the VM replica is not started alone: it is started together with VMs on which the VM replica is dependent. Starting a VM replica in conjunction with other VMs enables full functionality of applications running inside the VM replica and lets you run these applications just like in the production environment.
2. Virtual lab. The virtual lab is the isolated virtual environment in which the VM replica and VMs from the application group are started and tested.
3. SureBackup job. The SureBackup job is a task for VM replica verification process. You can run the SureBackup job manually or schedule it to run automatically by schedule.

# Replica Recovery Verification Tests

To verify a VM replica, Veeam Backup & Replication performs the same tests as for VM backup verification, except backup validation test. You can run predefined tests or perform your own tests against VMs. The predefined tests include the following ones:

- Heartbeat test
- Ping test
- Application test

For more information, see Backup Recovery Verification Tests.

# Application Group

You can add to the same application groups both VMs from backups and VMs from replicas. Keep in mind that all VMs from the application group must have at least one valid restore point created by the time the SureBackup job starts.

For more information, see Application Group.

# Virtual Lab Configuration

Veeam Backup & Replication offers three types of the virtual lab configuration for VM replica verification:

- Basic single-host virtual lab
- Advanced single-host virtual lab
- Advanced multi-host virtual lab

## Basic Single-Host Virtual Labs

The basic single-host virtual lab configuration can be used if your disaster recovery (DR) site is configured in the following way:

- All VM replicas that you want to verify are registered on the same ESXi host.
- All VM replicas that you want to verify are connected to the same network.

> **IMPORTANT**
>
> For this configuration type, the virtual lab must be created on the same ESXi host where VMs replicas are located. If you create the virtual lab on some other ESXi host, the SureBackup job will fail.

For the basic single-host virtual lab, Veeam Backup & Replication creates one virtual network that is mapped to the production network. Additionally, Veeam Backup & Replication automatically adds a number of new VMware objects on the ESXi host where the virtual lab is created:

- A resource pool
- A VM folder
- A standard vSwitch

The vSwitch is only used by the VMs started in the virtual lab. There is no routing outside the virtual lab to other networks.

Veeam Backup & Replication automatically configures all settings for the basic single-host virtual lab. The proxy appliance is also created and configured automatically and placed to the virtual lab folder and resource pool on the ESXi host.

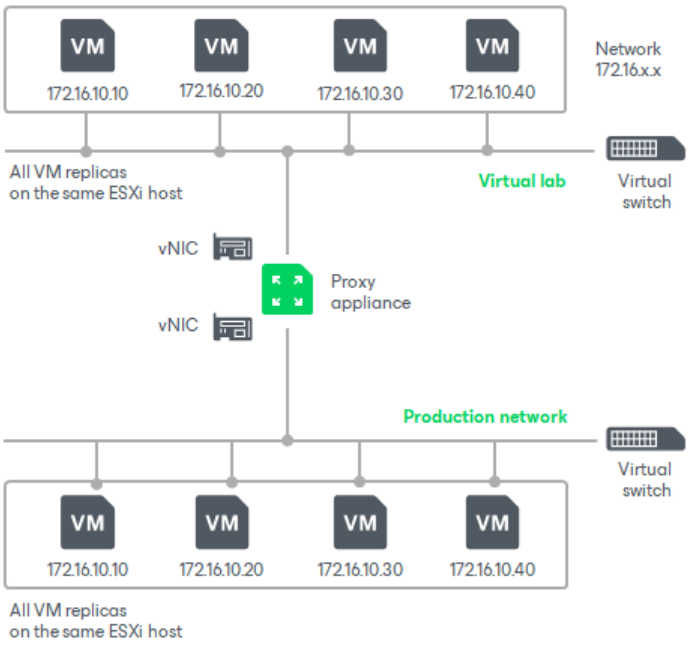


# Advanced Single-Host Virtual Labs

The advanced single-host virtual lab configuration can be used if your virtual environment is configured in the following way:

- All VM replicas that you want to verify are located on the same ESXi host.
- VM replicas you want to verify are connected to different networks.

> **IMPORTANT**
>
> For this configuration type, the virtual lab must be created on the same ESXi host where VMs replicas are located. If you create the virtual lab on some other ESXi host, the SureBackup job will fail.

In the advanced single-host virtual lab, Veeam Backup & Replication creates several virtual networks. The number of virtual networks corresponds to the number of production networks to which verified VM replicas are connected. Networks in the virtual lab are mapped to production networks.

Veeam Backup & Replication automatically adds a number of new VMware objects on the ESXi host where the virtual lab is created:

- A resource pool
- A VM folder
- A standard vSwitch

The vSwitch is only used by the VMs started in the virtual lab. There is no routing outside the virtual lab to other networks.

When you create an advanced single-host virtual lab, Veeam Backup & Replication configures basic settings for networks that are created in the virtual lab. You need to review these settings and manually adjust them.

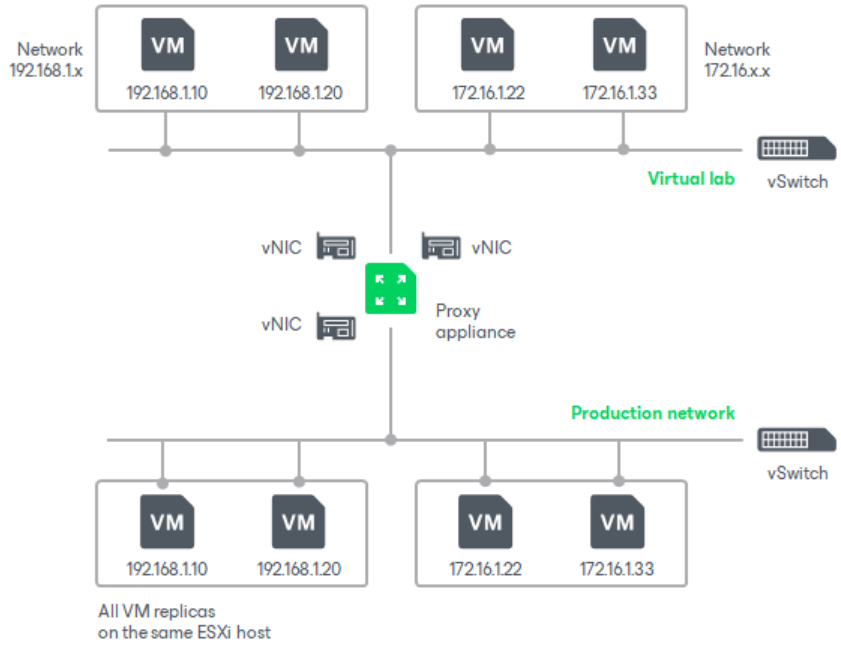


# Limitations of Single-Host Virtual Labs

If VM replicas are located on different hosts, you cannot use the single-host virtual lab configuration (basic or advanced). A single-host virtual lab uses standard vSwitches that have specific configuration limitations.

When you create or edit a virtual lab, Veeam Backup & Replication creates a new port group for each isolated network in the virtual lab. All VMs in the isolated network are added to this port group. Such configuration helps differentiate the traffic passing through the standard vSwitch to the isolated network in the virtual lab.

However, the standard vSwitch has a restriction: it is "limited" to one ESXi host. A standard vSwitch is configured on a specific ESXi host. The configuration of the standard vSwitch, such as information about port groups, resides on the ESXi host where the vSwitch is configured. Other ESXi hosts in the virtual environment do not have access to this information.

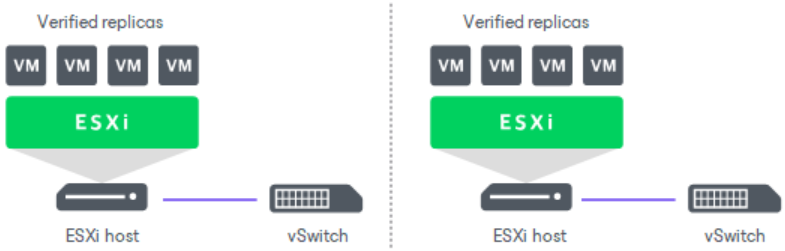


For this reason, the single-host configuration can only be used if all VM replicas are registered on the same ESXi host. If attempt to verify VM replicas registered on different ESXi hosts in the single-host virtual lab, VMs from different port groups will not be able to "see" each other and communicate with each other.

To overcome this limitation and verify VM replicas that are registered on different ESXi hosts, you can use advanced multi-host virtual labs.

# Advanced Multi-Host Virtual Labs

The advanced multi-host virtual lab configuration can be used if your disaster recovery (DR) site is configured in the following way:

- VM replicas that you want to verify are located on different ESXi hosts.
- VM replicas that you want to verify are connected to one or several networks.

The advanced multi-host virtual lab leverages the VMware Distributed vSwitch (DVS) technology. For more information, see VMware Docs.

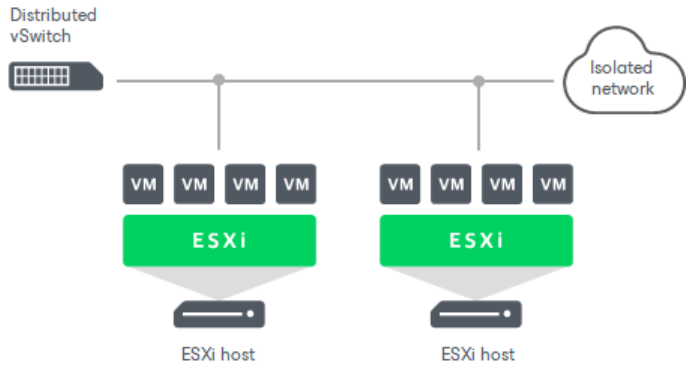


When you configure an advanced multi-host virtual lab, you must select an ESXi host on which the proxy server will be created and DVS on which Veeam Backup & Replication will create isolated networks.

Veeam Backup & Replication does not offer an option to automatically configure the DVS. The DVS must be preconfigured in your virtual environment correctly and allow traffic to traverse between all ESXi hosts assigned to the DVS. Failure to assign valid uplinks to the DVS will cause the job to fail.

> **IMPORTANT**
>
> DVS is limited to one datacenter. For this reason, all verified VM replicas and VM replicas from the application group must belong to the same datacenter. If VM replicas belong to different datacenters, you will be able to start them in the virtual lab but Veeam Backup & Replication will not be able to automatically verify them.

## Isolated Networks on DVS

For every isolated network in the virtual lab, Veeam Backup & Replication adds a new DVS port group to the DVS. The added port group is named after the isolated network.

The port groups created on the DVS must be isolated from the production environment. To isolate port groups, you can use one of the following methods:

1. **Connect DVS uplinks to the isolated network**. You can link the DVS that you plan to use for recovery verification to an external isolated network using uplink adapters. This operation must be performed manually by the backup administrator.

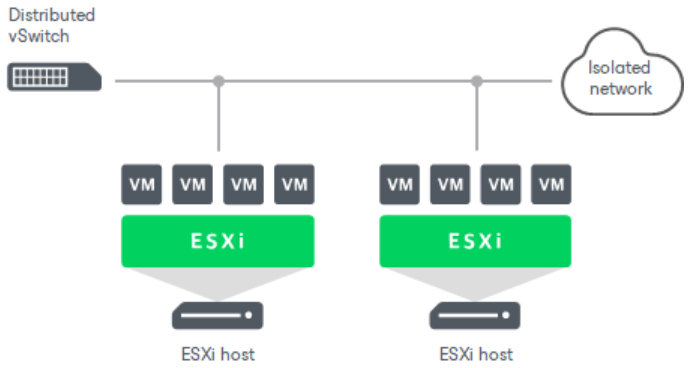


2. **Use VLAN tagging**. This method can be used only if your switch supports VLAN ID tagging. When you specify settings for isolated networks in the **New Virtual Lab** wizard, you can define different VLAN IDs for different isolated networks. Setting VLAN IDs restricts communication of VM replicas in the isolated network from the production environment.

**IMPORTANT**

If your network does not support VLAN ID tagging or the virtual lab is configured incorrectly, VM replicas will be started in the virtual lab but Veeam Backup & Replication will not be able to automatically verify them.

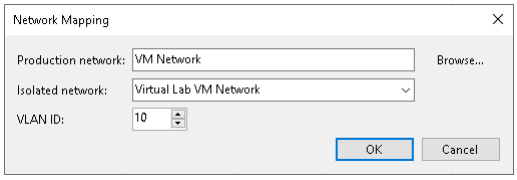


## Port Groups and VLAN IDs

You need to be extremely careful when specifying port group and VLAN ID settings for the advanced multi-host virtual lab.

# Port Groups in Advanced Multi-Host Virtual Labs

For the advanced multi-host virtual lab, Veeam Backup & Replication uses an existing DVS that was configured by the backup administrator beforehand. Veeam Backup & Replication creates a number of new port groups on the DVS, one per isolated network in the virtual lab.

When Veeam Backup & Replication creates a new port group, it performs a check for the DVS selected for the virtual lab:

- If a port group with the specified name already exists, Veeam Backup & Replication starts using it for the virtual lab. However, in this case, Veeam Backup & Replication will not be the owner of this port group.
- If a port group with the specified name does not exist, Veeam Backup & Replication creates it and becomes the owner of the created port group.

When a virtual lab is removed, Veeam Backup & Replication checks the ownership of the port group:

- If Veeam Backup & Replication is not the owner of the port group, the port group remains on the DVS. Veeam Backup & Replication stops using it.
- If Veeam Backup & Replication is the owner of the port group, Veeam Backup & Replication removes this port group from the DVS.

Several virtual labs can use the same port group. For this reason, you should be extremely careful when removing virtual labs. If Veeam Backup & Replication is the owner of the virtual lab and the port group is removed, other virtual labs using this port group may fail to start.

# VLAN IDs in Advanced Multi-Host Virtual Labs

A DVS port group has VLAN ID settings. If you select an existing port group for the virtual lab, you must specify its VLAN ID in the virtual lab settings.

- If VLAN ID settings are specified correctly, Veeam Backup & Replication will be able to configure the virtual lab and verify VM replicas in it.
- If VLAN ID settings are specified incorrectly, Veeam Backup & Replication will report an error informing that the selected port group exists but cannot be used due to incorrect VLAN ID settings.

# SureBackup Job for VM Replicas

You can verify VM replicas with the SureBackup job only in the **Full recoverability testing** mode. The SureBackup job aggregates all settings and policies of a recovery verification task, such as application group and virtual lab to be used, VM replicas that must be verified in the virtual lab and so on. The SureBackup job can be run manually or scheduled to be performed automatically.

When a SureBackup job runs, Veeam Backup & Replication first creates an environment for VM replica verification:

1. If antivirus software scan is enabled, Veeam Backup & Replication performs antivirus scan and use backup server as the mount server during this operation.
2. Veeam Backup & Replication starts the virtual lab.
3. In the virtual lab, it starts VMs from the application group or the linked job in the required order. VMs remain running until the verified VM replicas are booted and tested. If Veeam Backup & Replication does not find a successful VM replica or backup for any of VMs from the application group, the SureBackup job will fail.

When the virtual lab is ready, Veeam Backup & Replication starts VM replicas from the necessary restore point, tests and, depending on the specified settings, verifies them one by one or creates several streams and tests VM replicas simultaneously. If Veeam Backup & Replication does not find a successful restore point for any of verified VM replicas, verification of this VM replica fails, but the job continues to run.

By default, you can start and test up to three VM replicas at the same time. You can also increase the number of VMs to be started and tested simultaneously. Keep in mind that if these VMs are resource demanding, performance of the SureBackup job may decrease.

When the verification process is complete, VMs from the application group are powered off. Optionally, you can leave the VMs from the application group running to perform manual testing or enable user-directed application item-level recovery.

In some cases, the SureBackup job schedule may overlap the schedule of the replication job linked to it. The VM replica files may be locked by the replication job and the SureBackup job will be unable to verify such replica. In this situation, Veeam Backup & Replication will not start the SureBackup job until the replication job is over.

To overcome the situation of job overlapping, you may chain the replication and SureBackup jobs or define the timeout period for the SureBackup job. For more information, see Specifying Job Schedule.

> **NOTE**
>
> You can mix VM backups and replicas in the recovery verification job. For example, the application group may contain VMs that will be started from backup files and the job linked to the recovery verification job may be a replication job. Veeam Backup & Replication supports any type of a mixed scenario. Note that VMs that you verify with a SureBackup job must belong to the same platform — VMware or Hyper-V.

# SureBackup Job for VM Replicas Processing

The recovery verification process for VM replicas includes the following steps:

1. **Getting virtual lab configuration**. Veeam Backup & Replication gets information about configuration of the virtual lab where verified VM replicas must be started.
2. **Starting virtual lab routing engine**. Veeam Backup & Replication starts a proxy appliance. The proxy appliance is used as a gateway that provides access to VM replicas the virtual lab.

3. **Publishing**. Veeam Backup & Replication triggers a protective VMware snapshot for the verified VM replica.
4. **Reconfiguring**. Veeam Backup & Replication updates configuration files of the VM replica to connect the VM replica to the isolated network in the virtual lab.
5. **Configuring DC**. If the VM replica has the Domain Controller or Global Catalog role, the VM replica is reconfigured.
6. **Powering on**. Veeam Backup & Replication powers on the VM replica in the isolated network.
7. **Heartbeat test**. Veeam Backup & Replication checks whether the green or yellow VMware Tools heartbeat signal is coming from the VM replica or not. If the VM replica has no VMware Tools, the test is not performed and a notification is written to the session details.
8. **Running ping tests**. Veeam Backup & Replication checks if the VM replica responds to the ping requests or not. If the VM replica has no NICs and mapped networks for them and has no VMware Tools installed, the ping test is not performed and a notification is written to the session details.
9. **Application initialization**. Veeam Backup & Replication waits for applications installed in the VM replica, for example, Microsoft SQL Server, to start. The application initialization period is defined in the properties of the SureBackup job and by default is equal to 120 sec. Depending on the software installed in a VM, the application initialization process may require more time than specified in the SureBackup job settings. If applications installed in a VM are not initialized within the specified period of time, test scripts can be completed with errors. If such error situation occurs, you need to increase the Application initialization timeout value and start the job once again.
10. **Running test scripts**. Veeam Backup & Replication runs scripts to test whether the application installed in the VM replica is working correctly or not. If the VM replica has no VMware Tools installed and there are no NICs and mapped networks for them, Veeam Backup & Replication skips tests that use variables %vm_ip% and %vm_fqdn%, as the IP address of the VM cannot be determined. Test results are written to the session details. To define whether the script has completed successfully or not, Veeam Backup & Replication uses return codes. If the return code is equal to 0, the script is considered to complete successfully. Other values in the return code mean that the script has failed.
11. **Powering off**. After all tests have been performed, Veeam Backup & Replication powers off the verified VM replica.
12. **Unpublishing**. Veeam Backup & Replication deletes the protective VMware snapshot and rollbacks all changes made to the VM replica while it was running in the virtual lab.
13. **Stopping virtual lab engine**. Veeam Backup & Replication powers off the proxy appliance in the virtual lab.

# Stabilization Algorithm

To perform tests for a VM replica without errors, Veeam Backup & Replication needs to know that the VM replica is ready for testing. To determine this, Veeam Backup & Replication waits for the VM replica to reach a "stabilization point": — the moment when the VM replica booted and reports it is ready for tests. After the stabilization point has been reached, Veeam Backup & Replication can start heartbeat tests, ping tests and test scripts against the VM replica.

Veeam Backup & Replication establishes the stabilization point with the help of VMware parameters that it gets from the VM replica. Depending on the VM replica configuration, it uses one of the three algorithms:

- **Stabilization by IP**. This algorithm is used if the VM replica has VMware Tools installed, there are NICs and mapped networks for these NICs. In this case, Veeam Backup & Replication waits for an IP address of the VM replica for mapped networks that is sent by VMware Tools running in the VM replica. The sent IP address must be valid and must not change for a specific period of time.

- **Stabilization by heartbeat**. This algorithm is used if the VM replica has VMware Tools installed but there are no NICs and mapped networks for them. In this case, Veeam Backup & Replication waits for the green or yellow heartbeat signal from VMware Tools installed inside the VM replica.
- **Stabilization by Maximum allowed boot time**. This algorithm is used if the VM replica has neither VMware Tools installed, nor NICs and mapped networks for them. In this case, Veeam Backup & Replication will wait for the time specified in the **Maximum allowed boot time** field, which is considered to be a stabilization period for the VM replica. Once this time interval is exceeded, Veeam Backup & Replication considers that the VM replica is successfully booted and is ready for testing.

Once the stabilization point has been established, Veeam Backup & Replication runs ping, heartbeat tests and test scripts against the verified VM replica.

The stabilization process cannot exceed the value specified in the **Maximum allowed boot time** field. If the stabilization point cannot be determined within the **Maximum allowed boot time**, the recovery verification process will be finished with the timeout error. For this reason, you should be careful when specifying this value. Typically, a VM replica started by a recovery verification job requires more time to boot that a VM started regularly. When such error situation occurs, you need to increase the **Maximum allowed boot time** value and start the job again.

# On-Demand Sandbox

If you need to perform tests for production VMs, you can use an On-Demand Sandbox™. The On-Demand Sandbox is an isolated virtual environment where you can start one or more VMs from backups, VM replicas or VMs from storage snapshots. You can use the On-Demand Sandbox to perform the following tasks:

- Troubleshoot problems with VMs
- Test software patches and upgrades
- Install new software and so on

The On-Demand Sandbox uses a virtual lab — an isolated environment that is fully fenced off from the production environment. VMs started in the virtual lab remain in the read-only state. All changes made to VMs are written to redo logs (for VM backups and storage snapshots) or saved to delta files (for VM replicas). Redo logs and delta files are deleted after you finish working with the On-Demand Sandbox and power it off.

To create the On-Demand Sandbox, you must configure the following objects:

- Virtual lab in which VMs will be started. For more information, see Virtual Lab.
- Application group. The application group must include all VMs and VM replicas that you want to start in the On-Demand Sandbox. This can be one VM or a group of VMs that work together. For more information, see Application Group.
- SureBackup job. The virtual lab and application group must be linked to this job. For more information, see SureBackup Job.

# On-Demand Sandbox for Storage Snapshots

In On-Demand Sandbox, you can start VMs from snapshots existing on the production storage array. You can use On-Demand Sandbox to test VMs, troubleshoot issues, perform training and so on.

On-Demand Sandbox configuration where VMs from storage snapshots are started is similar to configuration of the regular On-Demand Sandbox. To start a VM from the storage snapshot in the isolated environment, you must configure the following objects:

- **Virtual lab**. The virtual lab must mirror the networking scheme of the production environment. You can configure a new virtual lab or use an existing virtual lab. Any type of the virtual lab configuration is supported: basic single-host, advanced single-host or advanced multi-host. For more information, see Virtual Lab in the Veeam Backup & Replication User Guide.
- **Application group**. The application group must contain one or several VMs that you want to start in the On-Demand Sandbox. You can select VMs from volumes or LUNs on the storage system. During the SureBackup job, Veeam Backup & Replication will detect the latest snapshot for this volume or LUN and start the VM from this snapshot. For more information, see Application Group in the Veeam Backup & Replication User Guide.
- **SureBackup job**. You must link the application group with VMs and virtual lab to the SureBackup job. For more information, see SureBackup Job. in the Veeam Backup & Replication User Guide.

## How On-Demand Sandbox for Storage Snapshots Works

To start a VM from a storage snapshot in On-Demand Sandbox, Veeam Backup & Replication needs to present this storage snapshot to an ESXi host as a datastore. To do this, Veeam Backup & Replication performs the following actions:

1. Veeam Backup & Replication detects the latest storage snapshot for the VM whose disks are located on the storage system.
2. Veeam Backup & Replication triggers the storage system to create a copy of the storage snapshot. The snapshot copy helps protect the storage snapshot from changes.

   To create a snapshot copy, Veeam Backup & Replication uses the same technology as for Data Recovery from Storage Snapshots. For more information, see the Data Recovery from Storage Snapshots section in the Storage System Snapshot Integration Guide.
3. The snapshot copy is presented as a new datastore to the ESXi host on which the virtual lab is registered.
4. Veeam Backup & Replication performs regular operations required for On-Demand Sandbox: reconfigures the VMX file, starts the VM, performs necessary tests for it and so on.

5. After you finish working with VMs and power off On-Demand Sandbox, Veeam Backup & Replication performs cleanup operations: powers off the VM and the proxy appliance in the virtual lab, unmounts the datastore from the ESXi host and triggers the storage system to remove the snapshot copy.

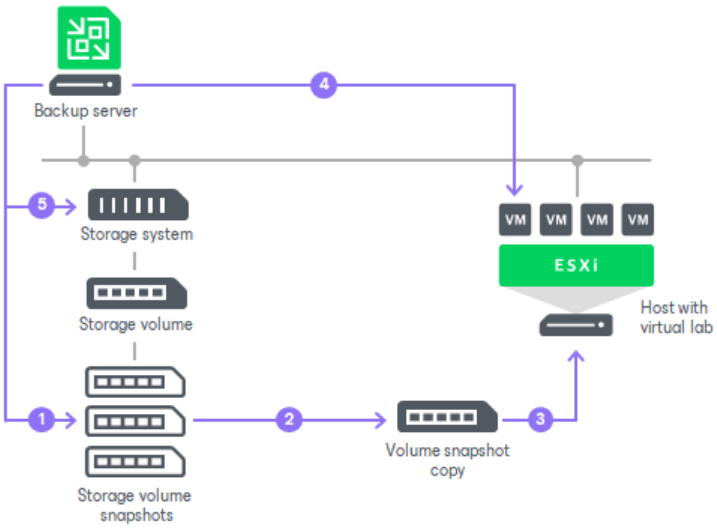


# Number of Mounted NFS Datastores

You can add to the application group several VMs that reside on different storage snapshots. In this case, Veeam Backup & Replication will trigger several snapshot copies (one per each storage snapshot) and present the equal number of datastores to the ESXi host.

The number of NFS datastores that can be mounted to the ESXi host is limited by VMware vSphere. If number of snapshot copies is great, Veeam Backup & Replication may fail to present all of them as datastores to the ESXi host. In this case, VMs in the application group will not be started and the SureBackup job will fail. For more information about limitations, see this VMware KB article.

To overcome this situation, Veeam Backup & Replication offers the mechanism of the snapshot copy re-mounting:

1. If Veeam Backup & Replication detects that there are not enough resources to mount a datastore, it displays a warning and offers you to free up resources on the ESXi host.
2. During the next 20 minutes, Veeam Backup & Replication attempts to mount the datastore with the time interval of 2 minutes.
3. If resources are freed and Veeam Backup & Replication manages to mount the datastore, VMs in the application group are started and the SureBackup job continues to run. If resources on the ESXi hosts are not freed within 20 minutes, the SureBackup job fails.

# Mixed Scenarios

You can start VMs from different sources in the On-Demand Sandbox:

- VM backups
- VM replicas
- VMs from storage snapshots

For example, you can add VMs from backups and VMs from storage snapshots to the same application group and link a replication job to the SureBackup job.

You cannot link jobs that trigger snapshots on storage arrays to the SureBackup job. This option is not supported.

| Type of Job/Object | SureBackup | SureReplica | SureSnap |
|---|---|---|---|
| Application group | ● | ● | ● |
| Linked job | ● | ● | ○ |

# Configuring On-Demand Sandbox

To configure the On-Demand Sandbox, perform the following steps:

1. Configure a virtual lab in which you plan to start VMs. For more information, see Creating Virtual Lab.
2. Configure an application group. The application group must contain all VMs that you plan to start in the On-Demand Sandbox and all VMs on which these VMs are dependent. For more information, see Creating Application Groups.
3. Configure a SureBackup job:
   a. Launch the **New SureBackup Job** wizard.
   b. At the **Virtual Lab** step of the wizard, select the configured virtual lab.
   c. At the **Application Group** step of the wizard, select the configured application group.
   d. Select the **Keep the application group running after the job completes** check box.
   e. Configure other job settings as required and save the job settings.

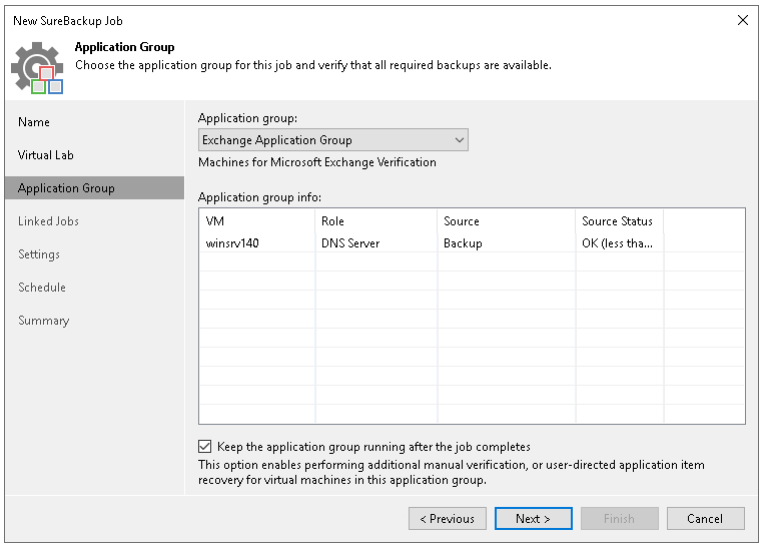


To start VMs in the On-Demand Sandbox, run the SureBackup job:

1. Open the **Home** view.
2. In the inventory pane, select **SureBackup**.
3. In the working area, right-click the configured SureBackup job and select **Start** or **Start to**.

Veeam Backup & Replication will start the virtual lab and power on VMs from the application group in the virtual lab. You will be able to connect to VMs and perform tests for them.

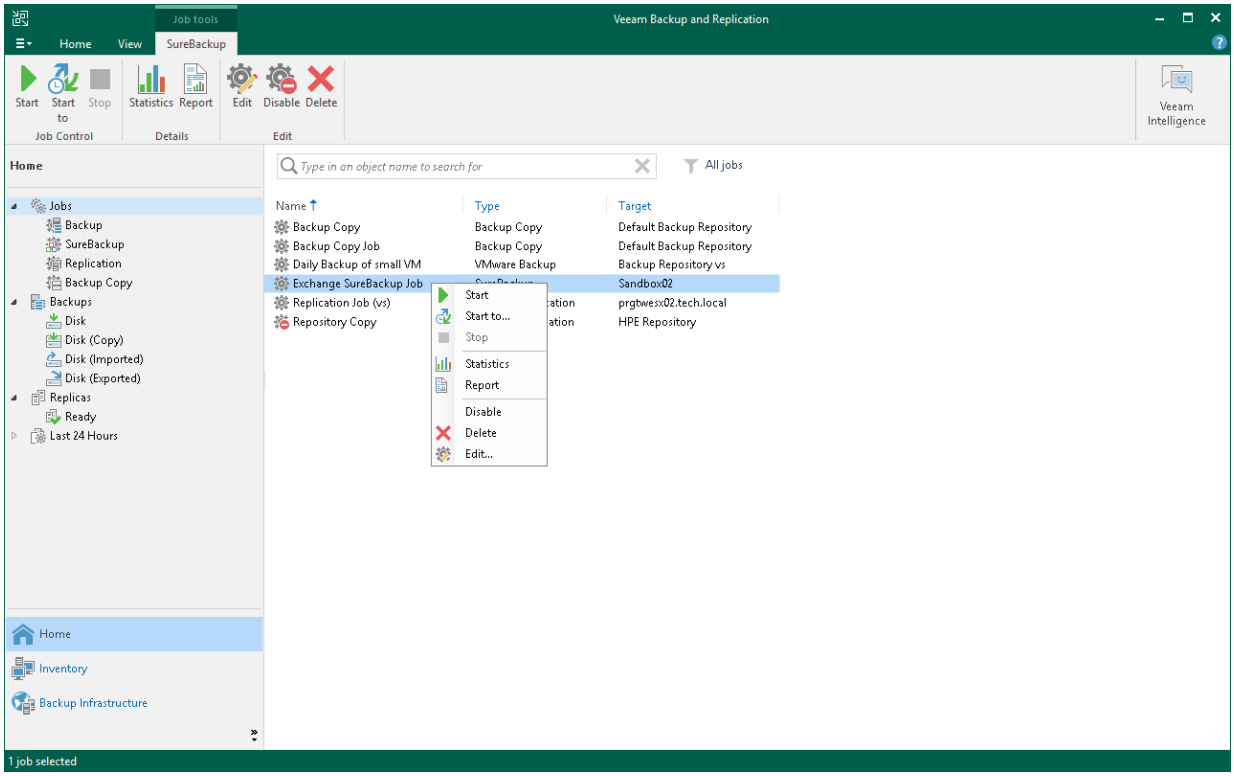

# Data Recovery

Veeam Backup & Replication offers the following types of recovery:

- VM recovery — to restore entire VMs to different data protection environments: to VMware vSphere, Hyper-V, Amazon EC2 and so on.
- Disk recovery — to recover and export disks.
- Item recovery — to recover VM files, guest OS files and folders, and application items.
- Veeam Data Integration API — to get the backup content using the iSCSI or FUSE protocol and analyze data stored in this backup.
- Secure restore — to scan data with antivirus software before restoring it to the production environment.
- Unstructured Data Recovery — to recover data previously backed up with file or object storage backup jobs.
- Restore from Tape — to recover data previously archived to tape.

**NOTE**

Veeam Backup & Replication provides backward compatibility: backups created with previous product versions can be restored with later product versions. However, backups created with later product versions cannot be restored with previous product versions.

# VM Recovery

VM recovery includes the following methods:

- Instant Recovery to VMware vSphere — to instantly recover workloads (VMs, EC2 instances, physical servers and so on) directly from compressed and deduplicated backup files as VMware vSphere VMs. When you perform Instant Recovery, Veeam Backup & Replication mounts recovered VM images to a host directly from backups stored on backup repositories.

  Instant Recovery helps improve recovery time objectives (RTO), minimize disruption and downtime of production workloads. However, Instant Recovery provides for VMs "temporary spares" with limited I/O performance. To provide the recovered VMs full I/O performance, you must finalize Instant Recovery — migrate the recovered VMs to production environment. If you do not want to migrate the recovered VM, you can stop publishing it. This removes the recovered VM.

  Use Instant Recovery for tier 1 VMs with little tolerance for business interruption and downtime. Besides disaster recovery matters, Instant Recovery can also be used for testing purposes.

- Instant Recovery to Microsoft Hyper-V — to instantly recover workloads directly from compressed and deduplicated backup files as Microsoft Hyper-V VMs. In many respects, this method is similar to Instant Recovery to VMware vSphere.

- Instant Recovery to Nutanix AHV — to instantly recover workloads directly from compressed and deduplicated backup files as Nutanix AHV VMs. Note that to restore to Nutanix AHV, you must install Nutanix AHV Plug-in on the backup server. For more information, see the Installation section in the Veeam Backup for Nutanix AHV User Guide.

- Entire VM restore — to recover entire VMs. When you recover VMs, you extract VM images from backups to the production storage. Entire VM restore takes more resources and time to complete than Instant Recovery but recovers VMs with full I/O performance. You also do not need to perform additional steps to finalize the recovery process.

  Use entire VM restore for VMs that require full I/O performance as soon as they are recovered and that tolerate some downtime.

- Staged restore — to run executable scripts for VMs before recovering them to the production environment. Staged restore is a part of entire VM restore.

  Use this option when you need to make sure that recovered VMs do not contain any personal or sensitive data.

- Restore to Amazon EC2 — to restore workloads of different data protection environments as EC2 instances.

- Restore to Microsoft Azure — to restore workloads of different data protection environments as Microsoft Azure VMs.

- Restore to Google Compute Engine (GCE) — to restore workloads of different data protection environments as Google VM instances.

- Restore to Nutanix AHV — to restore workloads of different data protection environments as Nutanix AHV VMs.

- Restore to Proxmox VE — to restore workloads of different data protection environments as Proxmox VE VMs.

- Restore to OLVM and RHV — to restore workloads of different data protection environments as OLVM and RHV VMs.

# Instant Recovery to VMware vSphere

With Instant Recovery to VMware vSphere, you can immediately recover different workloads (VMs, EC2 instances, physical servers and so on) as VMware vSphere VMs. Instant Recovery to VMware vSphere can be helpful, for example, if you want to migrate your infrastructure from one environment to another, or you want to recover your infrastructure in a matter of minutes but with limited performance.

During recovery, Veeam Backup & Replication runs workloads directly from compressed and deduplicated backup files. This helps improve recovery time objectives (RTO), minimize disruption and downtime of production workloads. The workloads are recovered in a matter of minutes.

When you perform Instant Recovery, Veeam Backup & Replication mounts workload images to a host directly from backups stored on backup repositories. This means that Veeam Backup & Replication creates fully functioning "temporary spares" with limited I/O performance. To provide full I/O performance, you must migrate these "temporary spares" to the production site. For more information, see Migration of Recovered VMs to Production Site.

Besides disaster recovery matters, Instant Recovery can also be used for testing purposes. Instead of extracting workload images to production storage to perform regular disaster recovery (DR) testing, you can run a workload directly from a backup file, boot it and make sure the guest OS and applications are functioning properly. For more information, see Finalizing Instant Recovery to VMware vSphere.

Instant Recovery supports bulk processing so you can immediately recover multiple workloads at once. If you perform Instant Recovery for several workloads, Veeam Backup & Replication uses the resource scheduling mechanism to allocate and use optimal resources required for Instant Recovery. For details, see Resource Scheduling.

## Supported Backup Types

You can recover workloads from the following types of backups:

- Backups of VMware vSphere virtual machines created by Veeam Backup & Replication

  You can also recover VMware vSphere VM data directly from storage snapshots.
- Backups of VMware Cloud Director virtual machines created by Veeam Backup & Replication
- Backups of Microsoft Hyper-V virtual machines created by Veeam Backup & Replication
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV
- Backups of Amazon EC2 instances created by Veeam Backup for AWS
- Backups of Microsoft Azure virtual machines created by Veeam Backup for Microsoft Azure
- Backups of Google Compute Engine VM instances created by Veeam Backup for Google Cloud
- Backups of oVirt VMs created by Veeam Backup for OLVM and RHV
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE

For details on how to restore Microsoft Hyper-V VMs, see the Instant Recovery to VMware vSphere section in the User Guide for Microsoft Hyper-V.

# How Instant Recovery Works

When Instant Recovery is performed, Veeam Backup & Replication uses the Veeam vPower technology to mount a workload image to an ESXi host directly from a compressed and deduplicated backup file. Since there is no need to extract the workload from the backup file and copy it to production storage, you can perform recovery from any restore point in a matter of minutes.

The image of the workload remains in read-only state to avoid unexpected modifications. By default, all changes to virtual disks that take place while a recovered VM is running are logged to auxiliary redo log files residing on the NFS server (backup server or backup repository). These changes are discarded as soon as the recovered VM is removed, or merged if you migrate the VM to the production site.

To improve I/O performance for a recovered VM, you can redirect VM changes to a specific datastore that is closer to the host where the VM resides. In this case, Veeam Backup & Replication will trigger a snapshot and will put it to the *Veeam IR* directory on the selected datastore, together with metadata files holding changes to the VM image.

# Migration of Recovered VMs to Production Site

To migrate the recovered VMs to the production storage, you can use one of the following relocation methods:

- Use Storage vMotion to quickly migrate the recovered VM to the production storage without any downtime. In this case, original VM data will be pulled from the NFS datastore to the production storage and consolidated with VM changes while the VM is still running. Storage vMotion, however, can only be used if you select to keep VM changes on the NFS datastore without redirecting them. Note that to use Storage vMotion, you need an appropriate VMware license.
- Use Quick Migration. In this case, Veeam Backup & Replication will perform a two-stage migration procedure — instead of pulling data from the vPower NFS datastore, it will recover the VM from the backup file on the production server, then move all changes and consolidate them with the VM data.

For more information on the relocation methods, see Quick Migration. For more information on how to launch the migration for workloads recovered with Instant Recovery, see Finalizing Instant Recovery to VMware vSphere.

# Performing Instant Recovery to VMware vSphere

With Instant Recovery, you can recover different workloads from backups and register them as VMware vSphere VMs. For the list of backups that you can use for recovery, see Supported Backup Types.

To perform Instant Recovery to VMware vSphere, use the **Instant Recovery to VMware** wizard.

> **NOTE**
>
> If you want to recover workloads as Hyper-V VMs, see Performing Instant Recovery to Hyper-V.

## Before You Begin

Before you perform Instant Recovery, consider the following:

- Prerequisites for Instant Recovery from storage snapshots are listed in the Data Recovery from Storage Snapshots section in the Storage System Snapshot Integration Guide.
- You can recover a workload from a backup that has at least one successfully created restore point.

- [For VMware vSphere version 8.0 and later] The backed-up VM for which you plan to perform Instant Recovery must not have disks on a vVol datastore and must not be located on an ESXi host version 8.0 or later. Otherwise, restore will fail.
- If you recover a workload to the production network, make sure that the original workload is powered off.
- Restore of CSV (Cluster Shared Volumes) is not supported. Cluster disks used as CSV are automatically excluded from restore.
- When you restore workloads other than VMware vSphere VMs, Veeam Backup & Replication assigns to the restored VMs the highest VM hardware version supported by the ESXi host to which you restore the workloads. For more information on ESXi hosts and compatible virtual machine hardware versions, see this VMware KB article.
- Consider the following for Linux workloads:
  - We strongly recommend having `dracut`, `mkinitrd` or `initramfs` installed on workloads that will be restored. Otherwise, they may not boot after restore.
  - Open the `/etc/fstab/` file and check that all filesystems are mounted using UUID. If any filesystems are mounted using block device name, the restored VM may not boot.
- If you want to scan recovered VM data for viruses, check the secure restore requirements and limitations.
- You must provide enough free disk space in vPower NFS datastore. The minimum amount of free space must equal the RAM capacity of the recovered VM plus 200MB. For example, if the recovered VM has 32 GB of virtual RAM, 32.2 GB of free space is required.

  By default, vPower NFS datastore is located in the `IRCache` folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`. The vPower NFS datastore is not used when you select to redirect virtual disk updates to a VMware vSphere datastore when configuring the job.
- You cannot use a vVol datastore as a destination for virtual disk updates.
- [For Veeam Agents] Restore of 4K native drives is not supported. For more information on VMware vSphere limitations, see this VMware KB article.
- [For Veeam Quick Migration with Smart Switch] In addition to the disk space mentioned above, you need to provide more disk space in vPower NFS datastore. The minimum amount of free space must equal the RAM capacity of the recovered VM.
- [For Nutanix AHV VMs] Instantly recoverd VM will have default virtual hardware settings: 2 CPU cores, 4GB RAM and one network adapter. If you want to change the default settings, turn off the VM and set the required virtual resources. Note that you must not switch off the instant recovery session before turning off the VM.

## Step 1. Launch Instant Recovery Wizard

To launch the **Instant Recovery to VMware** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select one of the following:
  - **VMware vSphere** > **Restore from backup** > **Entire VM restore** > **Instant recovery** > **Instant recovery to VMware vSphere** – if you want to recover VMware vSphere VMs from a VM backup created by Veeam Backup & Replication or you want to recover VMware vSphere VMs from storage snapshots.
  - **VMware Cloud Director** > **Restore from backup** > **VM restore** > **Entire VM restore** > **Instant recovery** > **Instant recovery to VMware vSphere** – if you want to recover VMware Cloud Director VMs from a VM backup created by Veeam Backup & Replication.
  - **Agent** > **Entire machine restore** > **Instant recovery** > **Instant recovery to VMware vSphere** – if you want to recover physical machines from a backup created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.
  - **AWS** > **Entire machine restore** > **Instant recovery** > **Instant recovery to VMware vSphere** – if you want to recover EC2 instances from a backup created by Veeam Backup for AWS.
  - **Azure IaaS backup** > **Entire machine restore** > **Instant recovery** > **Instant recovery to VMware vSphere** – if you want to recover Azure VMs from a backup created by Veeam Backup for Microsoft Azure.
  - **GCE backup** > **Entire machine restore** > **Instant recovery** > **Instant recovery to VMware vSphere** – if you want to recover VM instances from a backup created by Veeam Backup for Google Cloud.
  - **Nutanix backup** > **Entire machine restore** > **Instant recovery** > **Instant recovery to VMware vSphere** – if you want to recover VMs from backups created by Veeam Backup for Nutanix AHV.
  - **oVirt KVM** > **Entire machine restore** > **Instant recovery** > **Instant recovery to VMware vSphere** – if you want to recover VMs from backups created by Veeam Backup for OLVM and RHV.
  - **Proxmox VE** > **Entire machine restore** > **Instant recovery** > **Instant recovery to VMware vSphere** – if you want to recover VMs from backups created by Veeam Backup for Proxmox VE.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, select workloads that you want to recover and click **Instant Recovery** > **VMware vSphere** on the ribbon. Alternatively, you can right-click one of the selected workloads and select **Instant recovery** > **VMware vSphere**.

Alternatively, to recover VMware vSphere VMs from storage snapshots, you can open the **Storage Infrastructure** view. In the inventory pane, expand the storage system tree and select the necessary volume snapshot. In the working area, select the necessary VMs and click **Instant Recovery** on the ribbon.

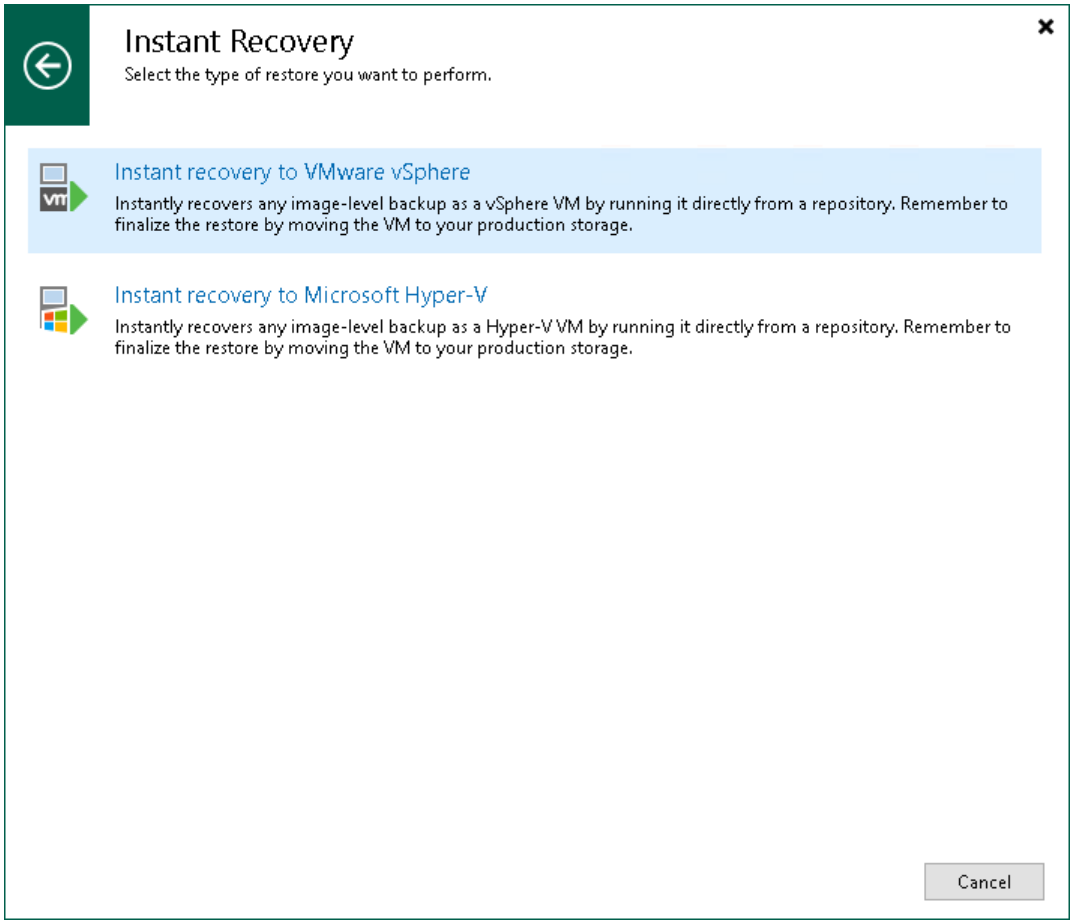

## Step 2. Select Workloads

At the **Machines** step of the wizard, select one or multiple workloads that you want to recover:

1. Click **Add**.
2. In the **Backup Browser** window, do the following:
   a. [For VMware vSphere VMs and VMware Cloud Director VMs] Select where to browse for VMs:
      - **From infrastructure** – browse the environment and select VMs or VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags) to recover. If you select a VM container, Veeam Backup & Replication will expand it to a plain VM list.

        When you add a VM to the list, Veeam Backup & Replication displays information about the most recent restore point in the **Restore point** column. If no restore point is available for the added VM, Veeam Backup & Replication displays a warning next to this VM.
      - **From backup** – browse existing backups or storage snapshots and select virtual machines.

   b. [For other workloads] In the list of backup jobs, expand a job and select workloads.
   c. Click **Add**.

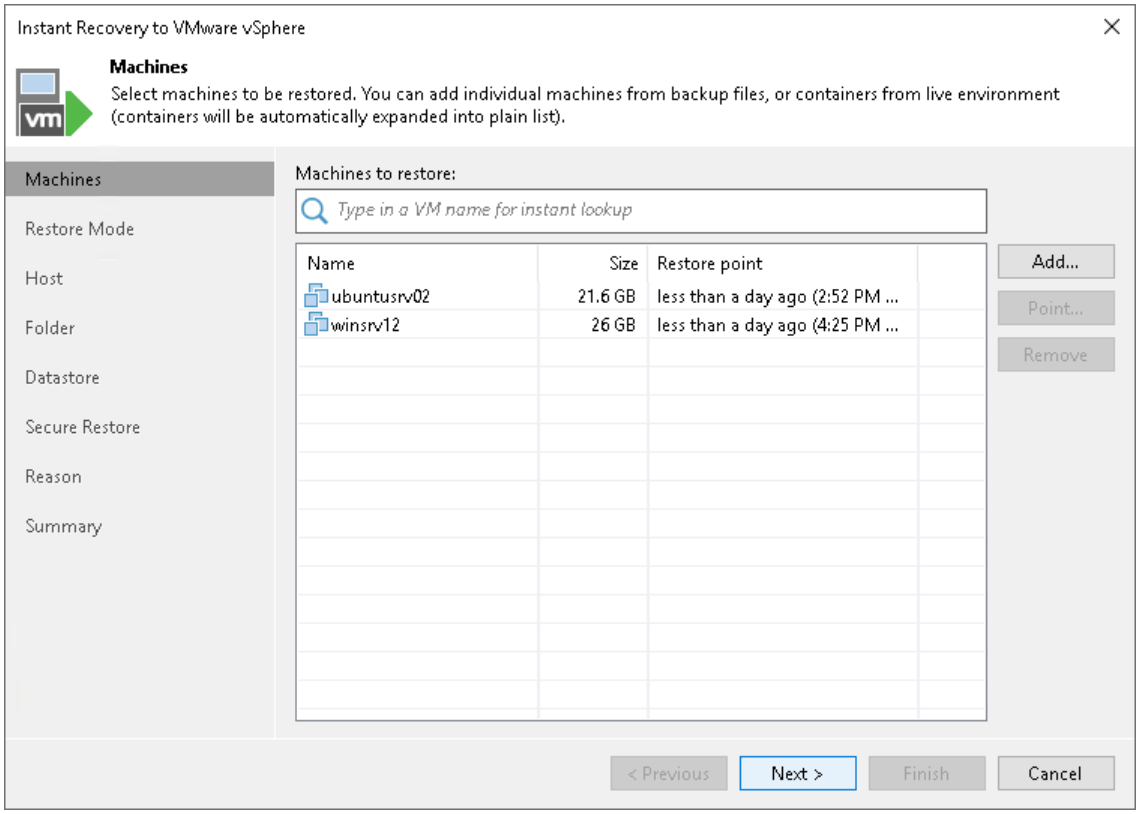

# Step 3. Select Restore Point

By default, Veeam Backup & Replication uses the latest valid restore point to recover workloads. You can recover a workload to an earlier state, if necessary:

1. In the **Machines to restore** list, select a workload.
2. Click **Point** on the right.
3. In the **Restore Points** window, select a restore point from which you want to recover the workload.

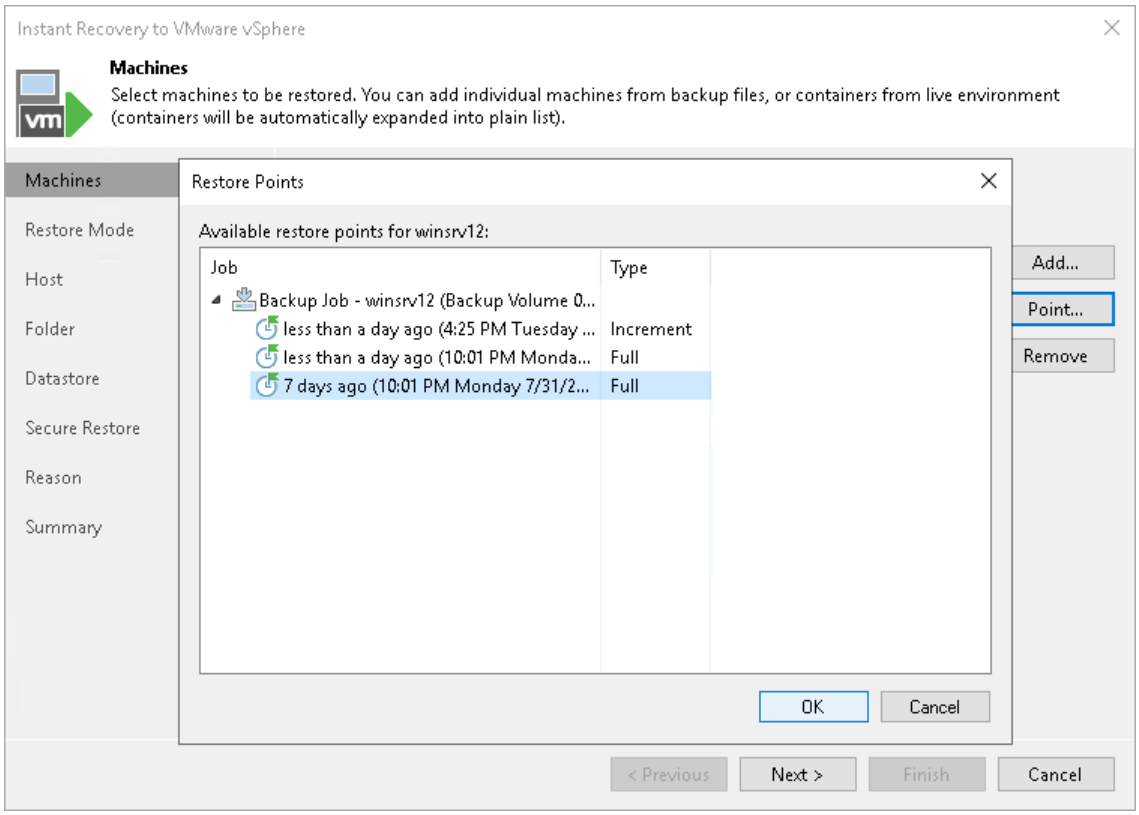

# Step 4. Select Restore Mode

This step is available only if you recover VMware vSphere VMs or VMware Cloud Director VMs.

At the **Restore Mode** step of the wizard, specify a destination for VM recovery and whether you want to recover VM tags:

1. Select a destination for recover:
   - **Restore to the original location** – select this option if you want to recover VMs with their initial settings and to their original location. If this option is selected, you will pass directly to the Reason step of the wizard.

     > **IMPORTANT**
     >
     > If you recover a VM with the original settings, and the original VM still exists in the virtual infrastructure, the original VM will be removed.

   - **Restore to a new location, or with different settings** – select this option if you want to recover VMs to a new location, or to any location but with different settings (such as VM location, network settings, format of recovered virtual disks and so on). If this option is selected, the **Instant Recovery to VMware** wizard will include additional steps for customizing VM settings.
2. If you want to recover tags that were assigned to the original VMs and assign them to the recovered VMs, select the **Restore VM tags** check box. Veeam Backup & Replication will recover the VMs with original tags if the following conditions are met:
   a. You recover VMs to their original location.
   b. The original VM tags are available on the source vCenter Server.

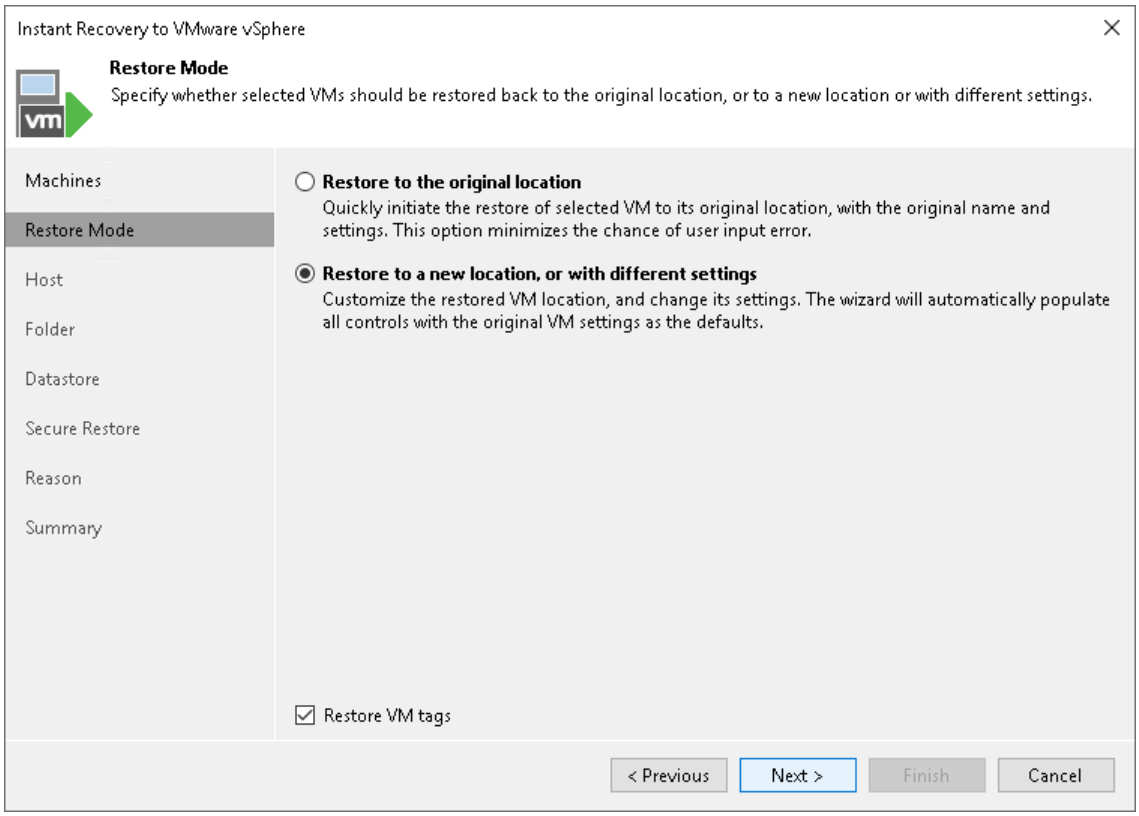

# Step 5. Specify Destination for Restored VMs

Specify a destination where the recovered VMs will reside. The destination settings differ depending on the number of workloads that you recover:

- Specifying Destination for One VM
- Specifying Destination for Multiple VMs

## Specifying Destination for One VM

The **Destination** step of the wizard is available if you recover one workload and recover it to a new location or with different settings.

At this step of the wizard, you configure destination settings such as the recovered VM name, target host, VM folder and so on. For workloads other than VMware vSphere VMs and VMware Cloud Director VMs, you must also configure a network mapping table. This table maps networks in the original site to networks in the target site. When the job starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update VM configuration files to replace the original networks with the specified networks in the target site. As a result, you will not have to re-configure network settings manually.

To specify the destination and network mapping settings for a recovered VM:

1. In the **Restored VM name** field, specify a name under which the workload will be recovered.
2. In the **Host** field, specify a host on which the VM will run.

   [For VMware vSphere VM recovery from storage snapshots] Veeam Backup & Replication will create a clone/virtual copy of the storage snapshot, mount it to the selected ESXi host and start the VM on this ESXi host.
3. In the **VM folder** field, specify a folder to which the recovered VM files will be placed.
4. In the **Resource pool** list, select a resource pool to which the VM will be placed.
5. [For workloads other than VMware vSphere VMs and VMware Cloud Director VMs] In the **Networks** section, configure the network mapping table. To configure one row of the table, select a network in the list and click **Choose**. The **Select Network** window displays all networks to which the target host or cluster is connected. In the list of networks, select a network to which the recovered VM will be connected instead of the original network.

6. Click the **Advanced** button and choose whether to preserve the BIOS UUID or generate a new BIOS UUID. By default, the BIOS UUID is preserved.

   We recommend that you select to generate a new BIOS UUID for the recovered VM to prevent conflicts if the original workload still resides in the production environment. The BIOS UUID change is not required if the original VM no longer exists, for example, it was deleted.

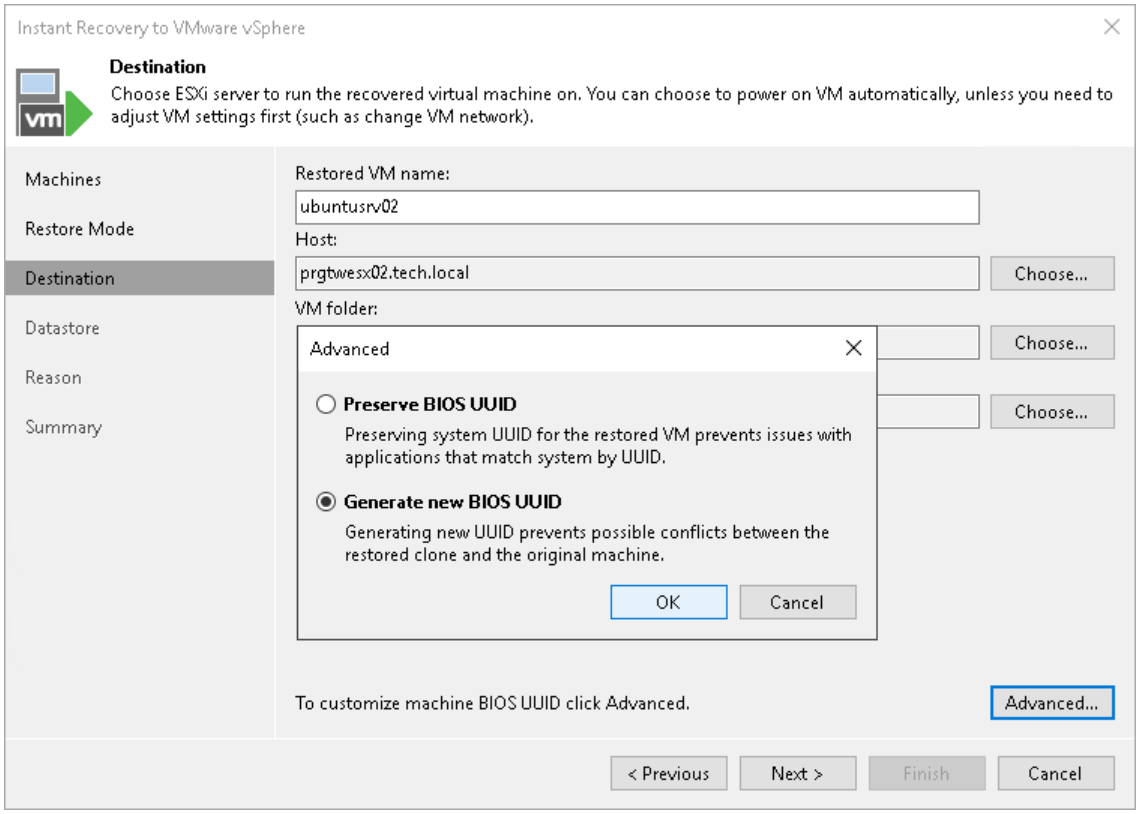


## Specifying Destination for Multiple VMs

The following steps are available if you recover multiple workloads and recover them to a new location or with different settings.

If you recover multiple workloads, specify the following settings for the destination:

1. At the Host step of the wizard, select a target host.
2. At the Folder step of the wizard, specify VM settings.
3. [For workloads other than VMware vSphere VMs and VMware Cloud Director VMs] At the Network step of the wizard, specify network settings.

# Selecting Target Host

At the **Host** step of the wizard, specify a target host and resource pool for recovered VMs:

1. In the list, select the necessary VMs and click the Host button.
2. From the virtual environment, select a standalone or clustered host where the selected VMs will be registered.

   [For VMware vSphere VM recovery from storage snapshots] Veeam Backup & Replication will create a clone/virtual copy of the storage snapshot, mount it to the selected ESXi host and start the VM on this ESXi host.

3. Select one or multiple VMs and click the **Pool** button.
4. In the list, select a resource pool where the selected VMs will be stored.

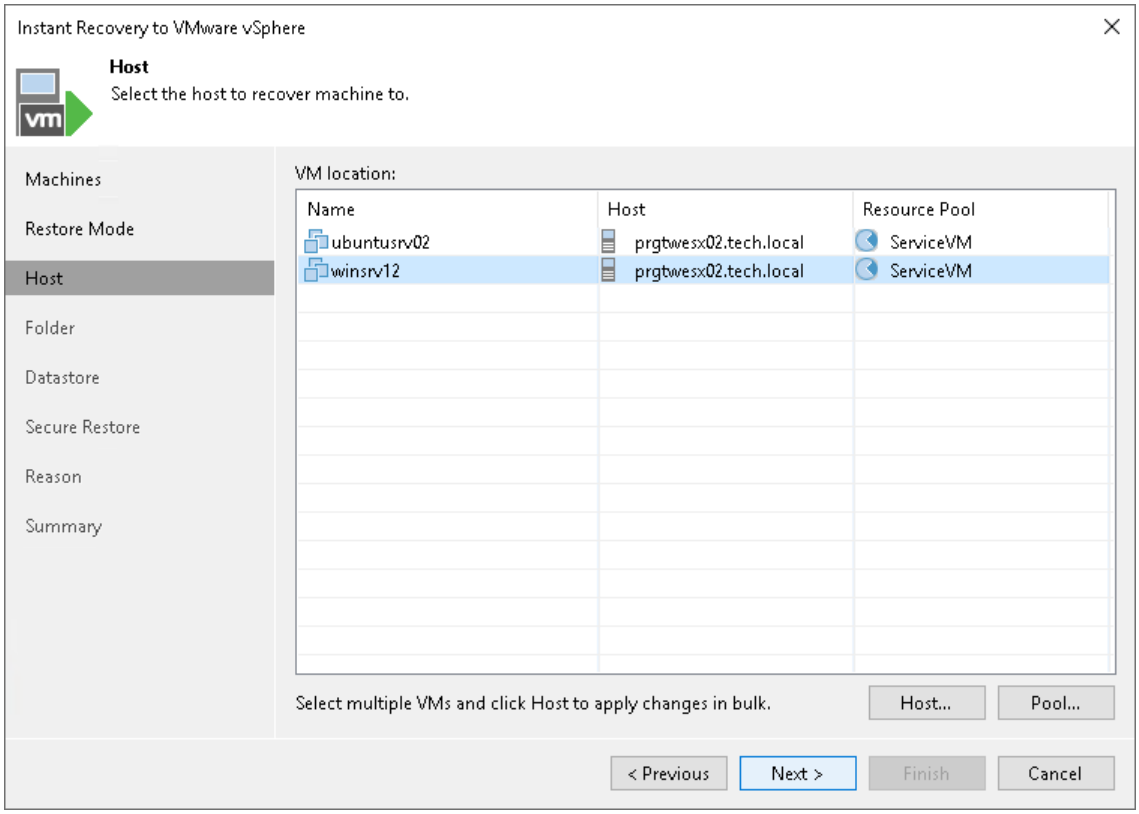


# Specifying VM Settings

For each recovered VM, you can change a VM name, BIOS UUID and folder where VM files must be stored. By default, the BIOS UUID is preserved.

We recommend that you specify a new name and generate a new BIOS UUID to prevent conflicts if the original workload still resides in the production environment. The name and BIOS UUID change is not required if the original workload no longer exists, for example, it was permanently deleted.

To specify a folder:

1. In the **VM settings** list, select the necessary VM. If you want to specify a folder for multiple VMs, select VMs that will be recovered to the same host.
2. Click the **Folder** button.
3. Select a folder where VM files must be stored.

To change a recovered VM name and BIOS UUID:

1. In the **VM settings** list, select one VM.
2. Click the **Customize** button.
3. In the **VM Settings** window, do the following:
   a. In the **Set name to** field, specify a new VM name.
   b. To add a prefix and suffix to the name specified in the **Set name to** field, select **Add prefix** and **Add suffix** check boxes.

c. In the **BIOS UUID** section, specify whether to preserve or generate a new BIOS UUID.

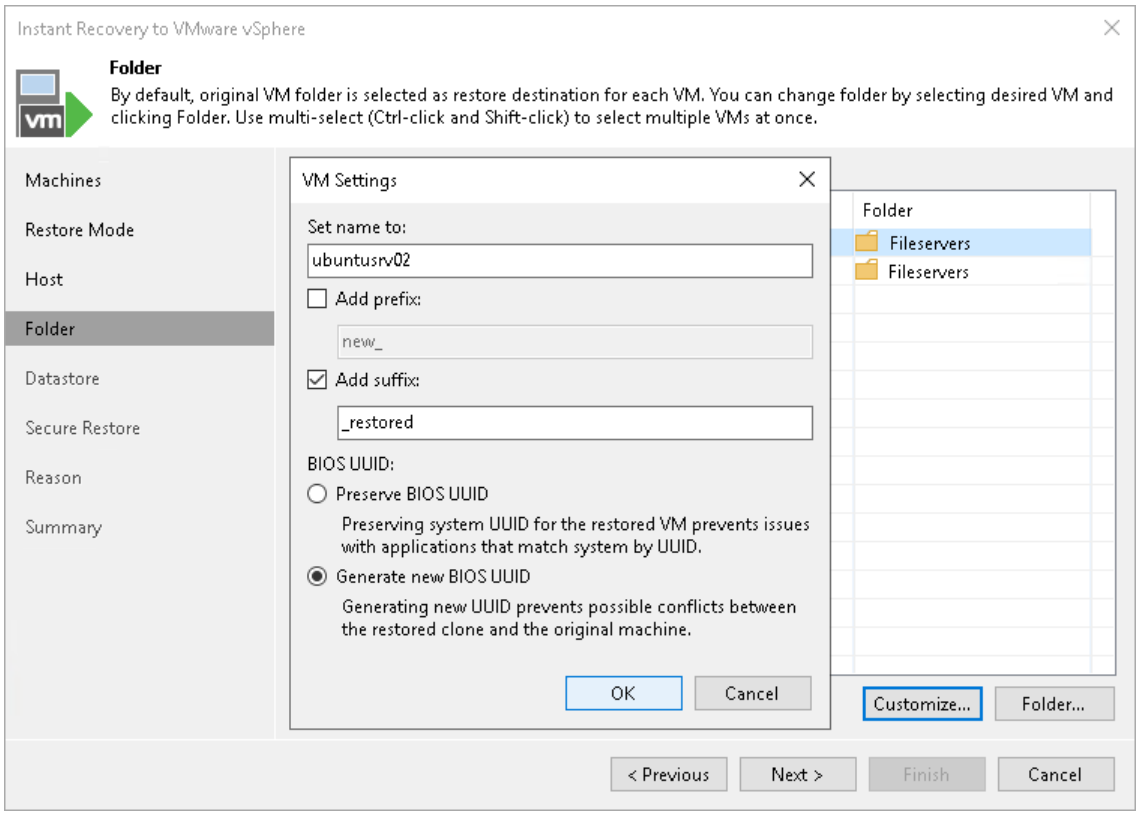


# Specifying Network Settings

This step is available if you recover workloads other than VMware vSphere VMs and VMware Cloud Director VMs.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the original site to networks in the target site (site where VMs will be recovered). When the job starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update VM configuration files to replace the original networks with the specified networks in the target site. As a result, you will not have to re-configure network settings manually.

To change networks to which the recovered VMs will be connected:

1. In the list, select one or multiple workloads and click the **Network** button.

   If a workload is connected to multiple networks, you can select a network to map and click **Network**.

2. The **Select Network** window displays all networks to which the target host or cluster is connected. In the list, select a network to which the recovered VM will be connected after recover.

If you do not want to connect a recovered VM to any virtual network, select the original workload and click **Disconnected**.

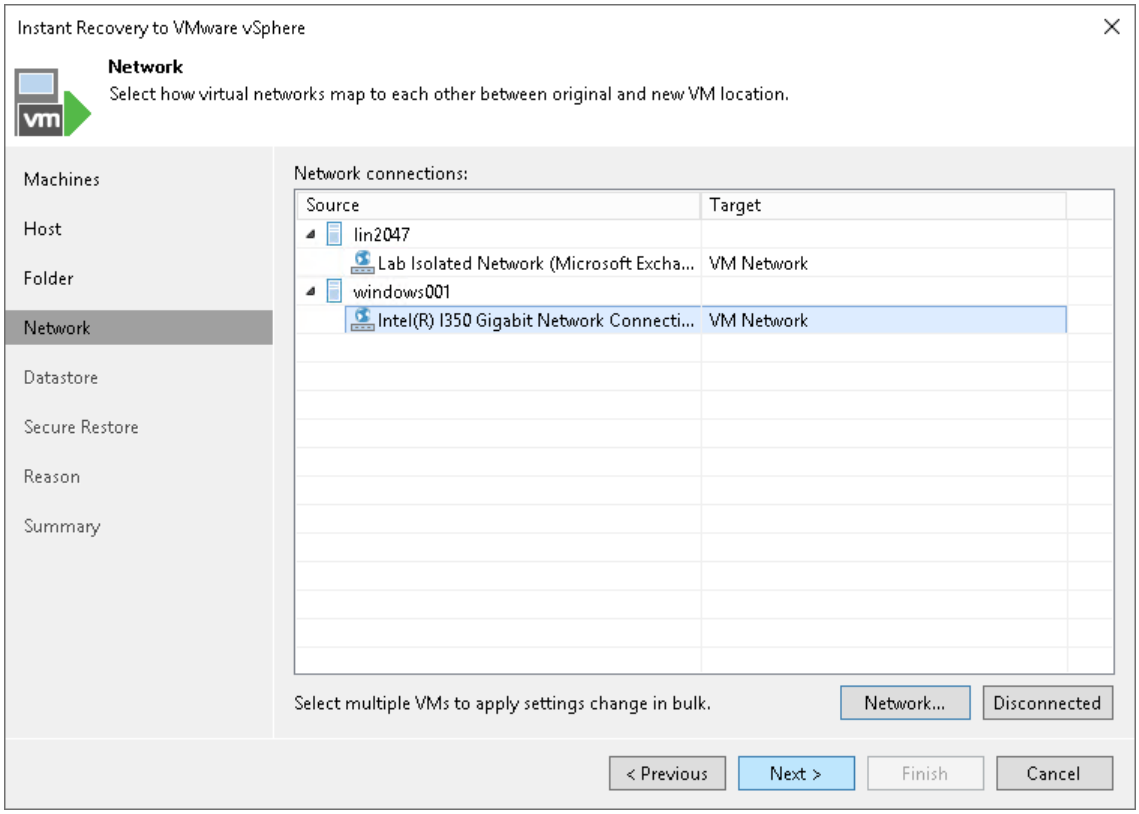

# Step 6. Select Destination for Virtual Disk Updates

This step is available if you recover workloads to a new location or with different settings. However, this step is not available if you recover VMware vSphere VMs from storage snapshots.

At the **Datastore** step of the wizard, you can select where to store redo logs when a VM is running from a backup. Redo logs are auxiliary files used to keep changes that take place while the recovered VM runs.

By default, redo logs are stored on the vPower NFS server. You can store redo logs on another datastore in the virtual environment. As soon as a recovery verification job completes, Veeam Backup & Replication deletes redo logs.

To redirect redo logs:

1. Select the **Redirect write cache** check box.
2. Click **Choose** and select a datastore from the list.

> **IMPORTANT**
>
> If the size of recovered VM disks is greater than 2 TB, you must not place redo logs on a VSAN datastore. Otherwise, Veeam Backup & Replication will fail to create a snapshot for the recovered VMs. For more information, see VMware Docs.

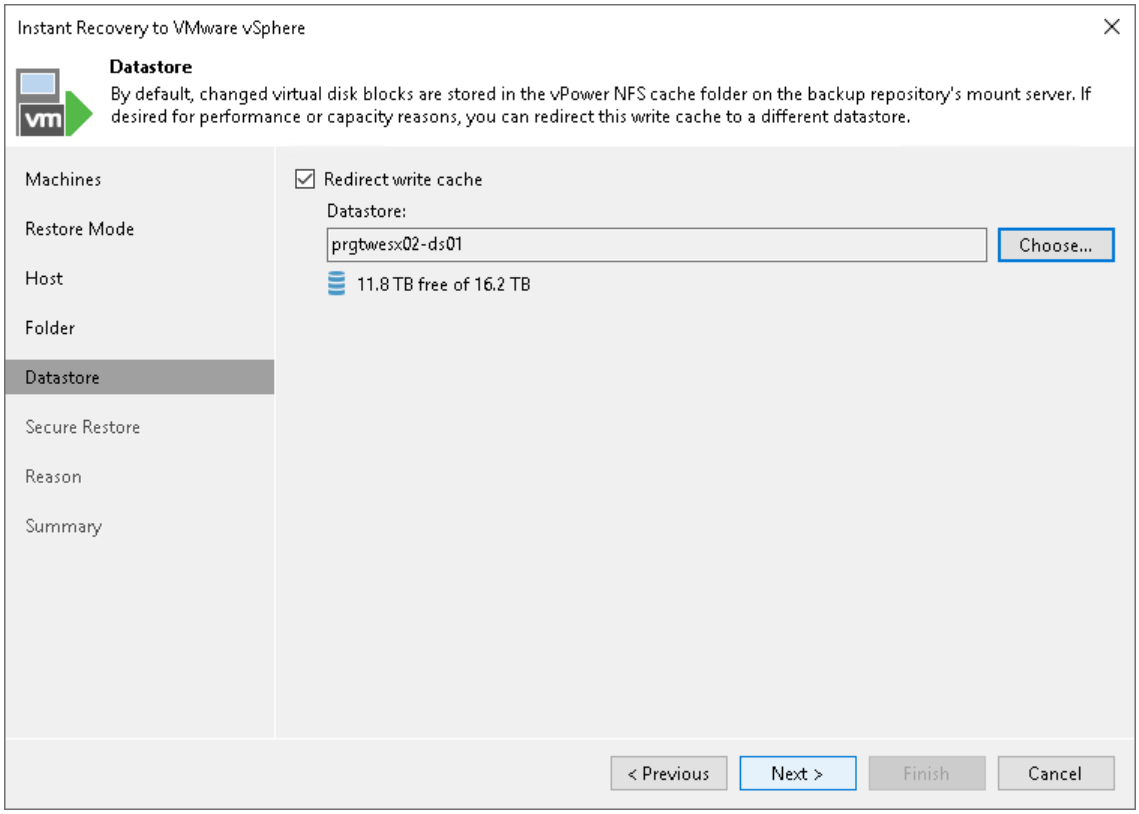

# Step 7. Configure Helper Appliance

This step is available if you recover workloads with Linux OS, recover them to a new location or with different settings, and VIX API is not available. However, this step is not available if you recover VMware vSphere VMs from storage snapshots.

Veeam Backup & Replication recovers Linux machines to a standalone host using a helper appliance. The helper appliance is an auxiliary Linux-based VM registered by Veeam Backup & Replication. The appliance is quite small — around 150 MB. It requires the same amount of RAM as the VM being restored and takes around 10 seconds to boot.

To configure the helper appliance:

1. [For multiple machines] In the **Network** list, expand a host and select one machine.
2. Click the **Configure** button.
3. In the **Network Settings** window, select a network for the helper appliance.
   a. Click the **Browse** button near the **Production network** field.
   b. In the **Select Network** window, Veeam Backup & Replication shows a list of networks to which the target host is connected. In this list, select a network to which the helper appliance must be connected.

      Consider that the backup server and the mount server must have access to the helper appliance over the network.
4. Specify IP addressing settings for the helper appliance and DNS server:
   a. Click **Configure**.
   b. Switch to the **IPv4** or **IPv6** tab depending on which addresses you want to configure. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
   c. Select the **Enable IPv4/IPv6 interface** check box.
   d. Configure IP settings for the helper appliance:
      - If you use a DHCP server in the network and want to obtain the IP address automatically, leave the **Obtain an IP address automatically** option selected.
      - To manually assign a specific IP address to the helper appliance, click **Use the following IP address** and specify the IP address settings.
   e. Configure IP settings for the DNS server:
      - If you use a DHCP server in the network and want to obtain the IP address automatically, leave the **Obtain DNS server address automatically** option selected.
      - To manually assign a specific IP address to the DNS server, click **Use the following DNS server address** and specify preferred and alternate addresses.

f. Click **OK**.

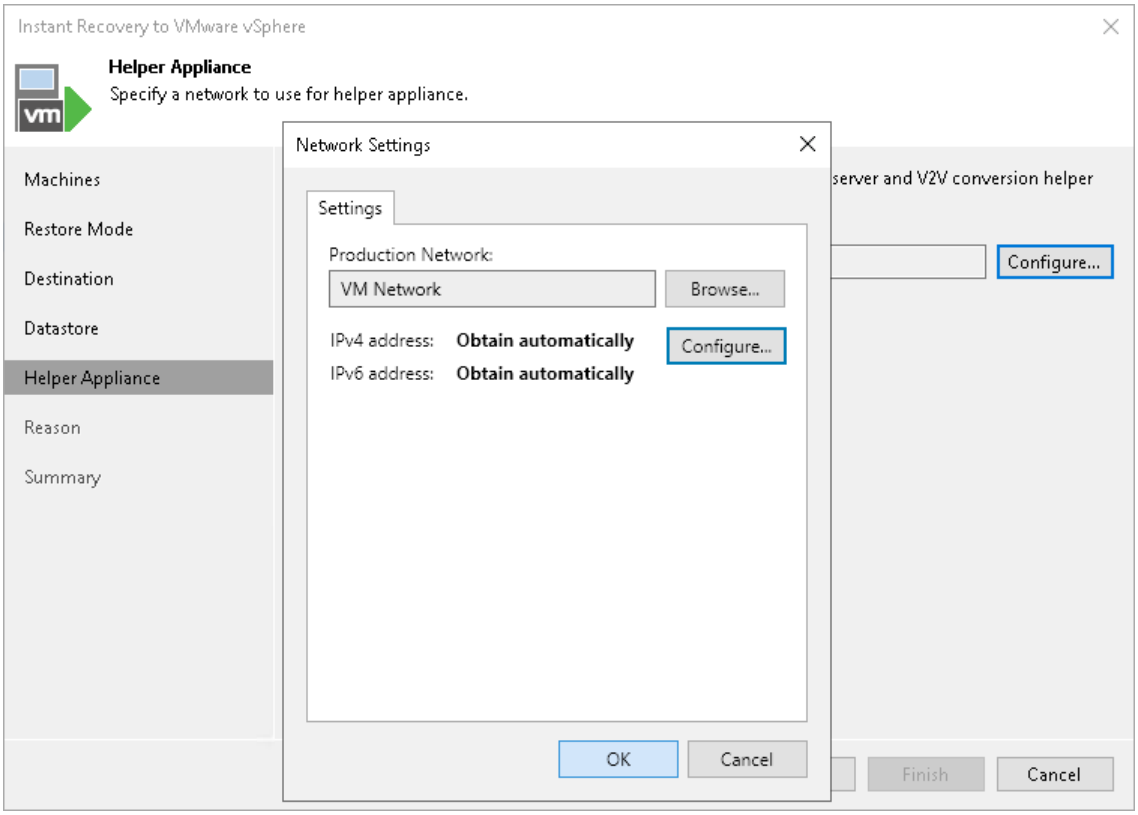
Instant Recovery to VMware vSphere
Helper Appliance
Specify a network to use for helper appliance.
Machines
Restore Mode
Destination
Datastore
Helper Appliance
Reason
Summary
server and V2V conversion helper
Configure...
Finish
Cancel
Network Settings
Settings
Production Network:
VM Network
Browse...
IPv4 address: Obtain automatically
IPv6 address: Obtain automatically
Configure...
OK
Cancel

# Step 8. Specify Secure Restore Settings

This step is available if you recover workloads with Microsoft Windows OS. However, this step is not available if you recover VMware vSphere VMs from storage snapshots.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Select which action Veeam Backup & Replication will take if scan finds a virus threat:
   - **Proceed to recovery but disable VM network adapters** — if you want to restore VMs with disabled network adapters (NICs).
   - **Abort VM recovery** — if you want to cancel the restore session.

3. Select the **Continue scanning all remaining files after the first occurrence** check box if you want to continue scanning VM data after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.

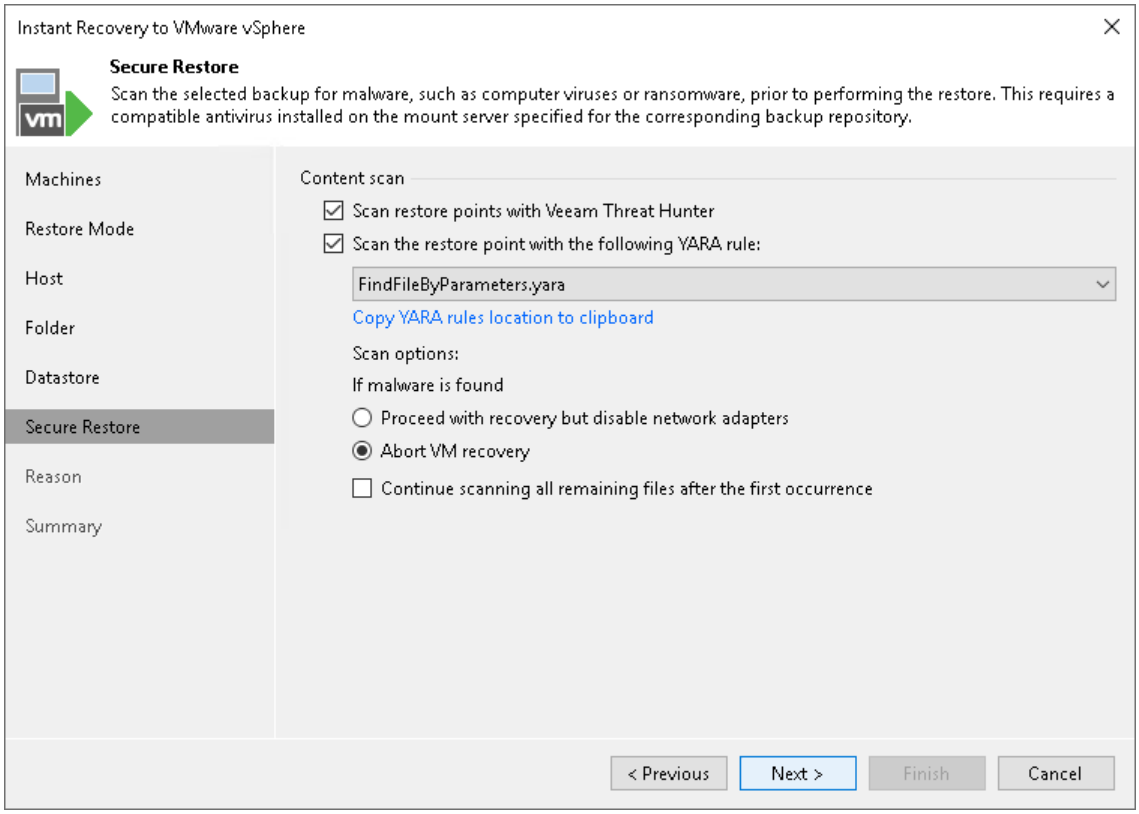

# Step 9. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for performing Instant Recovery. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

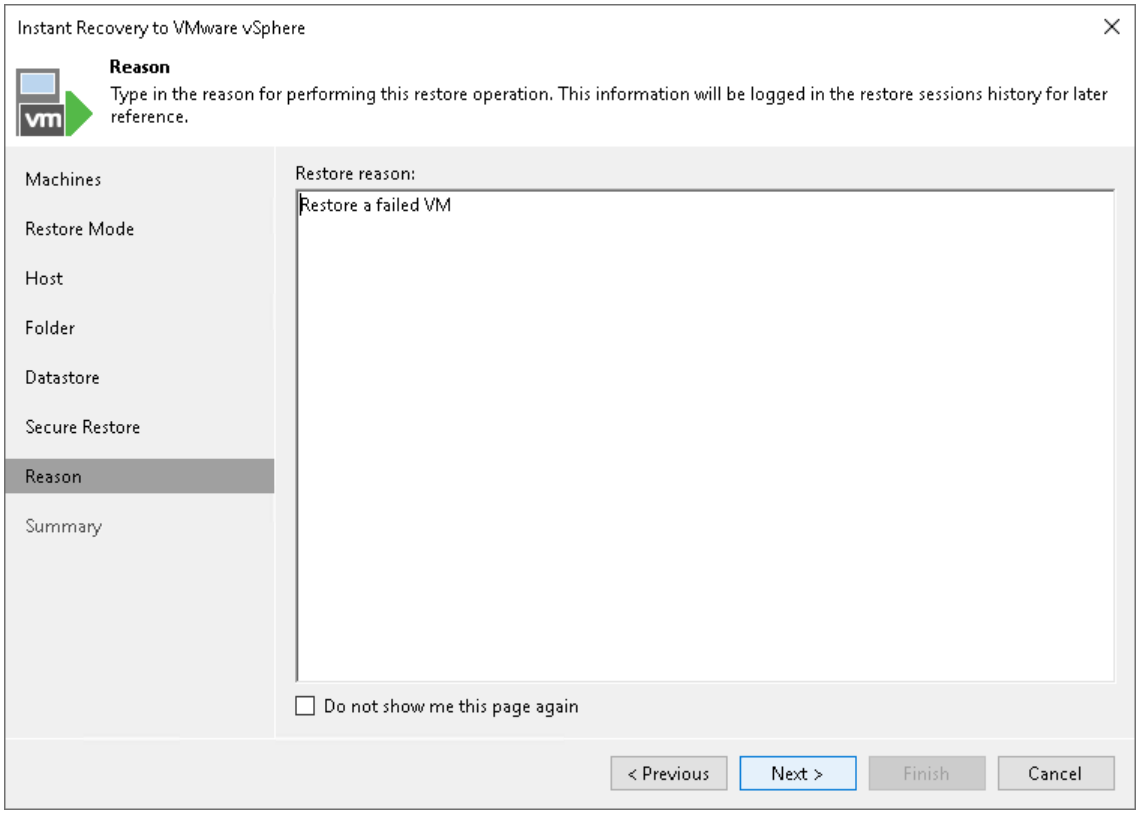

## Step 10. Verify Instant Recovery Settings

At the **Summary** step of the wizard, specify additional settings for Instant Recovery:

1. If you recover production workloads that have failed and want to recover them with initial network settings, select the **Connect VMs to network** check box.

   If you recover workloads for testing disaster recovery while the original workloads are still running, leave this check box unselected. Before you power on the recovered VMs, you must disconnect them from the production network and connect to a non-production network to avoid conflicts.

2. To start the VMs right after recovery, select the **Power on target VMs after restoring** check box. If you recover the workloads to the production network, make sure that the original workloads are powered off.

3. Check settings that you have specified for Instant Recovery and click **Finish**.

4. Check that the publishing process has started and click **Close**.

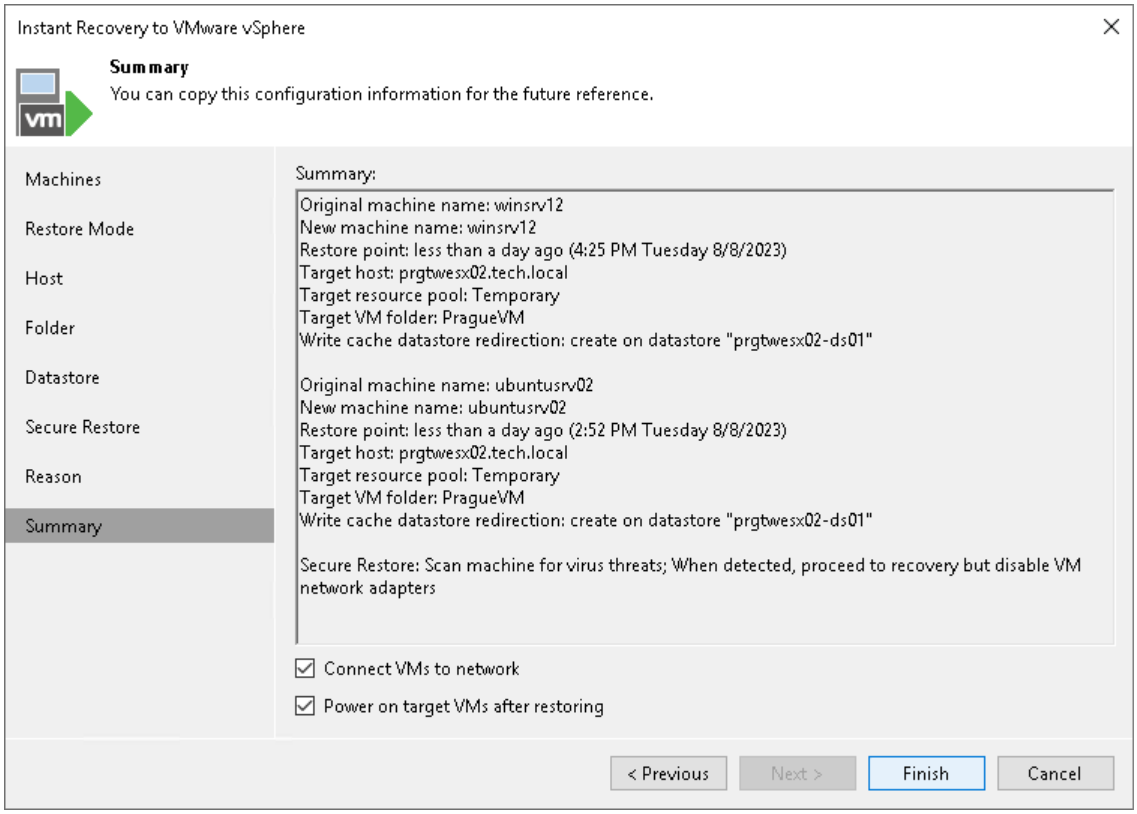


# What You Do Next

Finalizing Instant Recovery to VMware vSphere

# Finalizing Instant Recovery to VMware vSphere

After the VMs have been successfully recovered, you must finalize the process. For this, test the recovered VMs and decide whether to migrate them to production environment or stop publishing.

- Testing recovered VMs
- Migrating recovered VMs
- Stop publishing recovered VMs

# Testing Recovered VMs

To test the recovered VMs before you migrate them to production, you can launch the VMware Remote Console (VMRC) from the Veeam Backup & Replication console.

> **IMPORTANT**
>
> VMware Remote Console is not included as part of Veeam Backup & Replication installation and must be installed separately. For details, see Install the VMware Remote Console Application.

To open a VM console in Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click a VM and select **Open VM console**.

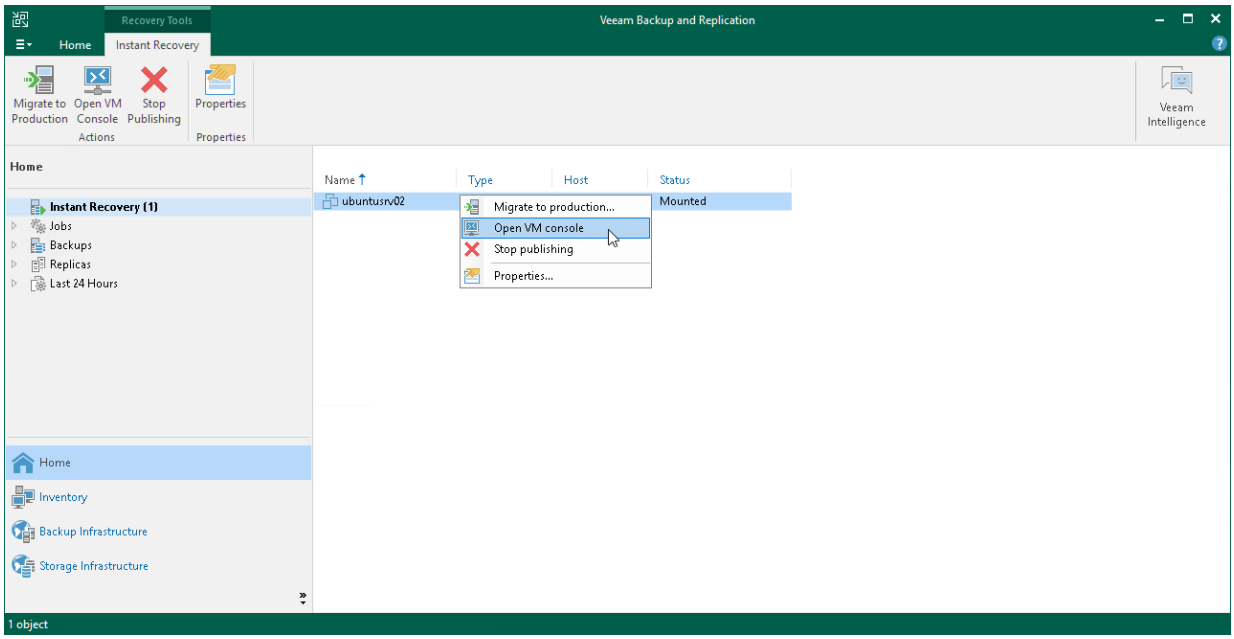


# Migrating Recovered VMs

If the VMs were recovered successfully, you can migrate them to the production environment.

To migrate a recovered VM to the production environment:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click a VM and select **Migrate to production**. Veeam Backup & Replication will launch the **Quick Migration** wizard. For more information on requirements and limitations for Quick Migration, see Migrating VMs.

   During migration, Veeam Backup & Replication will recover the VM from the backup file and additionally move all changes that were made while the VM was running from the backup in the Instant Recovery mode.

If you have launched Instant Recovery to a different location and you want to protect the recovered VM after migration finishes, you need to add the recovered VM to a backup job manually. If you have launched Instant Recovery to the original location, your actions depend on the method used for migration and whether the **Delete source VM files upon successful migration** check box is enabled in the migration wizard. For more information, see Finish Working with the Quick Migration Wizard.

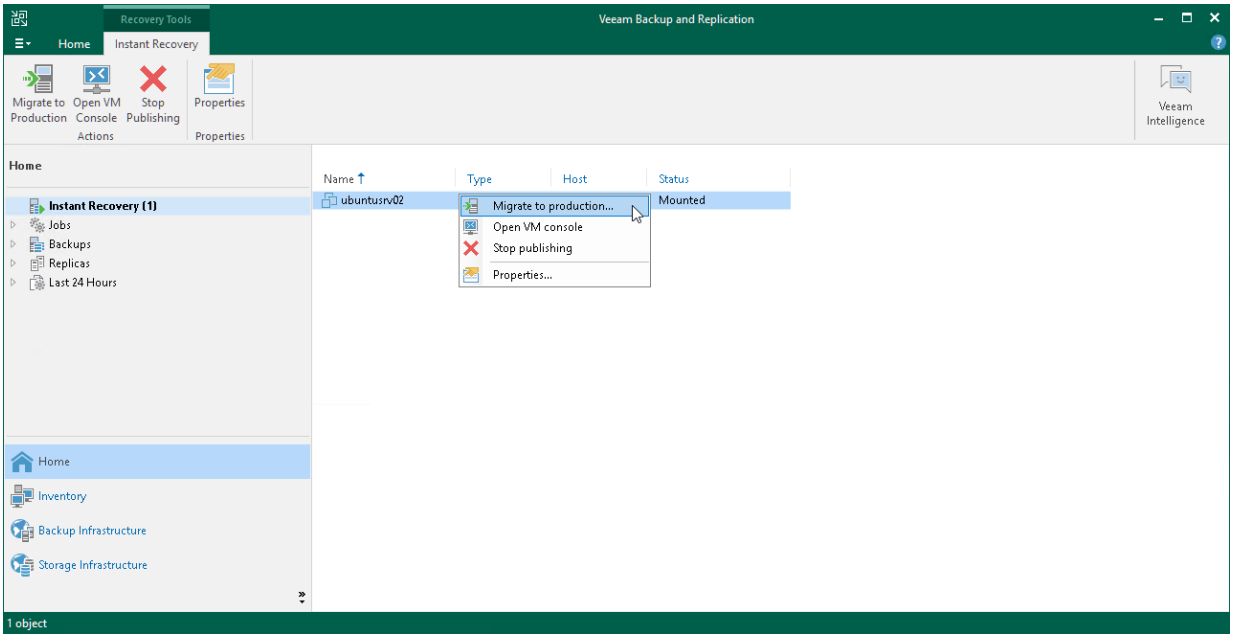


# Stop Publishing Recovered VMs

If your tests have failed, you can stop publishing the recovered VMs. This will remove the recovered VMs from the host that you selected as the destination for recovery. Note that all changes made in the recovered VMs will be lost.

> **IMPORTANT**
>
> [For restore to original location] Both the recovered and original VMs are removed if you stop publishing the recovered VM. This is because during restore to the original location, Veeam Backup & Replication removes the original VM.

To stop publishing a recovered VM:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.

3. In the working area, right-click a VM and select **Stop publishing**.

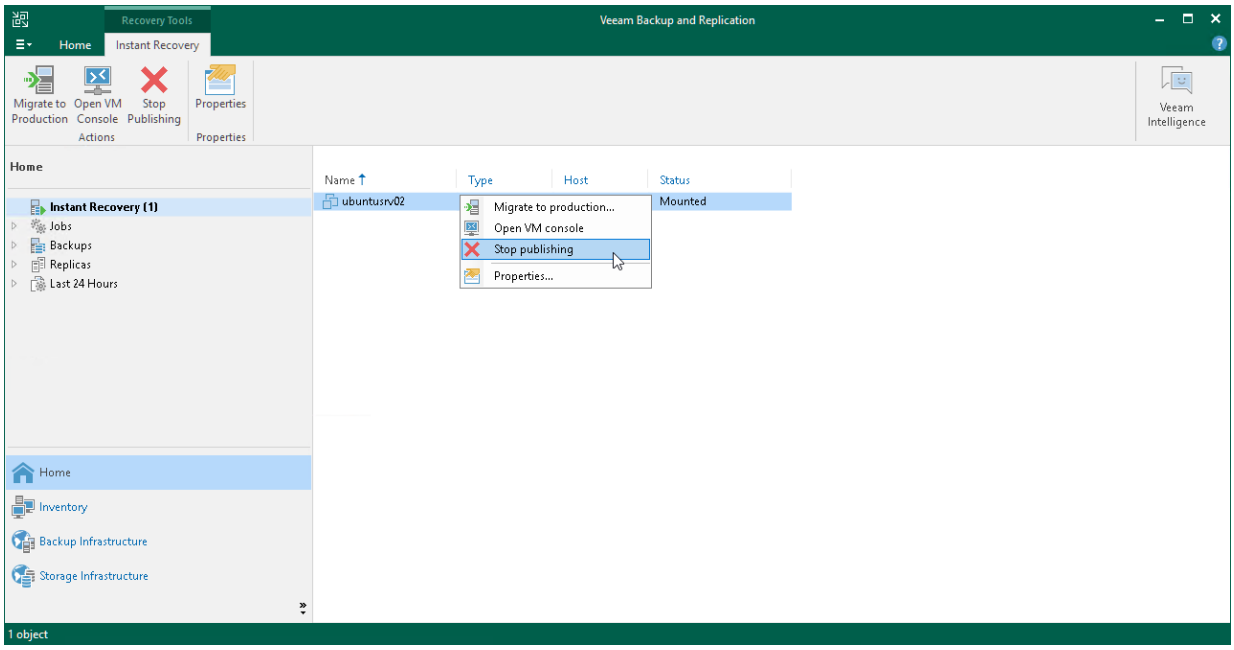

# Instant Recovery to Microsoft Hyper-V

With Instant Recovery to Microsoft Hyper-V, you can immediately recover different workloads (VMs, EC2 instances, physical servers and so on) as Microsoft Hyper-V VMs. Instant Recovery to Microsoft Hyper-V can be helpful, for example, if you want to migrate your infrastructure from one environment to another, or you want to recover your infrastructure in a matter of minutes but with limited performance.

During recovery, Veeam Backup & Replication runs workloads directly from compressed and deduplicated backup files. This helps improve recovery time objectives (RTO), minimize disruption and downtime of production workloads. The workloads are recovered in a matter of minutes.

When you perform Instant Recovery, Veeam Backup & Replication creates dummy VMs and mounts to VMs workload disks directly from backups stored on backup repositories. These dummy VMs have limited I/O performance. To provide full I/O performance, you must migrate the VMs to the production site. For more information, see Migration of Recovered VMs to Production Site.

Besides disaster recovery matters, Instant Recovery can also be used for testing purposes. Instead of extracting workloads to production storage to perform regular disaster recovery (DR) testing, you can run a workload directly from a backup file, boot it and make sure the guest OS and applications are functioning properly. For more information, see Finalizing Instant Recovery to Microsoft Hyper-V.

Instant Recovery supports bulk processing so you can immediately recover multiple workloads at once. If you perform Instant Recovery for several workloads, Veeam Backup & Replication uses the resource scheduling mechanism to allocate and use optimal resources required for Instant Recovery. For details, see Resource Scheduling.

## Supported Backup Types

You can recover workloads from the following types of backups:

- Backups of Microsoft Hyper-V virtual machines created by Veeam Backup & Replication
- Backups of VMware vSphere virtual machines created by Veeam Backup & Replication
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV
- Backups of Amazon EC2 instances created by Veeam Backup for AWS
- Backups of Microsoft Azure virtual machines created by Veeam Backup for Microsoft Azure
- Backups of Google Compute Engine VM instances created by Veeam Backup for Google Cloud
- Backups of oVirt VMs created by Veeam Backup for OLVM and RHV
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE

For details on how to restore Microsoft Hyper-V VMs and how restore works, see the Instant Recovery to Hyper-V section in the User Guide for Microsoft Hyper-V.

# How Instant Recovery Works

Instant Recovery is performed in the following way:

1. Veeam Backup & Replication reads the workload configuration from the backup file in the backup repository and creates a dummy VM with empty disks on the target host. Before creating the dummy VM, Veeam Backup & Replication analyzes the workload configuration and decides on the generation for this VM:
   - If Veeam Backup & Replication detects an EFI system partition, it creates a generation 2 VM.
   - If Veeam Backup & Replication detects BIOS boot partition, it creates a generation 1 VM.
   - If Veeam Backup & Replication detects at least one GPT partition, it creates a generation 2 VM.
   - In other cases, Veeam Backup & Replication creates a generation 1 VM.

   In other aspects, the created VM has the same settings as the workload in the backup file. If you select to preallocate disk space for the recovered VMs, Veeam Backup & Replication preallocates disk space at the beginning of the Instant Recovery process.
2. Veeam Backup & Replication initiates creation of a protective snapshot for the dummy VM and the VM is started. If the Instant Recovery process fails for some reason, the protective snapshot guarantees that no data is lost.
3. On the backup repository and on the target host, Veeam Backup & Replication starts a pair of Veeam Data Movers that are used to mount the VM disks from the backup file to the dummy VM.
4. On the target host, Veeam Backup & Replication starts a proprietary Veeam driver. The driver redirects requests to the file system of the recovered VM (for example, when a user accesses some application) and reads necessary data from the backup file in the backup repository using the pair of Veeam Data Movers that maintain the disk mount.

# Migration of Recovered VMs to Production Site

When you begin the migration process, Veeam Backup & Replication starts one more pair of Veeam Data Movers — one Veeam Data Mover on the backup repository and one on the target host. This pair of Veeam Data Movers copies data of the recovered VM from the backup repository to the target host in the background, and populates disks of the VM started on the target host.

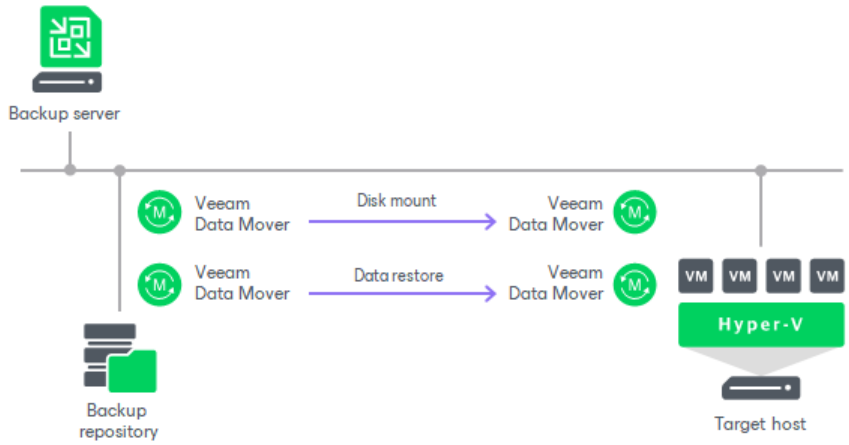


The driver on the target host knows which data has already been recovered permanently and does not redirect requests to such data, reading it directly from the disks of the recovered VM. Thus, performance of the instantly recovered VM will increase as more of the data is copied. When the VM is recovered completely, all Veeam Data Movers are stopped.

# Performing Instant Recovery to Microsoft Hyper-V

With Instant Recovery, you can recover different workloads from backups and register them as Microsoft Hyper-V VMs. For the list of backups that you can use for recovery, see Supported Backup Types.

To perform Instant Recovery, use the **Instant Recovery** wizard.

> **NOTE**
>
> If you want to recover workloads as VMware vSphere VMs, see Performing Instant Recovery to VMware vSphere.

## Before You Begin

Before you perform Instant Recovery to Microsoft Hyper-V, consider the following:

- You must add the Hyper-V target host to which you want to recover machines to your backup infrastructure.
- Make sure that the **Disable changed block tracking for this host** option is not selected for a host to which you plan to recover a workload. If this option is selected for the host, the driver required for work of Instant Recovery will be disabled. For more information, see the Configuring Connected Volumes section in the User Guide for Microsoft Hyper-V.
- You can recover a workload from a backup that has at least one successfully created restore point.
- If you recover a workload to the production network, make sure that the original workload is powered off to avoid conflicts.
- Veeam Backup & Replication restores disks of workloads other than Hyper-V in the dynamically expanded VHDX format. If the target host does not support the VHDX format, Veeam Backup & Replication uses the dynamically expanding VHD format.
- Consider the following for Linux workloads:
  - We strongly recommend having `dracut`, `mkinitrd` or `initramfs` installed on workloads that will be restored. Otherwise, they may not boot after restore.
  - Open the `/etc/fstab/` file and check that all file systems are mounted using UUID. If any filesystems are mounted using block device name, the restored VM may not boot.
- If you want to scan recovered VM data for viruses, check the secure restore requirements and limitations.
- On non-Microsoft Windows SMB3 storage, for example, Tintri, Veeam Backup & Replication may display the "*Failed to disable integrity bit on disk N*" warning during the restore process. You can ignore this warning for non-Microsoft Windows SMB3 storage.
- The recovered VM will have the same MAC address as the original workload. Therefore, if you recover the workload to the same Hyper-V host where the original workload is running, a MAC address conflict may occur. To overcome this situation, power off the original workload before you start the recovery process.
- [For Nutanix AHV VMs] The recovered VM will not be connected to a network. You must connect to the network manually.
- [For Nutanix AHV VMs, Amazon EC2 instances and Microsoft Azure virtual machines] Instantly recovered VM will have default virtual hardware settings: 2 CPU cores, 4GB RAM and one network adapter. If you want to change the default settings, turn off the VM and set the required virtual resources. Note that you must not switch off the instant recovery session before turning off the VM.

## Step 1. Launch Instant Recovery Wizard

To launch the **Instant Recovery to Hyper-V** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select one of the following:
  - 
  - **VMware vSphere** > **Restore from backup** > **Entire VM restore** > **Instant recovery to Microsoft Hyper-V** — if you want to recover VMware vSphere VMs from a VM backup created by Veeam Backup & Replication.
  - **Agent** > **Entire machine restore** > **Instant recovery** > **Instant recovery to Microsoft Hyper-V** — if you want to recover physical machines from a backup created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.
  - **AWS EC2 backup** > **Entire machine restore** > **Instant recovery** > **Instant recovery to Microsoft Hyper-V** — if you want to recover EC2 instances from a backup created by Veeam Backup for AWS.
  - **Azure IaaS** > **Entire machine restore** > **Instant recovery** > **Instant recovery to Microsoft Hyper-V** — if you want to recover Azure VMs from a backup created by Veeam Backup for Microsoft Azure.
  - **GCE backup** > **Entire machine restore** > **Instant recovery** > **Instant recovery to Microsoft Hyper-V** — if you want to recover VM instances from a backup created by Veeam Backup for Google Cloud.
  - **Nutanix backup** > **Entire machine restore** > **Instant recovery** > **Instant recovery to Microsoft Hyper-V** — if you want to recover VMs from backups created by Veeam Backup for Nutanix AHV.
  - **oVirt KVM** > **Entire machine restore** > **Instant recovery** > **Instant recovery to Microsoft Hyper-V** — if you want to recover VMs from backups created by Veeam Backup for OLVM and RHV.
  - **Proxmox VE** > **Entire machine restore** > **Instant recovery** > **Instant recovery to Microsoft Hyper-V** — if you want to recover VMs from backups created by Veeam Backup for Proxmox VE.

- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, select workloads that you want to recover and click **Instant Recovery** > **Microsoft Hyper-V** on the ribbon. Alternatively, you can right-click one of the selected workloads and select **Instant recovery** > **Microsoft Hyper-V**.

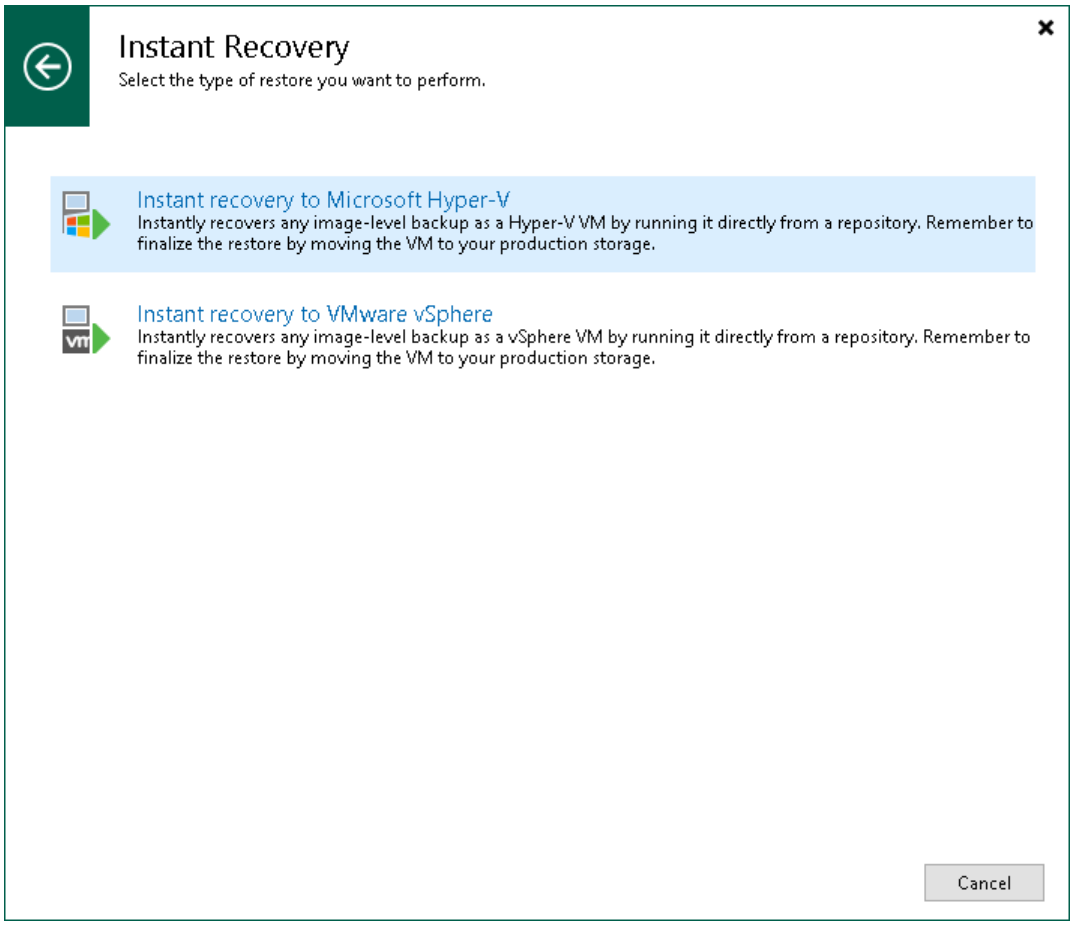

## Step 2. Select Workloads

At the **Machines** step of the wizard, select workloads that you want to recover:

1. Click **Add**.
2. In the **Backup Browser** window, do the following:
   a. [For VMware vSphere and VMware Cloud Director VMs] You can browse for necessary VMs using one of the following way:
      - **From infrastructure** – use this option to browse the virtual environment and select VMs or VM containers to recover. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list.
      - **From backup** – browse existing backups and select VMs under backup jobs.
   b. [For other workloads] In the list of backup jobs, expand a job and select workloads.
   c. Click **Add**.

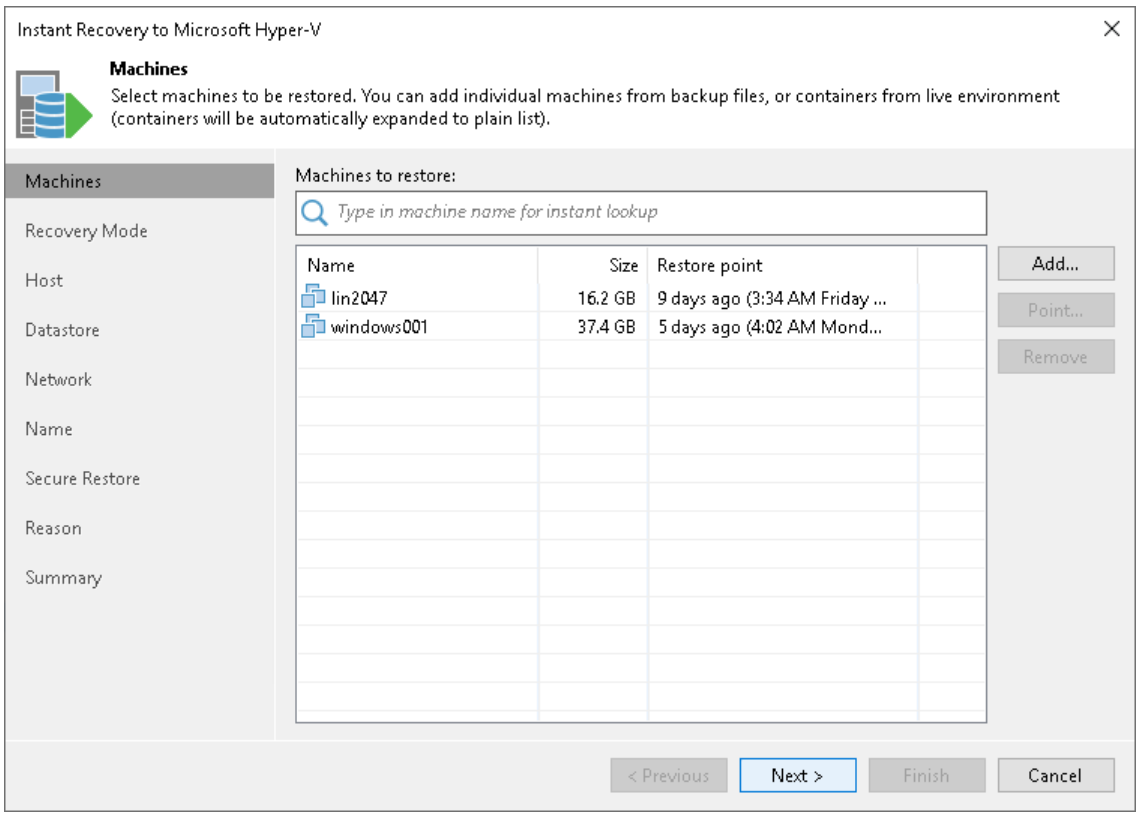

## Step 3. Select Restore Point

By default, Veeam Backup & Replication uses the latest valid restore points to recover the workloads. You can recover a workload to an earlier state, if necessary. If you have chosen to recover multiple workloads, you can select a restore point for each workload in the list.

To select a restore point:

1. In the **Machines to restore** list, select a workload.
2. Click **Point** on the right.
3. In the **Restore Points** window, select a restore point from which you want to recover the workload.

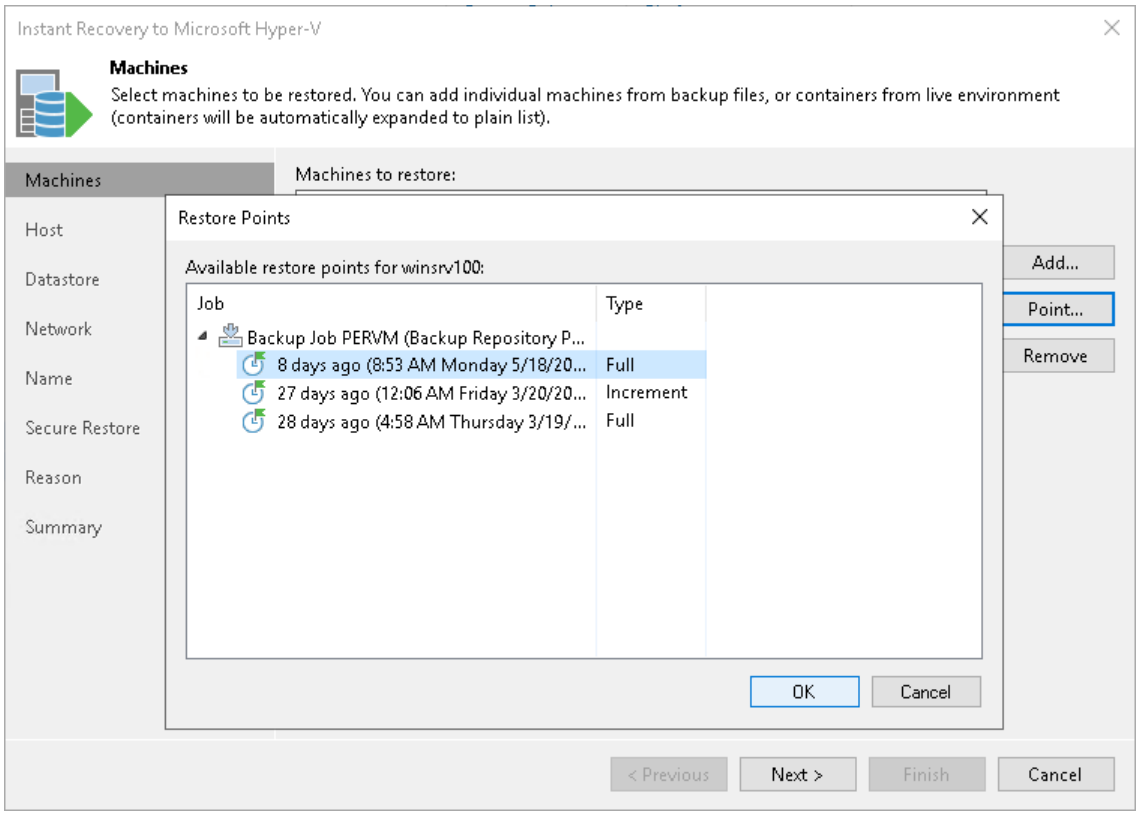

# Step 4. Select Recovery Mode

This step is available only if you recover Hyper-V VMs.

At the **Recovery Mode** step of the wizard, choose the necessary restore mode:

- Select Restore to the original location if you want to recover VMs with initial settings and to original location. If this option is selected, you will pass directly to the Reason step of the wizard.
- Select **Restore to a new location, or with different settings** if you want to recover VMs to a different location and with different settings (such as location, network settings and so on). If this option is selected, the **Instant Recovery to Hyper-V** wizard will include additional steps for customizing VM settings.

> **IMPORTANT**
>
> If you recover a VM to the original location, consider the following:
>
> - If the original VM still exists in the virtual infrastructure, the VM and its disks will be removed. Make sure that other VMs in the virtual environment do not use these disks.
> - The VM settings contain the ID of the VM group to which the machine belongs. To recover the VM to the original VM group, you must not delete the original VM group or change the hierarchy of its parent VM groups.

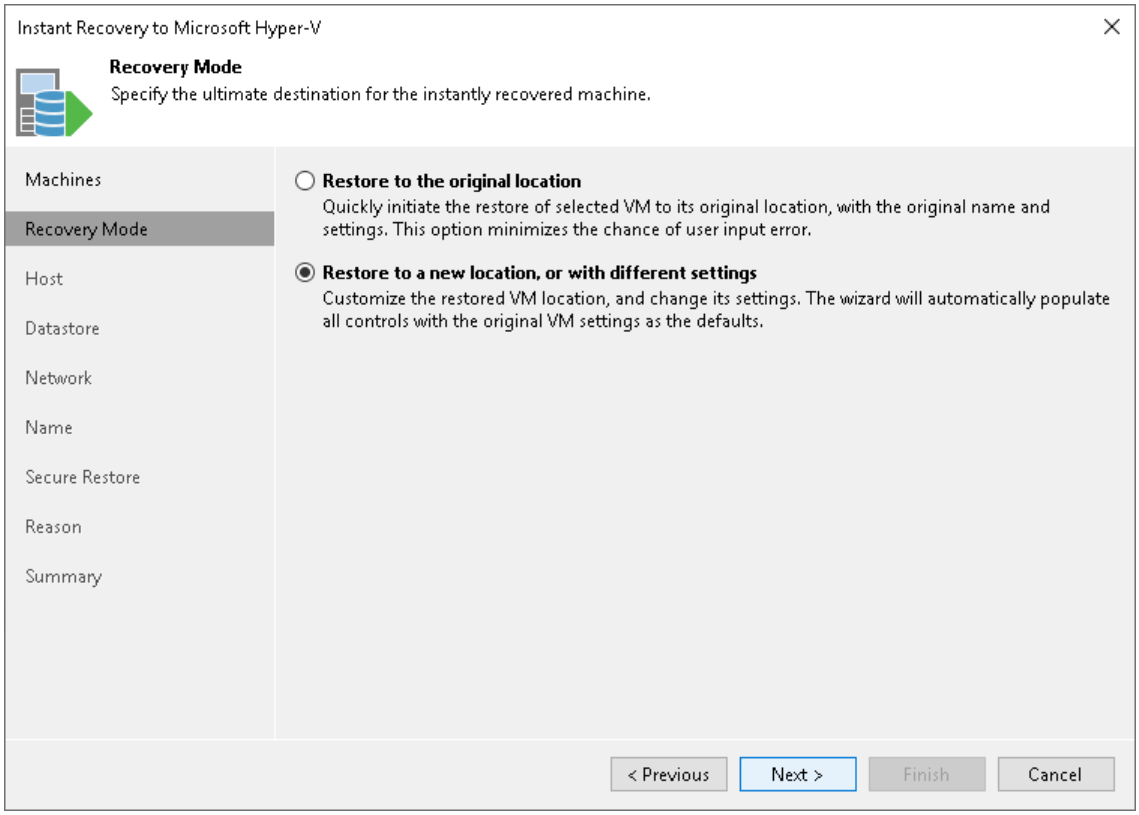

# Step 5. Select Target Host

This step is available if you recover workloads to a new location or with different settings.

At the **Host** step of the wizard, specify a target Hyper-V host or a Hyper-V cluster where you want to locate the recovered VMs:

1. In the **VM location** list, select the necessary VMs and click **Host**.
2. In the **Select Server** window, select a standalone host or cluster where the selected VM will be registered.
3. If you have selected a Hyper-V cluster, you can specify the cluster resource settings. Click **Resource** and select one of the following options in the **Cluster Resource Settings** window:
   - **Register VM as a cluster resource** – if you want to assign a cluster role to the recovered VM.
   - **Do not register VM as a cluster resource** – if you do not want to assign a cluster role to the recovered VM.

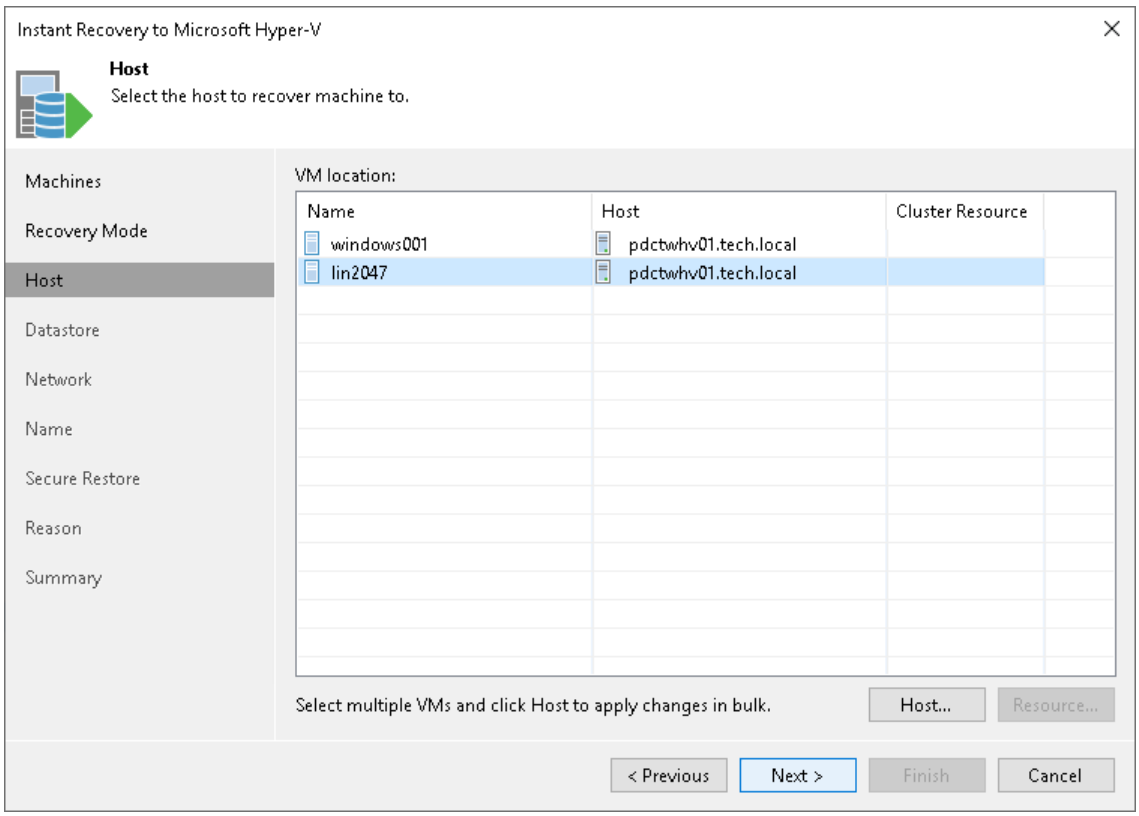

# Step 6. Select Target Datastore

This step is available if you recover workloads to a new location or with different settings.

At the **Datastore** step of the wizard, specify a path to the folder where VM configuration files and disks will be stored:

1. In the **Files location** list, select the workloads that will be recovered to the same host and click **Path**. Alternatively, you can expand a workload in the list and select individual files. Use this method if you want to place configuration and disk files to different locations.
2. In the **Select Folder** window, do one of the following:
   - Select an existing folder where VM files will be stored.
   - Create a new folder by clicking **New Folder** at the bottom of the window.
   - Type a path to an SMB3 shared folder in the search field at the bottom of the **Select Folder** window. The path must be specified in the UNC format, for example: *||172.16.11.38|Share01*.

     > **IMPORTANT**
     >
     > The host or cluster on which you register VMs must have access to the specified SMB3 shared folder. If you are using SCVMM 2012 or later, the server hosting the Microsoft SMB3 shared folder must be registered in SCVMM as a storage device. For more information, see Microsoft Docs.

3. Check that the **Allocate required disk space before migration** check box is selected if you want to preallocate disk space required for the recovered VM. Otherwise, clear the check box.

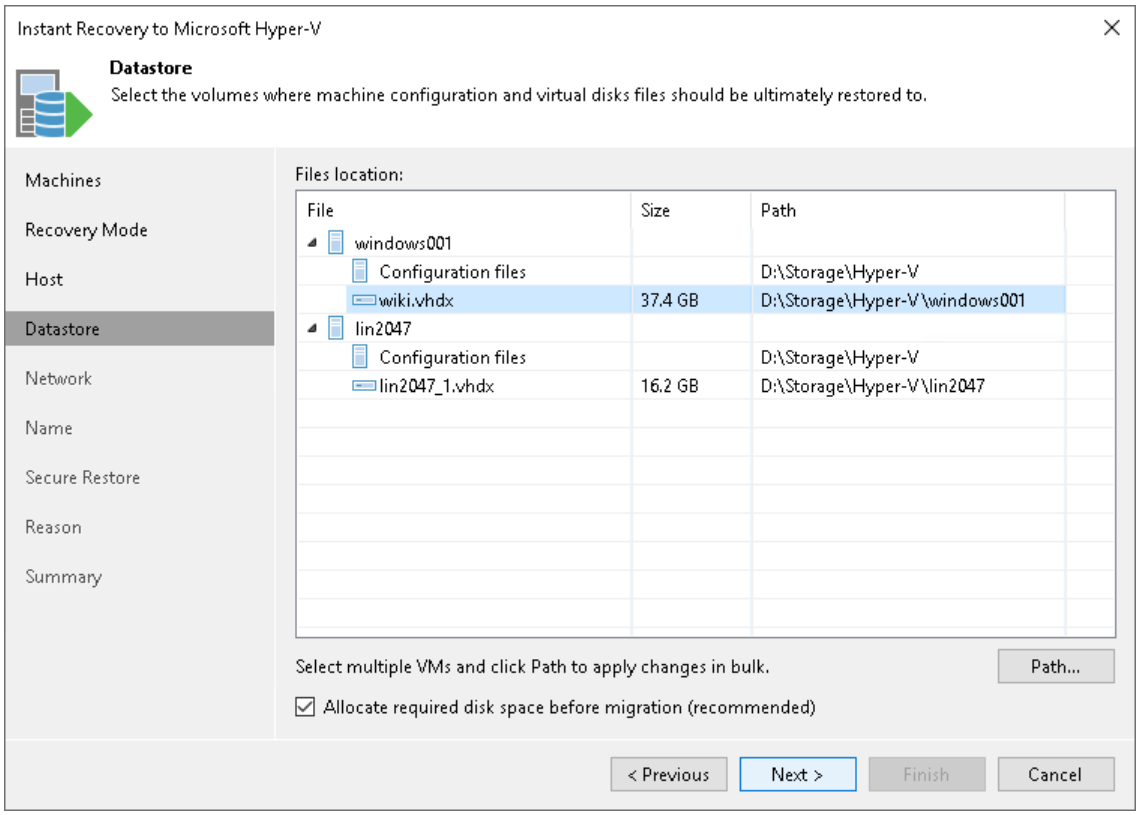

# Step 7. Specify Network Mapping

This step is available if you recover workloads to a new location or with different settings.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the original site to networks in the target site (site where VMs will be recovered). When the job starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update VM configuration files to replace the original networks with the specified networks in the target site. As a result, you will not have to re-configure network settings manually.

To specify network mapping settings:

1. In the **Network connections** list, select the necessary workloads and click **Network**.

   If a workload is connected to multiple networks, expand the workload, select a network to map and click **Network**.

2. In the **Select Network** window, select a network to which the selected workload must be connected after recovery.

If you do not want to connect the recovered VM to any virtual network, select the original workload in the list and click **Disconnect**.

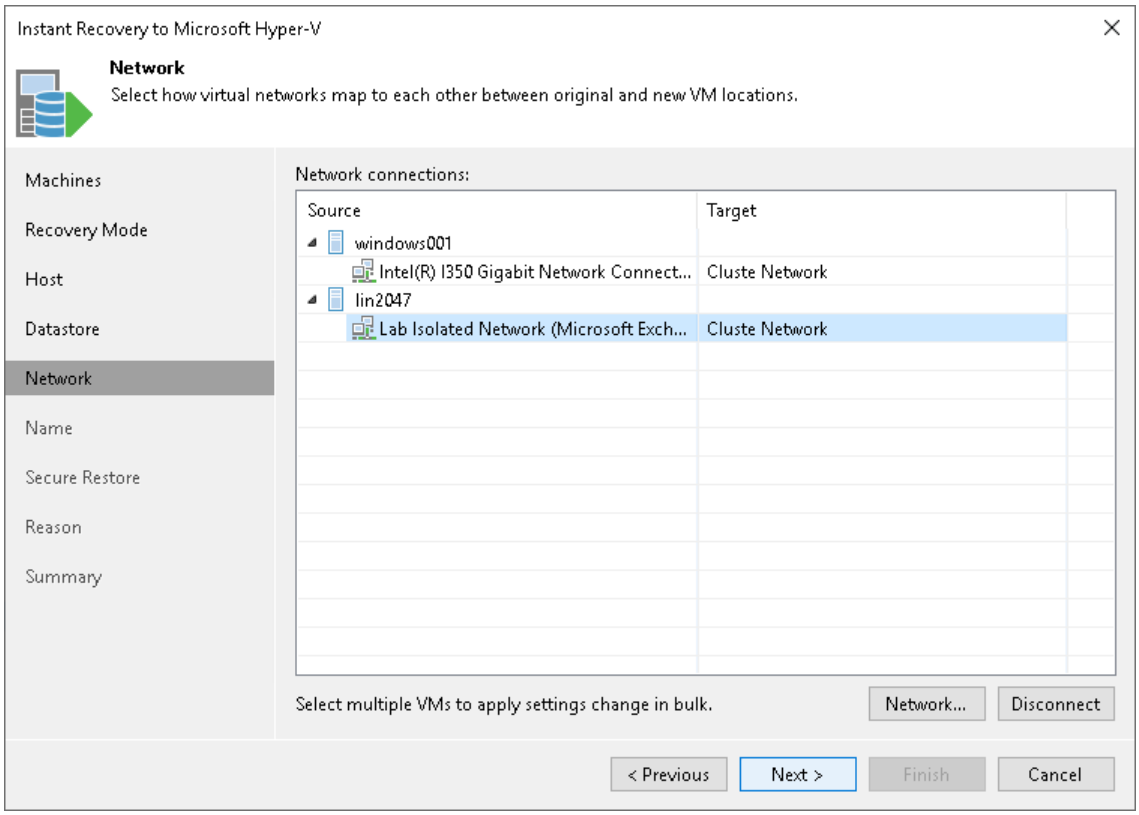

## Step 8. Change VM Name and UUID

This step is available if you recover workloads to a new location or with different settings.

At the **Name** step of the wizard, specify names under which VMs will be recovered and select whether you want to preserve VM UUIDs or change them. By default, Veeam Backup & Replication preserves the original names and UUIDs.

> **NOTE**
>
> We recommend that you specify a new name and generate a new UUID for a VM to prevent conflicts if the original workload still resides in the production environment. The name and UUID change is not required if the original workload no longer exists, for example, it was permanently deleted.

# Changing Names

To change a VM name:

1. In the **Virtual machines** list, select the necessary workloads and click **Name**.
2. In the **Change Name** section, enter a new name explicitly or specify a change name rule by adding a prefix or suffix to the original workload name.

Alternatively, you can change a VM name directly in the **Virtual machines** list. To do this, click the **New Name** field and enter the name to be assigned to the recovered VM.

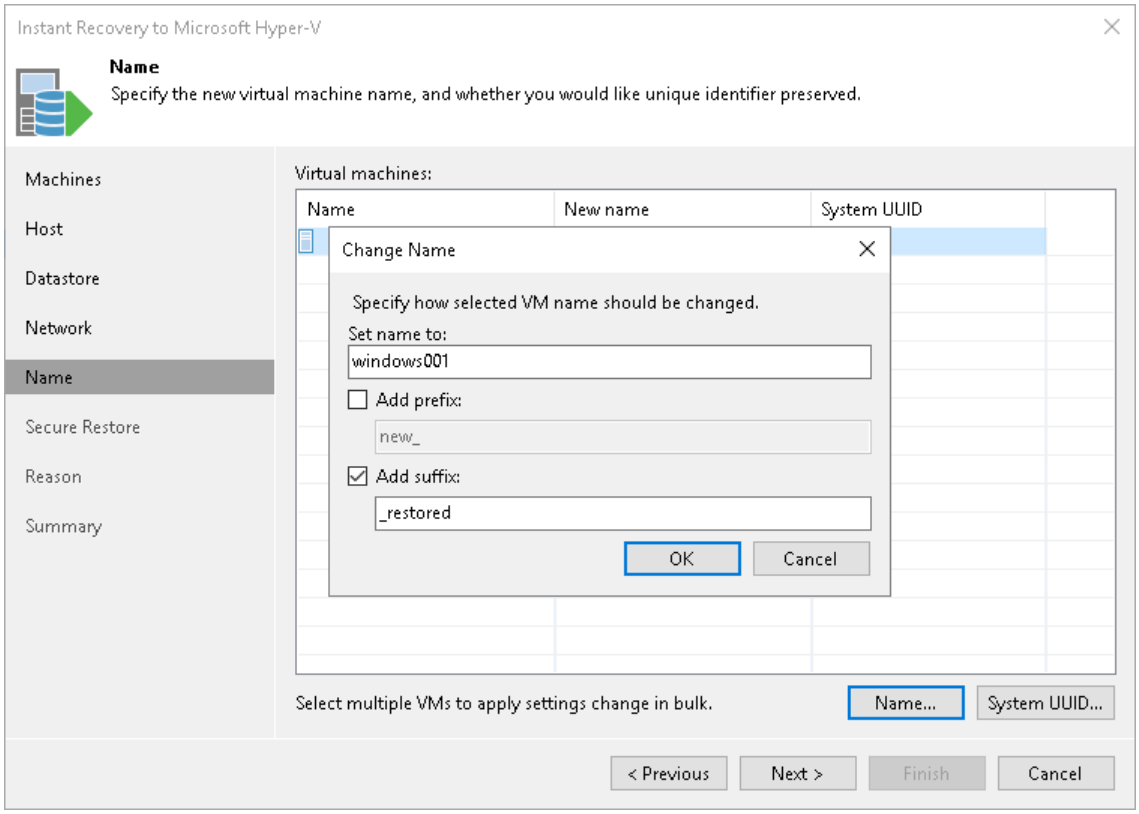


# Changing UUIDs

To change VM identification settings:

1. Select the necessary workloads in the list and click **System UUID**.

2. In the **BIOS UUID Settings** window, select to generate a new UUID.

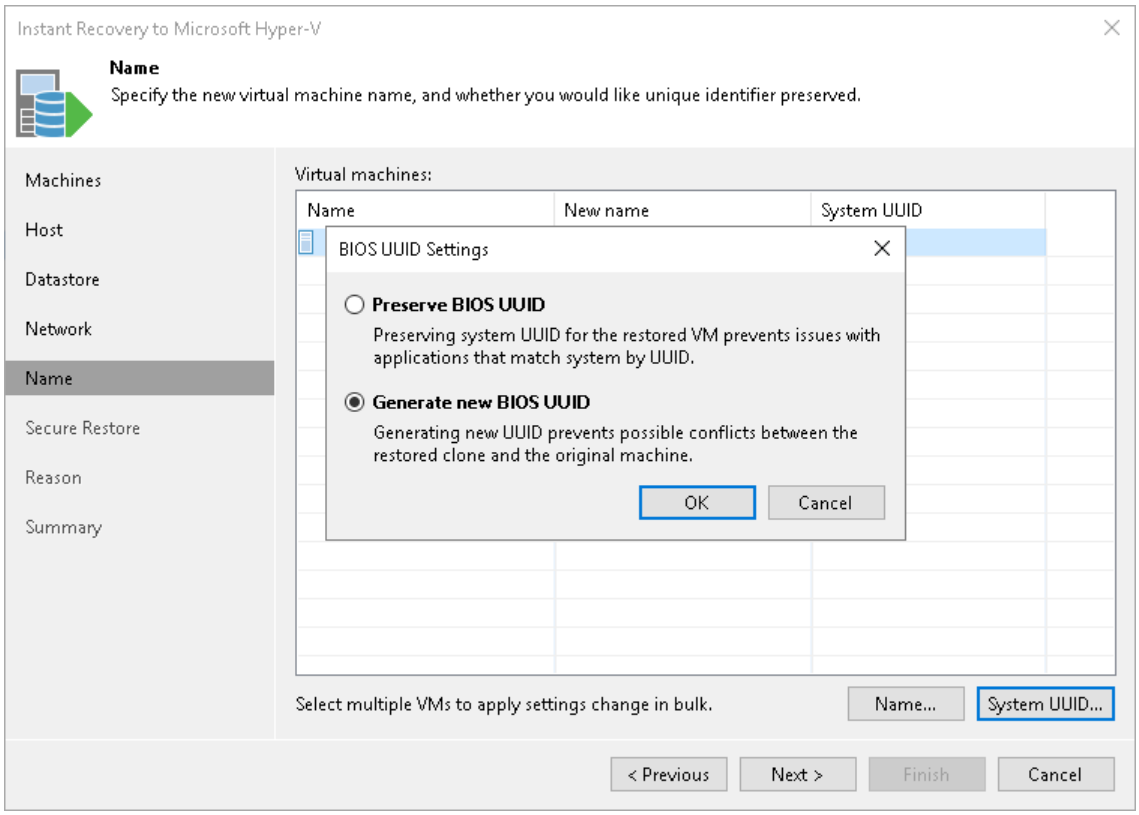

## Step 9. Configure Helper Appliance

This step is available if you recover workloads with Linux OS and recover them to a new location or with different settings.

Veeam Backup & Replication recovers Linux machines using a helper appliance. The helper appliance is an auxiliary Linux-based VM registered by Veeam Backup & Replication. The appliance is quite small — around 100 MB. It requires 1024 MB RAM and takes around 10 seconds to boot.

At the **Helper Appliance** step of the wizard, configure the helper appliance network settings:

1. [For multiple machines] In the **Network** list, expand a host and select one machine for which you want to configure the helper appliance.
2. Click the **Configure** button.
3. In the **Network Settings** window, select a network for the helper appliance.
   a. Click the **Browse** button to the right of the **Production network** field.
   b. In the **Select Network** window, Veeam Backup & Replication shows a list of networks to which the target host is connected. In this list, select a network to which the helper appliance must be connected. Click **OK**.

      Consider that the backup server and the mount server must have access to the helper appliance over the network.
   c. In the **VLAN ID** field, specify an ID of a VLAN where the helper appliance will reside.

      The 0 value means that the VLAN is not set.
4. Specify IP addressing settings for the helper appliance and DNS server:
   a. Click **Configure**.
   b. Switch to the **IPv4** or **IPv6** tab depending on which addresses you want to configure. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
   c. Select the **Enable IPv4/IPv6 interface** check box.
   d. Configure IP settings for the helper appliance:
      - If you use a DHCP server in the network and want to obtain the IP address automatically, leave the **Obtain an IP address automatically** option selected.
      - To manually assign a specific IP address to the helper appliance, click **Use the following IP address** and specify the IP address settings.
   e. Configure IP settings for the DNS server:
      - If you use a DHCP server in the network and want to obtain the IP address automatically, leave the **Obtain DNS server address automatically** option selected.
      - To manually assign a specific IP address to the DNS server, click **Use the following DNS server address** and specify preferred and alternate addresses.

f. Click **OK**.

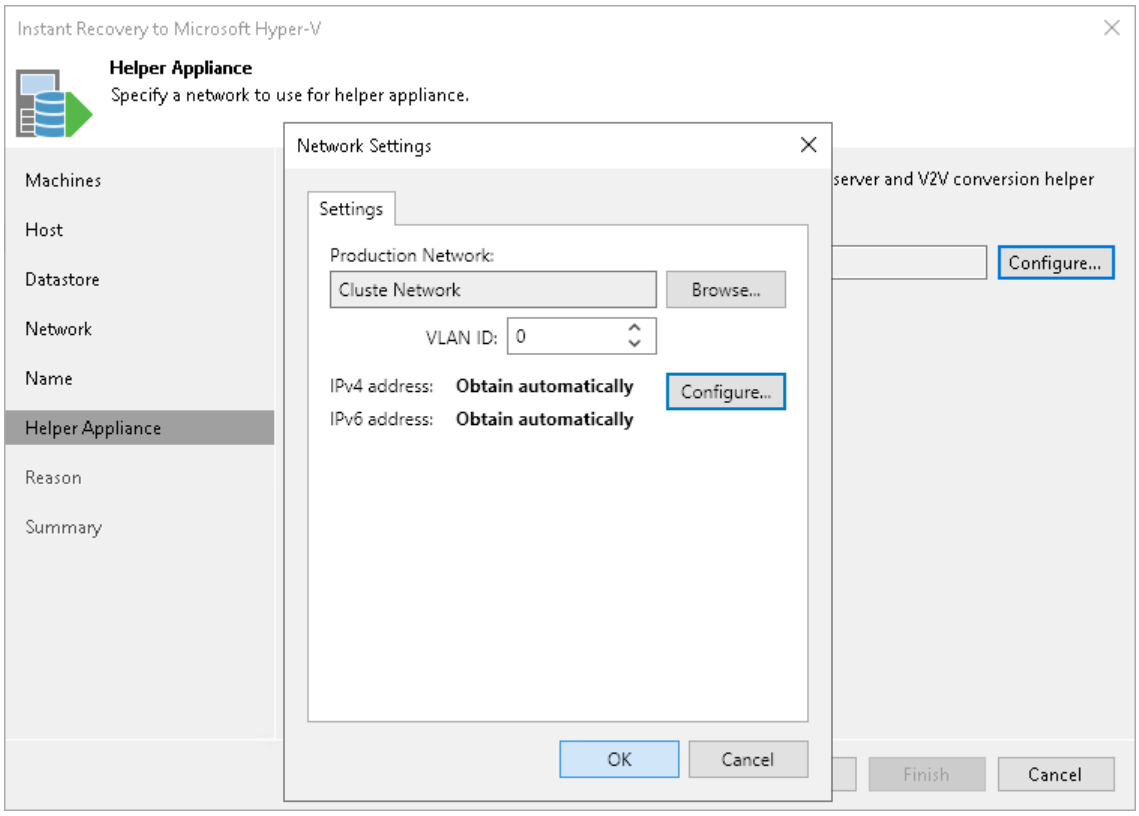
Instant Recovery to Microsoft Hyper-V
Helper Appliance
Specify a network to use for helper appliance.
Machines
Host
Datastore
Network
Name
Helper Appliance
Reason
Summary
server and V2V conversion helper
Configure...
Finish
Cancel
Network Settings
Settings
Production Network:
Cluste Network
Browse...
VLAN ID:
0
IPv4 address:
Obtain automatically
IPv6 address:
Obtain automatically
Configure...
OK
Cancel

# Step 10. Specify Secure Restore Settings

This step is available if you recover workloads with Microsoft Windows OS.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Select which action Veeam Backup & Replication will take if scan finds a virus threat:
   - **Proceed with recovery but disable network adapters** — if you want to restore the VM with disabled network adapters (NICs).
   - **Abort VM recovery** — if you want to cancel the restore session.

3. Select the **Scan the entire image** check box if you want to continue the VM data scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Antivirus Scan Results.

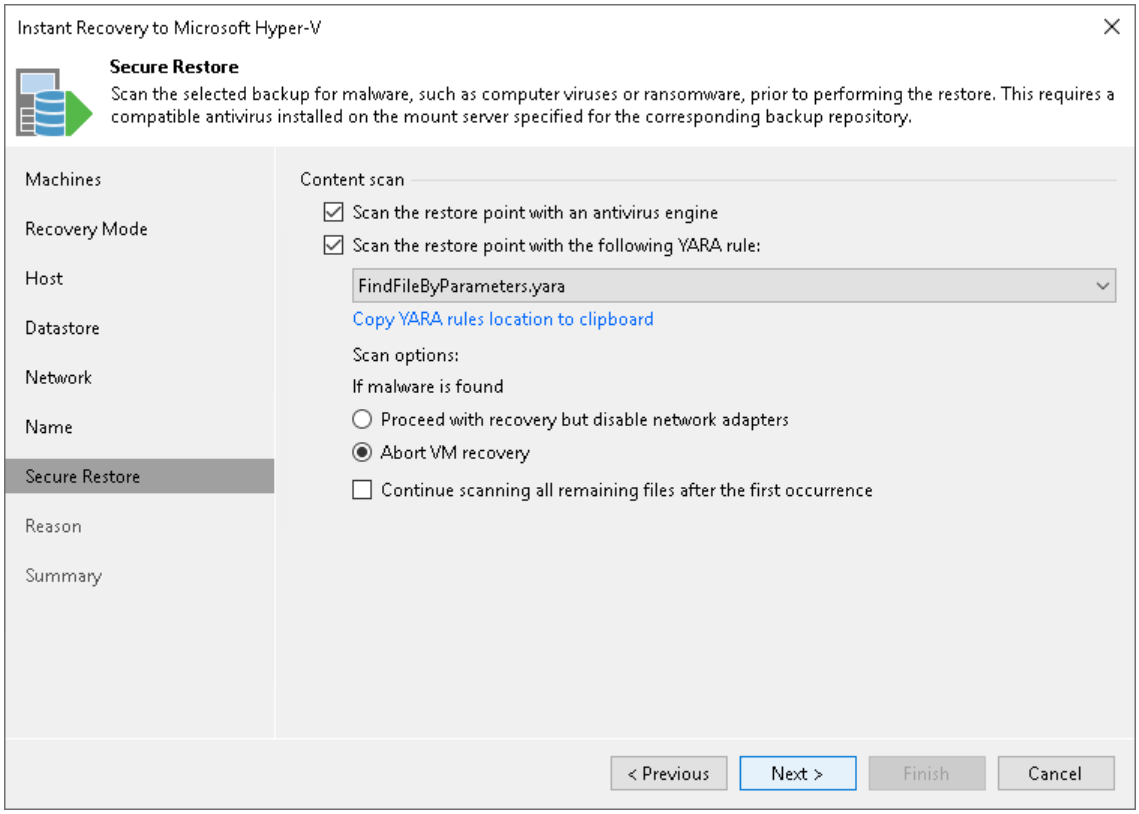

## Step 11. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for performing Instant Recovery of the workloads. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

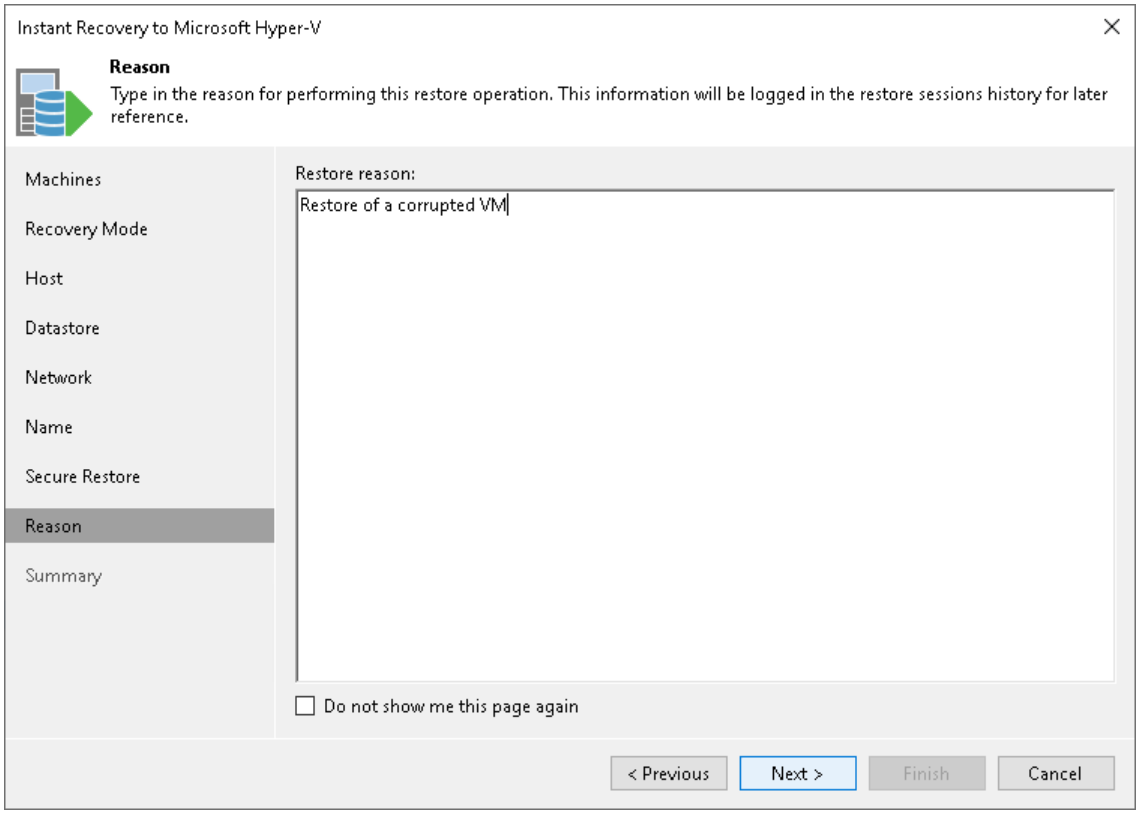

## Step 12. Verify Instant Recovery Settings

At the **Summary** step of the wizard, check settings of Instant Recovery and click **Finish**. If you want to start the recovered VMs on the target host, select the **Power on target VM after restoring** check box.

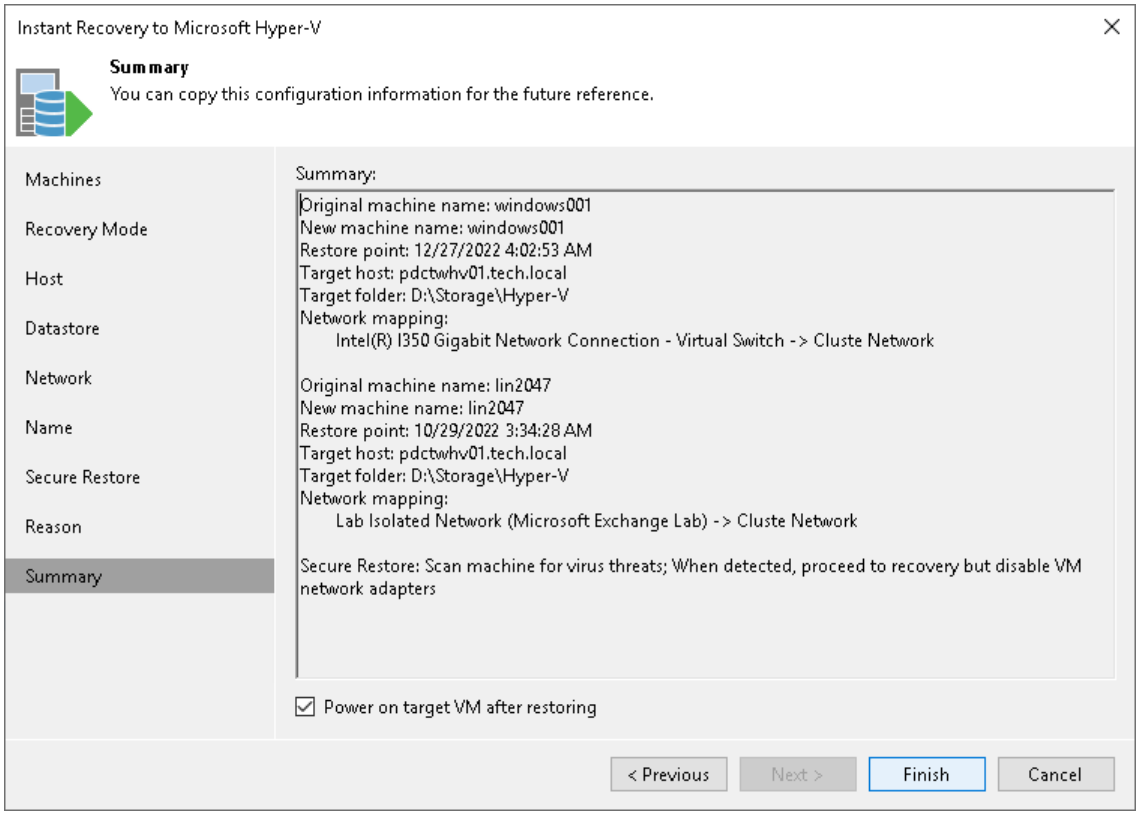


# What You Do Next

Finalizing Instant Recovery to Microsoft Hyper-V

# Finalizing Instant Recovery to Microsoft Hyper-V

After the VMs have been successfully recovered, you must finalize the process. For this, test the recovered VMs and decide whether to migrate them to production environment or stop publishing.

# Testing Recovered VMs

To test the recovered VMs before you migrate them to production, you can launch VM consoles from Veeam Backup & Replication or open the consoles in the Hyper-V client.

To launch a VM console from Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.

3. In the working area, right-click a VM and select **Open VM console**.

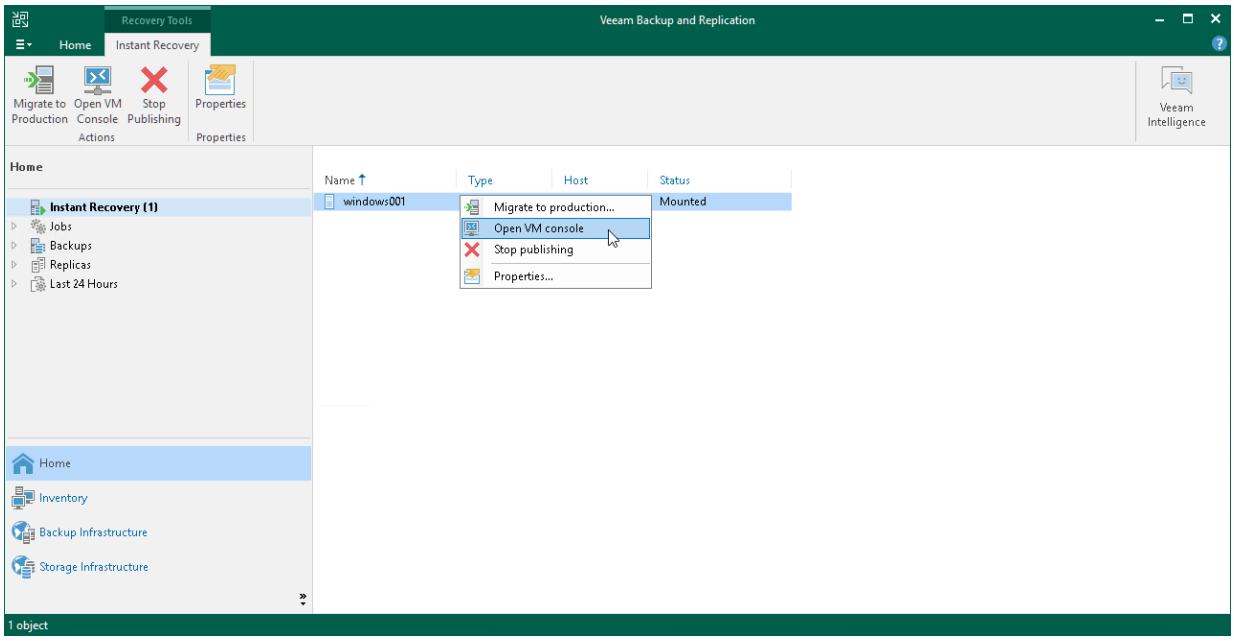


If the test fails, you can stop publishing the recovered VMs. For details, see Stop Publishing Recovered VMs.

# Migrating Recovered VMs

When Veeam Backup & Replication migrates VMs, it transfers VM disks data to the production storage that you have selected as a destination for the recovered VMs.

To migrate a recovered VM to production:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click a VM and select **Migrate to production**.

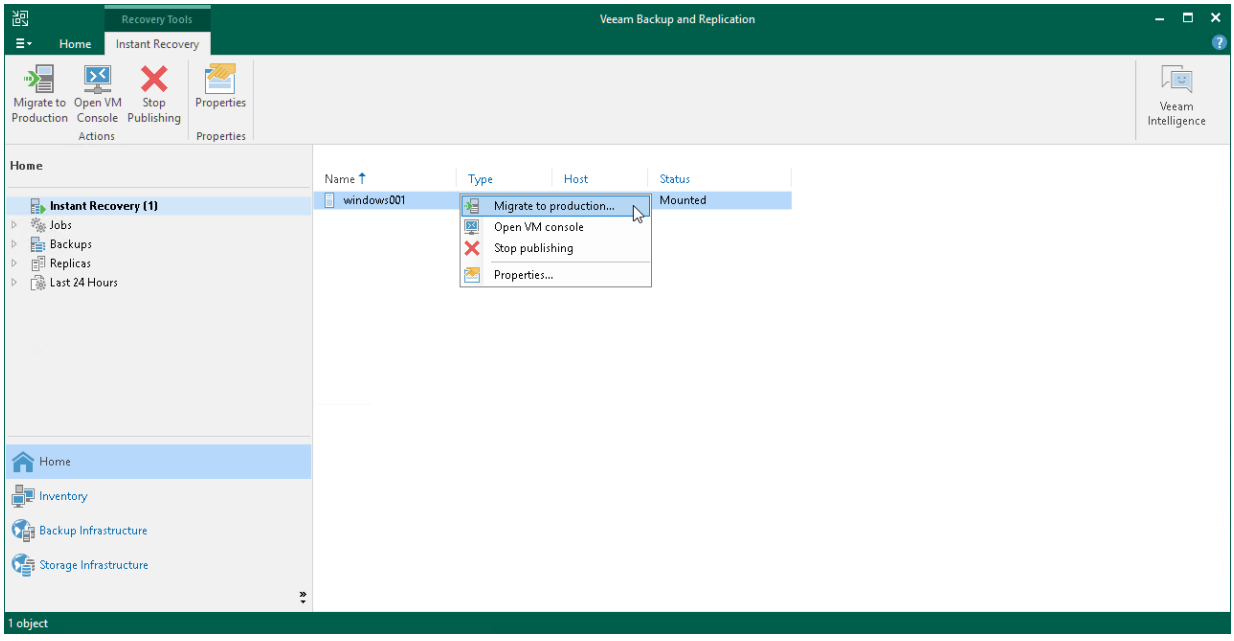

# Stop Publishing Recovered VMs

If you have ensured that the VM is working and you do not need it anymore, or your tests have failed, you can stop publishing the recovered VMs. This will remove the recovered VMs from the storage that you selected as the destination for recovery. Note that all changes made in the recovered VMs will be lost.

To remove a recovered VM:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click a VM and select **Stop publishing**.

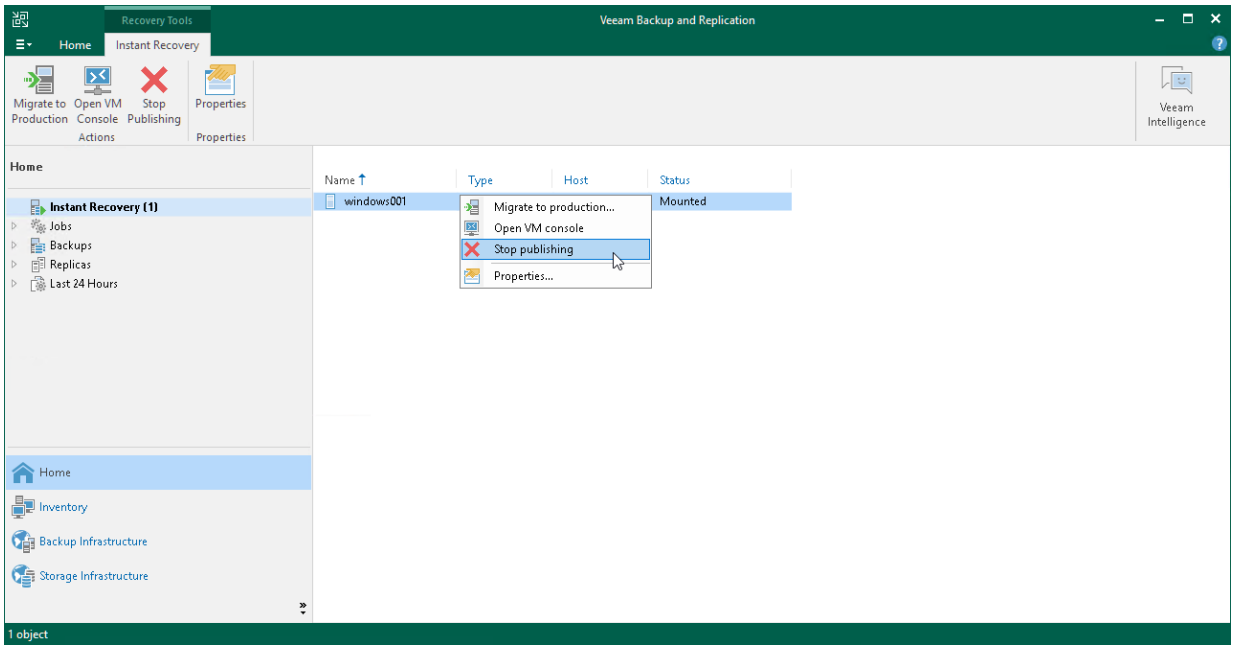

# Instant Recovery to Nutanix AHV

With Instant Recovery to Nutanix AHV, you can immediately recover different workloads as Nutanix AHV VMs. You can use Instant Recovery to Nutanix AHV to migrate the entire infrastructure or individual VMs to a Nutanix AHV cluster.

> **IMPORTANT**
>
> To restore to Nutanix AHV, you must install Nutanix AHV Plug-in on the backup server. To learn more, see the Installation section in the Veeam Backup for Nutanix AHV User Guide.

The procedure of restore to Nutanix AHV practically does not differ from the procedure described in the Performing Instant Recovery of Workloads to Nutanix AHV section in the Veeam Backup for Nutanix AHV User Guide.

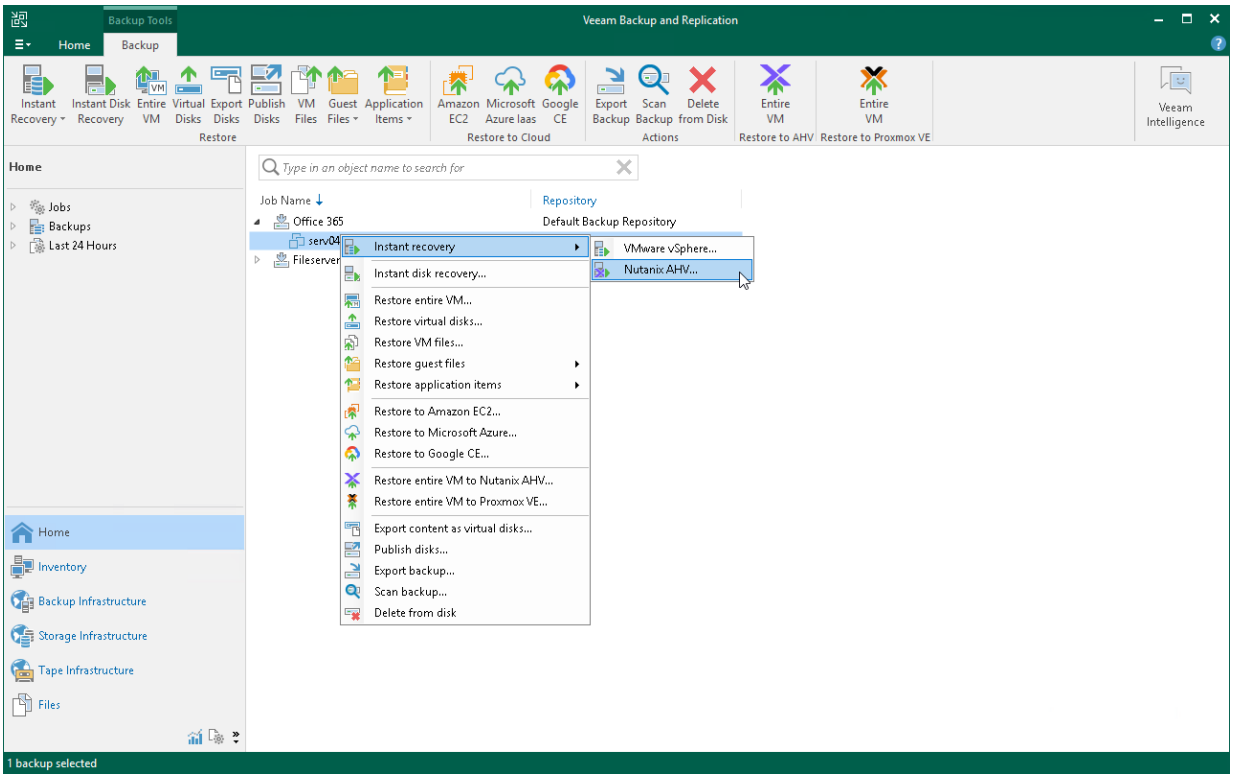

# Entire VM Restore

With Veeam Backup & Replication, you can restore an entire VM from a backup file to the latest state or to a previous point in time if the original VM fails.

When you restore an entire VM, Veeam Backup & Replication extracts the VM image from a backup to the production storage.Then Veeam Backup & Replication pulls the VM data from the backup repository to the selected storage, registers the VM on the chosen ESXi host and, if necessary, powers it on.

Entire VM restore enables full disk I/O performance while Instant Recovery provides a "temporary spare" for a VM as the vPower NFS throughput is limited.

## How Entire VM Restore Works

A VM can be restored to its original location or to a new location. When you restore a VM to its original location, Veeam Backup & Replication powers off the original VM and restores only those disks that are included in the backup file. All other disks remain unchanged.

When you restore a VM to a new location, you can specify new VM settings such as the new VM name, the host and datastore where the VM will reside, disk format (thin or thick provisioned) and network properties. Veeam Backup & Replication will change the VM configuration file and store the VM data to the location of your choice.

To perform entire VM restore, Veeam Backup & Replication uses one of the following transport modes:

- If the backup proxy is connected directly to the SAN fabric or has access to NFS datastores, Veeam Backup & Replication uses the Direct storage access transport mode. Veeam Data Movers, which are deployed on the backup repository and backup proxy, retrieve VM data from the backup file and put it directly to the necessary datastore.

  > **IMPORTANT**
  >
  > Veeam Backup & Replication can restore only thick VM disks using the Direct SAN access transport mode. For thin VM disks restore, Veeam Backup & Replication will use the Direct NFS access, Virtual appliance or Network transport modes. Alternatively, you can instruct Veeam Backup & Replication to restore VM disks as thick.

- If the backup proxy is virtualized and resides on the ESXi host to which the VM must be restored, Veeam Backup & Replication uses the Virtual appliance transport mode. The Virtual appliance mode utilizes VMware ESXi capabilities of HotAdding disks to the VM and thus eliminates the need to transfer the backup data across the network. Veeam Data Movers deployed on the backup repository and backup proxy retrieve VM data from the backup file and put it directly to the necessary datastore through the ESXi I/O stack.

- If the Direct storage access and Virtual appliance transport modes cannot be used, Veeam Backup & Replication uses the Network transport mode.

When Veeam Backup & Replication restores VMs, it performs a cyclic redundancy check (CRC) for the TCP traffic going between Veeam Data Movers installed on the backup repository and backup proxy. This happens provided that the multithreaded data transfer is enabled in the network traffic settings. Veeam Backup & Replication calculates and compares checksums for data blocks going from the source Veeam Data Mover and data blocks received on the target Veeam Data Mover. If the CRC check fails, Veeam Backup & Replication automatically re-sends data blocks without any impact on the restore job.

> **NOTE**
>
> If a VM has several VM disks, Veeam Backup & Replication restores VM disks in parallel.

# Quick Rollback

When you restore a full VM or VM hard disk to the original location, you can instruct Veeam Backup & Replication to perform quick rollback — incremental data restore. Instead of restoring an entire VM or VM disk from a backup file, Veeam Backup & Replication will recover only those data blocks that are necessary to revert the VM or VM disk to an earlier point in time. Quick rollback significantly reduces the recovery time and has little impact on the production environment.

For quick rollback, Veeam Backup & Replication uses the Changed Block Tracking technology. Veeam Backup & Replication gets information about the current VM state and compares it with the CBT information in the backup file. This way, Veeam Backup & Replication detects what data blocks must be transported back to the production datastore to rebuild the VM or VM disk to the necessary point in time.

It is recommended that you use quick rollback if you restore a VM or VM disk after a problem that has occurred at the level of the VM guest OS — for example, there has been an application error or a user has accidentally deleted a file on the VM guest OS. Do not use quick rollback if the problem has occurred at the VM hardware level, storage level or due to a power loss.

# Considerations and Limitations for Quick Rollback

- 
- VM or VM disk must be restored to its original location.
- CBT must be enabled for the VM and its disks.
- Quick rollback can be performed in the Direct NFS access, Virtual appliance, Network transport mode. The Direct SAN access transport mode cannot be used for quick rollback due to VMware limitations.
- After quick rollback, CBT will be disabled for the restored VM.
- Use quick rollback and VM guest OS file exclusion wisely. If you exclude specific files and folders from the VM guest OS during backup and use quick rollback to restore the VM or VM disk from such backup, Veeam Backup & Replication will restore only the content of the backup file. The excluded data will not be restored. For example, if you exclude `C:\Folder` from the backup, data in this folder will not be backed up and will not be available in the resulting backup file. After some time, data in `C:\Folder` may change but the folder will still not be backed up (since the job excludes this folder). For this reason, when you perform quick rollback, Veeam Backup & Replication will restore all data that have changed except the excluded `C:\Folder`.

# Restoring Entire VMs

To restore an entire VM, use the **Full VM Restore** wizard.

## Before You Begin

Before you restore a VM from a backup, consider the following:

- You can restore a VM from a backup that has at least one successfully created restore point.
- If you want to scan VM data for viruses, check the secure restore requirements and limitations.

- If you want to run an executable script for a VM, check the staged restore requirements and limitations.
- If you back up a VM with vRDM disks, Veeam Backup & Replication converts the disks into VMDK files. Thus, when you restore a VM with a vRDM disk, Veeam Backup & Replication restores this disk as a VMDK file. If you want to preserve the vRDM format for restored disks, use Quick Rollback. For more information, see Quick Rollback.
- [For restore to original location] If the original VM has vSphere Fault Tolerance (FT) enabled, you need to reenable FT manually after restore.
- When you restore a VM, consider the Virtual Hardware version compatibility. For more information, see this VMware KB article.

## Step 1. Launch Full VM Restore Wizard

To launch the **Full VM Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Entire VM restore** > **Entire VM restore**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and select the machine that you want to restore and click **Entire VM** on the ribbon. Alternatively, right-click the machine that you want to restore and select **Restore entire VM**.
- Double-click the VBK or VBM file (for example, in Microsoft Windows Explorer). In the displayed window, select the VM and click **Restore** > **Entire VM**.

  You can use this option if you perform restore on the backup server. You cannot use this option if you perform restore remotely over the Veeam Backup & Replication console.

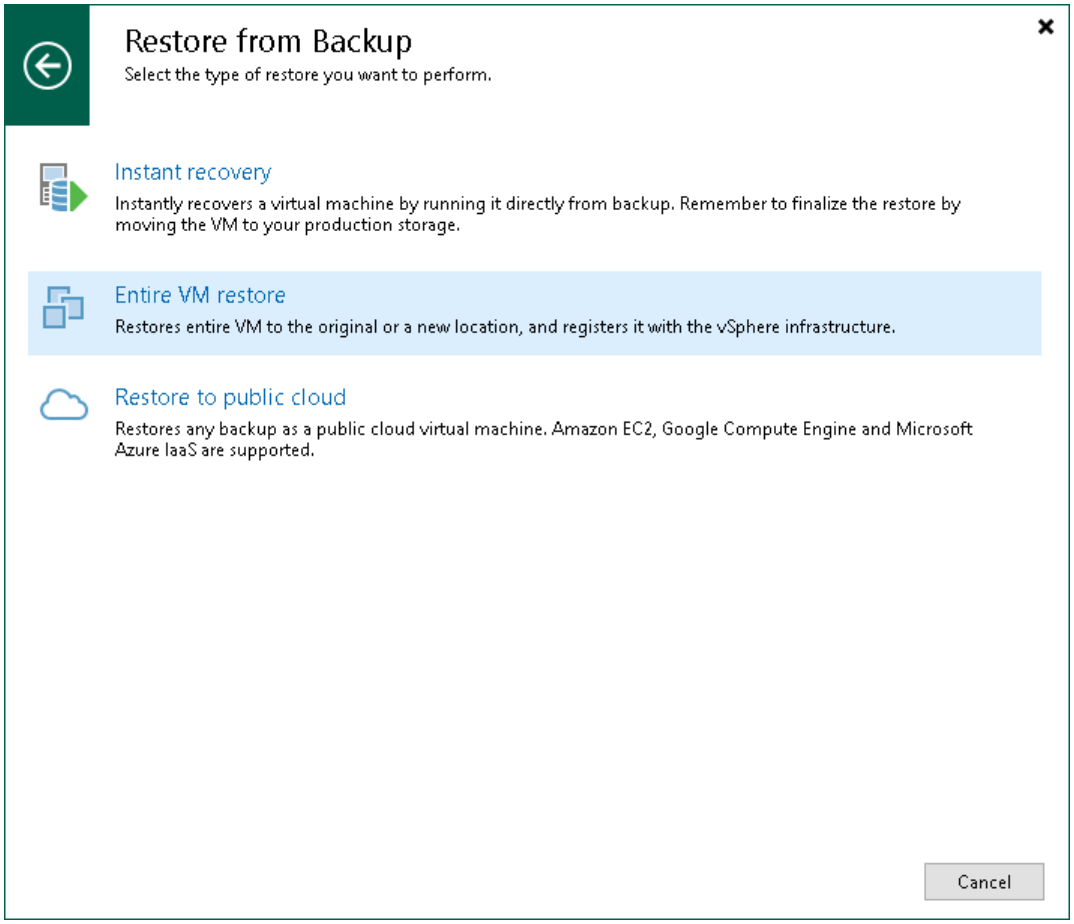

## Step 2. Select VMs

At the **Virtual Machines** step of the wizard, select VMs that you want to restore:

1. Click **Add VM**.
2. Select where to browse for VMs:
   - **From infrastructure** – browse the virtual environment and select VMs or VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags) to restore. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list.
   - **From backup** – browse existing backups and select VMs under backup jobs.

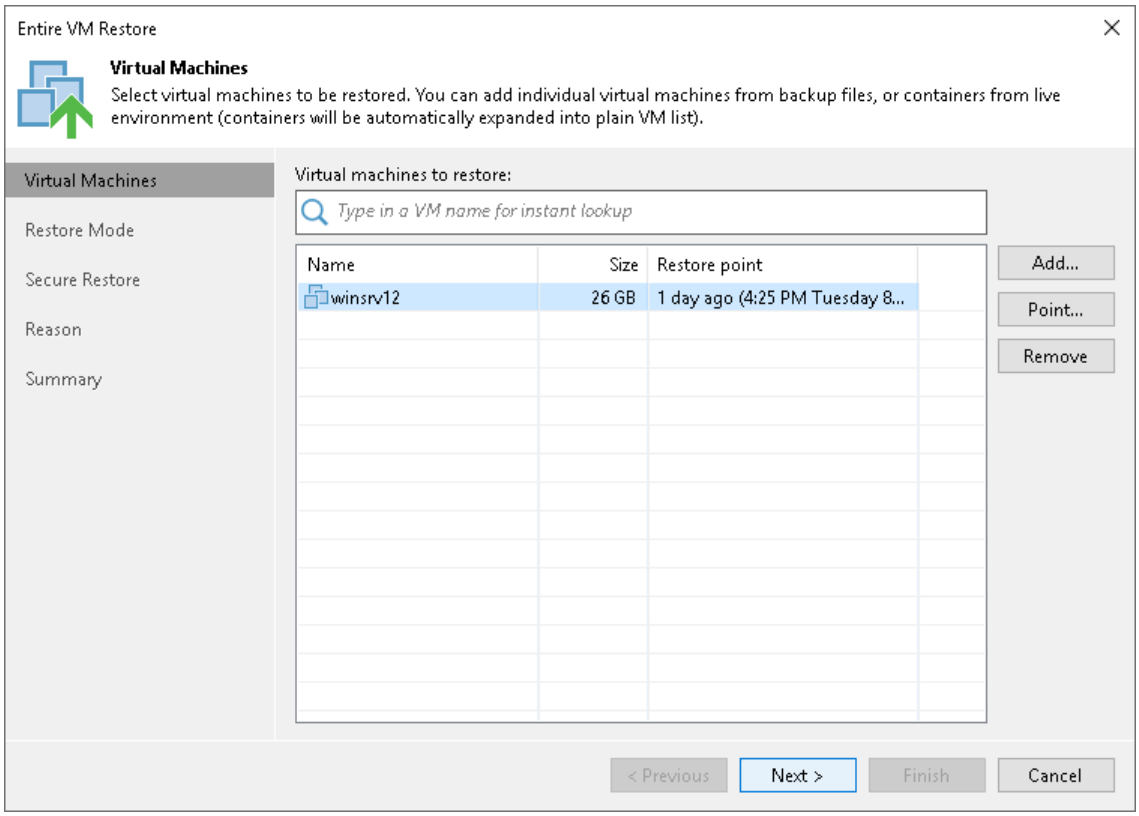

# Step 3. Select Restore Point

By default, Veeam Backup & Replication uses the latest valid restore point. However, you can restore the VM to an earlier state. If you have chosen to restore several VMs, you can select the necessary restore point for each VM in the list.

To select a restore point for a VM:

1. In the **Virtual machines to restore** list, select a VM.
2. Click **Point** on the right.
3. In the **Restore Points** window, select a restore point from which you want to restore the VM.

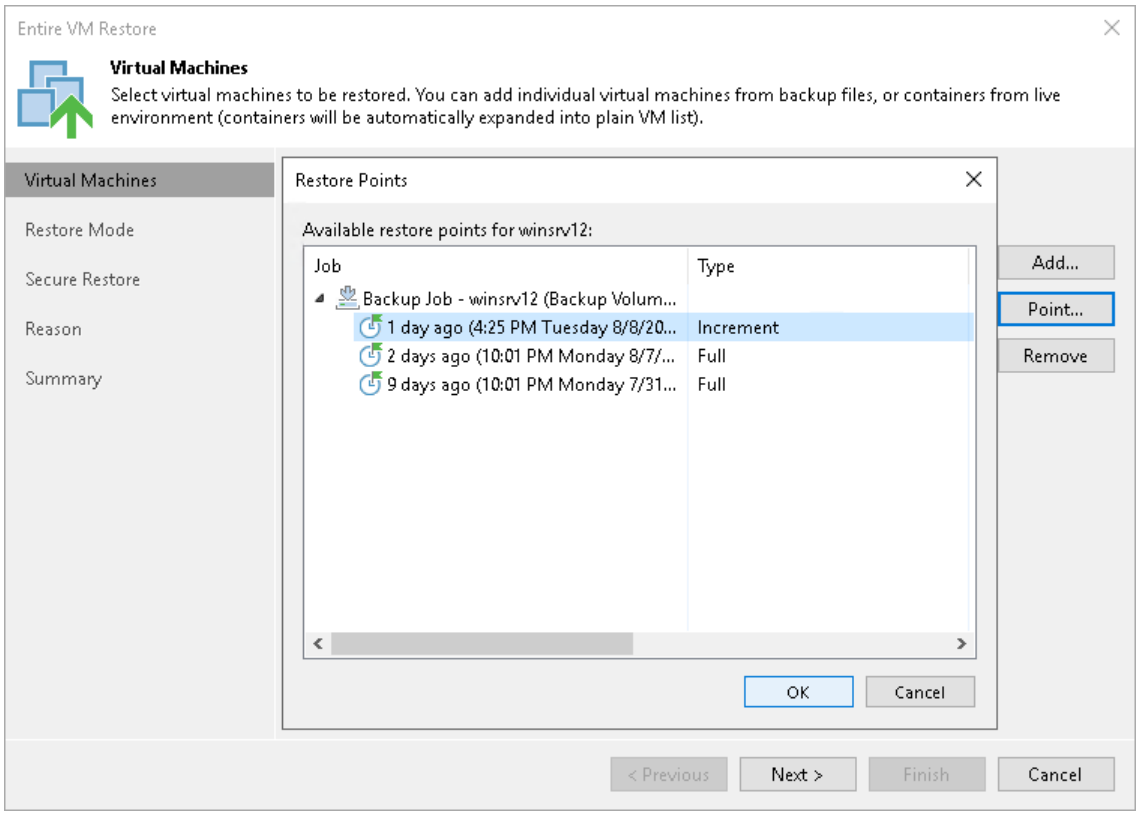

## Step 4. Select Restore Mode

At the **Restore Mode** step of the wizard, choose the necessary restore mode and backup proxy to transfer VM data:

1. Choose a restore mode:
   - Select **Restore to original location** to restore VMs with their initial settings and to their original location. If this option is selected, you will immediately pass to the Reason step of the wizard.

     During restore to the original location, Veeam Backup & Replication restores only those disks that are included in the backup file. This means that after the restore finishes, you do not have to update existing jobs which process the original VMs.
   - Select **Restore to a new location, or with different settings** to restore VMs to a different location or with different settings (such as VM location, network settings, format of restored virtual disks and so on). If this option is selected, the **Full VM Restore** wizard will include additional steps for customizing VMs settings.

     During restore to a new location, Veeam Backup & Replication creates new VMs. If you want to process the restored VMs, you must edit existing jobs or create new jobs to process the restored VMs. If you restore VMs with the same name and to the same folder as the original VMs, Veeam Backup & Replication deletes the original VMs. In this case, you must edit existing jobs to exclude original VMs from them.
   - Select **Staged restore** to run an executable script for VMs before restoring them to the production environment. If this option is selected, the **Full VM Restore** wizard will include an additional step for customizing staged restore settings.

     During staged restore to the original location, that is, when you leave the original settings on the next steps of the wizard, Veeam Backup & Replication removes the original VMs. However, Veeam Backup & Replication automatically updates the existing jobs to process the restored VMs and to exclude the original VMs.

     During staged restore to a new location, Veeam Backup & Replication creates new VMs. If you want to process the restored VMs, you need to edit existing jobs or create new jobs.
2. [For VM restore to the original location] Select the **Quick rollback** check box to perform incremental restore for the VM. Veeam Backup & Replication will query CBT to get data blocks that are necessary to revert the VM to an earlier point in time, and will restore only these data blocks from the backup. Quick restore significantly reduces the restore time and has little impact on the production environment.

   It is recommended that you enable this option if you restore a VM after a problem that occurred at the level of the VM guest OS: for example, there has been an application error or a user has accidentally deleted a file on the VM guest OS. Do not enable this option if the problem has occurred at the VM hardware level, storage level or due to a power loss.

   For more information on quick rollback, its requirements and limitations, see Quick Rollback.
3. Click the **Pick proxy to use** link to select backup proxies over which VM data must be transported to the target datastore. You can assign backup proxies explicitly or instruct Veeam Backup & Replication to automatically select backup proxies.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the target datastore and will automatically assign optimal proxy resources for processing VM data.

     During the restore process, VMs are processed simultaneously. Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use for writing data to target, current workload on these backup proxies, and selects the most appropriate resources for VMs processing.

- If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that will be used for restore. It is recommended that you select at least two proxies to ensure that VMs are recovered should one of backup proxies fail or lose its connectivity to the target datastore during restore.

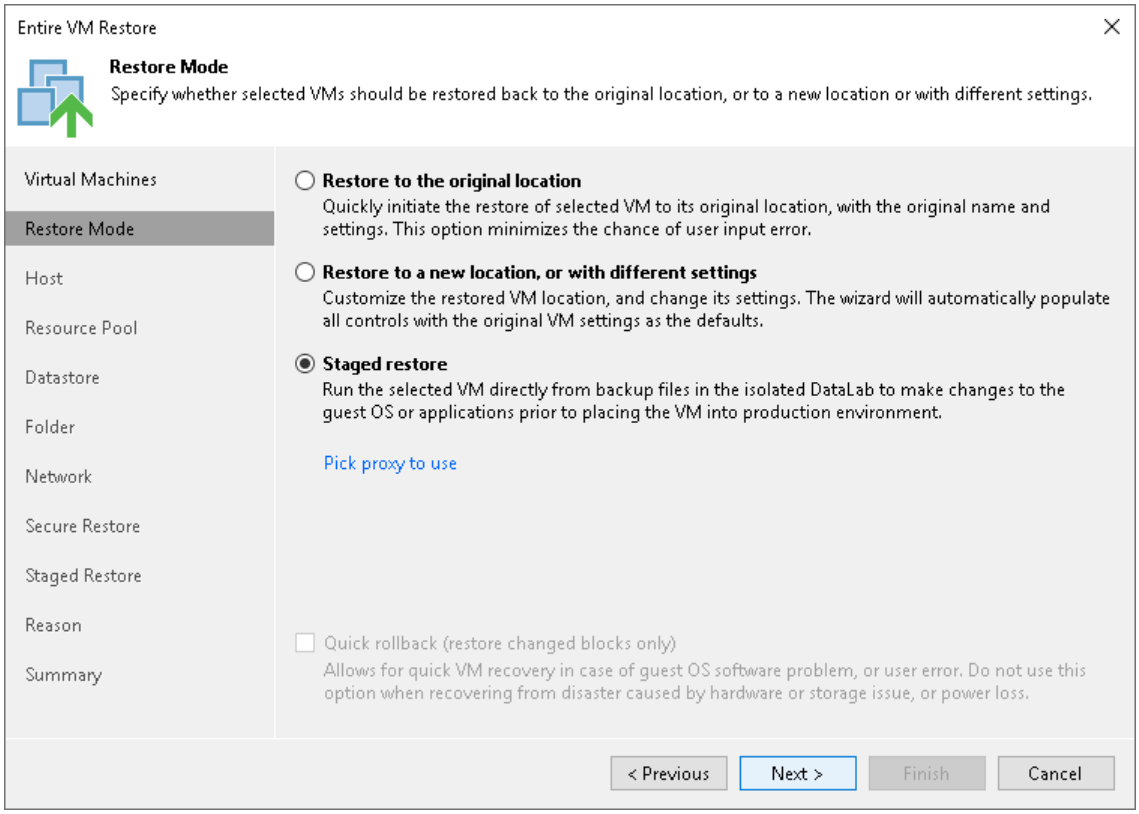


# Restoring Storage Policies

If the backed-up VM was associated with the storage policy, in the restore to original location scenario, Veeam Backup & Replication will associate the restored VM with this storage policy.

When you click **Next**, Veeam Backup & Replication will check storage policies in the virtual environment and compare this information with the information about the storage policy in the backup file. If the original storage policy has been changed or deleted, Veeam Backup & Replication will display a warning. You can select one of the following options:

- **Current** – the restored VM will be associated with the profile with which the original VM in the production environment is currently associated.
- **Default** – the restored VM will be associated with the profile that is set as default for the target datastore.
- **Stored** – the restored VM will be associated with the profile that was assigned to the original VM at the moment of backup, and whose information is stored in the backup file.

For more information, see Storage Policies.

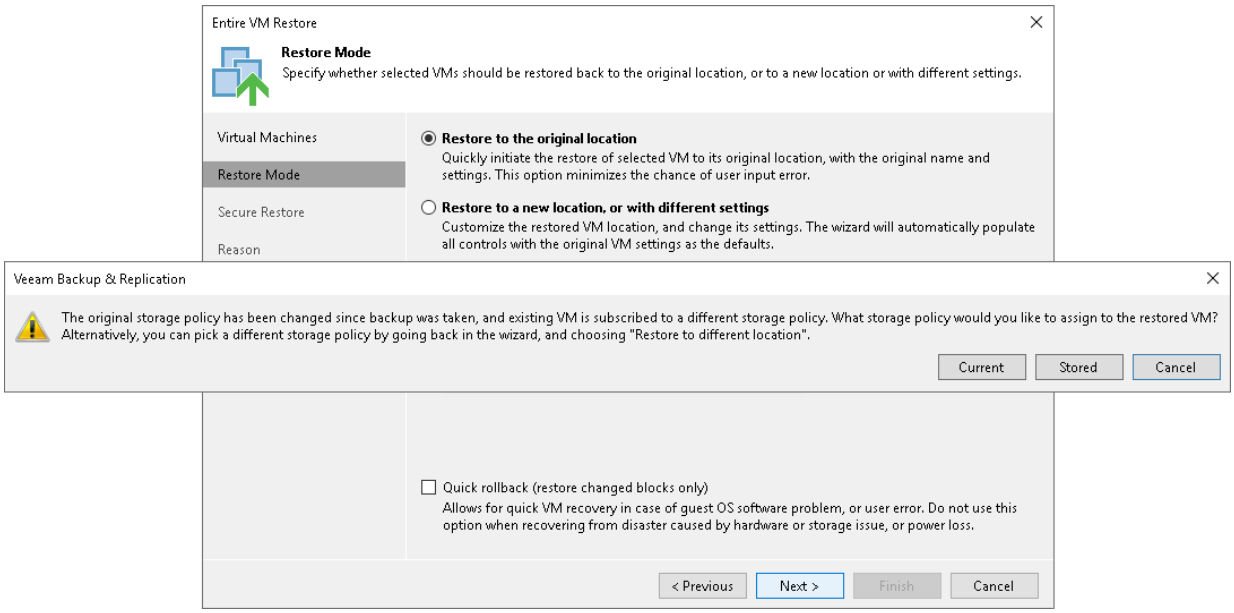

# Step 5. Select Target Host

The **Host** step of the wizard is available if you have chosen to change the location and settings for the restored VM.

To specify a target host:

1. Select the necessary VMs in the list and click **Host**.
2. Choose a host or cluster where the selected VMs must be registered.

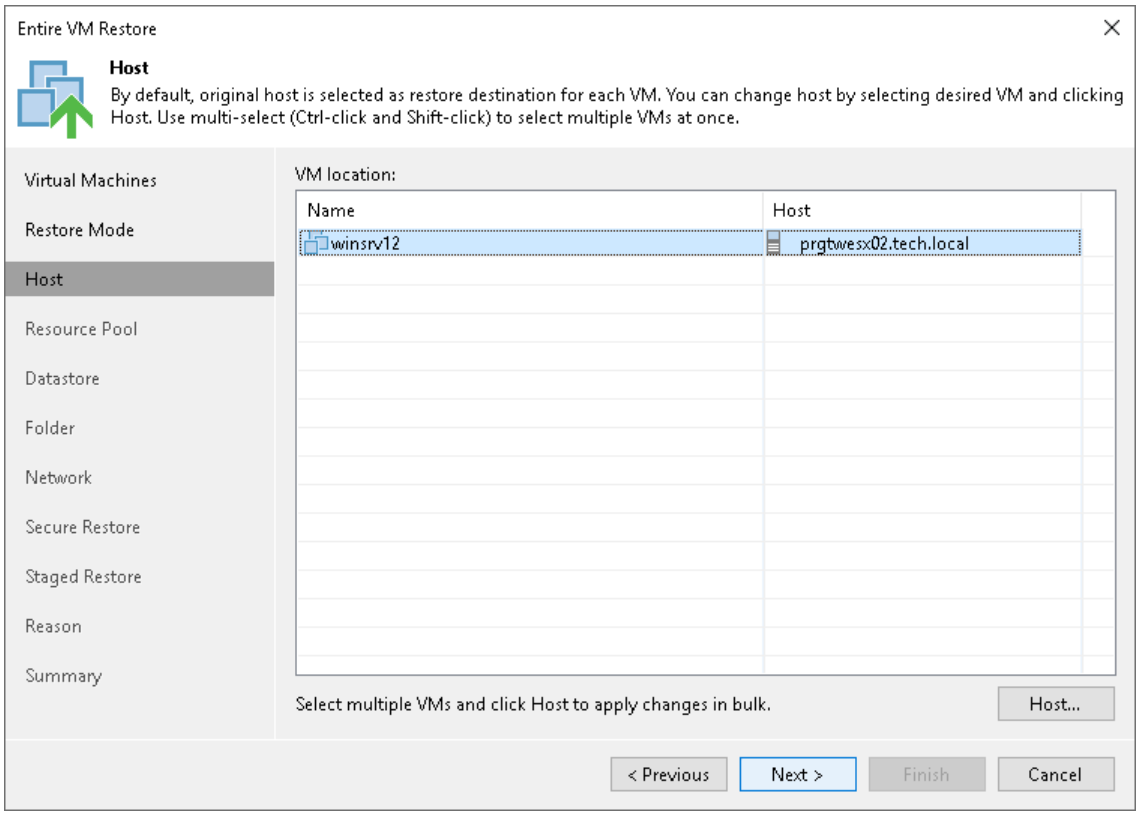

# Step 6. Select Target Resource Pool

The **Resource Pool** step of the wizard is available if you have chosen to change the location and settings for the restored VM.

To specify a destination resource pool:

1. Select the necessary VMs in the list and click **Pool**.
2. Select a resource pool to which the VMs must be placed.
3. If necessary, select a vApp in which the VMs must be included.

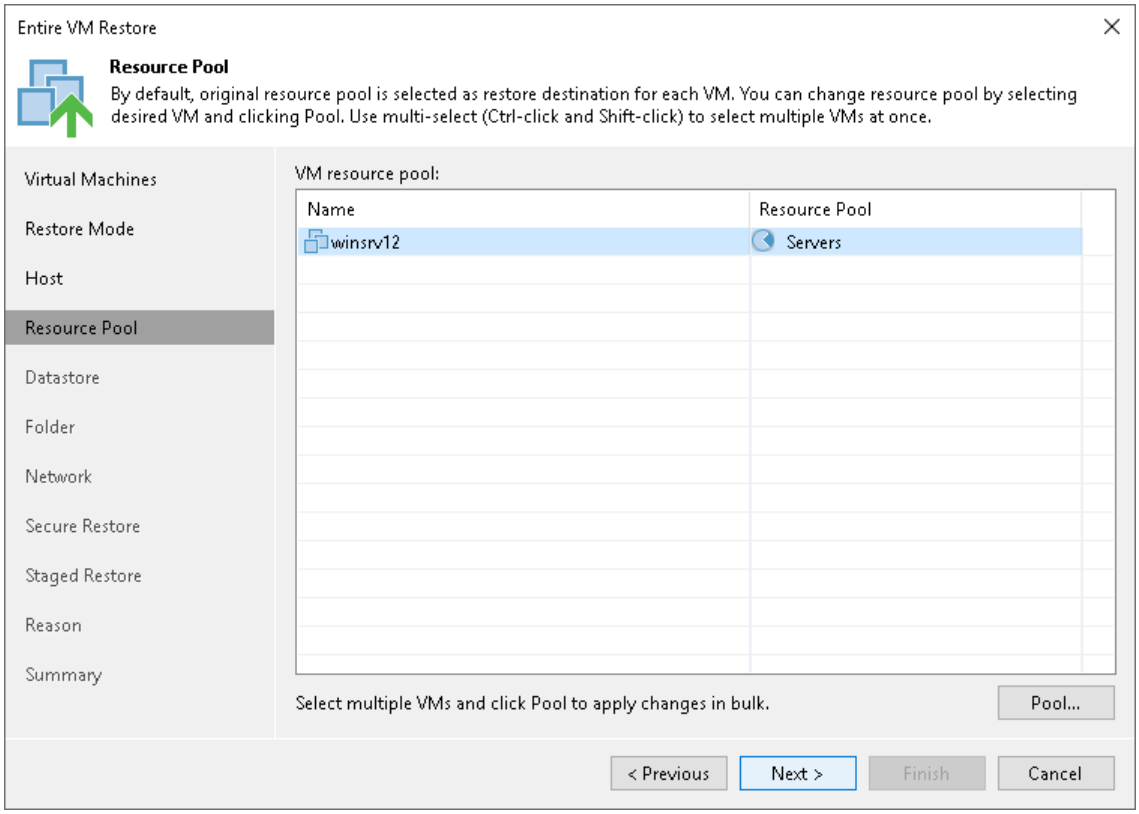

# Step 7. Select Target Datastore and Disk Type

The **Datastore** step of the wizard is available if you have chosen to change the location and settings for the restored VM.

You can place an entire VM to a particular datastore or choose to store configuration files and disk files of the restored VM in different locations.

1. To place VMs to a datastore, select the necessary VMs in the **Files location** list and click **Datastore**. In the **Disk Type Settings** window, select a datastore and click **OK**.

   If you want to place configuration and disk files to different datastores, expand a VM in the **File location** list and select individual files. Then specify the necessary datastore.

2. By default, Veeam Backup & Replication preserves the format of restored VM disks. To change the disk format, expand a VM in the **Files location** list, select the necessary disks and click **Disk Type**. In the **Disk Type Settings** section, choose the format that will be used to restore virtual disks of the VM: same as source, thin, thick lazy zeroed or thick eager zeroed. For more information about disk types, see VMware Docs. Click **OK**.

> **NOTE**
>
> Disk format change is supported only for VMs with Virtual Hardware version 7 or later.

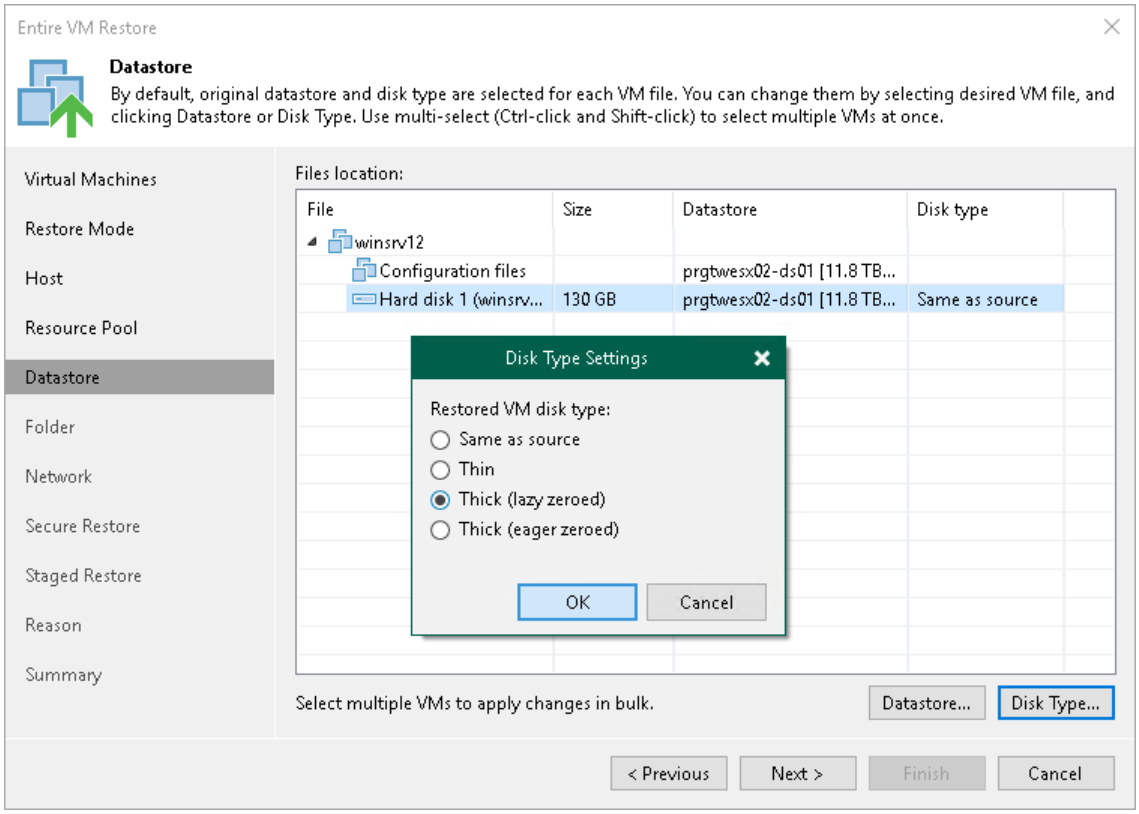

## Step 8. Select Target Folder and Change VM Settings

The **Folder** step of the wizard is available if you have chosen to change the location and settings for restored VMs.

At the **Folder** step of the wizard, specify a destination VM folder, choose whether you want to restore VM tags and change VM names. By default, Veeam Backup & Replication preserves the original names.

# Specifying Destination VM Folder

To specify a destination VM folder:

1. Select a VM in the list and click **Folder**.
2. Choose a folder to which the VM will be placed.

> **NOTE**
>
> Consider the following:
>
> - If you restore a VM to a standalone ESXi host that is not managed by the vCenter Server, you cannot select a destination folder: this option will be disabled.
> - During entire VM restore, Veeam Backup & Replication preserves the UUID of the original VM.

# Changing Names

To change a VM name:

1. Select a VM in the **VM folder** list and click **Name**.
2. In the **Change Name** section, enter a new name explicitly or specify a change name rule by adding a prefix or suffix to the original VM name.

Alternatively, you can change a VM name directly in the **VM folder** list. To do this, click the **New Name** field and enter the name to be assigned to the recovered VM.

# Restoring VM Tags

Select the **Restore VM tags** check box if you want to restore tags that were assigned to the original VM and to assign them to the restored VM. Veeam Backup & Replication will restore the VM with original tags if the following conditions are met:

- The VM is restored to its original location.

- The original VM tag is still available on the source vCenter Server.

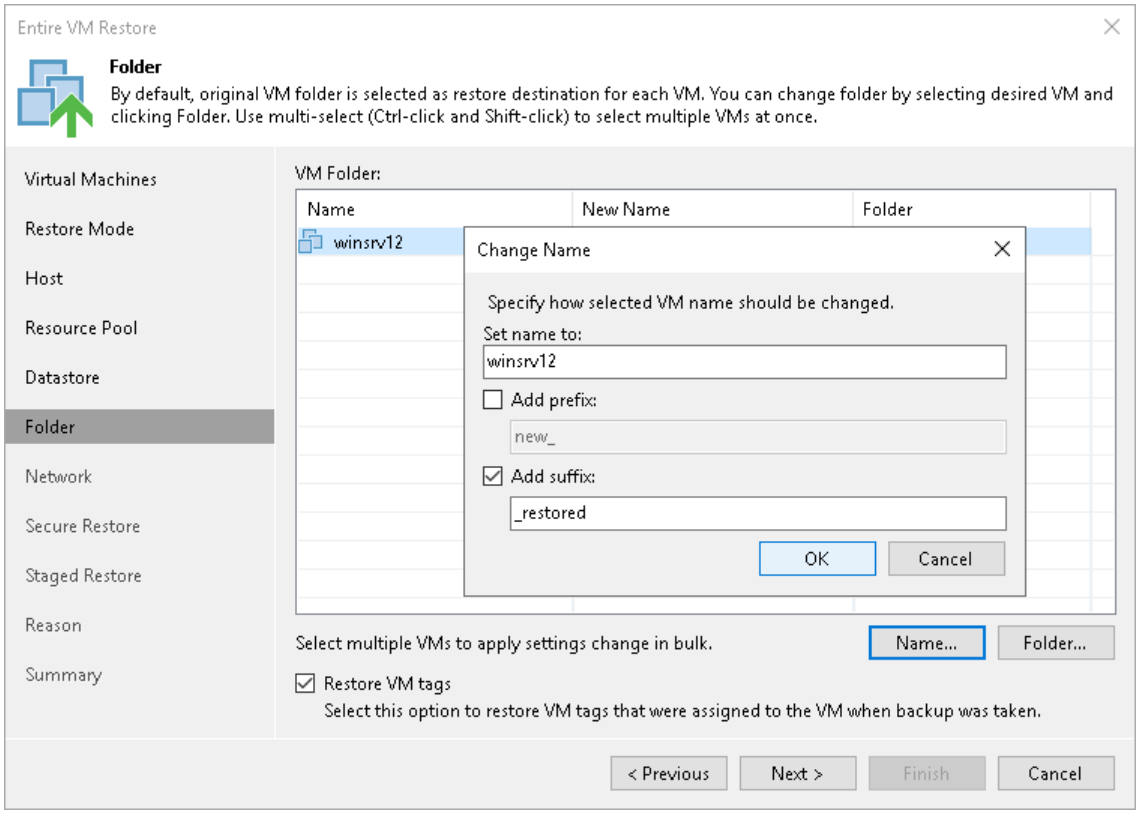

# Step 9. Specify Network Mapping

The **Network** step of the wizard is available if you have chosen to change the location and settings for the restored VM.

If you plan to restore a VM to a new location, for example, another site with a different set of networks, you can map source site networks to target site networks. Veeam Backup & Replication will use the network mapping table to update configuration files of the VM on the fly, during the restore process.

To change networks to which the restored VM will be connected:

1. Select a VM in the list and click **Network**. To apply changes in bulk, select several VMs in the list and click **Network**.

   If a VM is connected to multiple networks, expand the VM, select the network to map and click **Network**. The **Select Network** section displays all networks to which the target host or cluster is connected.

2. From the list of available networks, choose a network to which the VM must have access upon restore.

If you do not want to connect a restored VM to your virtual networks, select the VM in the list and click **Disconnected**.

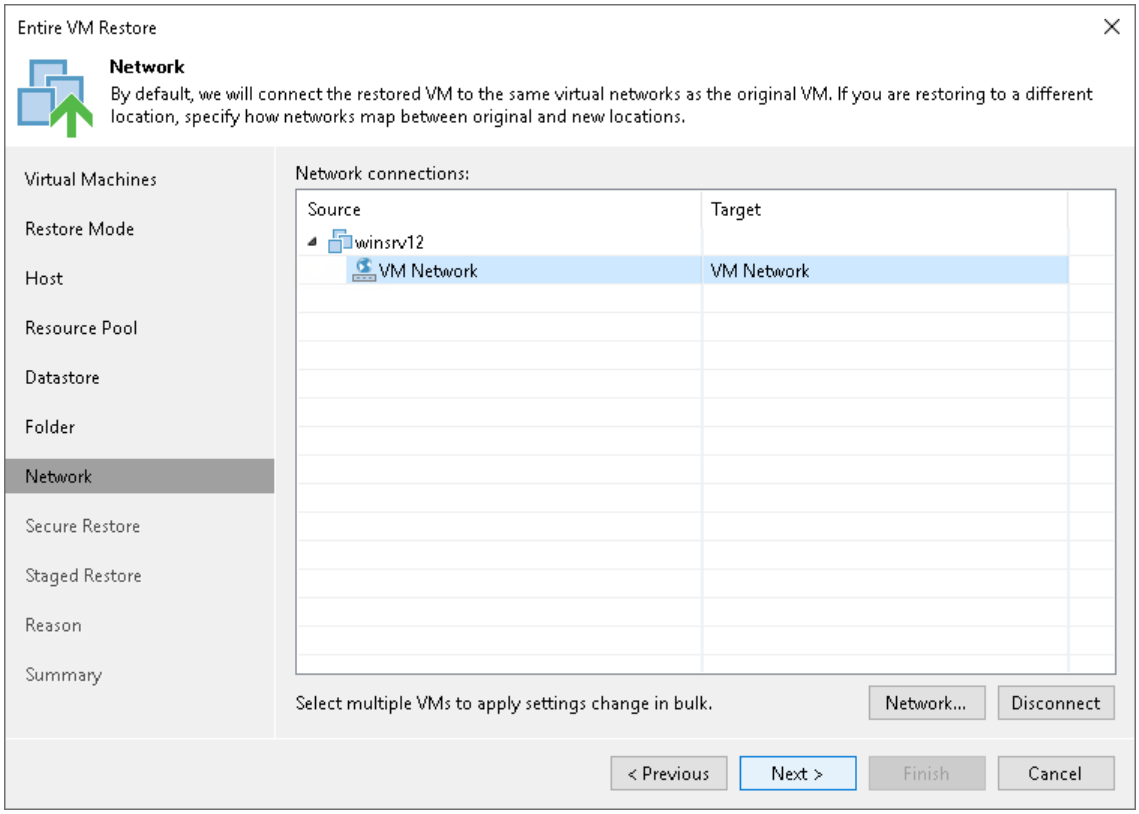

# Step 10. Specify Secure Restore Settings

The **Secure Restore** step of the wizard is available if you restore Microsoft Windows VMs.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Select which action Veeam Backup & Replication will take if scan finds a virus threat:
   - **Proceed with recovery but disable network adapters**. Select this action if you want to restore the machine with disabled network adapters (NICs).
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.

3. Select the **Continue scanning all remaining files after the first occurrence** check box if you want to continue machine scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.

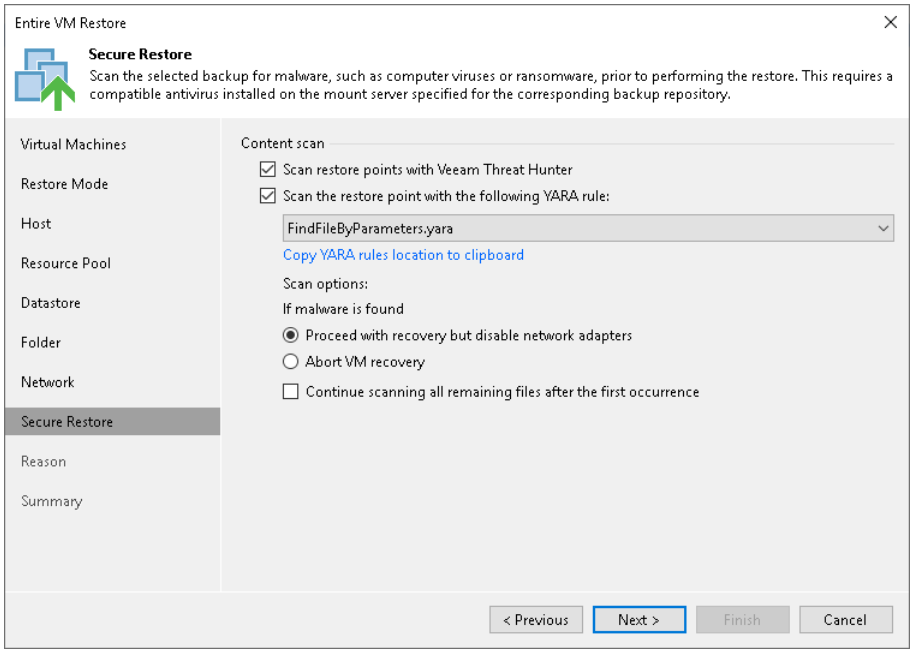

## Step 11. Specify Staged Restore Settings

The **Staged Restore** step of the wizard is available if you have chosen the **Staged restore** option at the Restore Mode step of the wizard.

Staged restore to run an executable script for VMs before restoring them to the production environment. For more information, see Staged Restore.

To specify staged restore settings:

1. From the **Virtual lab** list, select a virtual lab that will be used to start VMs. The list contains all virtual labs that are created or connected to the backup server.
2. From the **Application group** list, select an application group if script execution requires other VMs to be powered on. In the virtual lab during staged restore, Veeam Backup & Replication will start VMs from the selected application group in the required order. The **Application group** list contains all application groups that are created on the backup server. For more information, see Application Group.
3. On the right of the **Script** field, click **Browse** to choose the script from a local folder on the backup server.
4. From the **Credentials** list, select credentials for the account that has administrator privileges on VMs for which you want to run the script. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right of the **Credentials** field to add the credentials. For more information, see Credentials Manager.

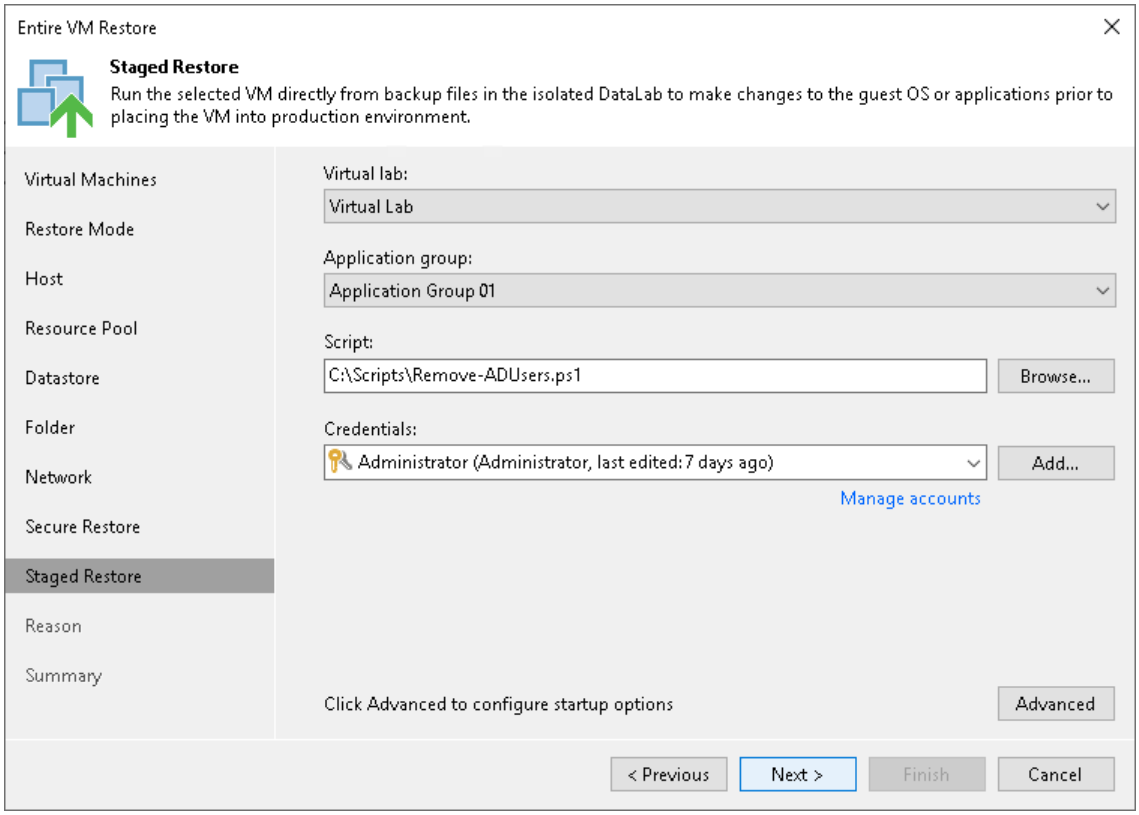


# VM Startup Settings

If you want to start VMs after recovery, perform the following steps:

1. Click **Advanced**.

2. In the **Memory** section, specify the amount of memory that you want to pre-allocate to a VM when it starts. The amount of pre-allocated memory is defined in percent. The percentage rate is calculated based on the system memory level available for the production VM. For example, if 4096 MB of RAM is allocated to the VM in the production environment and you specify 50% as a memory rate, 2048 MB of RAM will be allocated to the VM on startup.

3. In the **Startup time** section, specify the allowed boot time for the VM and timeout to initialize applications on the VM.

   Be careful when specifying the **Maximum allowed boot time** value. Typically, a VM started in a virtual lab requires more time to boot than a VM started in the production environment. If an application fails to be initialized within the specified interval of time, the recovery process fails with the timeout error. If such error occurs, you need to increase the **Maximum allowed boot time** value and perform VM restore again.

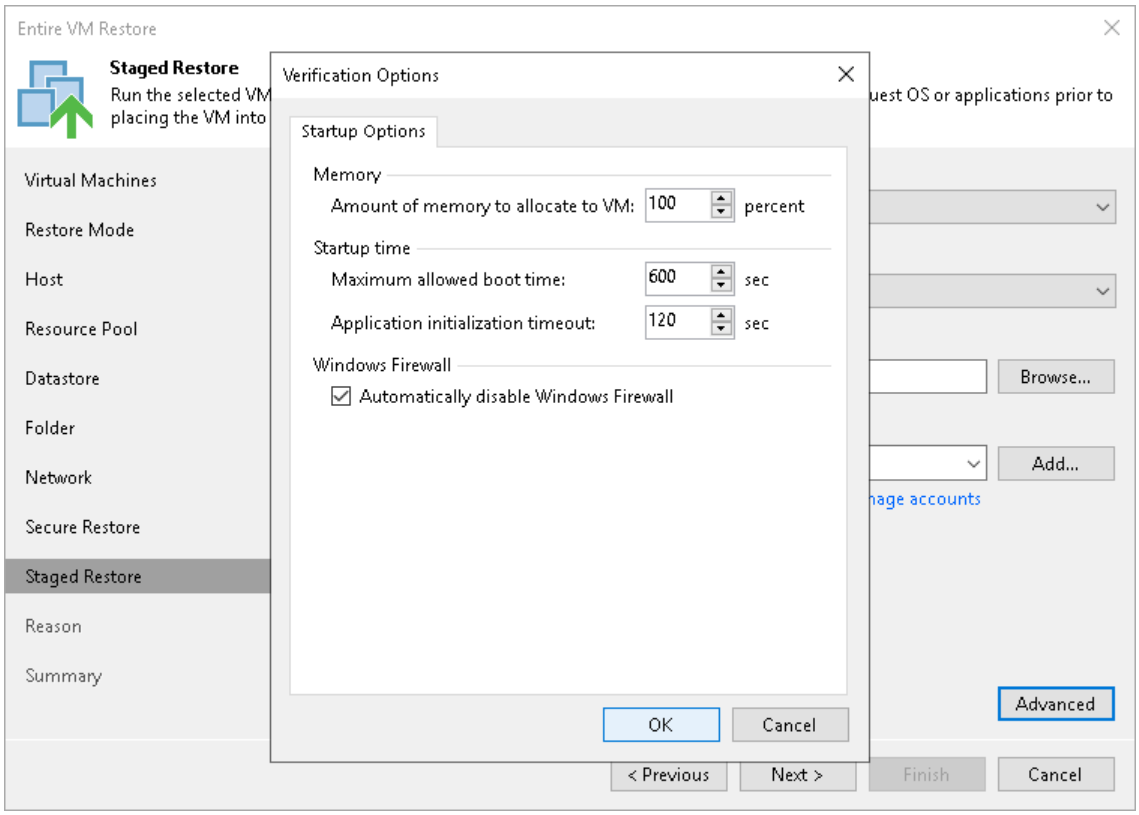

## Step 12. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the selected VMs. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

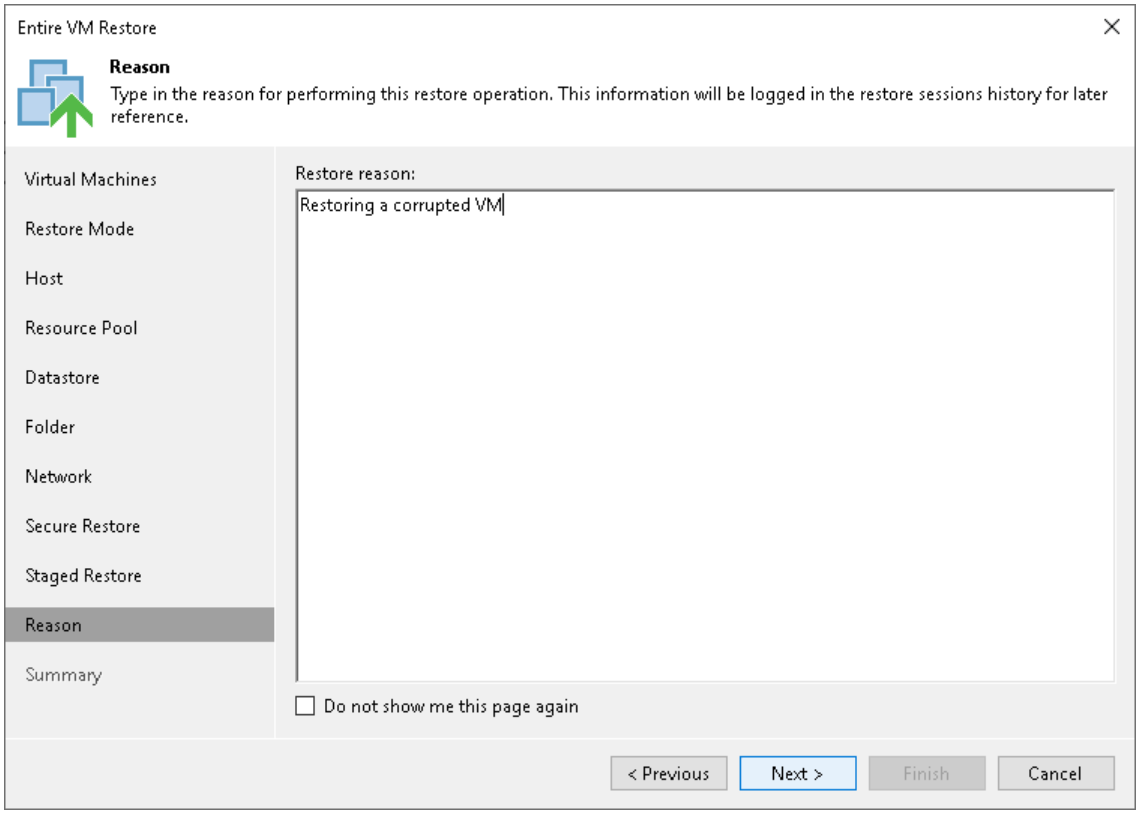

# Step 13. Verify Restore Settings

At the **Summary** step of the wizard, check restore settings and click **Finish**. If you want to start the recovered VMs on the target host, select the **Power on target VM after restoring** check box.

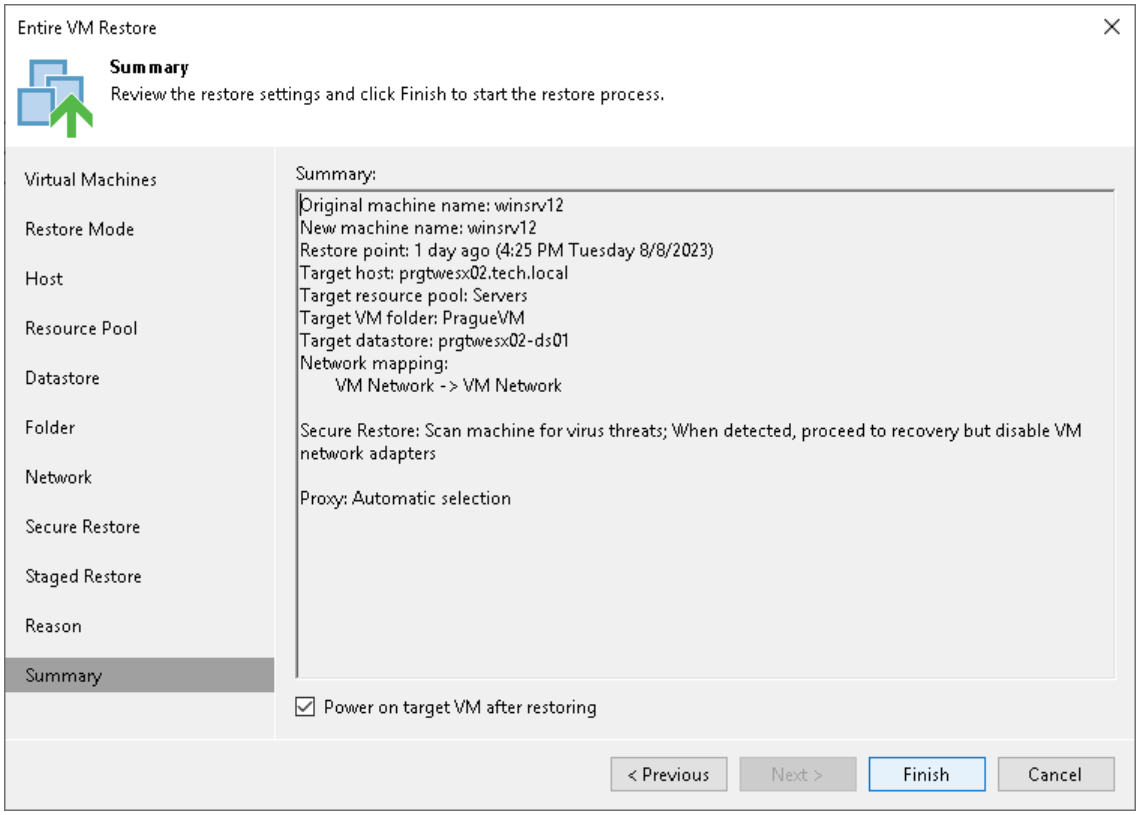

# Staged Restore

Staged restore allows you to run an executable script for VMs before recovering them to the production environment. Staged restore is a part of the entire VM restore operations. To perform staged restore, you must select the **Staged Restore** mode in the Full VM Restore wizard and specify staged restore settings.

Staged restore can help you ensure that recovered VMs do not contain any personal or sensitive data. For example, you can instruct Veeam Backup & Replication to run a Windows PowerShell script that removes Active Directory users:

```
$UserName = "John.Smith"
$ADUser = Get-ADUser -Filter 'Name -like $UserName'
if (!$ADUser)
   {
       [Environment]::Exit(1)
   }
Remove-ADUser -Identity $UserName -Confirm:$false
```

**NOTE**

The staged restore functionality is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.

## Requirements and Limitations for Staged Restore

Before you perform staged restore, check the following prerequisites:

- You must have a preconfigured virtual lab in your backup infrastructure. For more information, see Virtual Lab.
- Scripts that you plan to run must reside in a local folder on a backup server.
- If you plan to perform staged restore for several VMs within one restore session, make sure these VMs run OS of the same type: either Microsoft Windows or Linux. In the current version of Veeam Backup & Replication, you cannot specify credentials and scripts for each VM individually.
- When restoring VMs, Veeam Backup & Replication uses the Veeam Quick Migration method. vMotion and Storage vMotion methods cannot be used. For more information on the Quick Migration method, see Quick Migration.

## How Staged Restore Works

For staged restore, Veeam Backup & Replication uses a preconfigured virtual lab, an executable script located on the backup server, and credentials to connect to VMs and run the script. Veeam Backup & Replication performs staged restore in the following way:

1. In the virtual lab, Veeam Backup & Replication starts VMs directly from compressed and deduplicated backup files that reside in the backup repository. To achieve this, Veeam Backup & Replication uses the Veeam vPower NFS Service.

   If you selected to use an application group to run a script, Veeam Backup & Replication first starts VMs from the application group in the required order.

2. Veeam Backup & Replication copies the script from the backup server to VMs that you plan to restore.

   To connect to VMs, Veeam Backup & Replication uses credentials specified in staged restore settings.

3. Veeam Backup & Replication runs the copied script on every VM.

   To run the script, Veeam Backup & Replication uses the same technology as for pre-freeze and post-thaw scripts. For more information, see Pre-Freeze and Post-Thaw Scripts.

4. All VM changes that take place during script execution are written to VM delta files.

   By default, Veeam Backup & Replication stores delta files on the vPower NFS server. You can change the destination for VM delta files in virtual lab settings.

5. After the script execution is complete, Veeam Backup & Replication makes a safe shutdown of VMs in the virtual lab.

6. Veeam Backup & Replication restores VMs in a changed state to the production environment.

   To achieve that, Veeam Backup & Replication copies VM data from the backup repository and delta files to the target host using Veeam Quick Migration.

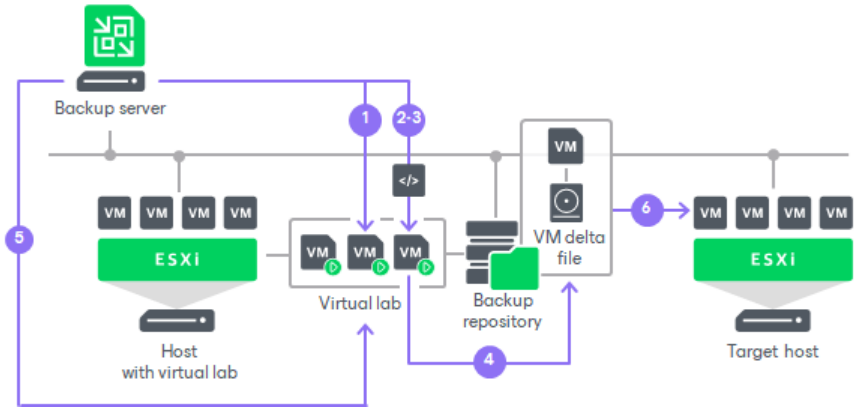

# Restore to Amazon EC2

Veeam Backup & Replication allows you to restore different workloads (VMs, Google VM instances, physical servers and so on) to Amazon Elastic Compute Cloud (Amazon EC2) as EC2 instances. An EC2 instance is a virtual machine in Amazon EC2 with a preconfigured combination of computing resources.

You can use Veeam Backup & Replication to perform the following operations:

- Restore workloads to Amazon EC2 from backups.
- Migrate workloads from the on-premises infrastructure to the cloud.
- Create a test environment in the cloud for troubleshooting, testing patches and updates, and so on.

## Supported Backup Types

You can restore workloads from the following types of backups:

- Backups of VMware vSphere or VMware Cloud Director virtual machines created by Veeam Backup & Replication.
- Backups of Microsoft Hyper-V virtual machines created by Veeam Backup & Replication.
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.

  Backups must be created at the entire machine level or volume level.
- Backups of Amazon EC2 instances created by Veeam Backup for AWS.
- Backups of Microsoft Azure virtual machines created by Veeam Backup for Microsoft Azure.
- Backups of Google Compute Engine VM instances created by Veeam Backup for Google Cloud.
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV.
- Backups of oVirt VMs created by Veeam Backup for OLVM and RHV.
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE.

## Helper Appliances

Helper appliance is an auxiliary Linux-based EC2 instance. It is used to upload backed-up data to Amazon EC2. Veeam Backup & Replication automatically deploys the helper appliance in Amazon EC2 only for the duration of the restore process and removes it immediately after that.

Depending on the type of backups you are restoring from and their location, the helper appliance may be required or optional. The helper appliance is required when you restore from:

- Backups of EC2 instances that are stored in external repositories.
- Backups of virtual and physical machines that are stored in object storage repositories.

The helper appliance is optional when you restore from backups of virtual and physical machines stored in backup repositories, or backups of EC2 instances copied to backup repositories with backup copy jobs. It is recommended, however, to use the helper appliance in scenarios where it is optional, as the helper appliance can significantly improve restore performance. You can specify the helper appliance settings at the Helper Appliance step of the **Restore to Amazon EC2** wizard.

# Requirements for Helper Appliance

When configuring a helper appliance, consider the following:

- To upload one machine disk to Amazon EC2, the helper appliance requires 1 GB RAM. Make sure that the type of EC2 instance selected for the helper appliance offers enough memory resources to upload all machine disks. Otherwise, the restore process may fail.
- To restore from a backup in an on-premises object storage repository, the helper appliance machine must have access to the source object storage repository. To provide access to object storage repository, you can use VPN or AWS Direct Connect.
- To restore from a backup in an Amazon S3 object storage, use the helper appliance that has the instance type with the maximum values for the Network Bandwidth (Gbps) and EBS Bandwidth (Mbps). In most cases, you can use c5.18xlarge. For more information on instance types, see Amazon EC2 Instance Types.
- A subnet and security group that you select for the helper appliance must meet the following requirements:
  - Auto-assignment of public IPv4 addresses must be enabled in the subnet. For more information on how to enable this option, see the AWS Documentation.
  - The subnet route table must contain a default route to an active AWS internet gateway. For more information on internet gateways and how to create route tables, see the AWS Documentation.
  - The subnet must have no network access control lists (ACLs) or a network ACL that allows inbound and outbound traffic on the ports listed in section Ports.
  - The security group must allow inbound and outbound traffic on the ports listed in section Ports.

# How Restore to Amazon EC2 Works

The workflow of the restore process depends on whether the helper appliance is used or not. For more information on the helper appliance, see Helper Appliances.

> **NOTE**
>
> If you use AWS Plug-in for Veeam Backup & Replication and plan to restore Amazon EC2 instances from restore points that were created using the appliance, you do not need to configure the helper appliance. Also, restore to Amazon EC2 works as described in the EC2 Instance Restore section in the Veeam Backup for AWS User Guide.

# Restoring to Amazon EC2 with Helper Appliance

If a helper appliance is required or you selected to use it for restore to Amazon EC2, the restore algorithm contains the following steps:

1. Veeam Backup & Replication creates a helper appliance in Amazon EC2.

   During the restore process, the helper appliance communicates with backup infrastructure components over the SSH protocol and the network redirector that is deployed on the helper appliance.
2. For every disk of a backed-up workload, Veeam Backup & Replication creates an empty EBS volume in Amazon EC2.
3. Veeam Backup & Replication hot-adds empty disks to the helper appliance and restores backed-up data to the EBS volumes.

4. [For backups of EC2 instances] Veeam Backup & Replication creates a target instance in Amazon EC2.
5. [For backups of EC2 instances] Veeam Backup & Replication detaches the EBS volumes from the helper appliance and attaches them to the target instance.
6. [For backups of other workloads] Veeam Backup & Replication requests Amazon to create snapshots of the attached EBS volumes. The snapshots are created in Amazon S3.
7. [For backups of other workloads] Veeam Backup & Replication requests Amazon to use the VM Import functionality to import a VM from EBS volume snapshots to Amazon EC2.
8. [For backups of other workloads] After the restore process is complete, Veeam Backup & Replication removes the snapshots created in Amazon S3.
9. Veeam Backup & Replication removes helper appliance from Amazon EC2.

# Restoring to Amazon EC2 without Helper Appliance

If you selected not to use a helper appliance for restore to Amazon EC2, the restore algorithm contains the following steps:

1. The backup server connects to a backup repository or a gateway server and locates the required backup.
2. The backup repository or gateway server uploads disks of the backed-up workload to Amazon S3.

   In Amazon S3, the uploaded disks are stored to the temporary bucket in the RAW format.
3. Veeam Backup & Replication imports the backed-up data from the temporary bucket in Amazon S3 to EBS volumes in Amazon EC2.
4. Veeam Backup & Replication creates a target instance in Amazon EC2 and attaches the EBS volumes to it.
5. After the import process is complete, Veeam Backup & Replication removes the temporary bucket from Amazon S3.

# AWS IAM User Permissions

To restore to Amazon EC2, it is recommended that the IAM user whose credentials you plan to use to connect to AWS has administrative permissions — access to all AWS actions and resources.

If you do not want to provide full access to AWS, you can grant to the IAM user a minimal set of permissions that will be sufficient for restore. To do that, create the following policy in the JSON format and attach it to the IAM user.

> **NOTE**
>
> The `iam: CreateRole` permission is required if you want to perform restore without helper appliances. This permission is used to create a service role named `vmimport` required for import to Amazon EC2. If you plan to restore workloads using helper appliances, you can omit the `iam: CreateRole` permission. However, restore without helper appliances will fail.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Action": [
   "ec2:DescribeInstances",
   "ec2:DescribeInstanceTypes",
```

```
      "ec2:RunInstances",
      "ec2:TerminateInstances",
      "ec2:StartInstances",
      "ec2:StopInstances",
      "ec2:ModifyInstanceAttribute",
      "ec2:DescribeImages",
      "ec2:ImportImage",
      "ec2:DeregisterImage",
      "ec2:DescribeVolumes",
      "ec2:CreateVolume",
      "ec2:ModifyVolume",
      "ec2:ImportVolume",
      "ec2:DeleteVolume",
      "ec2:AttachVolume",
      "ec2:DetachVolume",
      "ec2:GetEbsEncryptionByDefault",
      "ec2:CreateSnapshot",
      "ec2:DescribeSnapshots",
      "ec2:DeleteSnapshot",
      "ec2:DescribeSubnets",
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeSecurityGroups",
      "ec2:DescribeKeyPairs",
      "ec2:CreateKeyPair",
      "ec2:DeleteKeyPair",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeVpcs",
      "ec2:DescribeConversionTasks",
      "ec2:DescribeImportImageTasks",
      "ec2:DescribeVolumesModifications",
      "ec2:CancelImportTask",
      "ec2:CancelConversionTask",
      "ec2:CreateTags",
      "ec2:DescribeAccountAttributes",
      "ec2:DescribeDhcpOptions",
      "ec2:DescribeVpcAttribute",
      "iam:GetRole",
      "iam:CreateRole",
      "iam:PutRolePolicy",
      "iam:DeleteRolePolicy",
      "s3:CreateBucket",
      "s3:ListBucket",
      "s3:ListAllMyBuckets",
      "s3:DeleteBucket",
      "s3:PutObject",
      "s3:DeleteObject",
      "s3:GetBucketLocation",
      "s3:PutLifeCycleConfiguration",
      "s3:GetObject",
      "s3:RestoreObject",
      "s3:AbortMultiPartUpload",
      "s3:ListBucketMultiPartUploads",
      "s3:ListMultipartUploadParts"
    ],
    "Effect": "Allow",
    "Resource": "*"
  }]
}
```

Alternatively, you can attach the created policy to the IAM group or role to which the IAM user is assigned.

For information on how to create and attach a policy to an IAM user, see the Creating IAM Policies and Adding and Removing IAM Identity Permissions sections in the AWS IAM User Guide.

# Restoring to Amazon EC2

To restore workloads to Amazon EC2, use the **Restore to Amazon EC2** wizard.

> **NOTE**
>
> If you use AWS Plug-in for Veeam Backup & Replication and plan to restore Amazon EC2 instances from restore points that were created using the appliance, the steps of the restore wizard differ from the steps described in this guide. For more information, see the Restoring to Amazon section in the Veeam Backup for AWS User Guide.

## Before You Begin

Before you restore workloads to Amazon EC2, consider the following requirements and limitations:

- Check whether a helper appliance must be configured for restore. For more information, see Helper Appliances.
- Make sure that a user whose credentials you plan to use to connect to AWS has permissions to restore to Amazon EC2. For more information, see AWS IAM User Permissions.
- You must have a backup of the workload that you plan to restore to Amazon EC2.
- The backup server and repositories with workload backup files must have access to the internet.

  If backup files are located on deduplicating storage appliances or shared folder repositories, the internet connection is required for gateway servers that communicate with these repositories.
- If you use a cloud-init-based Linux distribution, we recommend that you use SSH keys on these distributions. If you use a password, it is blocked after restore for security reasons. To reset the password on the restored instance, use the technologies described in AWS Documentation.
- If you plan to restore a Linux- or Unix-based workload, check that Amazon EC2 supports the kernel version. If the version is not supported, consider changing the kernel version before backing up. For more information on kernel versions, see AWS Documentation.
- If you plan to restore workloads other than EC2 instances, check the supported OS, EC2 instance and file system types in the AWS Documentation.
- Check that the logical sector size of disks that you plan to restore equals 512 bytes. Contents of disks whose sector size is 4096 bytes will be unreadable in Amazon EC2.
- Veeam Backup & Replication does not support restoring disks encrypted by BitLocker, except for restoring from backups created by Veeam Agent for Microsoft Windows. For more information, see the Veeam Agent for Microsoft Windows User Guide.
- You cannot restore EC2 instances with encrypted EBS volumes using a helper appliance. To restore such instances, disable the use of the helper appliance. If your backups are stored on an object storage repository, you can disable the appliance with a registry value. For more information, contact Veeam Customer Support.

- [For restore of EC2 instances without helper appliance] If you restore workloads with more than five disks, check that the **Limit maximum concurrent tasks** option of the repository where the backups are stored is equal or less than the limit of the AWS ImportVolume service for concurrent tasks.
  Veeam Backup & Replication uses this service during the restore. For more information on the limit, see AWS Documentation.

- If you plan to assign AWS tags to the restored EC2 instance, check limitations for tags in the AWS Documentation.

- [For Veeam Backup & Replication server located in Amazon EC2] If you restore workloads from a backup stored in an object storage repository, we recommend using Private IPs to increase the transfer speed. For more information, see this Veeam KB article.

## Step 1. Launch Restore to Amazon EC2 Wizard

To begin the restore process, do one of the following.

- On the **Home** tab, click **Restore** and select the type of backups from which you want to restore:
  - VMware vSphere
  - VMware Cloud Director
  - Microsoft Hyper-V
  - Agent
  - AWS
  - Azure IaaS backup
  - GCE backup
  - Nutanix AHV
  - oVirt KVM
  - Proxmox VE

  In the displayed window, click **Entire VM restore** > **Restore to public cloud** > **Restore to Amazon EC2**.
- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary backup, select workloads that you want to restore and click **Restore to Amazon EC2** on the ribbon. Alternatively, you can right-click one of the workload that you want to restore and select **Restore to Amazon EC2**.

- Double-click a full backup file (VBK) or backup metadata file (VBM) in a file browser.
  Veeam Backup & Replication will start its console. In the **Backup Properties** window, select the necessary workload and click **Restore** > **Restore to Amazon EC2**.

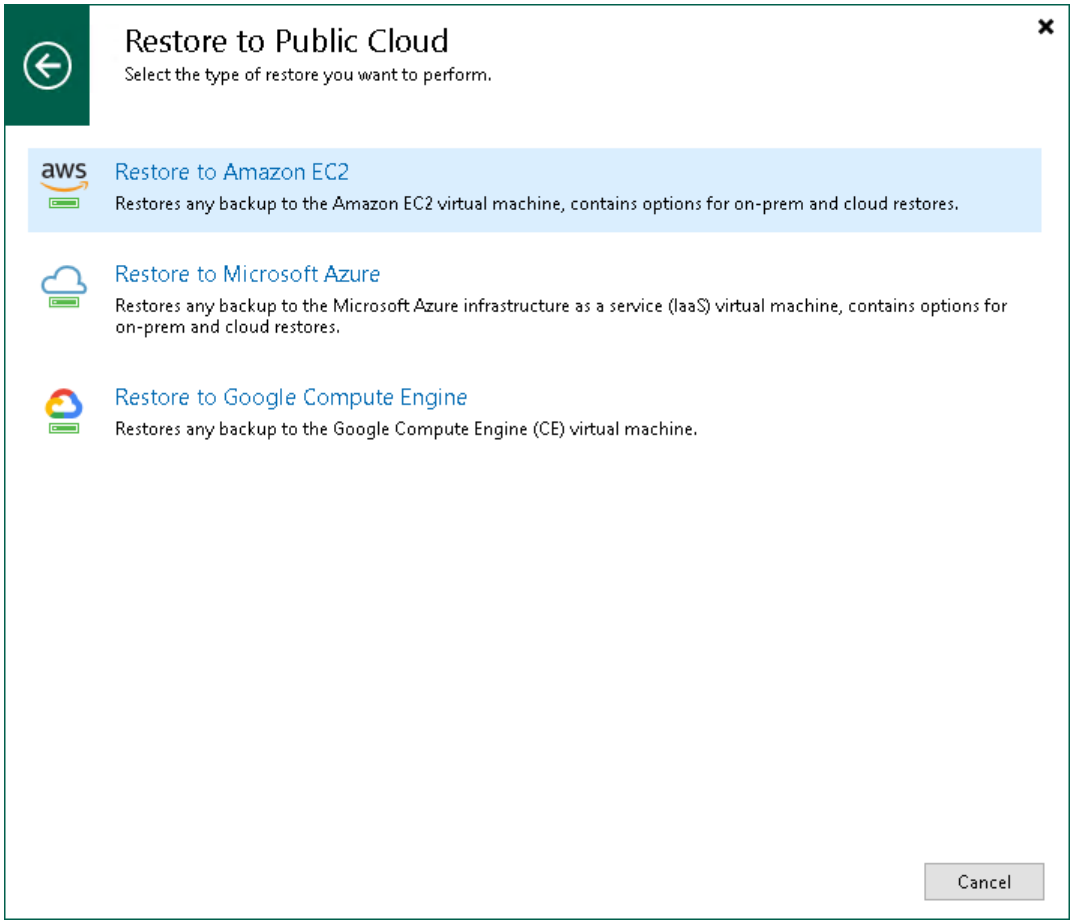

## Step 2. Select Workloads and Restore Points

At the **Machine** step of the wizard, specify workloads that you plan to restore and restore points to which you want to restore the workloads. By default, Veeam Backup & Replication restores workloads to the latest valid restore point in the backup chain.

# Selecting Workloads

To select workloads to restore:

1. On the right of the **Machine** list, click **Add**.
2. In the **Backup Browser** window, expand the necessary backup, select the workloads and click **Add**.

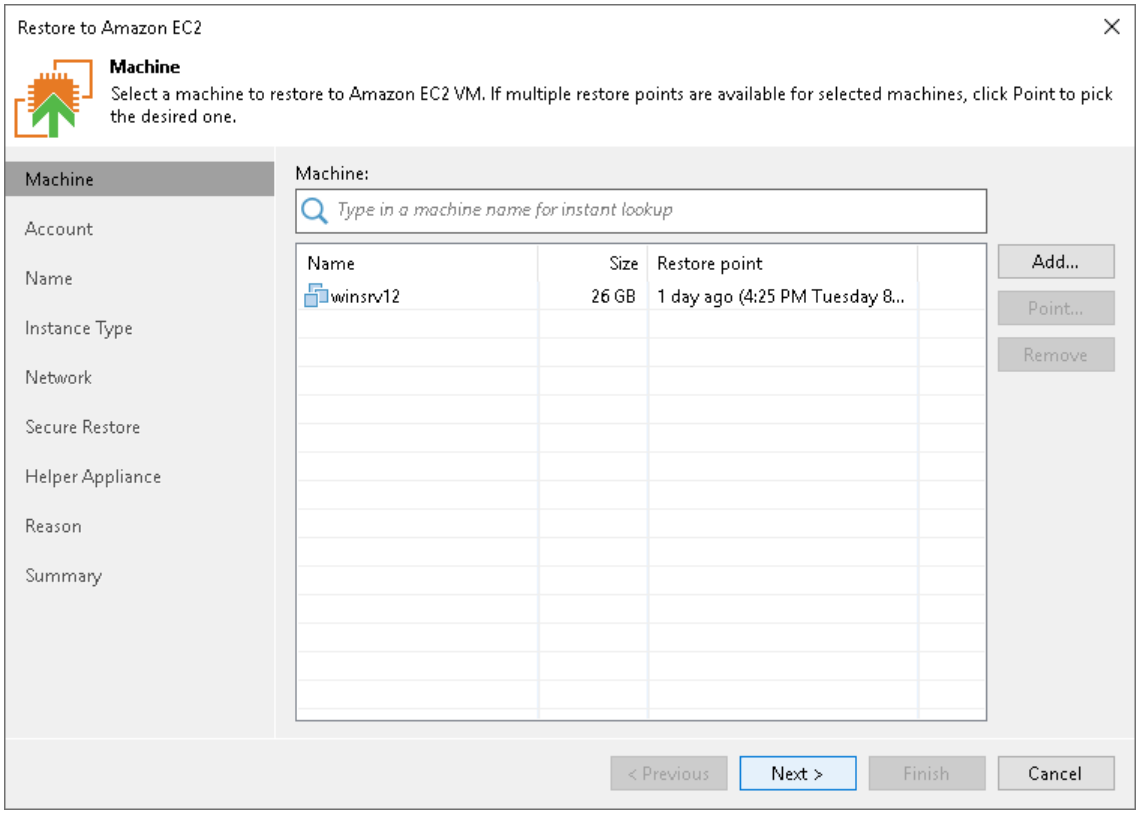


# Selecting Restore Point

To select a restore point for a workload:

1. In the **Machine** list, select a workload.
2. Click **Point** on the right.

3. In the **Restore Points** window, select a restore point to which you want to restore the workload.

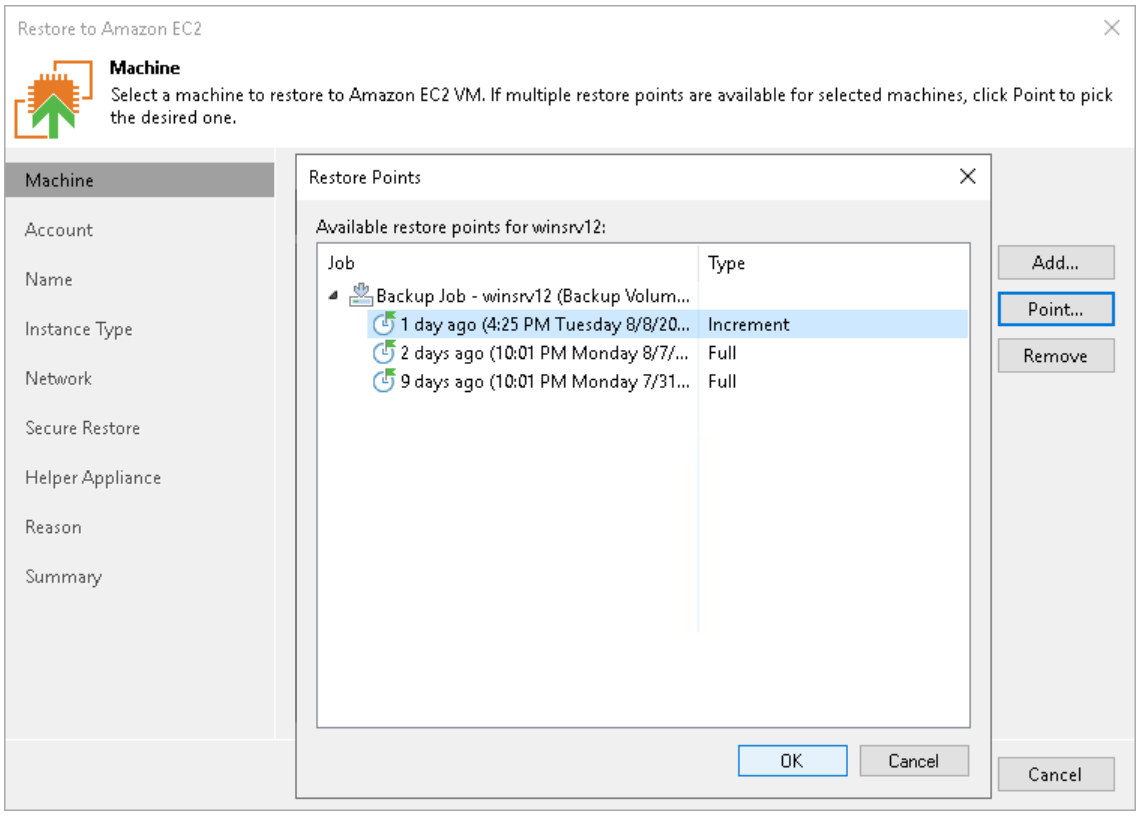

# Step 3. Specify Credentials and Region Settings

At the **Account** step of the wizard, specify AWS user credentials and region:

1. From the **AWS account** list, select credentials of a user account that will be used to connect to AWS. This user account must have permissions listed in section AWS IAM User Permissions.

   When you add AWS user credentials, Veeam Backup & Replication imports information about resources associated with this user. During the restore process, Veeam Backup & Replication accesses these resources and uses them to create a target instance in Amazon EC2.

   If you have not set up credentials beforehand in the Cloud Credentials Manager, click the **Manage accounts** link or click **Add** on the right to add the necessary credentials.

2. From the **AWS region** list, select the AWS region in which Veeam Backup & Replication will restore workloads as Amazon EC2 instances.

3. From the **Datacenter region** list, select the geographic region where Veeam Backup & Replication will restore the workloads.

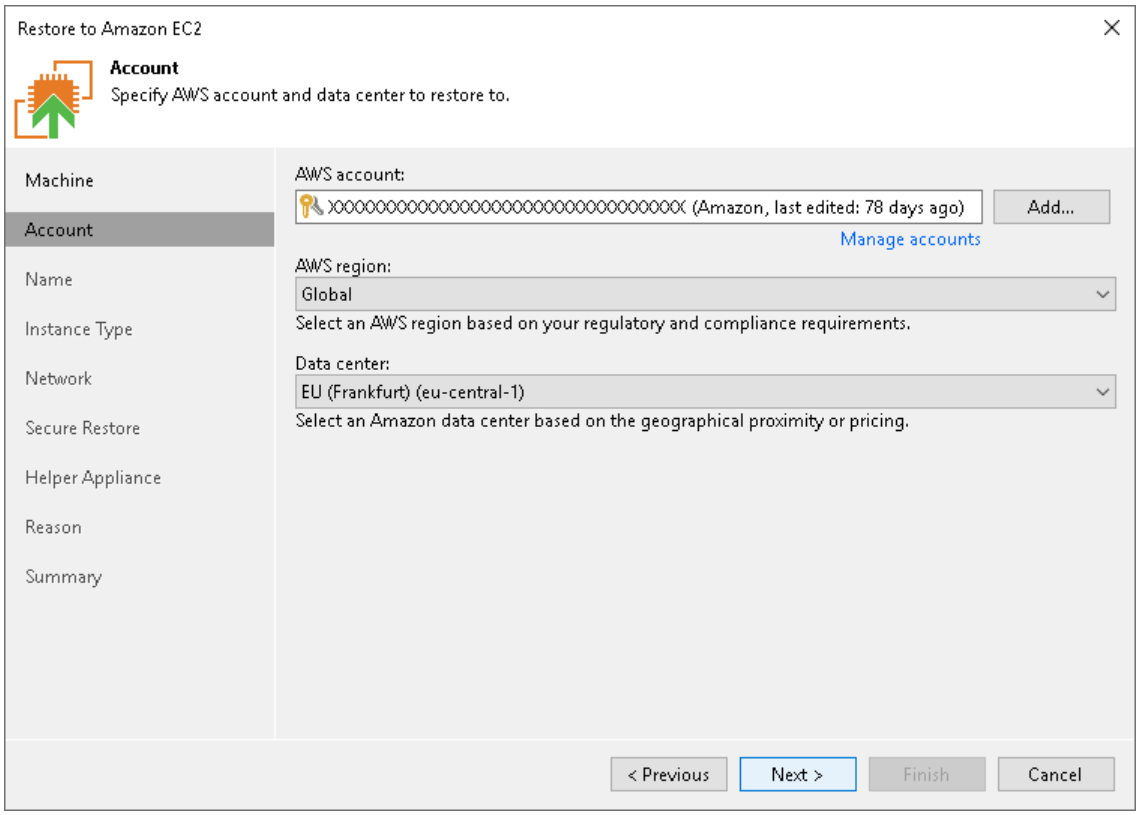

## Step 4. Specify Name and Tags

At the **Name** step of the wizard, you can specify names and manage AWS tags for the restored workloads. By default, Veeam Backup & Replication uses the original workload names and adds the **Original name** and **Restore start time** tags.

# Specifying New Name

To define a new name for a workload that will be restored:

1. In the **Machines to restore** list, select a workload and click **VM name**.
2. In the **Change Name** window, enter a new name explicitly or specify a change name rule — add a prefix and suffix to the original workload name.

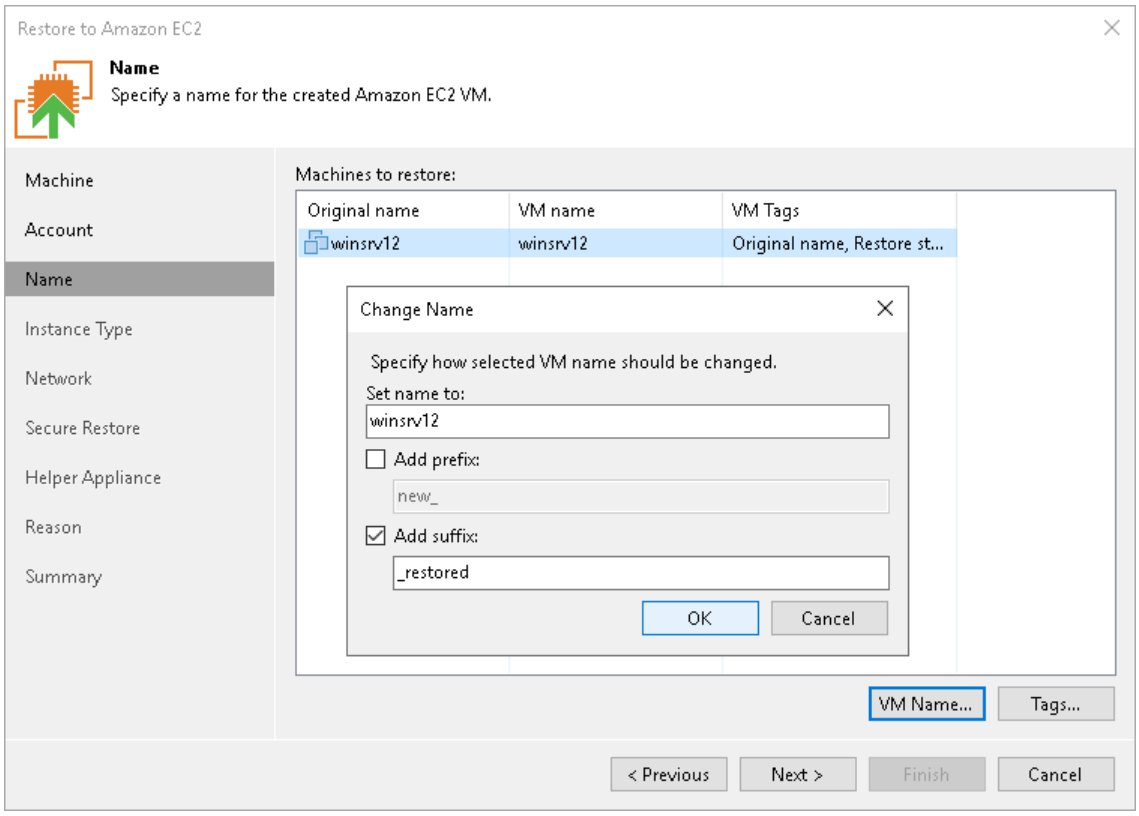


# Managing AWS Tags

You can use AWS tags to categorize instances in Amazon EC2. A tag is a label with metadata that includes two properties: a key and a value. For more information on AWS tags, see the AWS Documentation.

You can modify or delete these tags, or add new ones.

> **NOTE**
>
> If you restore a workload from backups of an Amazon EC2 instance, Veeam Backup & Replication displays tags that were assigned to this instance. You can modify or delete these tags as well.

## Adding Tag

To add a new tag:

1. In the **Machines to restore** list, select a workload and click **Tags**.
2. In the **Tags** window, click **Add**.
3. In the **EC2 VM Tag** window, specify the **Key** and **Value** properties.

   Note that you cannot add the tag with the *Name* key. Veeam Backup & Replication uses the *Name* tag to set the name for the restored EC2 instance in Amazon EC2.

## Modifying Tag

To modify a tag:

1. In the **Machines to restore** list, select a workload and click **Tags**.
2. In the **Tags** window, select the necessary tag and click **Edit**.
3. In the **EC2 VM Tag** window, edit the **Key** or **Value** properties.

## Deleting Tag

To delete a tag:

1. In the **Machines to restore** list, select a workload and click **Tags**.
2. In the **Tags** window, select the necessary tag and click **Remove**.

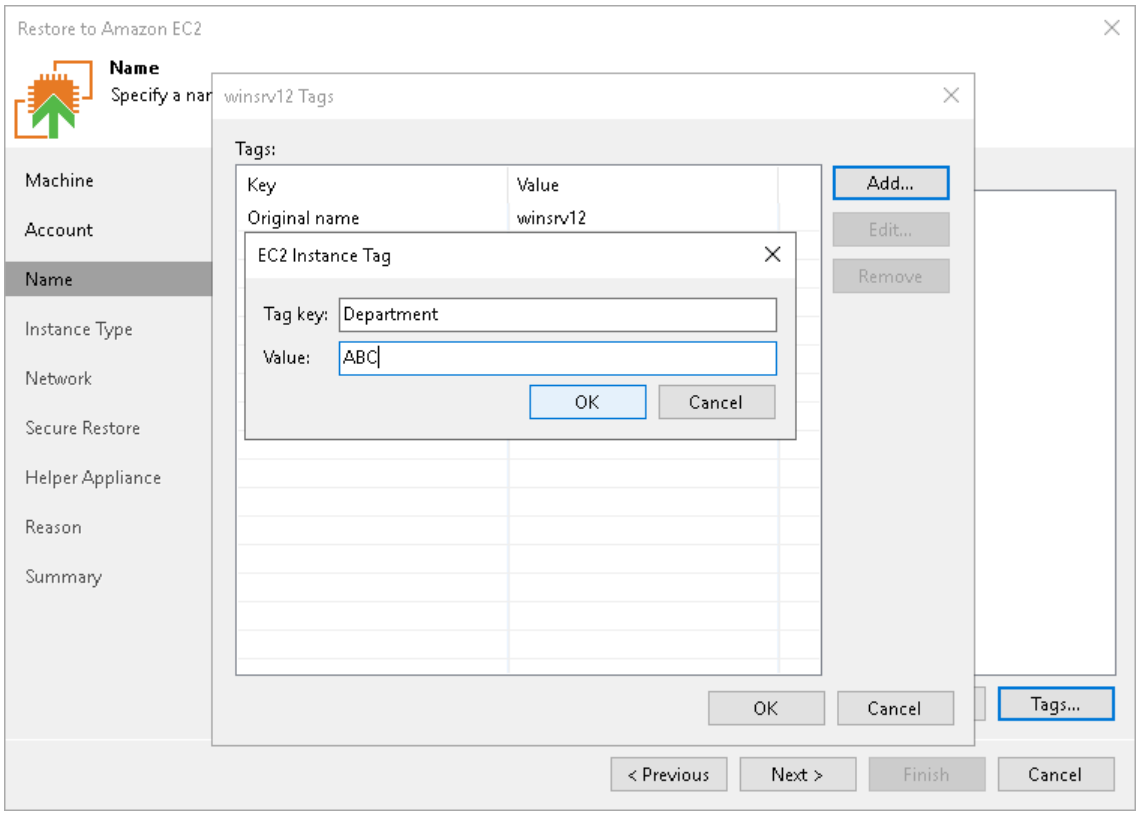

## Step 5. Specify Instance Type and Disks

At the **Instance Type** step of the wizard, you can configure the instance type for the restored workload, select which disks to restore and change their type. By default, Veeam Backup & Replication restores all disks as Amazon Elastic Block Store (EBS) volumes of the General Purpose SSD type. For information on types of EBS volumes, see the AWS Documentation.

# Selecting Instance Type

You can select the amount of computing resources that AWS will provision for your restored workload — an Amazon EC2 instance type. Each instance type offers a unique combination of CPU, memory, storage, and networking resources.

To select an instance type for a workload that will be restored:

1. In the **Virtual machines** list, select a workload and click **Edit**.
2. From the **EC2 instance type** list, select the instance type for the restored workload.

   Make sure that you select the right instance type that corresponds to the initial workload configuration. For information on instance types, see Amazon EC2 Instance Types.

   If you restore an EC2 instance from backups created by Veeam Backup for AWS, Veeam Backup & Replication will identify the type of a backed-up instance and select it by default.
3. From the **OS license** list, select the license policy that AWS will apply for software on the restored workload:
   - [For Linux workloads] The Bring Your Own License (BYOL) policy is used.
   - [For Microsoft Windows workloads] Select one of the following license policies:
     - **Provided by Amazon AWS**. Select this option if you want to obtain licenses for Microsoft software from AWS.

- **Bring Your Own License (BYOL)**. Select this option if you want to use your existing licenses for Microsoft software.

For more information on Microsoft software licensing in AWS, see Microsoft Licensing on AWS.

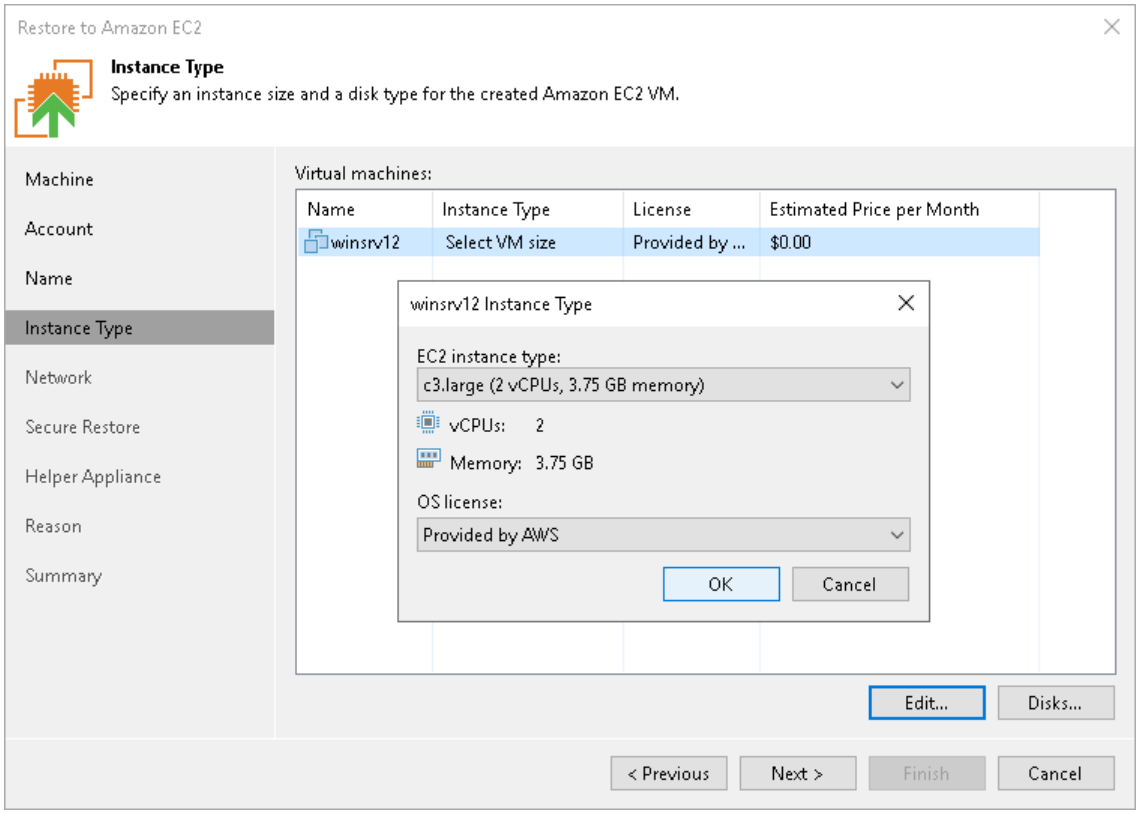


# Selecting Workload Disks to Restore and Changing Their Types

You can restore all disks or specific disks of a workload. You can also change disk types of the restored disks.

To select specific disks and change their type, do the following:

1. In the **Virtual machines** list, select a workload and click **Disks**.
2. In the **Disks To Restore** window, make sure that check boxes next to disks that you want to restore are selected. Clear check boxes next to disks that you do not want to restore.
3. Select a disk whose type you want to change and click **Disk type**.
4. In the **Select EC2 Disk Type** window, choose the disk type.

   If you selected the **Provisioned IOPS SSD (IO1)** type, you can also specify the maximum number of input/output operations per second (IOPS) for the volume. For more information on IOPS, see the AWS Documentation.

> **TIP**
>
> For your convenience, Veeam Backup & Replication uses the AWS Simple Monthly Calculator tools to estimate an approximate price per month for using a selected instance. The estimated price is calculated based on the instance type, license policy and disk configuration.

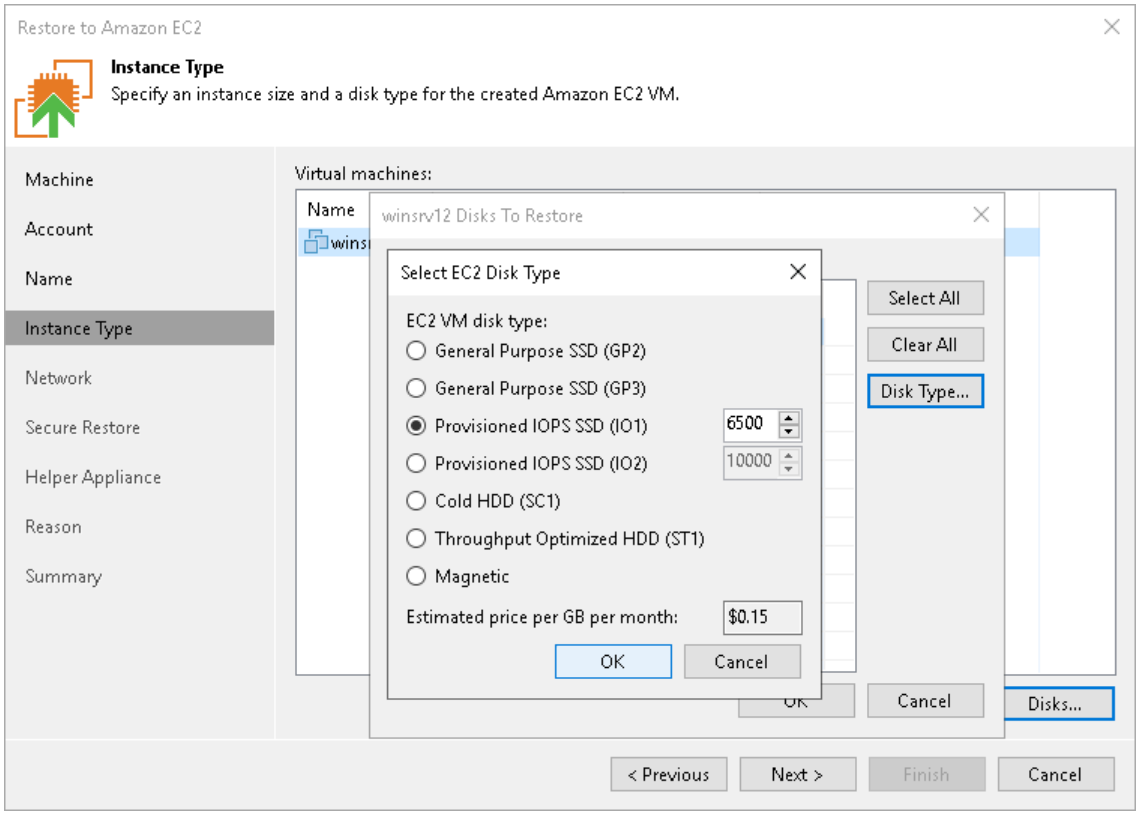
Restore to Amazon EC2
Instance Type
Specify an instance size and a disk type for the created Amazon EC2 VM.
Machine
Account
Name
Instance Type
Network
Secure Restore
Helper Appliance
Reason
Summary
Virtual machines:
Name
winsrv12 Disks To Restore
Select EC2 Disk Type
EC2 VM disk type:
General Purpose SSD (GP2)
General Purpose SSD (GP3)
Provisioned IOPS SSD (IO1)
6500
Provisioned IOPS SSD (IO2)
10000
Cold HDD (SC1)
Throughput Optimized HDD (ST1)
Magnetic
Estimated price per GB per month:
$0.15
OK
Cancel
Select All
Clear All
Disk Type...
Cancel
Disks...
< Previous
Next >
Finish
Cancel

# Step 6. Configure Network Settings

At the **Network** step of the wizard, you can select to which Amazon Virtual Private Cloud (Amazon VPC) the restored workload must be connected. You can also specify a subnet, and a security group — a virtual firewall for the restored EC2 instance. For more information on Amazon VPC, see the AWS Documentation.

To configure network settings for the restored workload, do the following:

1. From the **Amazon VPC** list, select the Amazon VPC where the restored workload will be launched.
2. From the **Subnet** list, select the subnet for the restored workload.
3. From the **Security group list**, select a security group that will be associated with your restored workload.
4. From the **Public IP address** list, select one of the following:
   - **Assign (restored VM will be accessible from the Internet)** — if you want to assign a public IP to the restored workload. For security reasons, make sure traffic filtration rules are properly configured in the security group.
   - **Do not assign (more secure)** — if you do not want to assign a public IP.

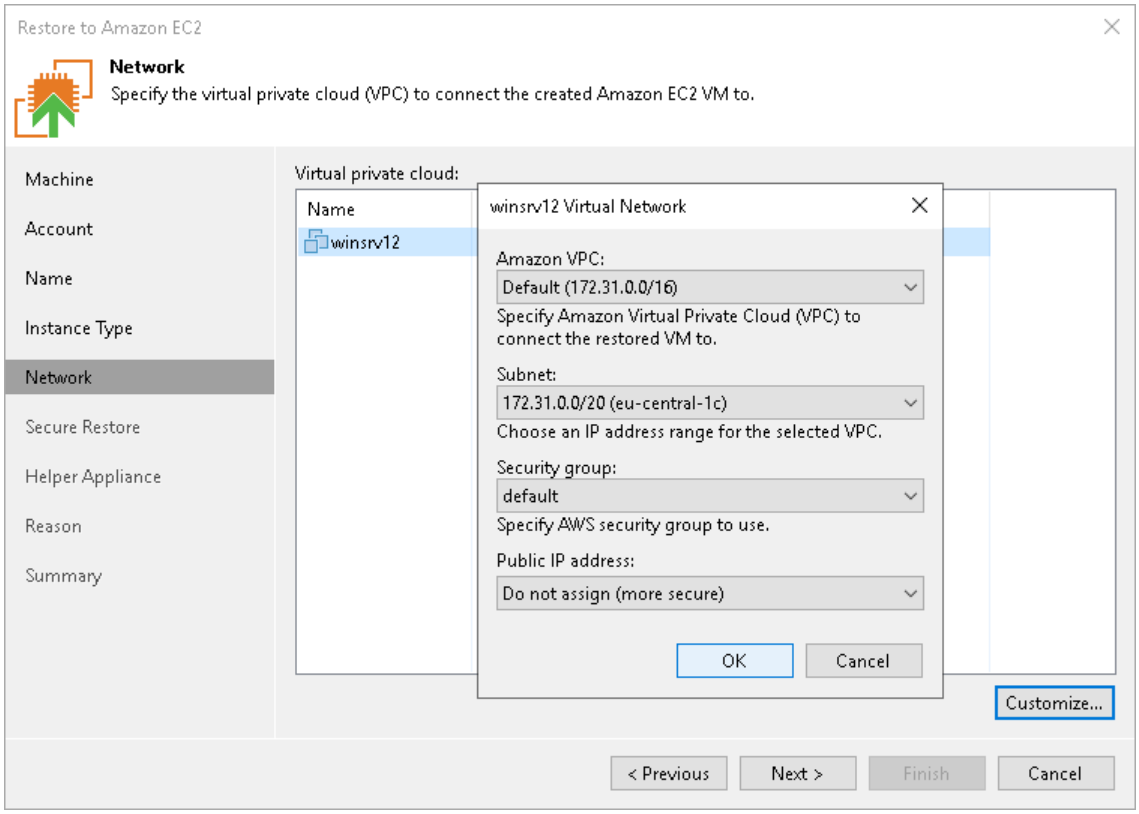

# Step 7. Specify Secure Restore Settings

This step is available if you restore Microsoft Windows-based workloads.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Select which action Veeam Backup & Replication will take if scan finds a virus threat:
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.
   - **Proceed with recovery but connect the VM to a different network**. Select this action if you want to restore the workload to a different AWS security group.

     Click the **Click to change** link to select the security group.

3. Select the **Continue scanning all remaining files after the first occurrence** check box if you want to continue workload scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Antivirus Scan Results.

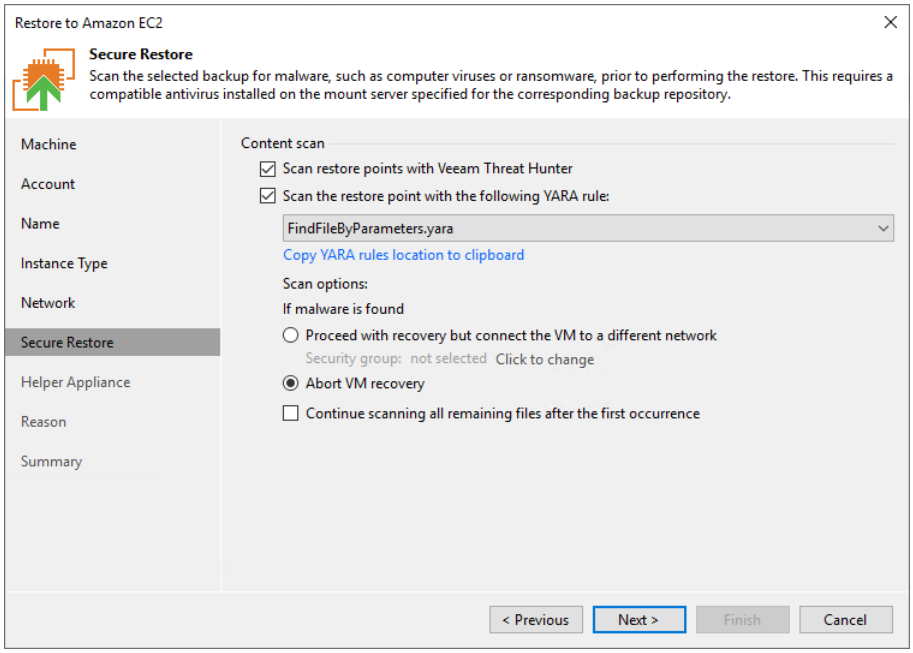

# Step 8. Configure Helper Appliance

At the **Helper Appliance** step of the wizard, you can specify helper appliance settings. A helper appliance is an auxiliary Linux-based instance used to upload disks of a backed-up workload to Amazon EC2. For more information on the helper appliance and requirements for it, see Helper Appliances.

To specify helper appliance settings, do the following:

1. Select the **Use the helper appliance** check box to enable the use of a helper appliance.
2. Click **Customize**.
3. From the **EC2 instance type** list, select the instance type for the helper appliance.
4. From the **Subnet** list, select the subnet for the helper appliance.
5. From the **Security group list**, select a security group that will be associated with the helper appliance.
6. In the **Redirector port** field, specify the port that Veeam Backup & Replication will use to route requests between the helper appliance and backup infrastructure components.

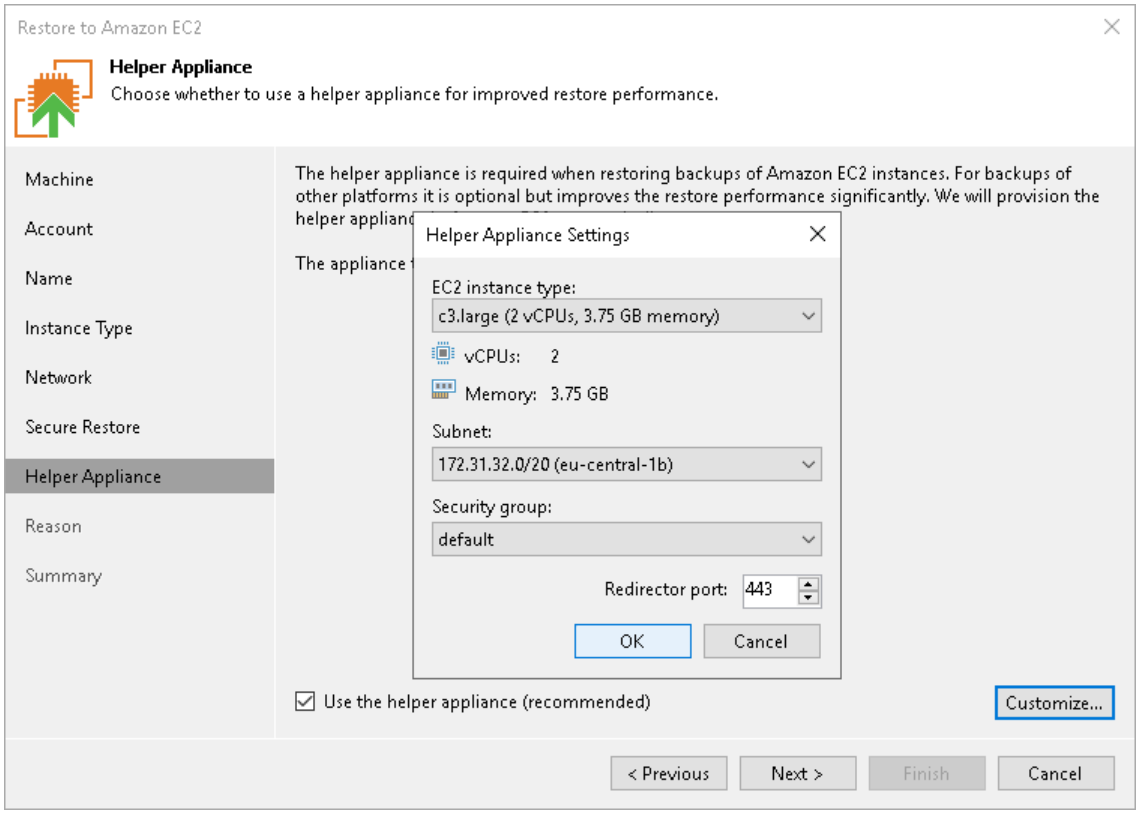

# Step 9. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the workload. The information you provide will be saved in the session history in Veeam Backup & Replication, and you can view it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

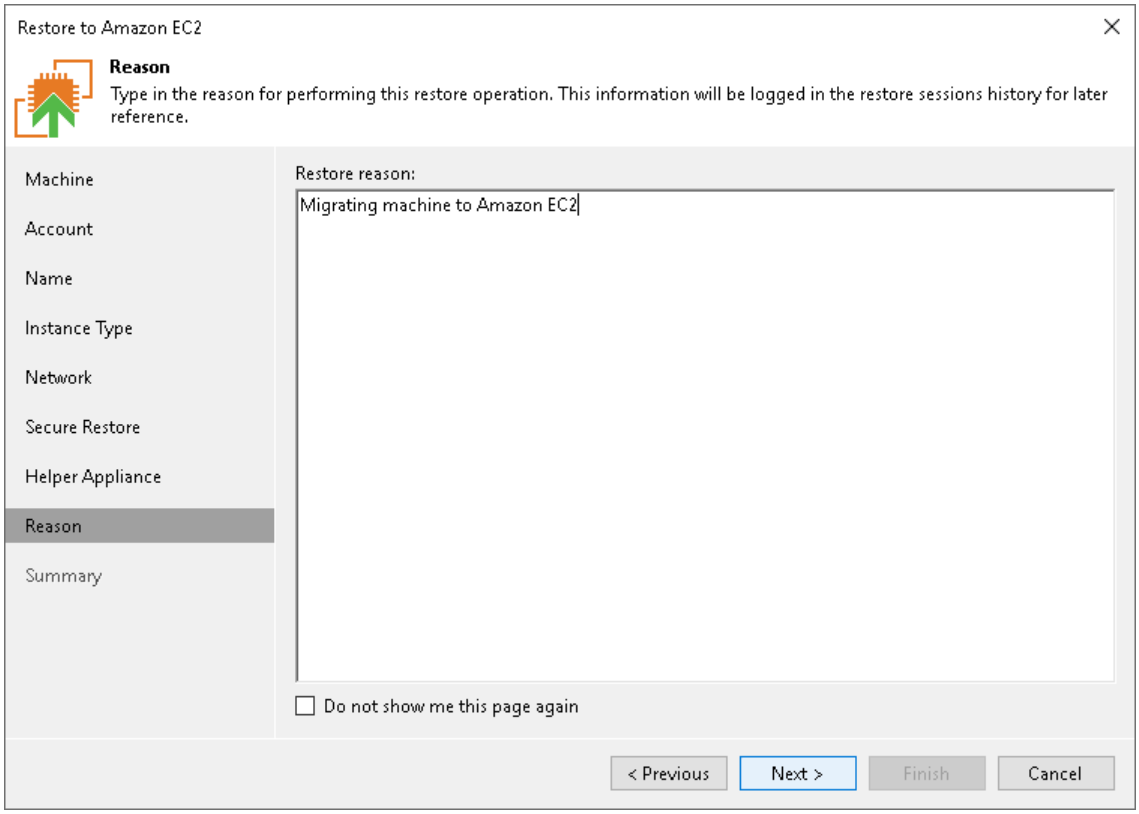

# Step 10. Verify Restore Settings

At the **Summary** step of the wizard, check the specified settings and click **Finish**. If you want to start the EC2 instance right after restore, select the **Power on VM after restoring** check box.

After the wizard closes, you can track the restore process in the **Restore Session** window. If you need to cancel the workload restore, click the **Cancel** restore task link.

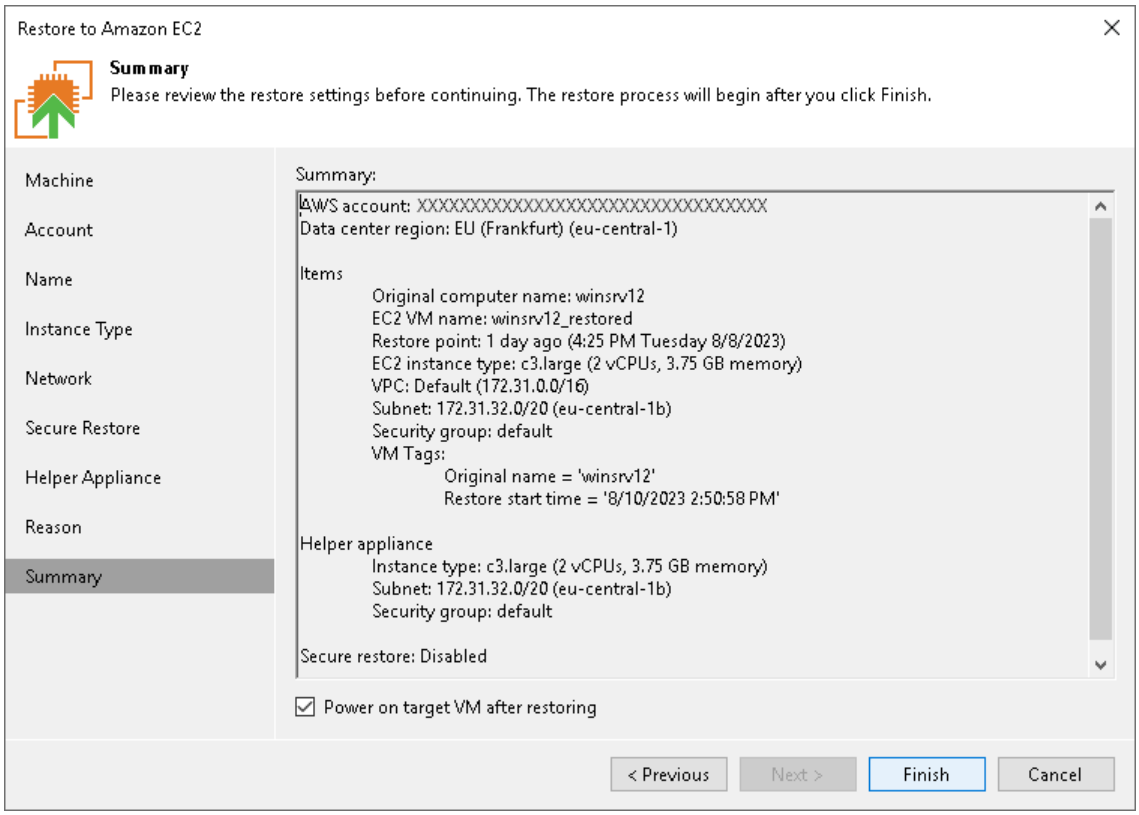

# Restore to Microsoft Azure

Veeam Backup & Replication allows you to restore different workloads (VMs, Google VM instances, physical servers and so on) from backups to Microsoft Azure.

You can use Veeam Backup & Replication to complete the following tasks:

- Restore workloads from Veeam backups to Microsoft Azure.
- Migrate workloads from the on-premises infrastructure to the cloud.
- Create a test environment in the cloud for troubleshooting, testing patches and updates and so on.

## Supported Backup Types

You can restore workloads from the following types of backups:

- Backups of Microsoft Windows and Linux VMs created by Veeam Backup & Replication. You can use backups of VMware vSphere VMs or VMware Cloud Director VMs.
- Backups of Microsoft Windows machines created by Veeam Agent for Microsoft Windows. Backups must be created at the entire machine level or volume level.
- Backups of Linux machines created by Veeam Agent for Linux. Backups must be created at the entire machine level or volume level.
- Backups of Amazon EC2 instances created by Veeam Backup for AWS.
- Backups of Microsoft Azure virtual machines created by Veeam Backup for Microsoft Azure.
- Backups of Google Compute Engine VM instances created by Veeam Backup for Google Cloud.
- Backups of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV.
- Backups of oVirt VMs created by Veeam Backup for OLVM and RHV.
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE.

You can restore a workload to the latest restore point or any previous restore point in the backup chain.

Veeam Backup & Replication employs the Microsoft Azure Resource Manager deployment model.
Veeam Backup & Replication supports batch restore — you can launch the restore process for several VMs at a time.

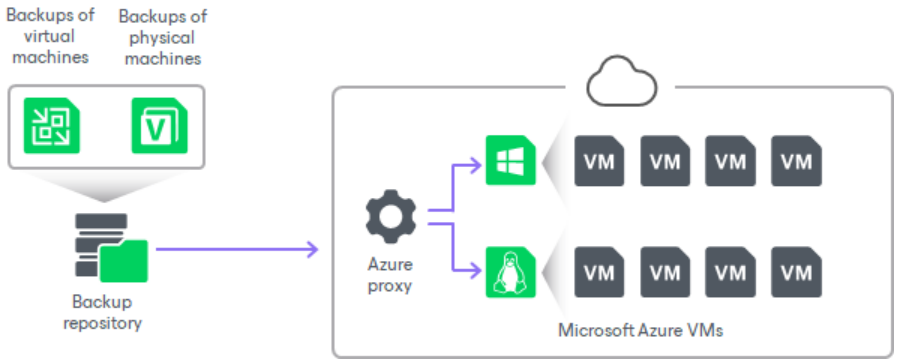

# VM Generation Support

Veeam Backup & Replication chooses the generation for the restored VM based on the boot loader used for the backed-up VM.

If a VM boots using the Unified Extensible Firmware Interface (UEFI) mode, Veeam Backup & Replication restores it as Generation 2. If the VM boots using the Basic Input/Output System (BIOS), Veeam Backup & Replication restores it as Generation 1. If both modes can be used, Veeam Backup & Replication restores the VM as Generation 2.

For more information about Microsoft Azure support for Generation 2 VMs, see Microsoft Docs.

# Considerations and Limitations for Restore to Microsoft Azure

When planning to restore workloads to Microsoft Azure, consider the following general limitations:

- Veeam Backup & Replication supports restore to Microsoft Azure for the following workloads:
  - Microsoft Windows workloads that run Windows Server 2008/Windows Vista and later.
  - Linux workloads (see the Supported Distributions & Versions section in Microsoft Docs).

    **IMPORTANT**

    We strongly recommend having `dracut`, `mkinitrd` or `initramfs` installed on Linux machines that will be restored to Azure. Otherwise, they may not boot after restore.
- Veeam Backup & Replication does not support restoring of disks whose logical sector size is 4096 bytes. Contents of such disks will be unreadable in Microsoft Azure.
- The restore to Microsoft Azure functionality does not support the Azure Hybrid Use Benefit program.
- Veeam Backup & Replication does not support restoring disks encrypted by BitLocker, except for restoring from backups created by Veeam Agent for Microsoft Windows. For more information, see the Veeam Agent for Microsoft Windows User Guide.
- You can use restore to Microsoft Azure for environments with ExpressRoute or site-to-site VPN connectivity to Microsoft Azure. In this case, Azure restore proxy appliances (former Azure proxies) and helper appliances must have private IP addresses. For more information, see this Veeam KB article.
- During restore, Veeam Backup & Replication uses Azure Blob SAS links to the target VM disks. That is why DNS must resolve these Azure Blob endpoints and the firewall rules must allow such traffic.
- If you use a cloud-init-based Linux distribution, we recommend that you use SSH keys on these distributions. If you use a password, it is blocked after restore for security reasons. To reset the password on the restored VM, you need to use the VMAccess extension. For more information, see Microsoft Docs.
- If the system disk of an initial workload uses the GPT partitioning scheme, the number of partitions on the disk cannot exceed 4. During restore such disk will be converted to a disk with the MBR partitioning scheme.
- When you select a storage account whose resources you want to use to store disks of the restored workload, consider the following:
  - Storage accounts with the zone-redundant storage (ZRS), geo-zone-redundant storage (GZRS) and geo-redundant storage (GRS) replication options are not supported. However, read-access geo-redundant storage (RA-GRS) option is supported. For details on replication options, see Microsoft Docs.

- If you plan to use a premium storage account and want to store unmanaged disks there, the restore speed for such disks will be limited to 30 MB/s (approximately).

- When you select a geographic region to which you want to restore workloads, consider that some regions are access restricted to support specific customer scenarios. For example, VMs cannot be created there. To be able to perform different actions in those regions, create a support request in the Azure portal.

  For the full list of access-restricted regions, see Microsoft Docs. The regions are marked with an asterisk *.

- The following applies if you restore Microsoft Azure VMs:

  - Microsoft has strict rules for naming Azure resources. For more information on resource naming requirements, see Microsoft Docs: Resource name rules and Resolve errors for reserved resource names.

  - Veeam Backup & Replication does not support the Azure VM sizes that support only NVMe storage controllers. If you restore a VM with this configuration, the VM will fail to boot. For more information on the unsupported sizes, see Microsoft Docs.

  - Veeam Backup & Replication does not support restore of Azure Ultra Disks.

  - [Unmanaged VM disks] Veeam Backup & Replication supports restoring disks that are equal to or less than 4093 GB. During the restore process, VM disks can increase in size up to 2 GB because of conversion, and Azure supports disk up to 4095 GB. For more information on all disk sizes that Azure supports, see Microsoft Docs.

  - [Managed VM disks] Veeam Backup & Replication supports restoring disks equal or less than 4093 GB for OS disks and equal to or less than 32765 GB for other disks. During the restore process, VM disks can increase in size up to 2 GB because of conversion. For more information on all managed disk sizes that Azure supports, see Microsoft Docs. For more information on OS disk size that Azure supports, see Microsoft Docs. Note that supported disk sizes for Azure and Veeam Backup & Replication differ.

  > **IMPORTANT**
  >
  > The price of a restored VM disk can become higher because of the increase in disk size during the restore process. For more information on pricing, see Managed Disks pricing and Unmanaged Disk and Page Blob pricing.

- [Azure Stack VMs] Veeam Backup & Replication supports restoring of managed and unmanaged disks equal to or less than 1021 GB. This is due to the following reasons: VM disks can increase in size up to 2 GB because of conversion during the restore process; Azure Stack supports disk up to 1023 GB. For more information on all disk sizes that Azure supports, see Microsoft Docs.

  You can change the maximum supported size for unmanaged VM disks with a registry value. For more information, contact Veeam Customer Support.

  > **IMPORTANT**
  >
  > The price of a restored VM disk can become higher because of the increase in disk size during the restore process. For more information on pricing, see Azure Stack Hub Pricing.

- [For restore from backups created by Veeam Agent for Microsoft Windows] Workloads from a backup that contains a failover cluster are restored as separate VMs, not as a cluster. Shared cluster disks of these VMs are restored as regular disks.

# Requirements and Limitations for Generation 2 VM Support

Consider the following requirements and limitations:

- When you select VM sizes at the Specify VM Size step of the **Restore to Azure** wizard, make sure that the selected size is compatible with Generation 2 VMs.
- Generation 2 VMs support only managed disks. Thus, you will need to select the *managed* storage type from the **Storage type** list at the Specify VM Size step of the **Restore to Azure** wizard.

# How Restore to Microsoft Azure Works

The restore process differs for Microsoft Windows and Linux workloads.

> **NOTE**
>
> If you use Veeam Backup for Microsoft Azure and plan to restore Microsoft Azure VMs from restore points that were created using the appliance, the restore works as described in the Performing VM Restore section in Veeam Backup for Microsoft Azure User Guide.

# Restore of Microsoft Windows Workloads

To restore a Microsoft Windows workload, Veeam Backup & Replication performs the following steps:

1. If you use an Azure restore proxy appliance (former Azure proxy) for restore, Veeam Backup & Replication powers on the Azure restore proxy appliance. For more information about the Azure restore proxy appliance, see Managing Azure Restore Proxy Appliances.
2. Veeam Backup & Replication converts disks of a backed-up workload to the VHD format and uploads converted disks to Blob storage in Microsoft Azure.
3. Veeam Backup & Replication mounts uploaded disks to the backup server.
4. Veeam Backup & Replication enables storage controller drivers needed to boot the VM in Microsoft Azure.
5. Veeam Backup & Replication enables remote desktop connections on the restored VM.
6. Veeam Backup & Replication configures Windows Firewall rules to allow incoming remote desktop connections.
7. Veeam Backup & Replication prepares Microsoft Azure agent installation on the restored VM. Installation will start after the VM will be powered on.
8. Veeam Backup & Replication unmounts the uploaded disks from the backup server.
9. If you use an Azure restore proxy appliance for restore, Veeam Backup & Replication powers off the Azure restore proxy appliance after a timeout.

10. Veeam Backup & Replication registers a Microsoft Azure VM with the prepared workload disks. After the registration process is complete, the Microsoft Azure VM is powered on immediately. Then the Microsoft Azure agent is installed on the VM.

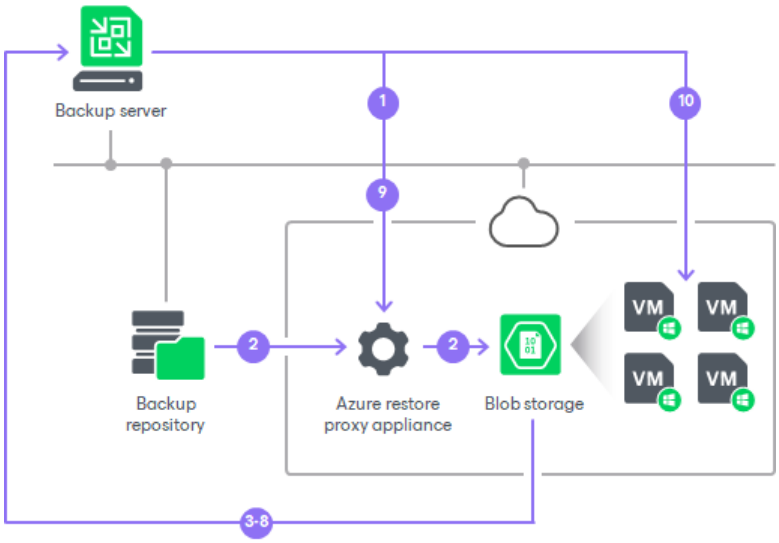


# Restore of Linux Workloads

For restore of Linux workloads, Veeam Backup & Replication uses a helper appliance. The helper appliance is a Linux-based VM in Microsoft Azure registered by Veeam Backup & Replication. During the restore process, Veeam Backup & Replication mounts disks of a backed-up workload to the helper appliance to prepare disks for restore. For more information, see Managing Helper Appliances.

To restore a Linux workload, Veeam Backup & Replication performs the following steps:

1. If you use an Azure restore proxy appliance (former Azure proxy) for restore, Veeam Backup & Replication powers on the Azure restore proxy appliance. For more information about the Azure restore proxy appliance, see Managing Azure Restore Proxy Appliances.
2. Veeam Backup & Replication converts disks of a backed-up workload to the VHD format and uploads converted disks to Blob storage in Microsoft Azure.
3. Veeam Backup & Replication mounts uploaded disks to the helper appliance that resides in the location to which you restore the Linux workload.
4. Veeam Backup & Replication starts the helper appliance with mounted disks.
5. Veeam Backup & Replication makes configuration changes needed for the VM to boot in Microsoft Azure.
6. Veeam Backup & Replication unmounts prepared disks from the helper appliance and powers off the helper appliance.
7. If you use an Azure restore proxy appliance for restore, Veeam Backup & Replication powers off the Azure restore proxy appliance after a timeout.

8. Veeam Backup & Replication registers a Microsoft Azure VM with the prepared workload disks. After the registration process is complete, the VM is powered on immediately.

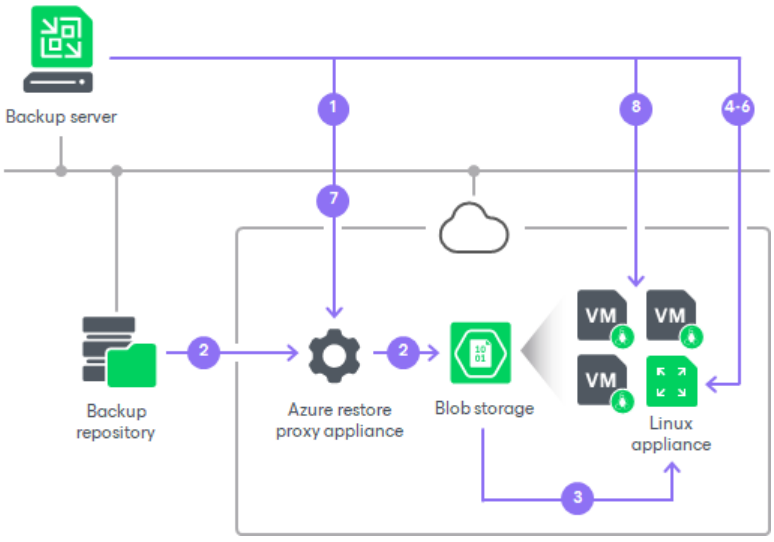


# Configuring Components and Accounts for Restore

Before you restore workloads, you must first add an account to be used for restore and then configure the required components:

1. Add a restore account:
   - Add a Microsoft Azure Compute account if you want to restore workloads to Microsoft Azure.
   - Add a Microsoft Azure Stack Hub Compute account if you want to restore workloads to Azure Stack Hub.

   These accounts must have specific built-in Azure roles (the roles are listed in the sections about adding the accounts). If you do not want to use built-in roles, you can create a custom role with granular permissions. For more information, see Creating Custom Role for Azure and Azure Stack Hub Accounts.
2. [For restore of Linux workloads] Configure helper appliances in Microsoft Azure.

   **NOTE**

   [If you have not changed default credentials for helper appliance] Before you configure the helper appliances, we recommend you to change the default credentials used during the helper appliance deployment. For more information on how to do this, see Changing Credentials for Helper Appliances.
3. [For restore process speed-up] Configure an Azure restore proxy appliance.

**NOTE**

If you use Veeam Backup for Microsoft Azure and plan to restore VMs from restore points that were created using the appliance, you do not need to configure the helper appliance and Azure restore proxy appliance (former Azure proxy). Also, restore to Microsoft Azure works as described in the Performing VM Restore section in Veeam Backup for Microsoft Azure User Guide.

## Creating Custom Role for Azure and Azure Stack Hub Accounts

Granular permissions differ depending on whether you create an Azure Stack Hub account, or Azure Compute account using a new Microsoft Entra ID (formerly Azure Active Directory) application, or Azure Compute account using an existing Microsoft Entra application.

> **NOTE**
>
> This section describes permissions required for Veeam Backup & Replication to perform tasks. If you need to perform other tasks, for example create virtual networks, add the required permissions for those tasks manually.
>
> Instead of granular permissions, you can use built-in roles. For more information, see Permissions.

# Permissions for Azure Compute Account (Existing Application)

If you plan to add an Azure Compute account using an existing Microsoft Entra ID (formerly Azure Active Directory) application (select the **Use the existing account** option at the Access Type step of the wizard), and you do not want to use built-in Azure roles, you can create a custom role with granular permissions:

1. In the Azure Portal, go to subscription properties and open Access control (IAM).
2. Create a custom role from a JSON file as described in Microsoft Docs. Use the following JSON. In the **assignableScopes** field, specify your subscription ID.

   > JSON — Permissions for Existing Application

```
{
   "properties": {
       "roleName": "
Veeam Restore Operat
or",
       "description"
: "Permissions neede
d for an application
for an Azure Compute
Account",
       "assignableSc
opes": [
           "/subscri
ptions/your_subscrip
tion_ID_here"
       ],
       "permissions"
: [
           {
               "acti
ons": [
                   "
Microsoft.Storage/st
orageAccounts/listke
ys/action",
                   "
Microsoft.Storage/st
orageAccounts/read",
                   "
Microsoft.Storage/st
orageAccounts/queueS
ervices/queues/delet
e",
                   "
Microsoft.Storage/st
orageAccounts/queueS
ervices/queues/read"
,
                   "
Microsoft.Storage/st
orageAccounts/queueS
ervices/queues/write
",
                   "
Microsoft.Network/lo
cations/checkDnsName
Availability/read",
                   "
Microsoft.Network/vi
rtualNetworks/read",
                   "
Microsoft.Network/vi
rtualNetworks/subnet
s/join/action",
                   "
Microsoft.Network/pu
blicIPAddresses/read
",
                   "
Microsoft.Network/pu
```

```
blicIPAddresses/writ
e",
                    "
Microsoft.Network/pu
blicIPAddresses/dele
te",
                    "
Microsoft.Network/pu
blicIPAddresses/join
/action",
                    "
Microsoft.Network/ne
tworkInterfaces/read
",
                    "
Microsoft.Network/ne
tworkInterfaces/writ
e",
                    "
Microsoft.Network/ne
tworkInterfaces/dele
te",
                    "
Microsoft.Network/ne
tworkInterfaces/join
/action",
                    "
Microsoft.Network/ne
tworkSecurityGroups/
read",
                    "
Microsoft.Network/ne
tworkSecurityGroups/
write",
                    "
Microsoft.Network/ne
tworkSecurityGroups/
delete",
                    "
Microsoft.Network/ne
tworkSecurityGroups/
join/action",
                    "
Microsoft.Compute/lo
cations/vmSizes/read
",
                    "
Microsoft.Compute/lo
cations/usages/read"
,
                    "
Microsoft.Compute/vi
rtualMachines/read",
                    "
Microsoft.Compute/vi
rtualMachines/write"
,
                    "
Microsoft.Compute/vi
```

```
rtualMachines/delete
",
                    "
Microsoft.Compute/vi
rtualMachines/start/
action",
                    "
Microsoft.Compute/vi
rtualMachines/deallo
cate/action",
                    "
Microsoft.Compute/vi
rtualMachines/instan
ceView/read",
                    "
Microsoft.Compute/vi
rtualMachines/extens
ions/read",
                    "
Microsoft.Compute/vi
rtualMachines/extens
ions/write",
                    "
Microsoft.Compute/vi
rtualMachines/conver
tToManagedDisks/acti
on",
                    "
Microsoft.Compute/di
sks/read",
                    "
Microsoft.Compute/di
sks/write",
                    "
Microsoft.Compute/di
sks/beginGetAccess/a
ction",
                    "
Microsoft.Compute/di
sks/delete",
                    "
Microsoft.Compute/di
sks/endGetAccess/act
ion",
                    "
Microsoft.Resources/
checkResourceName/ac
tion",
                    "
Microsoft.Resources/
subscriptions/resour
ceGroups/read",
                    "
Microsoft.Resources/
subscriptions/resour
ceGroups/write",
                    "
Microsoft.Resources/
subscriptions/locati
ons/read",
```

```
                    "
Microsoft.Marketplac
eOrdering/offerTypes
/publishers/offers/p
lans/agreements/read
",
                    "
Microsoft.Marketplac
eOrdering/offerTypes
/publishers/offers/p
lans/agreements/writ
e"
                ],
                "notA
ctions": [],
                "data
Actions": [
                    "
Microsoft.KeyVault/v
aults/keys/encrypt/a
ction",
                    "
Microsoft.KeyVault/v
aults/keys/decrypt/a
ction",
                    "
Microsoft.KeyVault/v
aults/keys/read",
                    "
Microsoft.Storage/st
orageAccounts/queueS
ervices/queues/messa
ges/delete",
                    "
Microsoft.Storage/st
orageAccounts/queueS
ervices/queues/messa
ges/read",
                    "
Microsoft.Storage/st
orageAccounts/queueS
ervices/queues/messa
ges/write",
                    "
Microsoft.Storage/st
orageAccounts/queueS
ervices/queues/messa
ges/process/action"
                ],
                "notD
ataActions": []
            }
        ]
    }
}
```

3. Assign the created role to the required Microsoft Entra application. For details, see the Manage access to Azure resources using RBAC and the Azure portal section in the RBAC for Azure resources documentation.
4. At the Account Type step of the **Microsoft Azure Compute Account** wizard, select **Use existing account.**

5. At the Subscription step of the wizard, specify the Azure Microsoft Entra application with the assigned role.

# Permissions for Azure Compute Account (New Application)

If you plan to add an Azure Compute account using a new Microsoft Entra ID (formerly Azure Active Directory) application (select the **Create a new account** option at the Subscription step of the wizard), and you do not want to use built-in Azure roles, you can create a custom role with granular permissions:

1. In the Azure Portal, go to subscription properties and open Access control (IAM).
2. Create a custom role from a JSON file as described in Microsoft Docs. Use the following JSON. In the **assignableScopes** field, specify your subscription ID.

   > JSON — Permissions for New Application

```
{
   "properties": {
       "roleName": "
Veeam Register Azure
Compute Account usin
g new Microsoft Entr
a application",
       "description"
: "Permissions neede
d for a user to add
an Azure Compute Acc
ount based on new Mi
crosoft Entra applic
ation",
       "assignableSc
opes": [
           "/subscri
ptions/00000000-0000
-0000-0000-000000000
000"
       ],
       "permissions"
: [
           {
               "acti
ons": [
                   "
Microsoft.Authorizat
ion/roleDefinitions/
read",
                   "
Microsoft.Authorizat
ion/roleAssignments/
read",
                   "
Microsoft.Authorizat
ion/roleAssignments/
write"
               ],
               "notA
ctions": [],
               "data
Actions": [],
               "notD
ataActions": []
           }
       ]
   }
}
```

3. Assign the created role to the required Microsoft Entra user. For details, see the Manage access to Azure resources using RBAC and the Azure portal section in the RBAC for Azure resources documentation.
4. At the Account Type step of the **Microsoft Azure Compute Account** wizard, select **Create a new account.**
5. At the **Subscription** step, configure the account as described in section Creating New Entra ID Application. On the Microsoft Azure device authentication page, specify an Microsoft Entra user account with the assigned role.

NOTE

The described permissions are required for assigning a role on the subscription level for the registered application. Also, privileges to register applications are required. For more information, see Permissions.

# Permissions for Azure Stack Hub Compute Account (Existing Application)

If you plan to add an Azure Stack Hub account using an existing Microsoft Entra application (select the **Use the existing account** option at the Subscription step of the wizard), and you do not want to use built-in Azure roles, you can create a custom role with granular permissions:

1. In the Azure Stack Hub management portal, go to subscription properties and open Access control (IAM).
2. Create a custom role from a JSON file as described in Microsoft Docs. Use the following JSON. In the **assignableScopes** field, specify your subscription ID.

   > JSON — Permissions for Existing Application

```
{
   "properties": {
       "roleName": "
Veeam Restore Operat
or",
       "description"
: "Permissions neede
d for an application
for an Azure Compute
Account",
       "assignableSc
opes": [
           "/subscri
ptions/your_subscrip
tion_ID_here"
       ],
       "permissions"
: [
           {
               "acti
ons": [
                   "
Microsoft.Storage/st
orageAccounts/listke
ys/action",
                   "
Microsoft.Storage/st
orageAccounts/read",
                   "
Microsoft.Network/lo
cations/checkDnsName
Availability/read",
                   "
Microsoft.Network/vi
rtualNetworks/read",
                   "
Microsoft.Network/vi
rtualNetworks/subnet
s/join/action",
                   "
Microsoft.Network/pu
blicIPAddresses/read
",
                   "
Microsoft.Network/pu
blicIPAddresses/writ
e",
                   "
Microsoft.Network/pu
blicIPAddresses/dele
te",
                   "
Microsoft.Network/pu
blicIPAddresses/join
/action",
                   "
Microsoft.Network/ne
tworkInterfaces/read
",
```

```
                    "
Microsoft.Network/ne
tworkInterfaces/writ
e",
                    "
Microsoft.Network/ne
tworkInterfaces/dele
te",
                    "
Microsoft.Network/ne
tworkInterfaces/join
/action",
                    "
Microsoft.Network/ne
tworkSecurityGroups/
read",
                    "
Microsoft.Network/ne
tworkSecurityGroups/
write",
                    "
Microsoft.Network/ne
tworkSecurityGroups/
delete",
                    "
Microsoft.Network/ne
tworkSecurityGroups/
join/action",
                    "
Microsoft.Compute/lo
cations/vmSizes/read
",
                    "
Microsoft.Compute/lo
cations/usages/read"
,
                    "
Microsoft.Compute/vi
rtualMachines/read",
                    "
Microsoft.Compute/vi
rtualMachines/write"
,
                    "
Microsoft.Compute/vi
rtualMachines/delete
",
                    "
Microsoft.Compute/vi
rtualMachines/start/
action",
                    "
Microsoft.Compute/vi
rtualMachines/deallo
cate/action",
                    "
Microsoft.Compute/vi
rtualMachines/extens
ions/read",
```

```
                    "
Microsoft.Compute/vi
rtualMachines/extens
ions/write",
                    "
Microsoft.Resources/
checkResourceName/ac
tion",
                    "
Microsoft.Resources/
subscriptions/resour
ceGroups/read",
                    "
Microsoft.Resources/
subscriptions/resour
ceGroups/write",
                    "
Microsoft.Resources/
subscriptions/locati
ons/read"
                ],
                "notA
ctions": [],
                "data
Actions": [],
                "notD
ataActions": []
            }
        ]
    }
}
```

3. Assign the created role to the required Microsoft Entra application. For details, see the Manage access to Azure resources using RBAC and the Azure portal section in the RBAC for Azure resources documentation.
4. At the Account Type step of the **Microsoft Azure Compute Account** wizard, select **Use existing account.**
5. At the Subscription step of the wizard, specify the Microsoft Entra application with the assigned role.

# Permissions for Azure Stack Hub Compute Account (New Application)

If you plan to add an Azure Stack Hub account using a new Microsoft Entra ID (formerly Azure Active Directory) application (select the **Create a new account** option at the Subscription step of the wizard), and you do not want to use built-in Azure roles, you can create a custom role with granular permissions:

1. In the Azure Stack Hub management portal, go to subscription properties and open Access control (IAM).
2. Create a custom role from a JSON file as described in Microsoft Docs. Use the following JSON. In the **assignableScopes** field, specify your subscription ID.

   > JSON — Permissions for New Application

```
{
   "properties": {
       "roleName": "
Veeam Register Azure
Compute Account usin
g new Microsoft Entr
a application",
       "description"
: "Permissions neede
d for a user to add
an Azure Compute Acc
ount based on new Mi
crosoft Entra applic
ation",
       "assignableSc
opes": [
           "/subscri
ptions/00000000-0000
-0000-0000-000000000
000"
       ],
       "permissions"
: [
           {
               "acti
ons": [
                   "
Microsoft.Authorizat
ion/roleDefinitions/
read",
                   "
Microsoft.Authorizat
ion/roleAssignments/
read",
                   "
Microsoft.Authorizat
ion/roleAssignments/
write"
               ],
               "notA
ctions": [],
               "data
Actions": [],
               "notD
ataActions": []
           }
       ]
   }
}
```

3. Assign the created role to the required Microsoft Entra user. For details, see the Manage access to Azure resources using RBAC and the Azure portal section in the RBAC for Azure resources documentation.
4. At the Account Type step of the **Microsoft Azure Compute Account** wizard, select **Create a new account.**
5. At the **Subscription** step, configure the account as described in section Creating New Entra ID Application. On the Microsoft Azure device authentication page, specify an Microsoft Entra user account with the assigned role.

> **NOTE**
>
> The described permissions are required for assigning a role on the subscription level for the registered application. Also, privileges to register applications are required. For more information, see Permissions.

# Managing Helper Appliances

Helper appliances are Linux-based VMs in Microsoft Azure registered by Veeam Backup & Replication. Helper appliances are required to restore Linux workloads to Microsoft Azure. During the restore process, Veeam Backup & Replication mounts disks of a backed-up workload to a helper appliance to prepare disks for restore.

To instruct Veeam Backup & Replication to deploy helper appliances, you must configure them as described in section Configuring Helper Appliances. When deploying the helper appliance, Veeam Backup & Replication analyzes the VM sizes available in the restore region and selects the cheapest and smallest size suitable for using the helper appliance.

Helper appliances are persistent. After the restore process finishes, helper appliances get powered off and remain in Microsoft Azure. The appliances remain in the powered off state until you start a new restore process. Note that Microsoft Azure will bill you for storing helper appliances disks in the storage account. To remove a helper appliance, follow the instruction provided in Removing Helper Appliances.

## Configuring Helper Appliances

Before you configure helper appliances, consider the following:

- Veeam Backup & Replication uses its built-in credentials record to work with all helper appliances. For security reasons, we recommended that you change a password for this account before you set up the helper appliances. Changing credentials is required only once. For more information, see Changing Credentials for Helper Appliances.
- If you plan to restore Linux workloads to different locations, you must configure a helper appliance in each location to which workloads will be restored.

Helper appliances are configured when you add a Microsoft Azure compute or Microsoft Azure Stack account. For more information, see the **Helper Appliance** step description in the Microsoft Azure Compute Accounts or Microsoft Azure Stack Hub Compute Accounts section.

## Changing Credentials for Helper Appliances

By default, Veeam Backup & Replication uses its built-in credentials record to work with all helper appliances in Microsoft Azure and Azure Stack Hub. You can find this credential record in the credentials manger in the Veeam Backup & Replication console: the *root* account with "Azure helper appliance credentials" in the **Description** column.

For security reasons, we recommend you to change the password for this credentials record before you set up helper appliances.

> **IMPORTANT**
>
> When you change the password in the built-in credentials record, you must re-deploy all existing helper appliances in Microsoft Azure and Azure Stack Hub. To redeploy appliances, you must remove all configured appliances and then configure them once again.

To change the password in the credentials record for the helper appliances:

1. From the main menu, select **Credentials and Passwords** > **Datacenter Credentials**.
2. In the **Manage Credentials** window, click the built-in credentials record for the helper appliances.
3. Click **Edit**.
4. In the **Password** field, specify a new password.
5. Click **OK** to save changes.

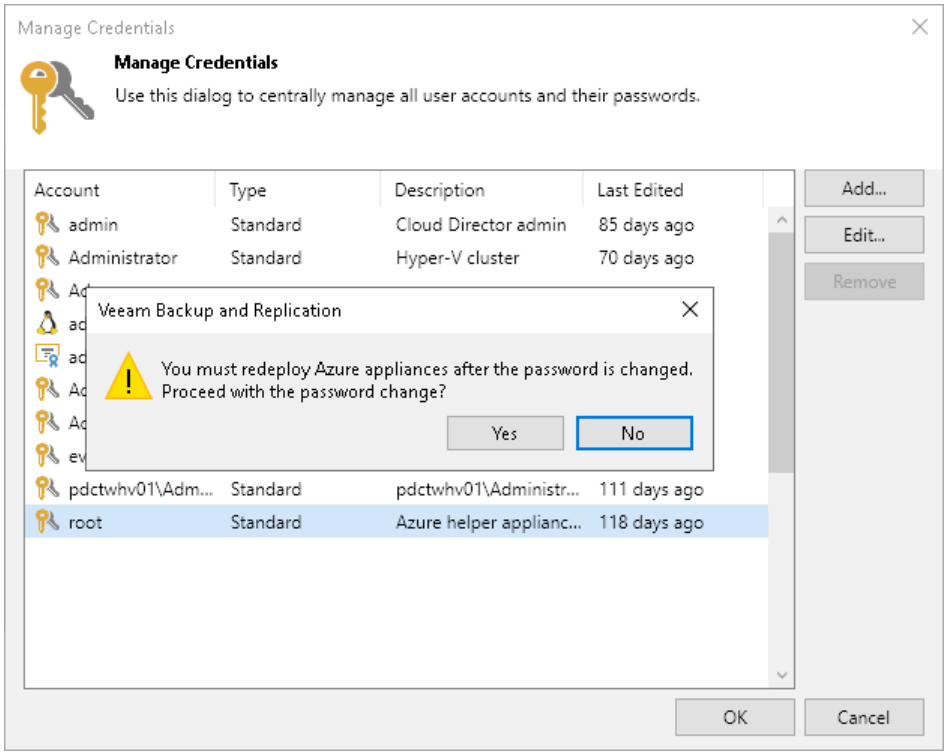


## Removing Helper Appliances

If you no longer need to restore Linux workloads to Azure or Azure Stack Hub, you can remove helper appliances:

1. From the main menu, select **Credentials & Passwords** > **Cloud Credentials**.
2. In the accounts list, select the Azure account and click **Edit**.
3. Pass to the **Helper Appliance** step of the **Microsoft Azure Compute Account** wizard.
4. In the **Helper appliances** list, select the helper appliance and click **Remove**.

> **IMPORTANT**
>
> Do not clear the **Enable restore of Linux-based computers** check box at the **Subscription** step of the wizard before removing helper appliances. If you clear the check box, the **Microsoft Azure Compute Account** wizard will not display the **Helper Appliance** step. Helper appliances themselves will remain in Microsoft Azure. Clear the **Enable restore of Linux-based computers** check box only after you remove the helper appliances.
>
> You can remove all helper appliances using the Deploy-VBRAzureLinuxRestoreAppliance PowerShell cmdlet.

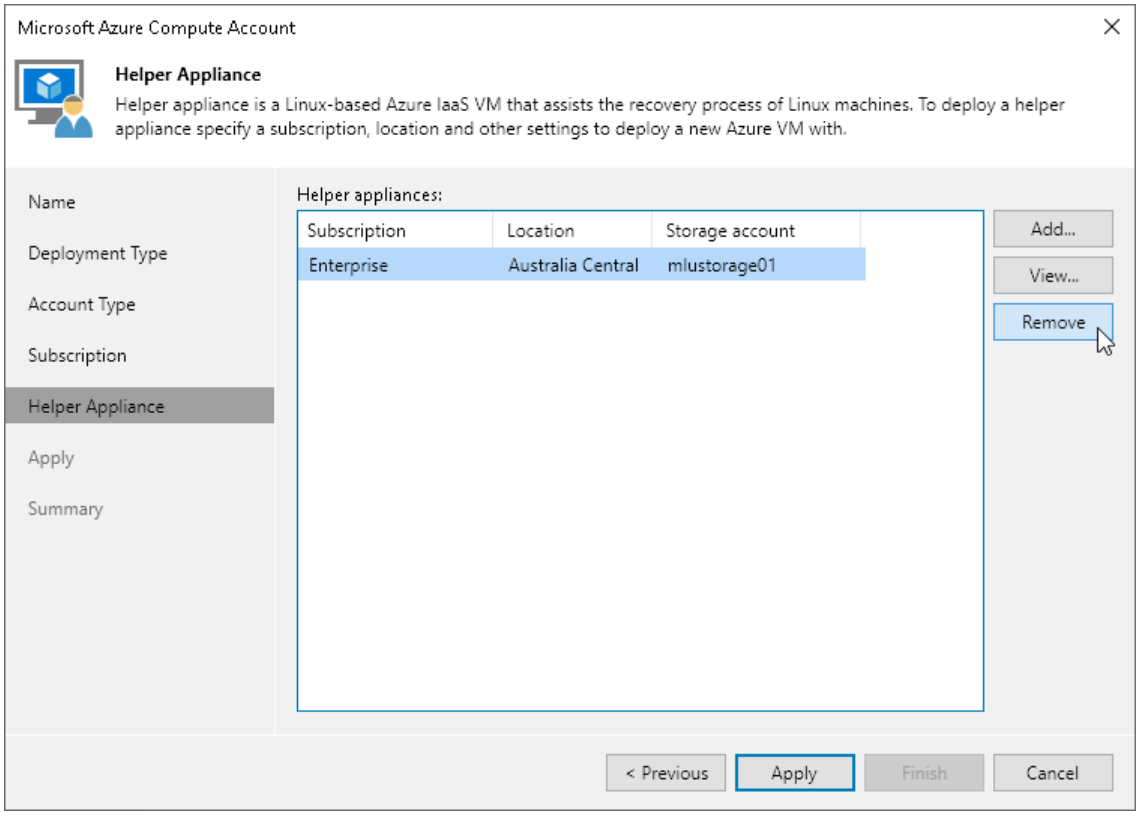


# Managing Azure Restore Proxy Appliances

Azure restore proxy appliances (former Azure proxies) are Windows-based VMs over which Veeam Backup & Replication transports VM disk data to Blob storage. Veeam Backup & Replication uses Azure restore proxy appliances during restore of Windows-based and Linux-based workloads.

Azure restore proxy appliances help speed up the restore process especially if you restore workloads to a distant location or the network connection is slow. Veeam components installed on an Azure restore proxy appliance compress and deduplicate disk data, which helps reduce network traffic.

Although Azure restore proxy appliances are optional, we recommend you to configure them. Azure restore proxy appliances do not require a lot of resources but can significantly improve restore performance. Configure an Azure restore proxy appliance in a location to which you plan to restore workloads or close to this location. If you plan to restore workloads to different locations, configure at least one Azure restore proxy appliance in each location.

The process of Azure restore proxy appliance deployment takes some time. We recommend you to configure the Azure restore proxy appliance in advance, before you start the restore process. To configure an Azure restore proxy appliance, use the **New Azure Restore Proxy Appliance** wizard as described in section Configuring Azure Restore Proxy Appliances. Veeam Backup & Replication will deploy a Microsoft Windows Server machine in Microsoft Azure and will assign the role of the Azure restore proxy appliance to this machine. You can then instruct Veeam Backup & Replication to use the Azure restore proxy appliance for restore tasks.

The Azure restore proxy appliance is persistent. After the restore process finishes, the Azure restore proxy appliance gets powered off and remains in Microsoft Azure. The Azure restore proxy appliance remains in the powered off state until a new restore process is started. Note that Microsoft Azure will bill you for storing Azure restore proxy appliance disks in the storage account. To remove an Azure restore proxy appliance, follow the instruction provided in Removing Azure Restore Proxy Appliances.

## Configuring Azure Restore Proxy Appliances

To configure an Azure restore proxy appliance, use the **New Azure Restore Proxy Appliance** wizard.

## Before You Begin

Before you configure an Azure restore proxy appliance, check the following prerequisites:

- You must add information about your Microsoft Azure Compute account or Microsoft Azure Stack Hub Compute account to Veeam Backup & Replication. For more information, see Adding Microsoft Azure Compute Accounts or Adding Azure Stack Hub Compute Accounts.
- You must configure the following objects in Microsoft Azure beforehand:
  - Storage account whose resources will be used to deploy the Azure restore proxy appliance.
  - Networks to which you plan to connect the Azure restore proxy appliance.

  For storage accounts and network configuration, you must use the same deployment model that you plan to use for Azure restore proxy appliance creation.

> **IMPORTANT**
>
> When you deploy Azure restore proxy appliance for Azure Stack Hub, make sure that Windows Server 2019 is available in Azure marketplace.

## Step 1. Launch New Azure Restore Proxy Appliance Wizard

To launch the **New Azure Restore Proxy Appliance** wizard, do one of the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Proxies** node and select **Add Proxy**. Alternatively, you can click **Add Proxy** on the ribbon.
3. In the **Add Backup Proxy** window, select **Azure proxy appliance**.

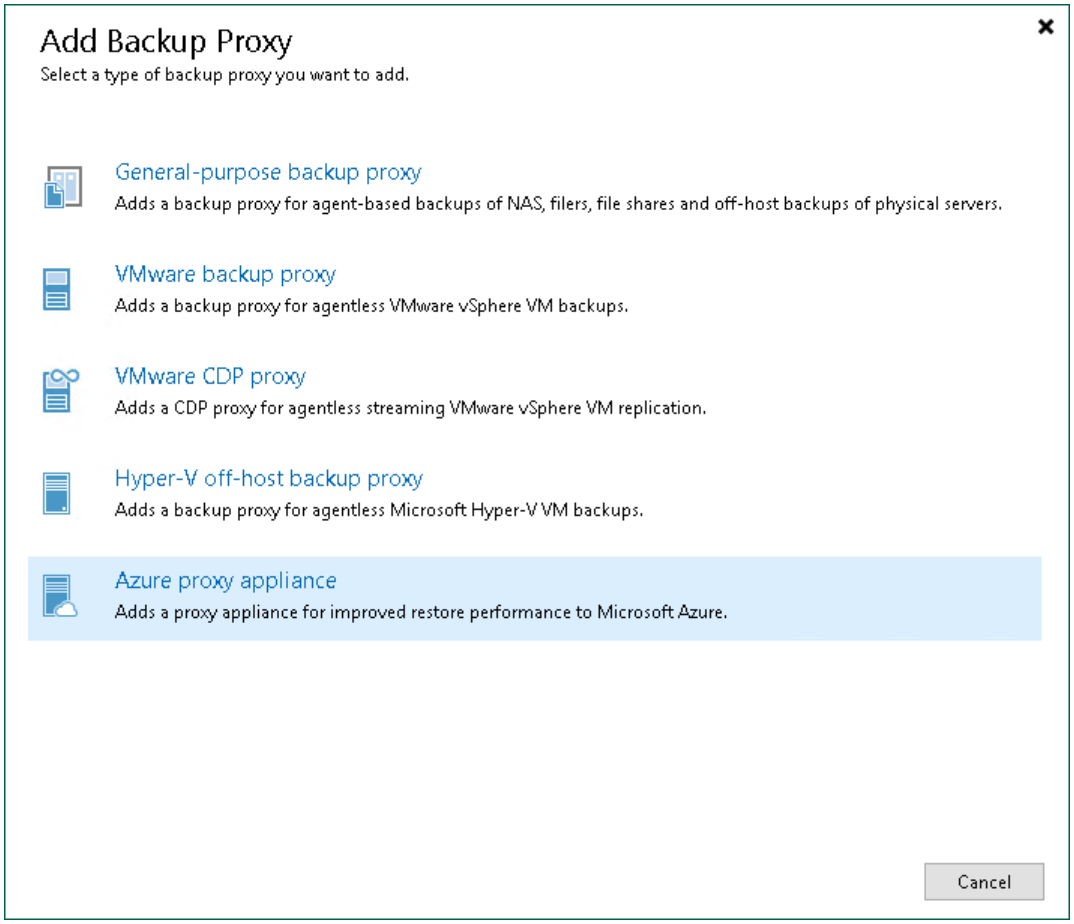

## Step 2. Specify Azure Restore Proxy Appliance Name

At the **Name** step of the wizard, specify a name and description for the Azure restore proxy appliance:

1. In the **Name** field, specify a name for the Azure restore proxy appliance. The name must meet the following requirements:
   - The name must not be longer than 15 characters.
   - The name must contain only alphanumeric characters and hyphens.
   - The name must start with a letter and end with a letter or number.
   - The name must not contain only numeric characters.
   - The name must not contain special characters: *`!@#$%^&*()+=_[]{}|;:.'",<>/?*.
2. In the **Description** field, provide a description for the Azure restore proxy appliance.
3. At the **Max concurrent tasks** field, specify the number of tasks that the Azure restore proxy appliance must handle in parallel. If the **Max concurrent tasks** value is exceeded, the Azure restore proxy appliance will not start a new task until one of current tasks finishes.

   Veeam Backup & Replication creates one task per one workload disk. By default, Azure restore proxy appliance handles 4 concurrent tasks.

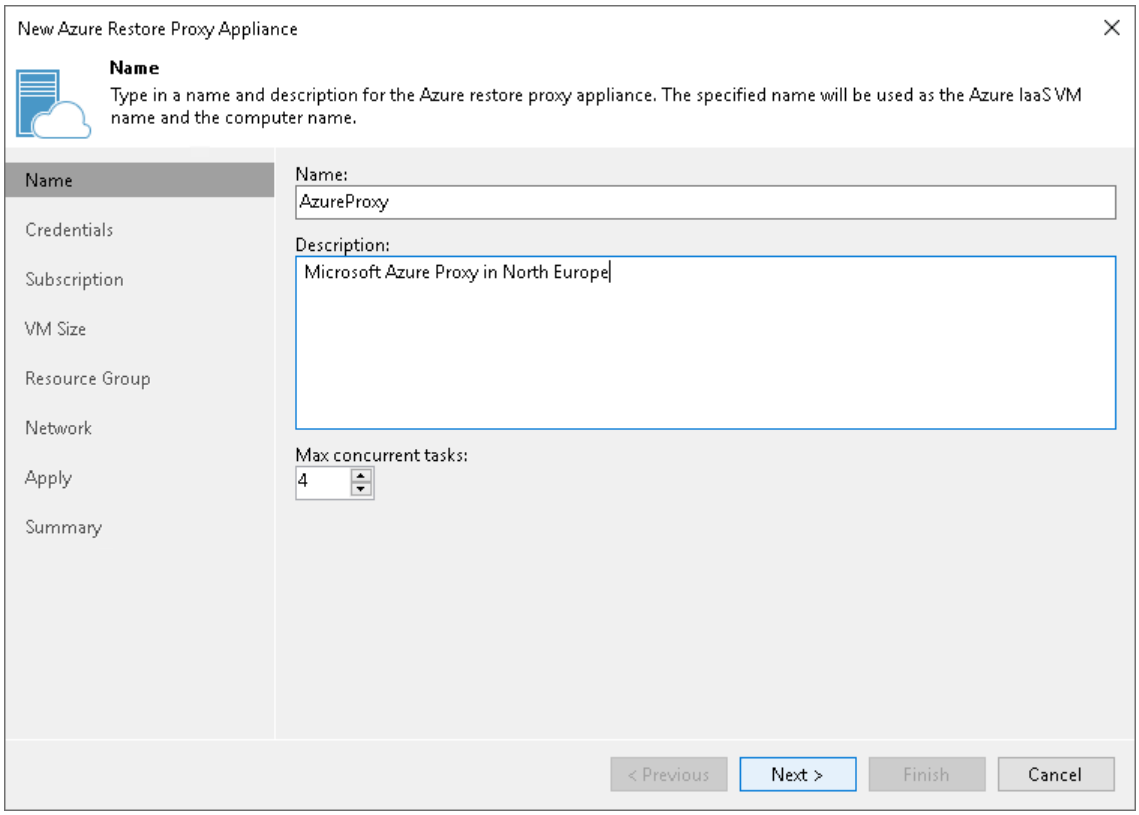

## Step 3. Specify Credentials and Transport Port

At the **Credentials** step of the wizard, specify credentials of the local administrator account on the Azure restore proxy appliance and define the transport port:

1. From the **Credentials** list, select credentials of a user that will be assigned the Local Administrator permissions on the Azure restore proxy appliance.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add the credentials. For more information, see Credentials Manager.

   > **IMPORTANT**
   >
   > Consider the following:
   >
   > - You cannot use reserved names such as 'administrator', 'admin', 'user', 'abc@123', 'P@$$w0rd' and so on as a user name and password of the local administrator account.
   > - You must specify the user name without a domain or Microsoft Azure machine name.
   > - The password must be at least 8 characters long, and must contain at least 1 uppercase character, 1 lowercase character, 1 numeric character and 1 special character.

2. In the **Traffic port** field, specify a port over which Veeam Backup & Replication will control components installed on the Azure restore proxy appliance and transport workload disks data to Blob storage. The port must be accessible from the backup server and backup repository that stores backups.

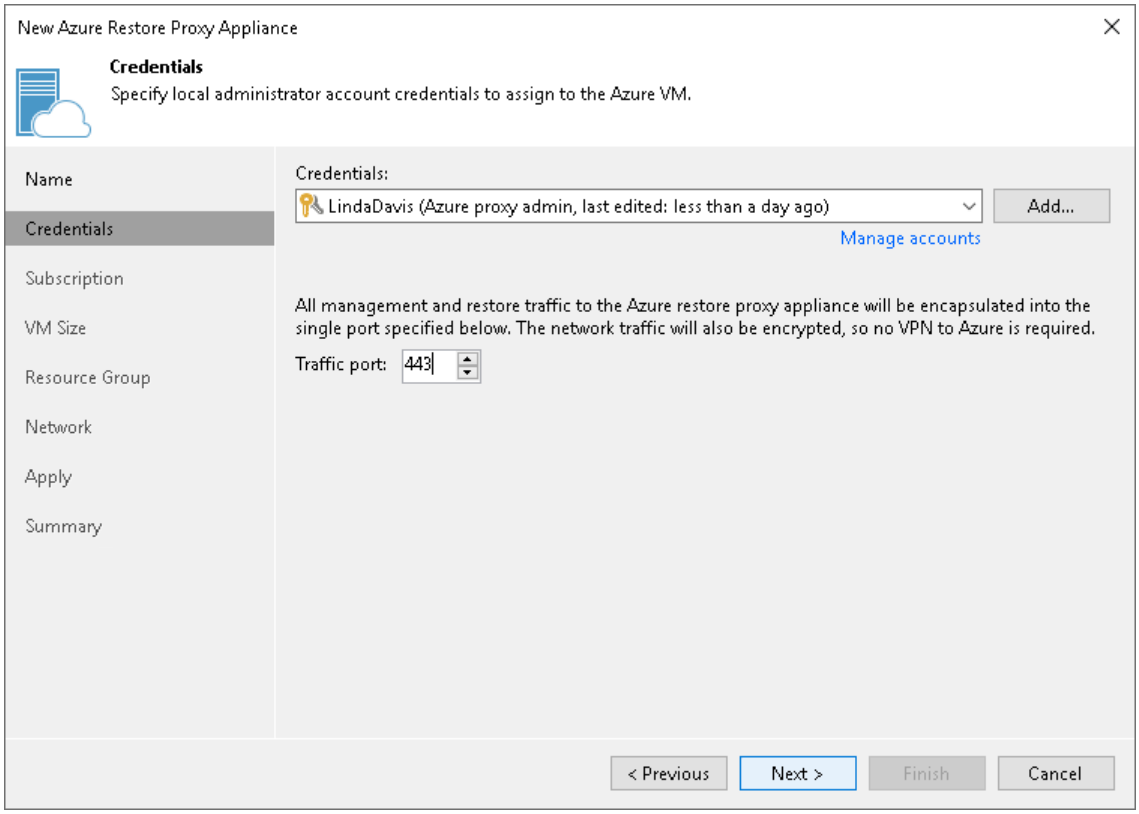

## Step 4. Select Subscription and Location

At the **Subscription** step of the wizard, select a subscription and location for the Azure restore proxy appliance:

1. From the **Subscription** list, select a subscription whose resources you want to use to deploy the Azure restore proxy appliance. The subscription list contains all subscriptions associated with Azure or Azure Stack Hub Compute accounts that you added to Veeam Backup & Replication.
2. From the **Locations** list, select a geographic region to which you want to place the Azure restore proxy appliance. Make sure that you select a geographic region with which at least one storage account of the subscription is associated.

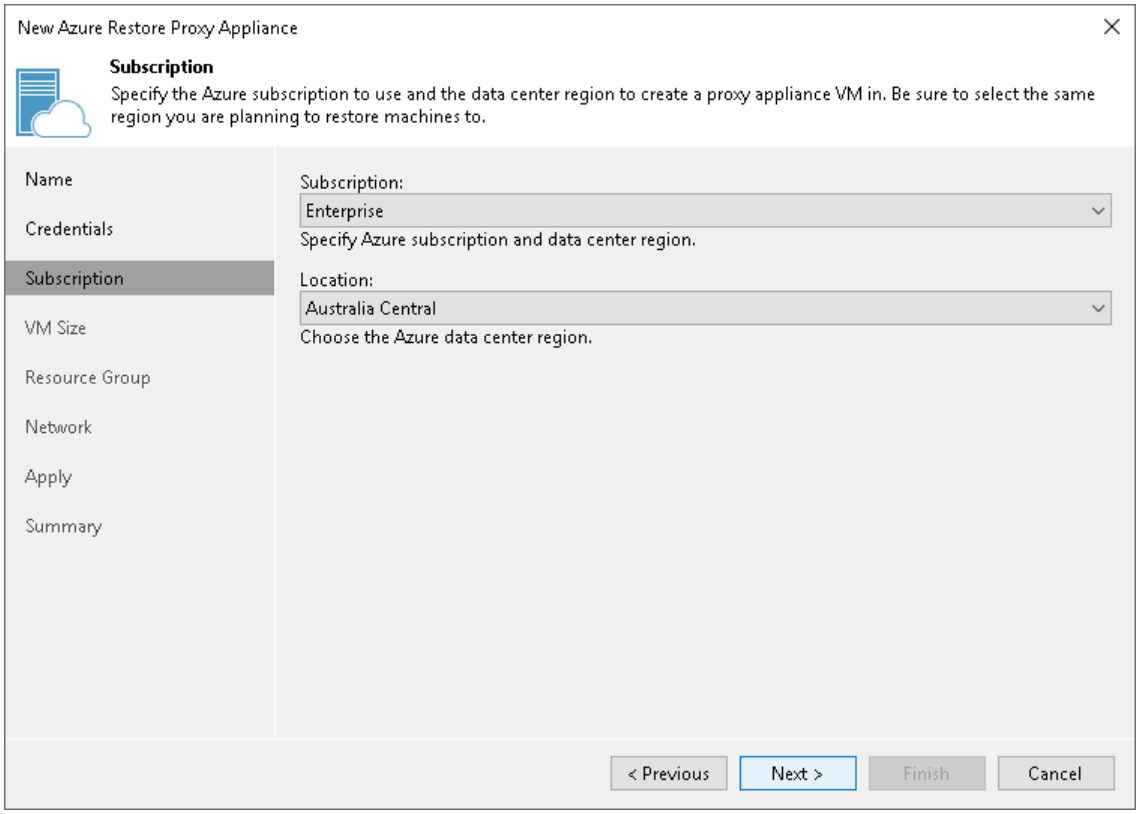

## Step 5. Select VM Size

At the **VM size** step of the wizard, you can select the size for the Azure restore proxy appliance VM and specify which storage account you want to use to deploy the Azure restore proxy appliance VM:

1. From the **Size** list, select the size for the Azure restore proxy appliance.

   The default size is *Standard_F4s_v2*. If *Standard_F4s_v2* is not available, we recommend that you select a similar compute optimized VM size of the previous generation — F-series. For example, *Standard_F4*, *Standard_F4s*. These sizes are sufficient to transport VM disks data to Blob storage. If necessary, you can select a greater size for the Azure restore proxy appliance.

   > **NOTE**
   >
   > Azure restore proxy appliance VMs created in Veeam Backup & Replication version prior 10a have smaller sizes — *Basic_A2*. We recommend you to change sizes of such proxies to the sizes listed in this point. This will enhance the performance of restore to Azure.
   >
   > You can change VM sizes in Microsoft Azure Portal or deploy new proxies with the required sizes in the Veeam Backup & Replication.

2. From the **Storage account** list, select a storage account where Veeam Backup & Replication will store components required for Azure restore proxy appliance deployment. After restore, the components will be removed from the storage account.

   [For Azure Stack Hub] Veeam Backup & Replication will store disks of the Azure restore proxy appliance in the selected storage account. The storage account must be compatible with the VM size you select.

   The list of storage accounts will contain only general purpose storage accounts. Blob storage accounts will not be displayed in the list of subscriptions. For more information about account types, see Microsoft Docs.

   > **NOTE**
   >
   > You cannot use a storage account with the ZRS or GZRS replication option for the Azure restore proxy appliance. For details, see Microsoft Docs.
   >
   > **TIP**
   >
   > Microsoft Azure subscriptions have default limits on the number of CPU cores. Make sure that the VM size you select does not exceed limits of the subscription.

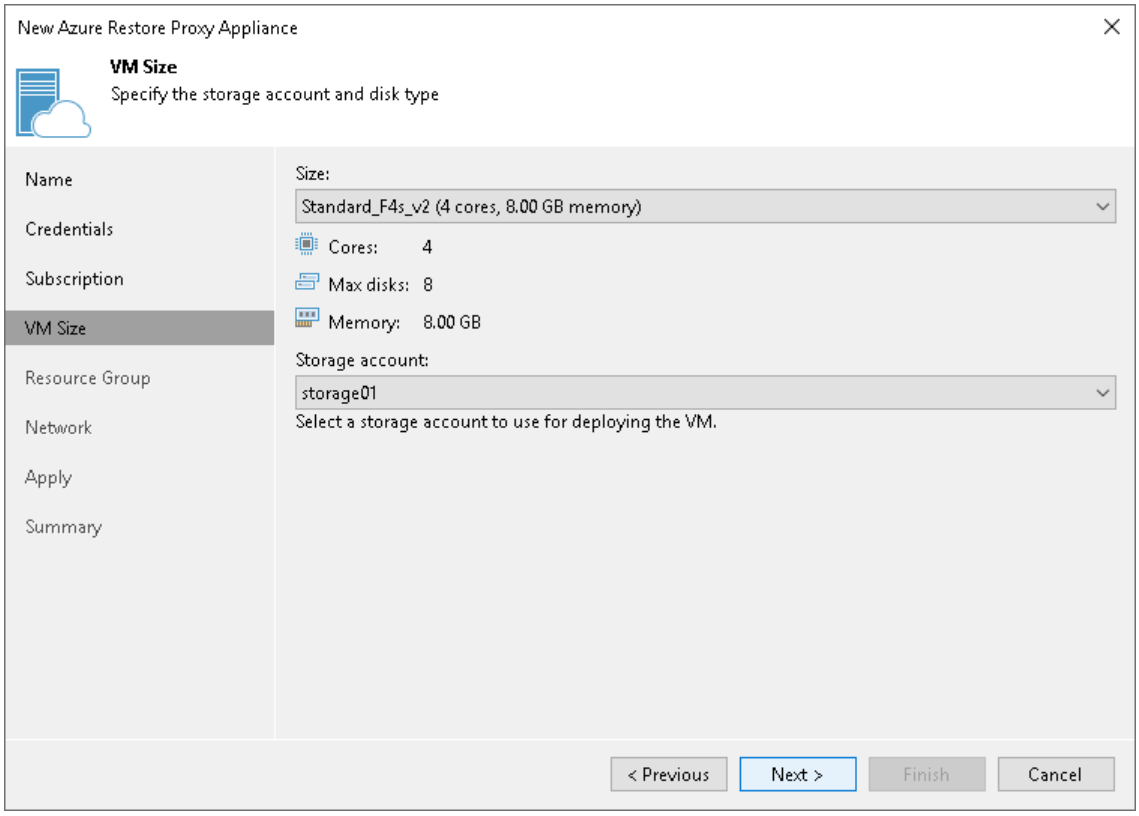
New Azure Restore Proxy Appliance
VM Size
Specify the storage account and disk type
Name
Credentials
Subscription
VM Size
Resource Group
Network
Apply
Summary
Size:
Standard_F4s_v2 (4 cores, 8.00 GB memory)
Cores: 4
Max disks: 8
Memory: 8.00 GB
Storage account:
storage01
Select a storage account to use for deploying the VM.
< Previous
Next >
Finish
Cancel

## Step 6. Select Resource Group

At the **Resource Group** step of the wizard, you specify the resource group to which the Azure restore proxy appliance must be placed and configure DNS name label:

1. You can place the Azure restore proxy appliance to an existing or new resource group:
   - Select **Place VM into the existing resource group** to place the Azure restore proxy appliance to an existing resource group. From the drop-down list, select the necessary resource group.
   - Select **Create a new resource group** to create a dedicated resource group for the Azure restore proxy appliance. In the **Name** field, enter a name for the new resource group. The resource group name can be up to 64 characters long and can contain only alphanumeric, underscore and hyphen characters.

     For the new resource group, Veeam Backup & Replication automatically creates a network security group, dynamic public IP and network interface.
2. In the **DNS name** label field, enter a name of the dynamic public IP. The DNS name label can be up to 80 characters long, and can contain only alphanumeric, dash and underscore characters. For more information, see Microsoft Docs.

> **TIP**
>
> Microsoft Azure subscriptions have default limits on the number of resource groups. If you decide to create a new resource group, make sure that you do not exceed limits of the subscription.

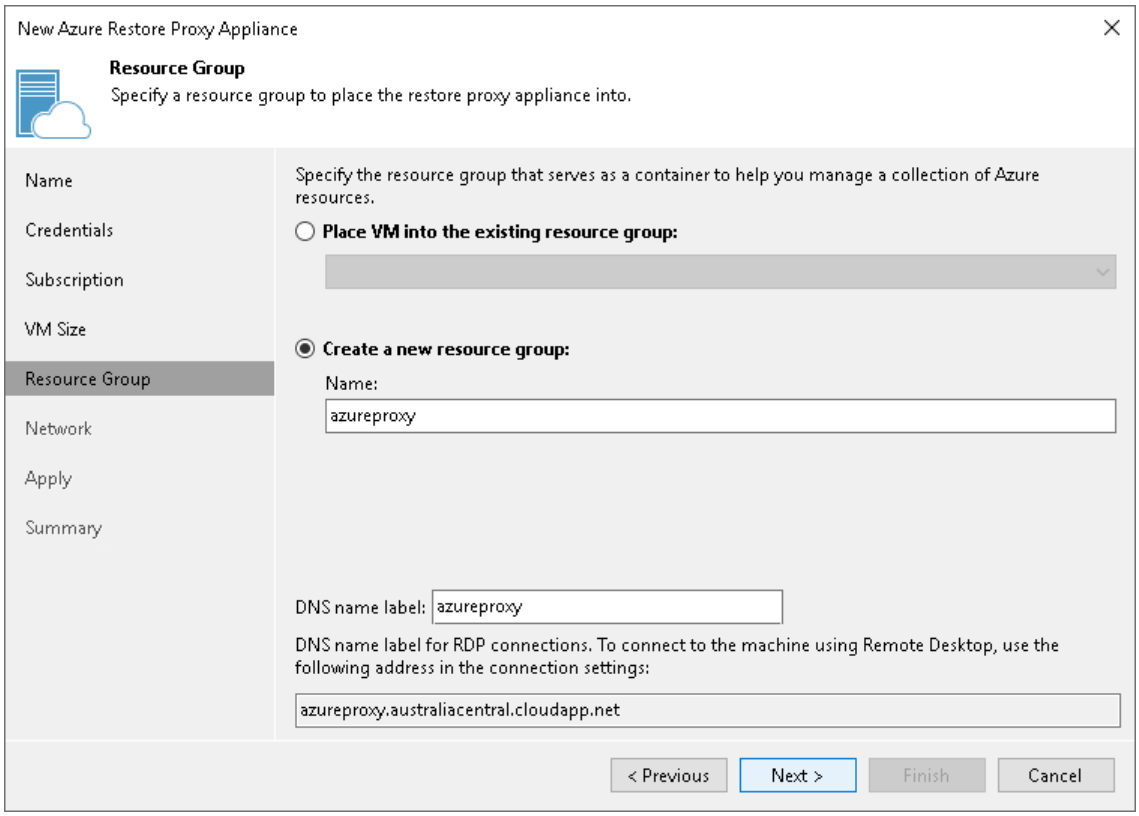

## Step 7. Select Virtual Network

At the **Network** step of the wizard, you select to which network and subnet the Azure restore proxy appliance will be connected.

> **IMPORTANT**
>
> If you want to restore from backups in an on-premises object storage repository, the selected virtual network must have access to the source object storage repository. To provide access to object storage repositories, you can use VPN or Azure ExpressRoute. For more information, see this Veeam KB article.

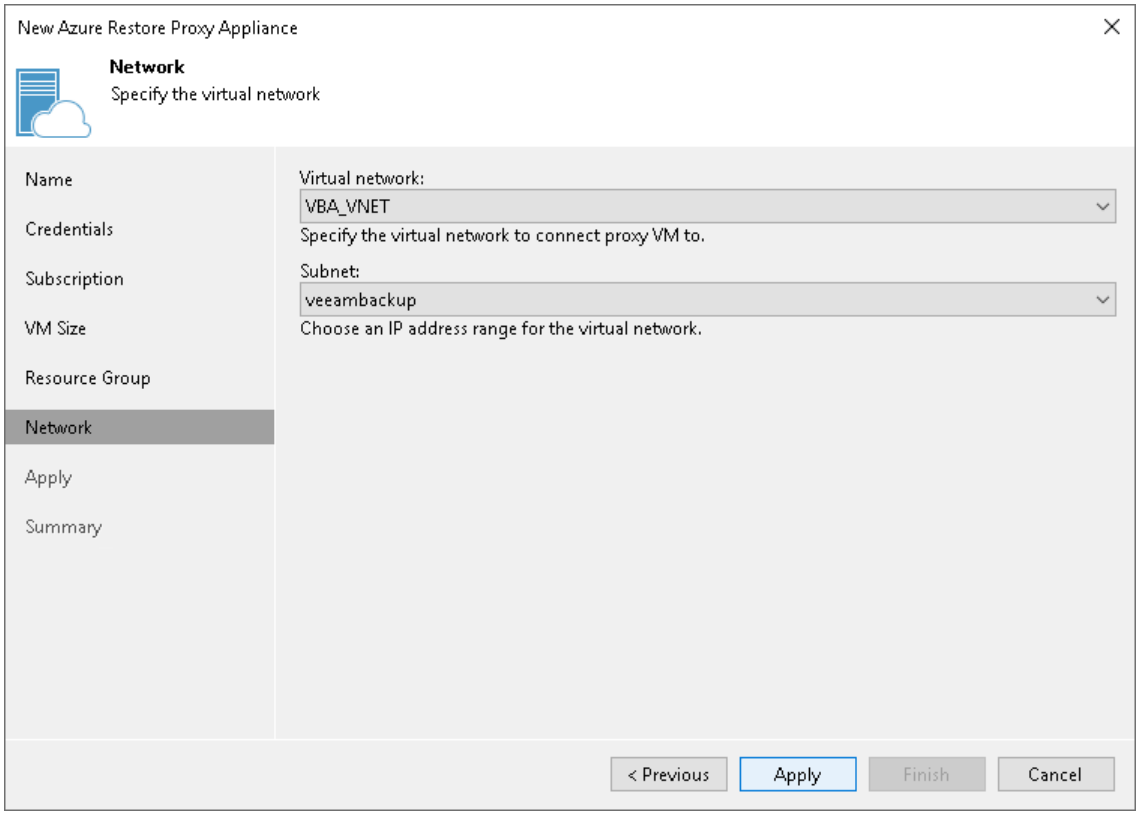

## Step 8. Start Deployment

At the **Apply** step of the wizard, Veeam Backup & Replication deploys the Azure restore proxy appliance with the specified settings. You can view the deployment progress in the real-time mode. When the configuration process is over, click **Next**. At the **Summary** step of the wizard, click **Finish** to close the wizard.

> **TIP**
>
> The Azure restore proxy appliance deployment may take several minutes. You can close the **New Azure Restore Proxy Appliance** wizard and continue working with Veeam Backup & Replication while the helper appliance is being deployed. To view the deployment progress, open the **History** view, in the inventory pane select **System**, and double-click the task of the helper appliance deployment in the working area.

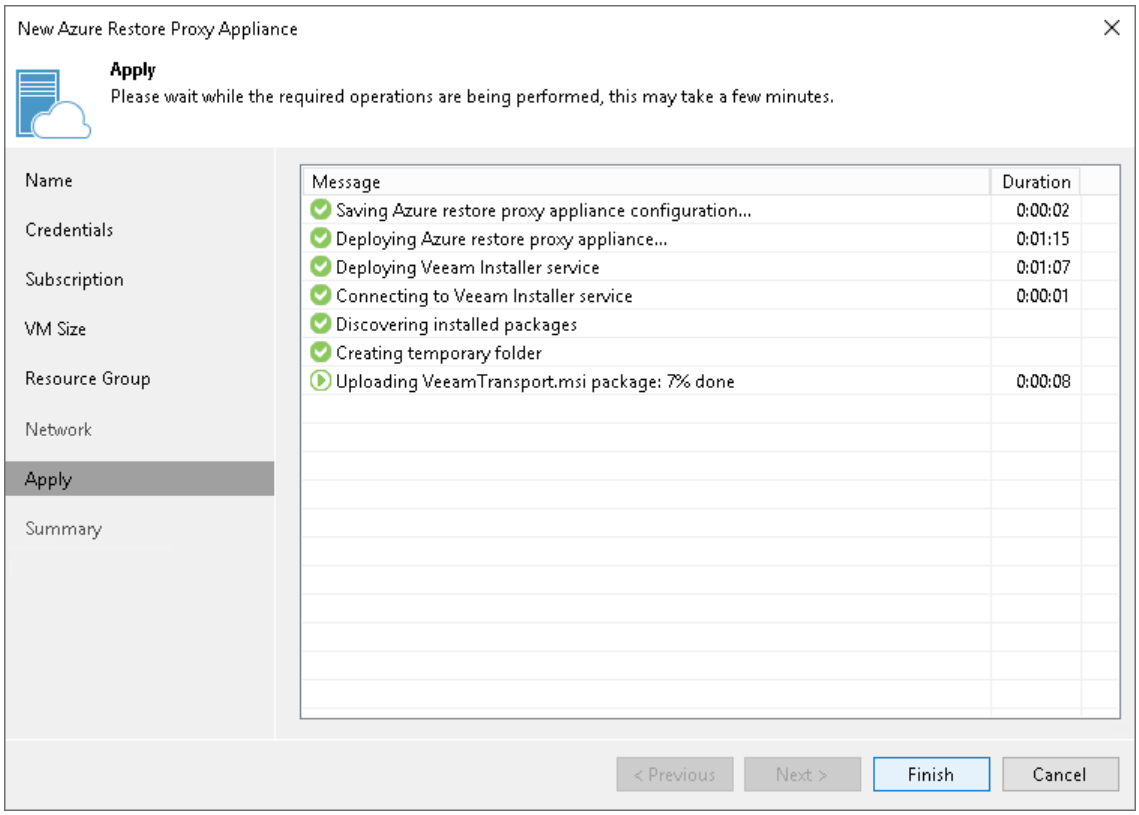


## Removing Azure Restore Proxy Appliances

Veeam Backup & Replication does not provide a possibility to edit settings of deployed Azure restore proxy appliances. If you want to change Azure restore proxy appliance configuration, remove the Azure restore proxy appliance and create a new Azure restore proxy appliance.

To remove an Azure restore proxy appliance, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, right-click the Azure restore proxy appliance and select **Remove**.

> **IMPORTANT**
>
> If you want to remove an Azure or an Azure Stack Hub Compute account from Veeam Backup & Replication, you must remove all Azure restore proxy appliances first.

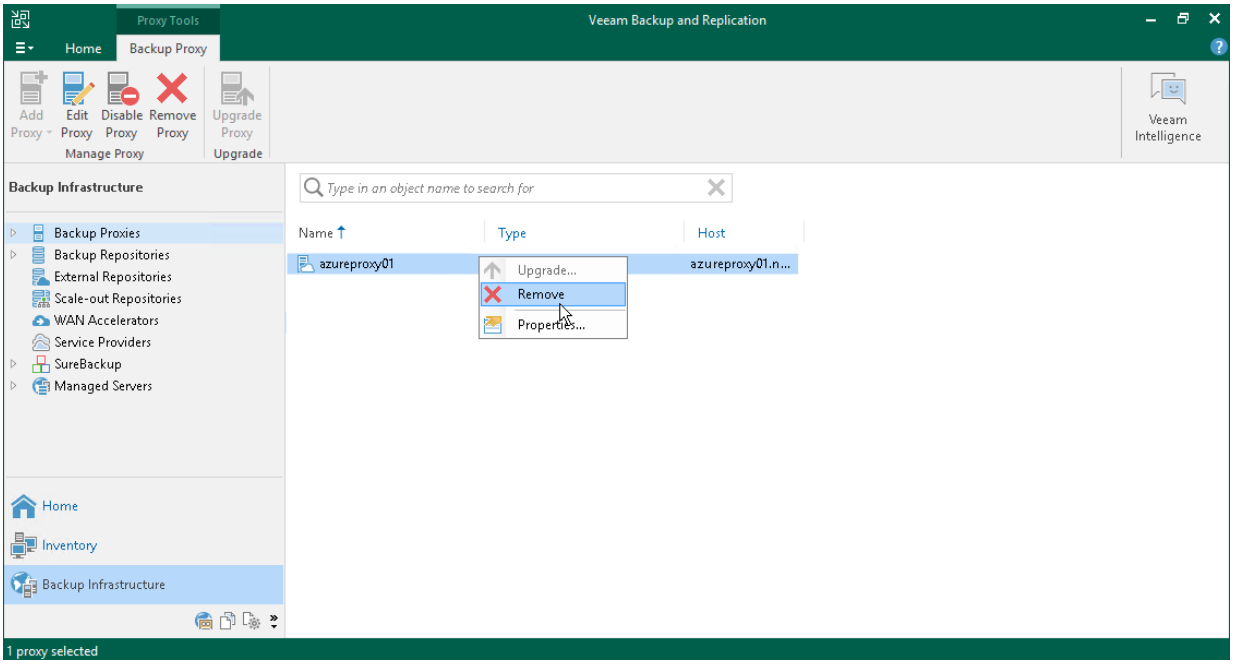


# Restoring to Microsoft Azure

Before you restore workloads to Microsoft Azure, you must configure an account to be used for restore and required components. For more information, see Configuring Components and Accounts for Restore. Then use the **Restore to Azure** wizard to restore the workloads.

> **NOTE**
>
> If you use Veeam Backup for Microsoft Azure and plan to restore Microsoft Azure VMs from restore points that were created using the appliance, the steps of the restore wizard differ from the steps described in this guide. For more information, see the Performing Entire VM Restore section in the Veeam Backup for Microsoft Azure User Guide.

## Before You Begin

Before you begin restore to Microsoft Azure, check the following prerequisites:

- Check limitations listed in section Considerations and Limitations for Restore to Microsoft Azure.
- You must create a backup of the workload that you want to restore in Microsoft Azure. For the list of supported backups, see Restore to Microsoft Azure.
- A backup chain from which you plan to restore a workload must reside in a backup repository added to the backup infrastructure. You can also import a backup to the Veeam Backup & Replication console. For more information, see Importing Backups Manually.
- You must configure the following objects in Microsoft Azure beforehand:
  - Storage account whose resources you plan to use to store disks of the restored workload.
  - Networks to which you plan to connect the restored workload.
- Make sure that you configured all the required components and accounts in Veeam Backup & Replication as described in section Configuring Components and Accounts for Restore.

- [For speeding up restore from Capacity Tier] It is strongly recommended to use Azure restore proxy appliance when you restore from backups residing on a Capacity Tier. For more information on Azure restore proxy appliances, see Managing Azure Restore Proxy Appliances.
- You must set up correct time on the backup server. Otherwise you may not be able to add a Microsoft Azure Compute account or Microsoft Azure Stack Hub Compute account to Veeam Backup & Replication, or the restore process may fail.

## Step 1. Launch Restore to Azure Wizard

To launch the **Restore to Azure** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select the type of backups from which you want to restore:
  - VMware vSphere
  - VMware Cloud Director
  - Microsoft Hyper-V
  - Agent
  - AWS
  - GCE backup
  - Nutanix backup
  - oVirt KVM
  - Proxmox VE

  In the displayed window, click **Entire machine restore** > **Restore to public cloud** > **Restore to Microsoft Azure**.
- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary backup, select workloads that you want to restore and click **Restore to Microsoft Azure** on the ribbon. Alternatively, you can right-click one of the workloads that you want to restore and select **Restore to Microsoft Azure**.

- Double-click a full backup file (VBK) or backup metadata file (VBM) in a file browser.
  Veeam Backup & Replication will start its console. In the **Backup Properties** window, select the necessary workload and click **Restore** > **Restore to Microsoft Azure**.

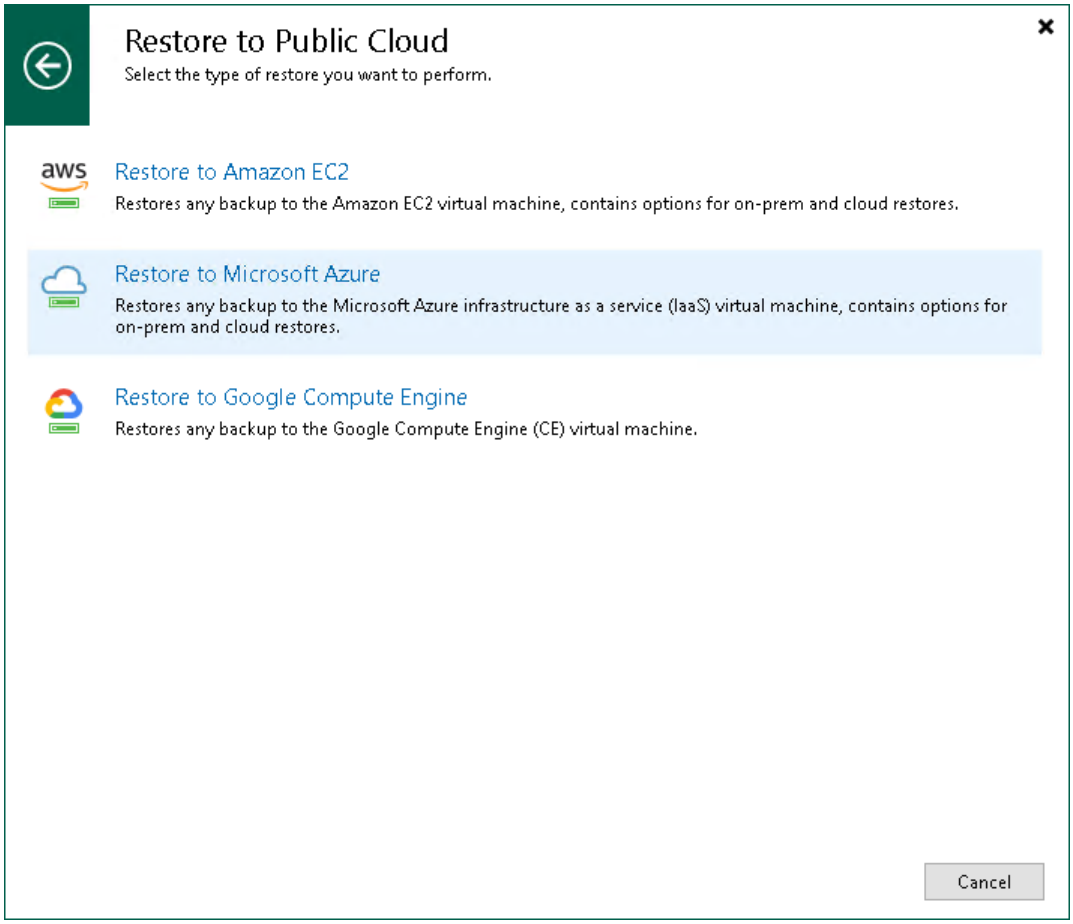

## Step 2. Select Workloads and Restore Points

At the **Machine** step of the wizard, specify workloads that you want to restore and specify restore points to which you want to restore the workloads. By default, Veeam Backup & Replication restores workloads to the latest valid restore point in the backup chain.

# Selecting Workloads

To select workloads to restore:

1. On the right of the **Machine** list, click **Add**.
2. In the **Backup Browser** window, expand the necessary backup, select workloads and click **Add**.

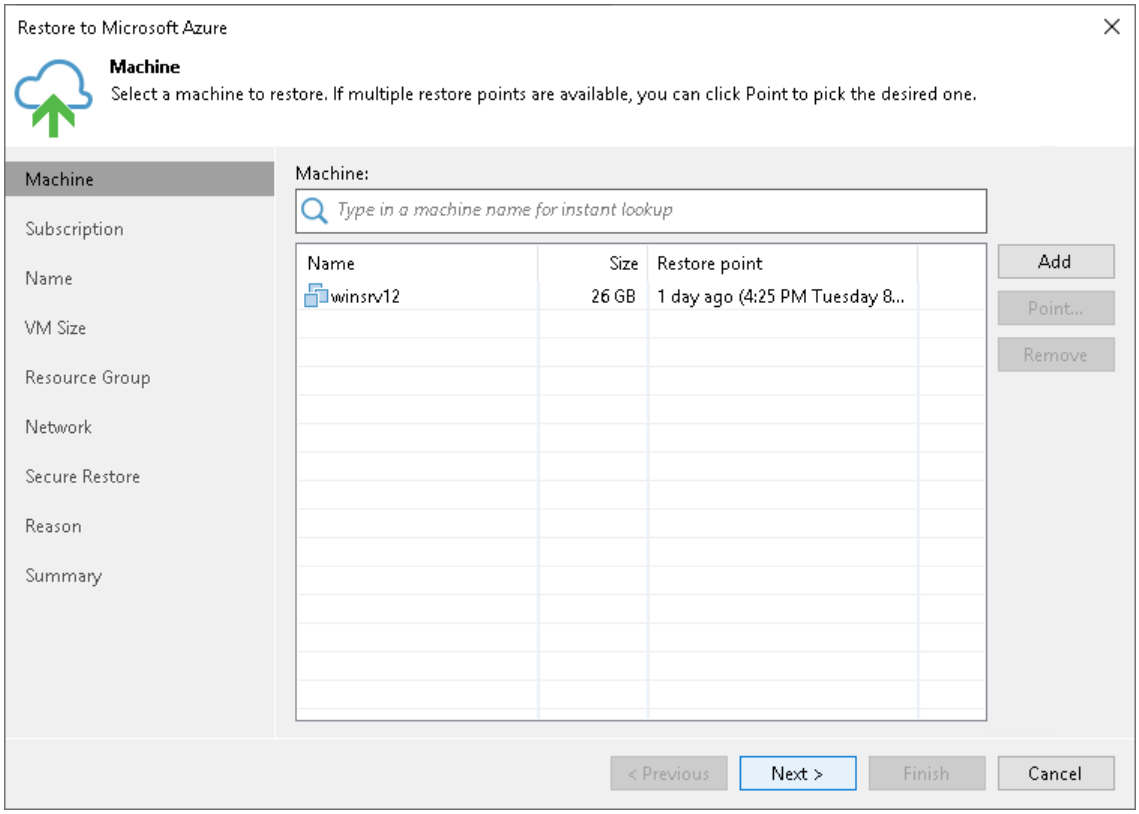


# Selecting Restore Points

To select a restore point for a workload, do the following:

1. In the **Machine** list, select a workload.
2. Click **Point** on the right.

3. In the **Restore Points** window, select a restore point to which you want to restore the workload.

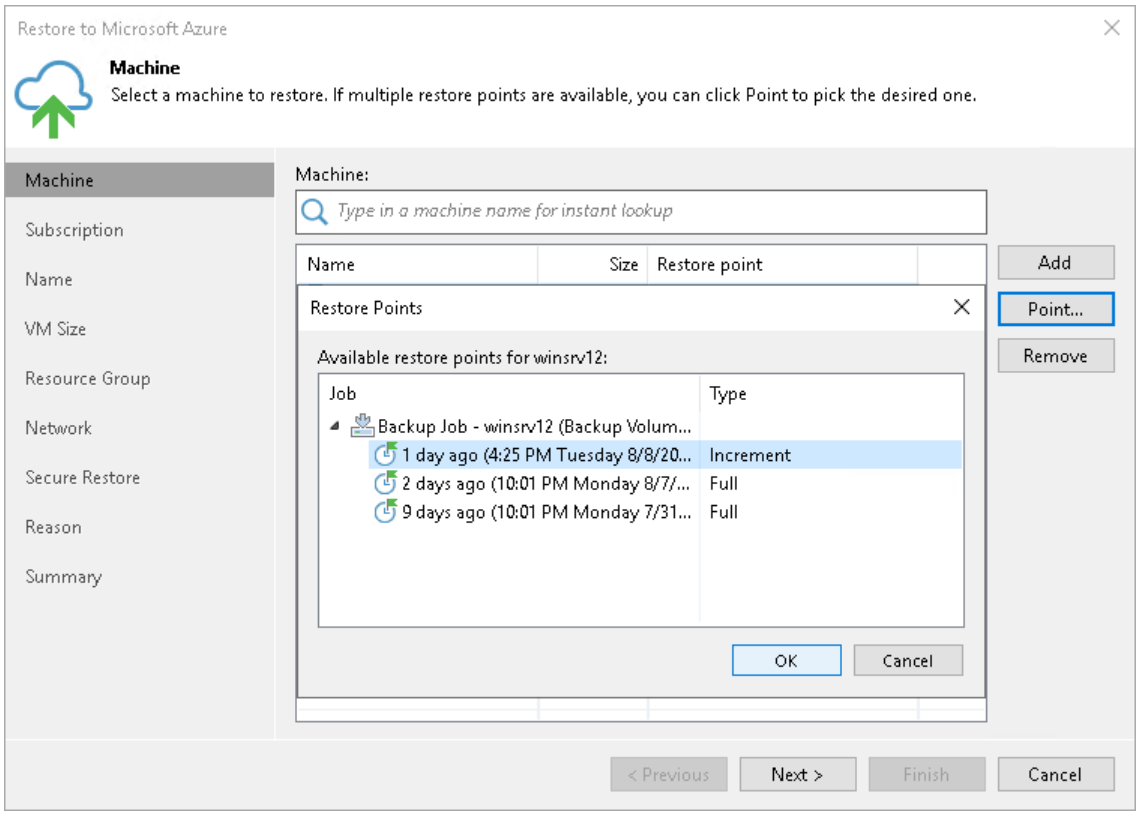

## Step 3. Select Subscription and Location

At the **Subscription** step of the wizard, select a subscription, location for the restored workloads and define how workload data must be transported to Microsoft Azure or Azure Stack Hub:

1. From the **Subscription** list, select a subscription whose resources you want to use. The subscription list contains all subscriptions associated with the Azure compute or Azure Stack Hub Compute accounts that you have added to Veeam Backup & Replication.
2. From the **Locations** list, select a geographic region to which you want to place the restored workloads. Make sure that you select a geographic region with which at least one storage account of the subscriptions is associated.
3. If you are restoring the workloads to a distant location and want to speed up the restore process, select the **Use a proxy appliance** check box. From the drop-down list, select an Azure restore proxy appliance.

   It is recommended that you configure the Azure restore proxy appliance in the same location to which you plan to restore the workload. For more information, see Managing Azure Restore Proxy Appliances.

> **IMPORTANT**
>
> [For restore of Linux workloads] You must have a preconfigured helper appliance in the location to which you restore Linux workloads. If the appliance is not configured, Veeam Backup & Replication will display the **Microsoft Azure Compute Account** wizard so that you can configure the appliance in the selected location.

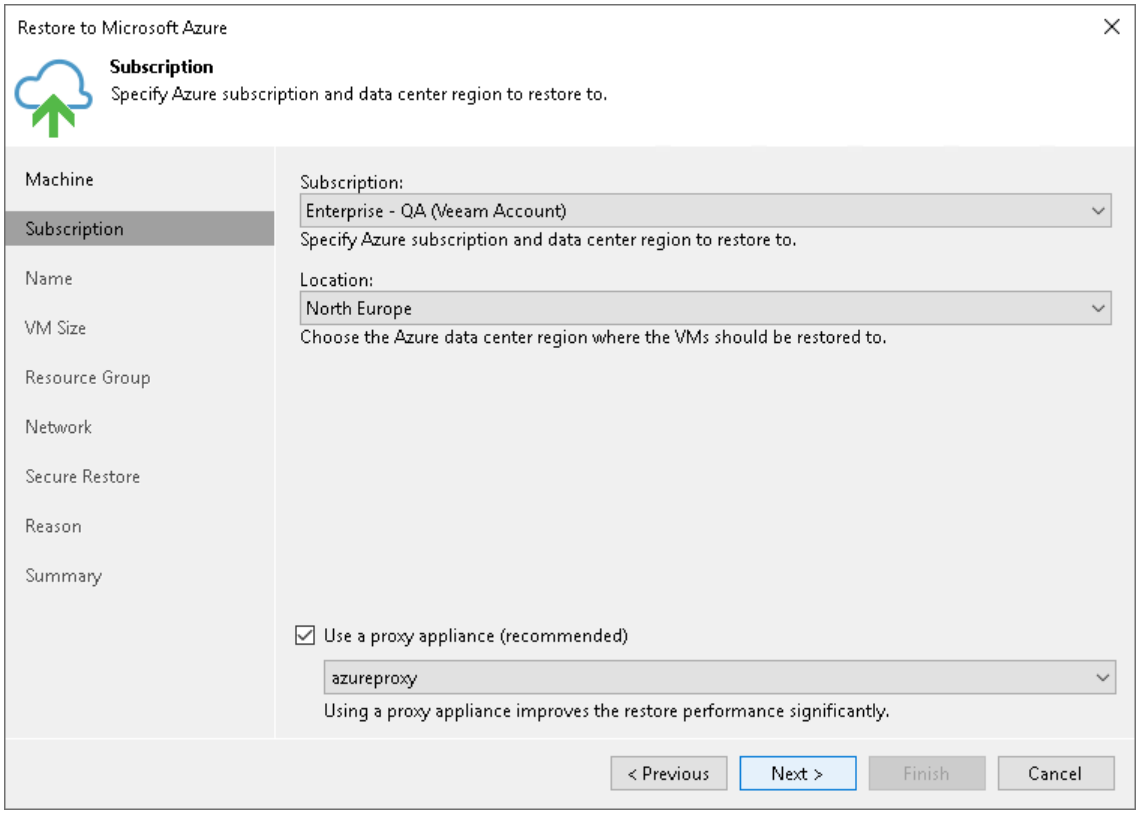

## Step 3. Specify Names and Tags

At the **Name** step of the wizard, you can specify new names for the restored workloads and assign tags to them. By default, Veeam Backup & Replication restores workloads with their original names.

# Specifying New Names

To define a new name for a workload:

1. In the **Machines to restore** list, select a workload and click **VM Name**.
2. In the **Change Name** window, enter a new name explicitly or specify a change name rule — add a prefix and suffix to the original workload name.

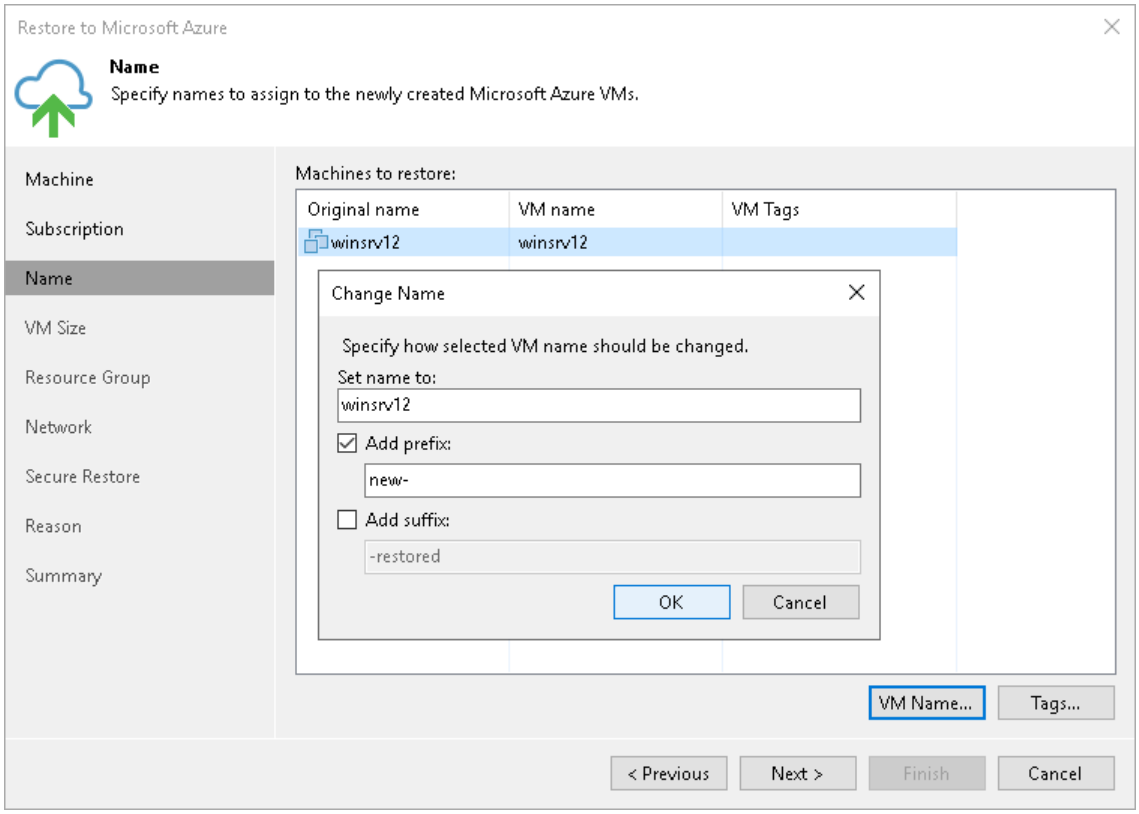


# Adding Tags

To add a tag for the restored workloads:

1. In the **Machines to restore** list, select a workload and click **Tags**.
2. In the **Tags** window, click **Add**.

3. In the **Microsoft Azure VM Tag** window, specify the tag name and its value.

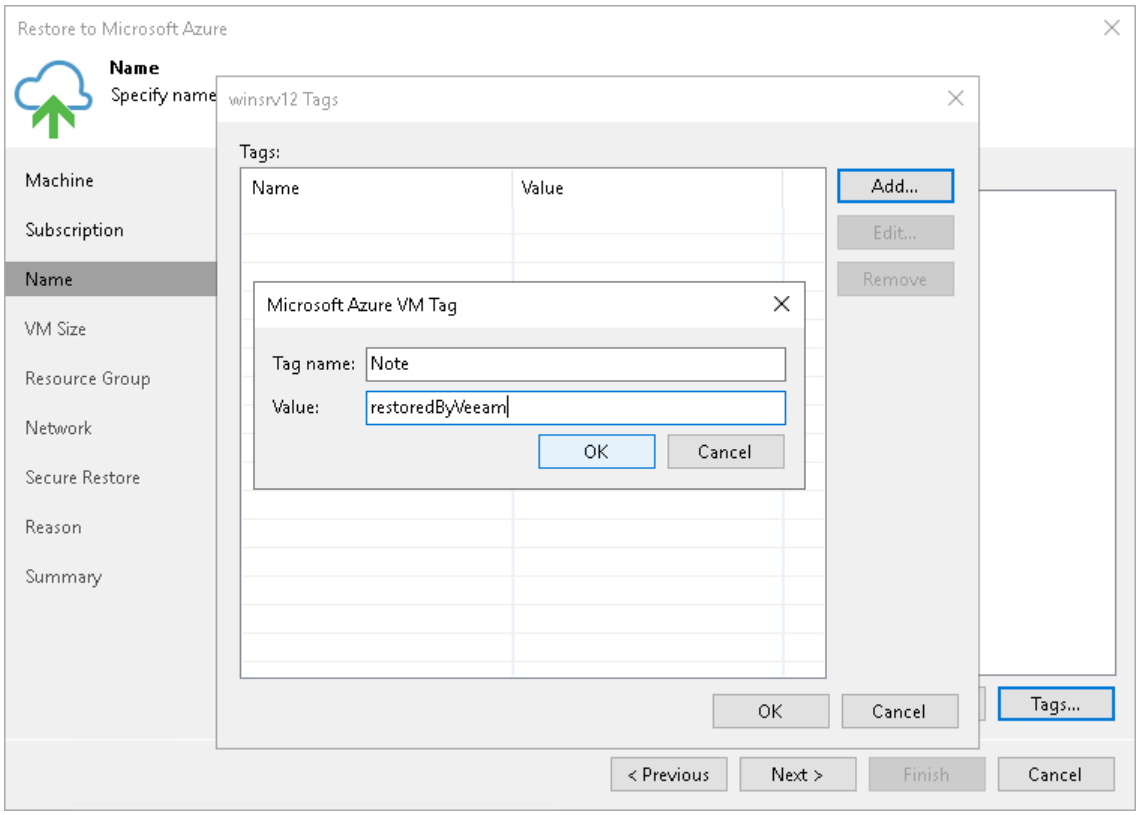

## Step 4. Specify VM Size and Disks

At the **VM Size** step of the wizard, you can select VM sizes, storage accounts where to store disks of the restored workloads, select disks to restore and change their type. By default, Veeam Backup & Replication selects the smallest size that can support the number of disks for the restored workload and restores all workload disks.

# Selecting VM Size and Storage Account

To select a size and storage account:

1. In the **Azure VM Configuration** list, select a workload and click **Edit**.
2. From the **Size** list, select a size for the restored workload.

   Make sure that you select the right workload size that corresponds to the initial workload configuration. The size affects the number of CPU cores, memory and disk resources that will be allocated to the restored workload. For more information, see Microsoft Docs.
3. From the **Storage type** list, select the disk type. Note the limitations for disk sizes in Considerations and Limitations for Restore to Microsoft Azure.
4. From the **Storage account** list, select a storage account whose resources you want to use to store disks of the restored workload. The storage account must be compatible with the workload size you select. Note the limitations for storage accounts in Considerations and Limitations for Restore to Microsoft Azure.

   The list of storage accounts contains only general purpose storage accounts. Blob storage accounts are not be displayed in the list of subscriptions. For more information about account types, see Microsoft Docs.

   If you select a premium storage account, make sure that the restored workload size is compatible with the selected account.

> **NOTE**
>
> Microsoft Azure subscriptions have default limits on the number of CPU cores. Make sure that the restored workload size that you select does not exceed limits of the subscription.

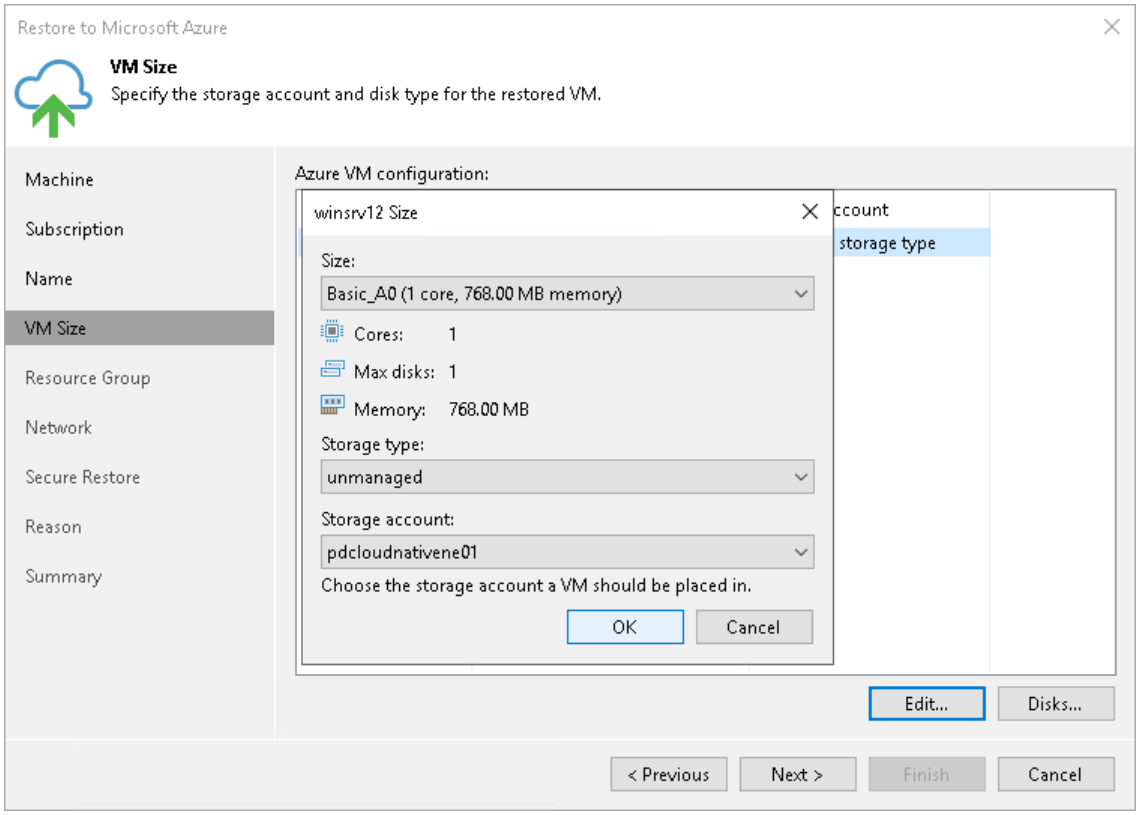


# Selecting Disks to Restore and Changing Their Types

You can restore all disks or specific disks of a workload. You can also change disk types of the restored disks if you have selected the managed storage type.

To select specific disks and change their type, do the following:

1. In the **Azure VM Configuration** list, select a workload and click **Disks**.
2. In the **Disks to restore** window, check that check boxes next to disks that you want to restore are selected. Clear check boxes next to disks that you do not want to restore.
3. [For managed storage type] Select a disk and click **Disk Type**. In the **Select Azure VM Disk Type** window, select the necessary type. For more information on disk types, see Microsoft Docs.

> **IMPORTANT**
>
> The selected disk type must be compatible with the selected workload size.

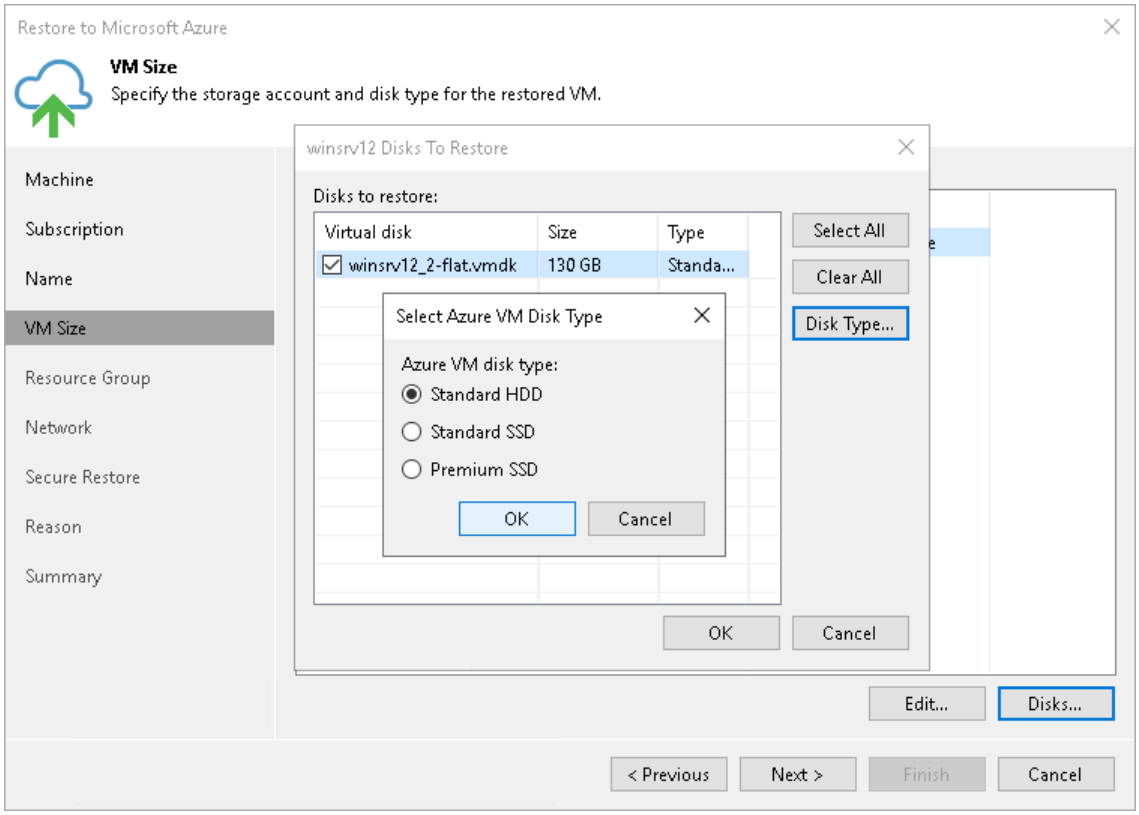
Restore to Microsoft Azure
VM Size
Specify the storage account and disk type for the restored VM.
Machine
Subscription
Name
VM Size
Resource Group
Network
Secure Restore
Reason
Summary
winsrv12 Disks To Restore
Disks to restore:
Virtual disk
Size
Type
winsrv12_2-flat.vmdk
130 GB
Standa...
Select All
Clear All
Disk Type...
Select Azure VM Disk Type
Azure VM disk type:
Standard HDD
Standard SSD
Premium SSD
OK
Cancel
OK
Cancel
Edit...
Disks...
< Previous
Next >
Finish
Cancel

# Step 5. Specify VM Name and Resource Group

At the **Resource Group** step of the wizard, you can select resource groups for the restored workloads for them. By default, Veeam Backup & Replication creates a new resource group for the restored workloads and places them to it.

To change a resource group to which a workload will be restored:

1. In the **Resource group** list, select a workload and click **Group**.
2. In the **VM Resource Group** window, select the necessary option for the workload:
   - Select **Place VM into the existing resource group** if you want to place the workload to an existing resource group. Then from the drop-down list, select the necessary resource group.
   - Select **Create a new resource group** if you want to create a dedicated resource group for the restored workload. In the **Name** field, enter a name for the new resource group.

     In the new resource group, Veeam Backup & Replication automatically creates a network security group, a dynamic public IP and network interface.

**NOTE**

Microsoft Azure subscriptions have default limits on the number of resource groups. If you decide to create a new resource group, make sure that you do not exceed limits of the subscription.

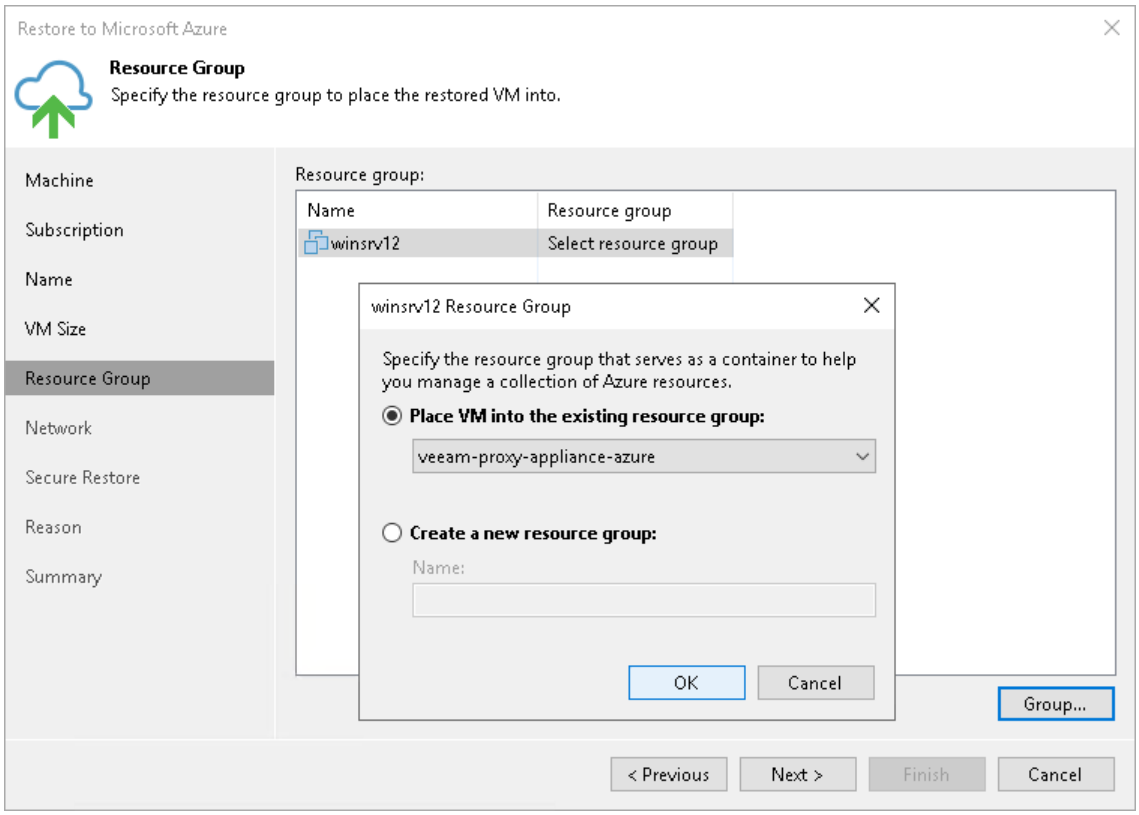

## Step 6. Configure Network Settings

At the **Network** step of the wizard, select to which networks and subnets the restored workloads will be connected. By default, Veeam Backup & Replication creates a new security group for the restored workloads. The default security group rules are listed in Default Security Group Rules.

# Configuring Network

To define network settings for a workload, do the following:

1. In the **Virtual network** list, select a workload and click **Network**.
2. From the **Virtual network** drop-down list, select a network to which the workload must be connected.
3. From the **Subnet** drop-down **list**, select a subnet for the workload.
4. In the **Public IP** field, specify whether to assign a public IP to the workload. You have two options:
   - **Assign (restored VM will be accessible from the Internet).** The public IP will be assigned to the restored workload. For security reasons, make sure traffic filtration rules are properly configured in the security group.
   - **Do not assign (more secure).** The public IP will not be assigned. You can assign a public IP later in the settings of the restored workload.

**NOTE**

Veeam Backup & Replication can connect a workload only to one virtual network. If necessary, you can manually configure additional network connections in Microsoft Azure after the workload is restored.

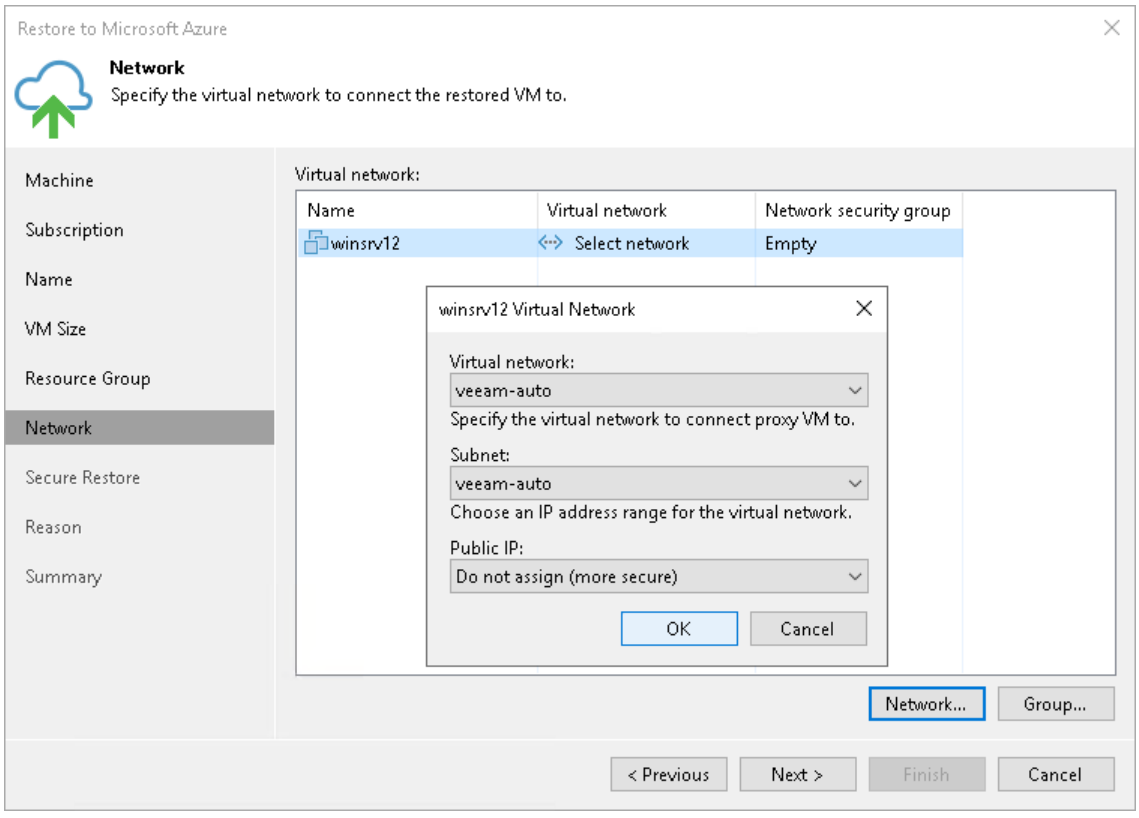

# Configuring Security Group

To change a security group to which a workload will be restored:

1. From the **Virtual network** list, select a workload and click **Group**.
2. Select the network security group from the **Network security group** list.

   If you leave the *Empty* value, Veeam Backup & Replication will create a new network security group. The rules for this group are described in Default Security Group Rules.

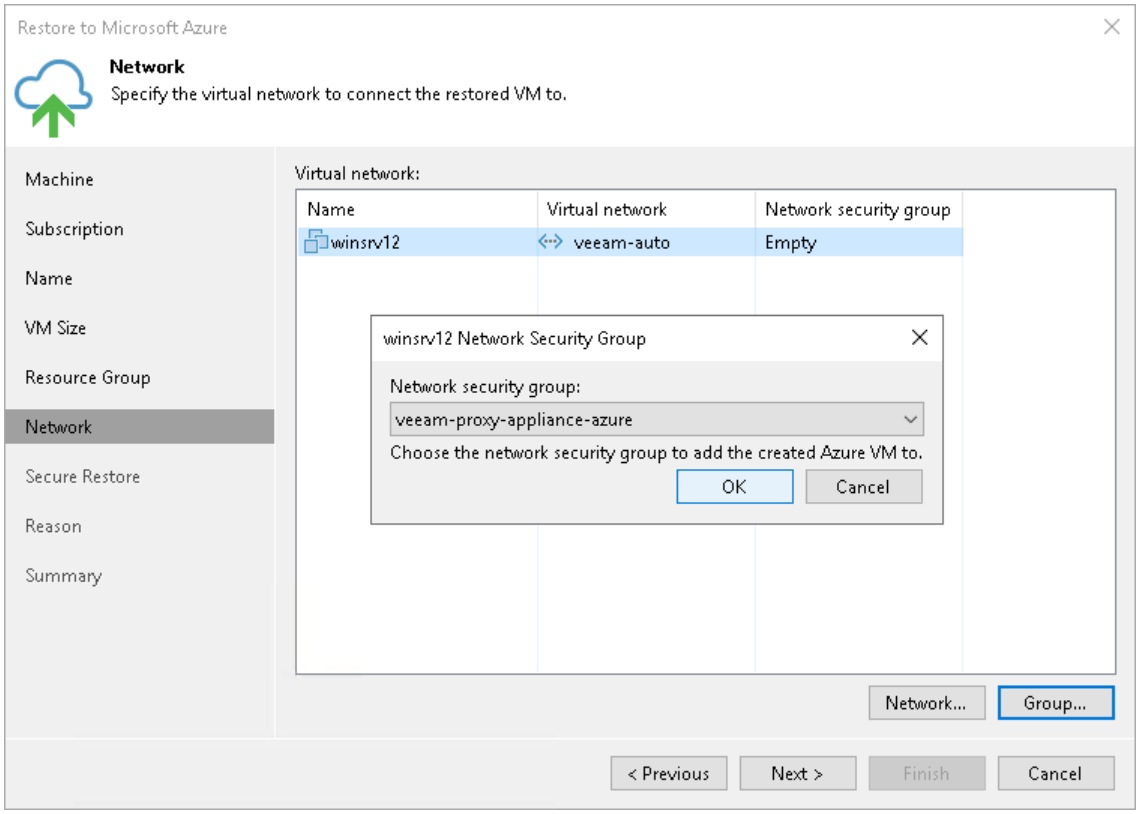


# Default Security Group Rules

When you leave the *Empty* value while configuring the security group, Veeam Backup & Replication creates a security group with the following rules.

| Priority | Name | Port | Protocol | Source | Destination | Action |
|---|---|---|---|---|---|---|
| **Inbound security rules** | | | | | | |
| 300 | SSH<br><br>**Note:** This rule applies if you restore Linux VMs. | 22 | TCP | Any | Any | Allow |

| Priority | Name | Port | Protocol | Source | Destination | Action |
|---|---|---|---|---|---|---|
| 300 | RemoteDesktop<br>**Note**: This rule applies if you restore Microsoft Windows VMs. | 3389 | TCP | Any | Any | Allow |
| 65000 | AllowVnetInBound | Any | Any | VirtualNetwork | VirtualNetwork | Allow |
| 65001 | AllowAzureLoadBalancerInBound | Any | Any | AzureLoadBalancer | Any | Allow |
| 65500 | DenyAllInBound | Any | Any | Any | Any | Deny |
| **Outbound security rules** | | | | | | |
| 65000 | AllowVnetOutBound | Any | Any | VirtualNetwork | VirtualNetwork | Allow |
| 65001 | AllowInternetOutBound | Any | Any | Any | Internet | Allow |
| 65500 | DenyAllOutBound | Any | Any | Any | Any | Deny |

# Step 7. Specify Secure Restore Settings

This step is available if you restore Microsoft Windows workloads.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Select which action Veeam Backup & Replication will take if scan finds a virus threat:
   - **Abort VM recovery**. Select this action if you want Veeam Backup & Replication to cancel the restore session.
   - **Proceed recovery but connect VM to a different network**. Select this action if you want to restore the workload to a different Microsoft Azure virtual network.

     Click the **Click to change** link to select the virtual network.

3. Select the **Continue scanning all remaining files after the first occurrence** check box if you want to continue workload scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Antivirus Scan Results.

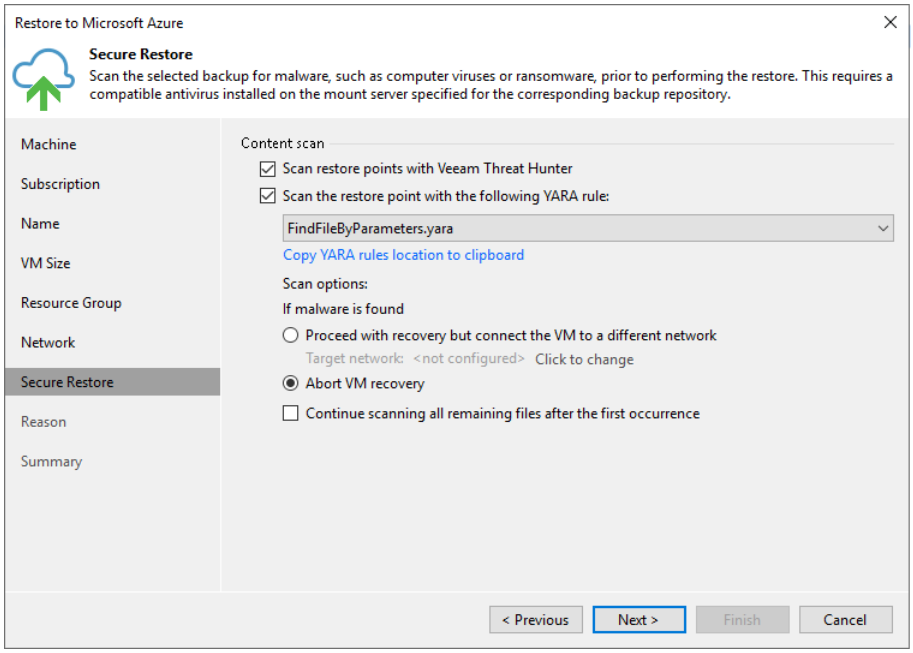

# Step 8. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the workload. The information you provide will be saved in the session history in Veeam Backup & Replication, and you can view it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

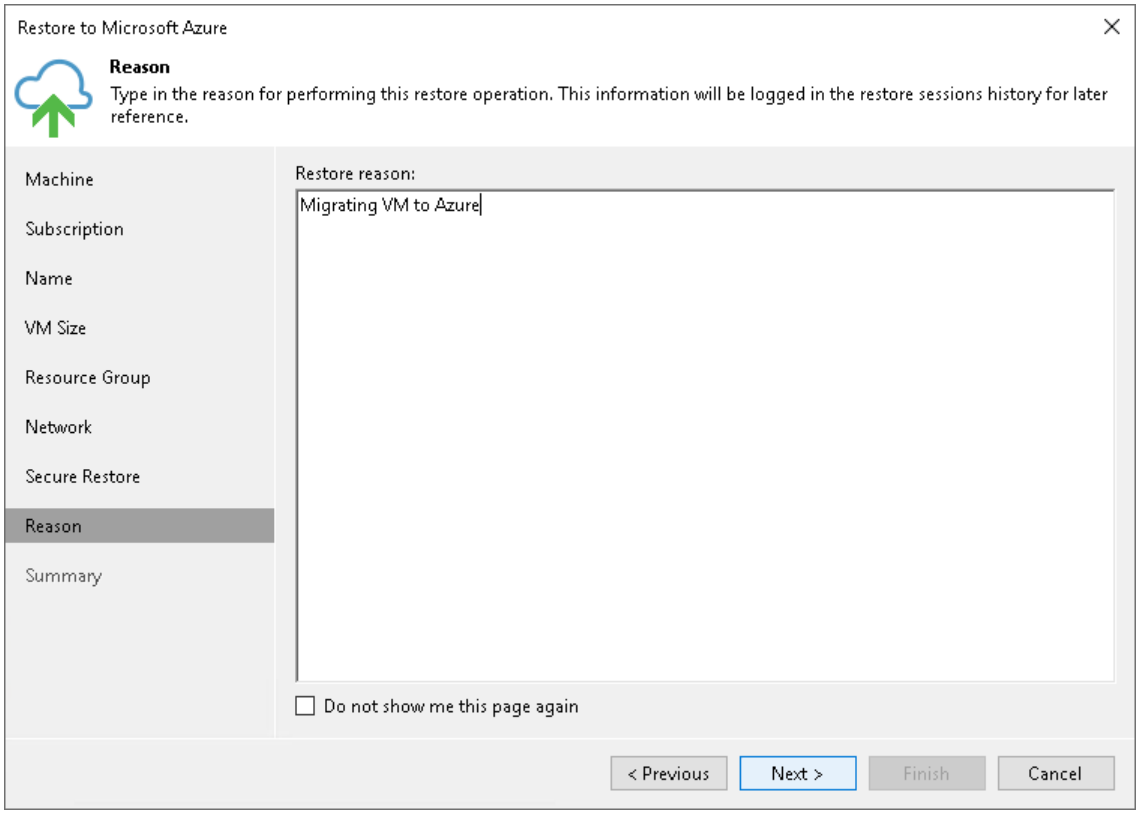

# Step 9. Verify Restore Settings

At the **Ready to Restore** step of the wizard, check the specified settings and click **Finish**. If you want to start the Azure VM right after restore, select the **Power on VM after restoring** check box.

You can trace the restore process in the **Restore Session** window. If you need to cancel the workload restore, click the **Cancel** restore task link.

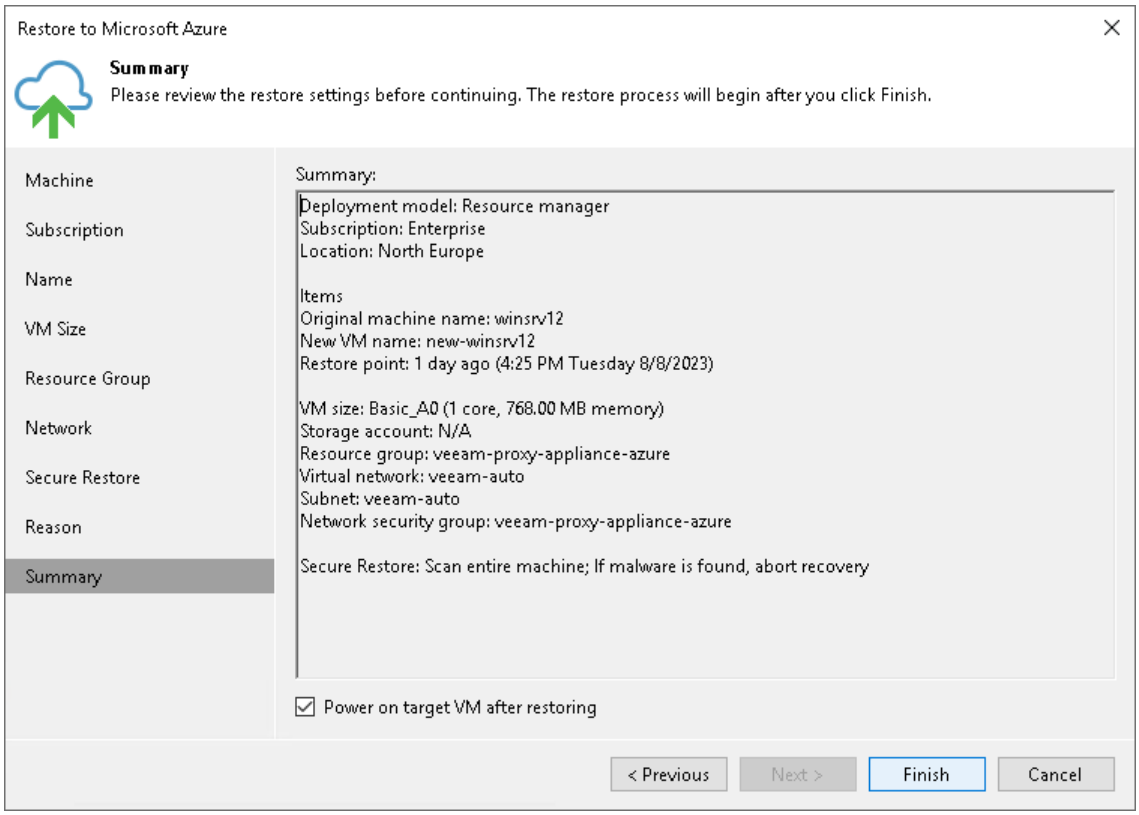

# Restore to Google Compute Engine

Veeam Backup & Replication allows you to restore different workloads (VMs, Google VM instances, physical servers and so on) to Google Compute Engine as VM instances. A VM instance is a virtual machine in Google Compute Engine with a preconfigured combination of computing resources.

You can use Veeam Backup & Replication to perform the following operations:

- Restore machines to Google Compute Engine from backups.
- Migrate machines from the on-premises infrastructure to the cloud.
- Create a test environment in the cloud for troubleshooting, testing patches and updates, and so on.

## Supported Backup Types

You can restore machines from the backups of the following types:

- Backups of VMware vSphere or VMware Cloud Director VMs created by Veeam Backup & Replication.
- Backups of Microsoft Hyper-V VMs created by Veeam Backup & Replication.
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux.

  Backups must be created at the entire machine level or volume level.
- Backups of Google Compute Engine VM instances created by Veeam Backup for Google Cloud.
- Backups of Amazon EC2 VM instances created by Veeam Backup for AWS.
- Backups of Microsoft Azure VMs created by Veeam Backup for Microsoft Azure.
- Backups of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV.
- Backups of oVirt VMs created by Veeam Backup for OLVM and RHV.
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE.

## Helper Appliances

Helper appliance is an auxiliary Linux-based VM instance. It is used to upload backed-up data to Google Compute Engine. Veeam Backup & Replication automatically deploys the helper appliance in Google Compute Engine only for the duration of the restore process and removes it immediately after that.

Depending on the type of backups you are restoring from and their location, the helper appliance may be required or optional. The helper appliance is required when you restore from:

- Backups of Google Compute Engine VM instances that are stored in external repositories.
- Backups of virtual and physical machines that are stored in object storage repositories.

The helper appliance is optional when you restore from backups of virtual and physical machines stored in backup repositories, or backups of Google Compute Engine virtual machines copied to backup repositories with backup copy jobs. It is recommended, however, to use the helper appliance in scenarios where it is optional, as the helper appliance may significantly improve restore performance. You can specify the helper appliance settings at the Helper Appliance step of the **Restore to Google Compute Engine** wizard.

# Requirements for Helper Appliance

When configuring a helper appliance, consider the following:

- If you want to restore from backups in an on-premises object storage repository, the helper appliance machine must have access to the source object storage repository. To provide access to object storage repositories, you can use VPN or Google Dedicated Interconnect. For more information, see the Google Cloud documentation.
- To upload one machine disk to Google Compute Engine, the helper appliance requires 1 GB RAM. Make sure that the type of Google Compute Engine instance selected for the helper appliance offers enough memory resources to upload all machine disks. Otherwise, the restore process may fail.
- The VPC route table must contain a route from the IP address of the Veeam Backup & Replication server to an active Google Cloud internet gateway. For more information on internet gateways and how to create route tables, see the Google Cloud documentation.
- Check that OS Login is disabled for the project where you plan to recover VM instances. For more information on how to configure OS Login, see the Google Cloud documentation. If you want to have OS Login enabled, use restore without the helper appliance.

# How Restore to Google Compute Engine Works

The workflow of the restore process depends on whether the helper appliance is used or not. For more information on the helper appliance, see Helper Appliances.

> **NOTE**
>
> If you use Google Cloud Plug-in for Veeam Backup & Replication and plan to restore Google Compute Engine virtual machines from restore points that were created using the appliance, you do not need to configure the helper appliance. Also, restore to Google Compute Engine works as described in the Performing Instance Restore section in the Veeam Backup for Google Cloud User Guide.

# Restoring to Google Compute Engine with Helper Appliance

If the helper appliance is used for restore to Google Compute Engine, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication creates a helper appliance in Google Compute Engine.

   During the restore process, the helper appliance communicates with backup infrastructure components over the SSH protocol and the network redirector that is deployed on the helper appliance.
2. For every disk of a backed-up workload, Veeam Backup & Replication creates a disk in Google Compute Engine.
3. Veeam Backup & Replication hot-adds empty disks to the helper appliance and restores backed-up data to the disks.
4. Veeam Backup & Replication creates a target instance in Google Compute Engine.
5. Veeam Backup & Replication detaches the disks from the helper appliance and attaches them to the target instance.
6. After the restore process is complete, Veeam Backup & Replication removes the helper appliance from Google Compute Engine.

# Restoring to Google Compute Engine without Helper Appliance

If the helper appliance is not used for restore to Google Compute Engine, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication uploads disks of a backed-up workload to Google Cloud Storage bucket.

   In Google Cloud Storage bucket, the uploaded disks are stored to the temporary bucket in the RAW format.

2. Veeam Backup & Replication imports the backed-up data from the temporary bucket in Google Cloud Storage to disks in Google Compute Engine.

3. Veeam Backup & Replication creates a target instance in Google Compute Engine and attaches disks to the target instance.

4. After the import process is complete, Veeam Backup & Replication removes the temporary bucket from Google Cloud Storage.

# Google Compute Engine IAM User Permissions

To enable restore of workloads to Google Compute Engine, do the following:

1. Grant the following roles to the IAM user whose credentials you plan to use to connect to Google Compute Engine:

   - Compute Admin role (roles/compute.admin)

     To avoid granting the Compute Admin role to the IAM user Compute Engine service account for security reasons, you can create a custom role with the following Compute Engine IAM permissions and grant it instead:

```
compute.addresses.list
compute.disks.create
compute.disks.delete
compute.disks.get
compute.disks.use
compute.disks.useReadOnly
compute.firewalls.create
compute.firewalls.delete
compute.firewalls.list
compute.globalOperations.get
compute.images.create
compute.images.delete
compute.images.get
compute.images.useReadOnly
compute.instances.attachDisk
compute.instances.create
compute.instances.delete
compute.instances.detachDisk
compute.instances.get
compute.instances.getGuestAttributes
compute.instances.list
compute.instances.setLabels
compute.instances.setMetadata
compute.instances.setTags
compute.instances.stop
compute.machineTypes.list
compute.networks.get
compute.networks.list
compute.networks.updatePolicy
compute.projects.get
compute.regions.list
compute.subnetworks.get
compute.subnetworks.list
compute.subnetworks.use
compute.subnetworks.useExternalIp
compute.zoneOperations.get
compute.zones.get
compute.zones.list
```

- Cloud Build Editor role (roles/cloudbuild.builds.editor)
- Project IAM Admin role (roles/resourcemanager.projectIamAdmin)
- Storage Admin role (roles/storage.admin)
- Storage HMAC Key Admin (roles/storage.hmacKeyAdmin)
- Viewer role (roles/viewer)

For more information, see the *Prerequisites for importing and exporting VM images* section in the Google Cloud documentation.

2. Make sure that the Cloud Build API is enabled. Then grant the following roles to the Cloud Build service account in Google Compute Engine:
   - Compute Admin role (roles/compute.admin)

     To avoid granting the Compute Admin role to the Cloud Build service account for security reasons, you can use the custom role that you created for the IAM user Compute Engine service account and grant it instead.

- Service Account Token Creator role (roles/iam.serviceAccountTokenCreator)
- Service Account User role (roles/iam.serviceAccountUser)
- [Optional: to export or import images that use shared VPCs] Compute Network User role (roles/compute.networkUser)

For more information, see the *Prerequisites for importing and exporting VM images* section in the Google Cloud documentation.

# Restoring to Google Compute Engine

To restore workloads to Google Compute Engine, use the **Restore to Google Compute Engine** wizard.

> **NOTE**
>
> If you use Google Cloud Plug-in for Veeam Backup & Replication and plan to restore Google Compute Engine virtual machines from restore points that were created using the appliance, you do not need to configure the helper appliance. Also, restore to Google Compute Engine works as described in the Performing Instance Restore section in the Veeam Backup for Google Cloud User Guide.

## Before You Begin

Before you restore workloads to Google Compute Engine, consider the following requirements and limitations:

- Check whether a helper appliance must be configured for restore. For more information, see Helper Appliances.
- The backup server and repositories with workload backup files must have access to the internet.

  If backup files are located on deduplicating storage appliances or shared folder repositories, the internet connection is required for gateway servers that communicate with these repositories.
- If you use a cloud-init-based Linux distribution, we recommend that you use SSH keys on these distributions. If you use a password, it is blocked after restore for security reasons. To reset the password on the restored VM, use the technologies described in Google Cloud Documentation.
- You must have a backup of the workload that you plan to restore to Google Compute Engine.
- Make sure that the Cloud Build API is enabled. For more information on enabling the Cloud Build API and other requirements for importing virtual disks into Google Compute Engine, see the Google Cloud Documentation.
- Make sure the IAM service account that you plan to use to restore workloads to Google Compute Engine has permissions to restore to Google Compute Engine. For more information, see Google Compute Engine IAM User Permissions.
- If you restore workloads from backups of virtual and physical machines (non-Google Compute Engine virtual machines), check the supported operating systems and their differences from standard images in the Google Cloud documentation.
- Check that the logical sector size of disks that you plan to restore is less than 4096 bytes. Contents of disks whose logical sector size is 4096 bytes will be unreadable in Google Compute Engine.
- Veeam Backup & Replication does not support restoring disks encrypted by BitLocker, except for restoring from backups created by Veeam Agent for Microsoft Windows. For more information, see the Veeam Agent for Microsoft Windows User Guide.

- If you plan to assign Google labels to the restored workload, check limitations for labels in the Google Cloud documentation.

## Step 1. Launch Restore to Google Compute Engine Wizard

To begin the restore process, do one of the following.

- On the **Home** tab, click **Restore** and select the type of backups from which you want to restore:
  - VMware vSphere
  - VMware Cloud Director
  - Microsoft Hyper-V
  - Agent
  - AWS
  - Azure IaaS backup
  - GCE backup
  - Nutanix AHV
  - oVirt KVM
  - Proxmox VE

  In the displayed window, click **Entire VM restore** > **Restore to public cloud** > **Restore to Google Compute Engine**.
- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary backup, select workloads that you want to restore and click **Restore to Google CE** on the ribbon. Alternatively, you can right-click one of the workloads that you want to restore and select **Restore to Google CE**.

- Double-click a full backup file (VBK) or backup metadata file (VBM) in a file browser.
  Veeam Backup & Replication will start its console. In the **Backup Properties** window, select the necessary workload and click **Restore** > **Restore to Google CE**.

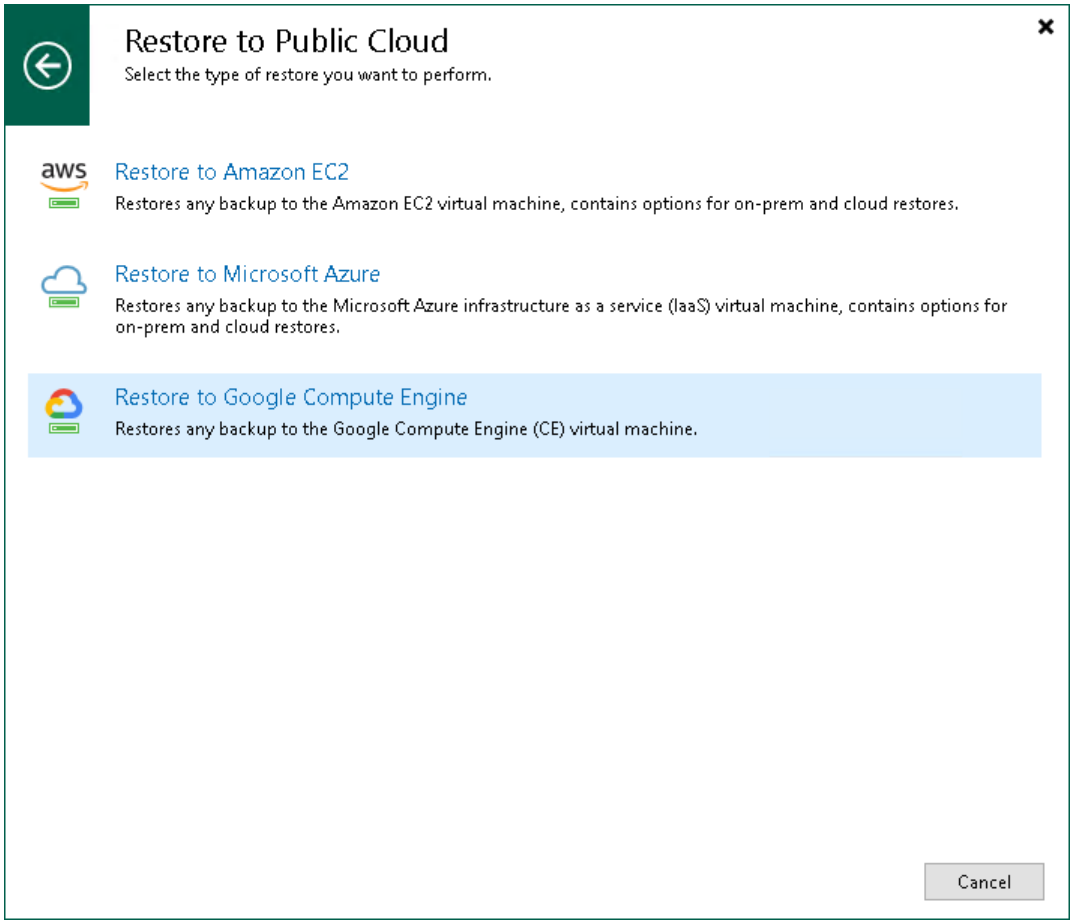

## Step 2. Select Workloads and Restore Points

At the **Virtual Machines** step of the wizard, specify the workload that you plan to restore and specify a restore point to which you want to restore the workload. By default, Veeam Backup & Replication restores workloads to the latest valid restore point in the backup chain.

# Selecting Workloads

To select workloads to restore:

1. On the right of the **Virtual machines to restore** list, click **Add**.
2. In the **Backup Browser** window, expand the necessary backup, select workloads and click **Add**.

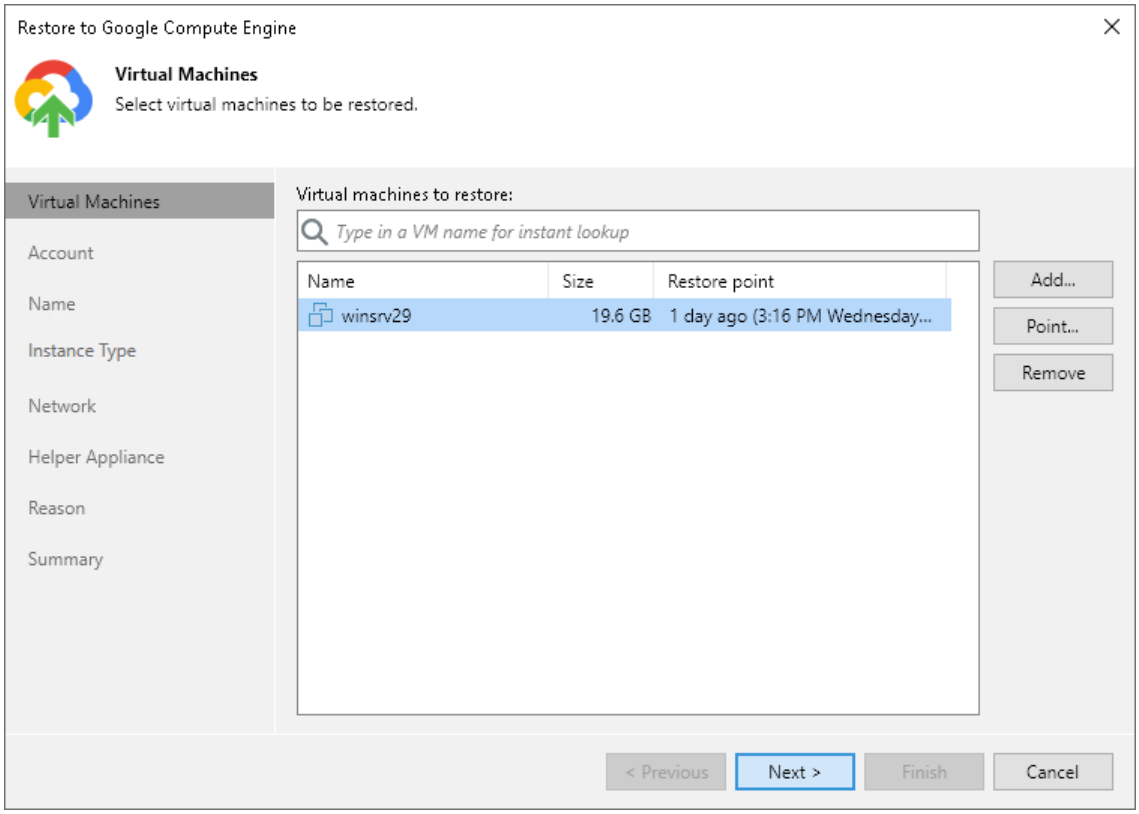


# Selecting Restore Points

To select a restore point, for a workload do the following:

1. In the **Virtual machines to restore** list, select a workload.
2. Click **Point** on the right.

3. In the **Restore Points** window, select a restore point to which you want to restore the workload.

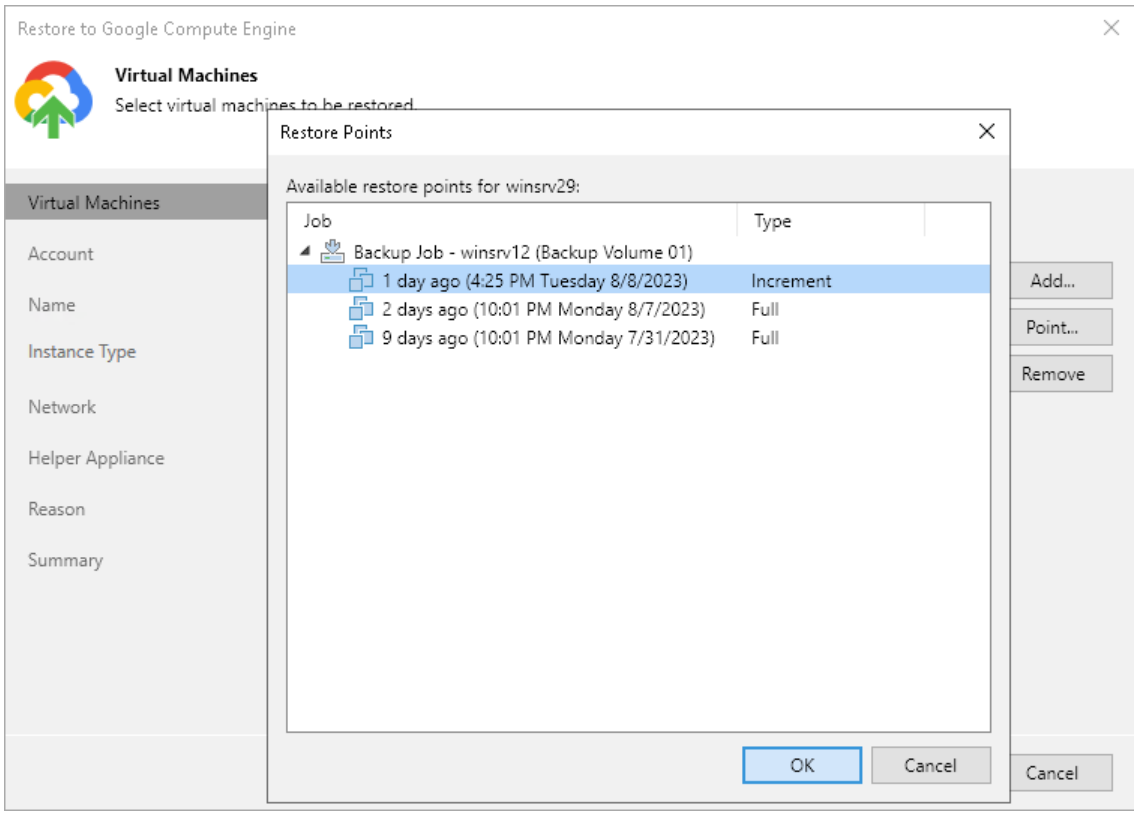

## Step 3. Specify Credentials and Datacenter Settings

At the **Account** step of the wizard, specify a Google Cloud service account, datacenter and availability zone to use for restore:

1. From the **Google Cloud service account** list, select user credentials to connect to Google Compute Engine.

   When you add credentials of the Google Cloud service account, Veeam Backup & Replication imports information about resources associated with this service account. During the restore process, Veeam Backup & Replication accesses these resources and uses them to create a target VM instance in Google Compute Engine.

   If you have not set up credentials beforehand in the Cloud Credentials Manager, click the **Manage accounts** link or click **Add** on the right to add the necessary credentials, as described in section Google Cloud Service Accounts.

2. From the **Datacenter** list, select the Google Cloud datacenter where Veeam Backup & Replication will restore your workload as a VM instance.

3. From the **Zone** list, select the availability zone inside the Google Cloud datacenter where the restored workload will reside.

If you restore a Google Compute Engine VM instance from a backup created by Veeam Backup for Google Cloud to the same Google Cloud region where the instance is placed, after you click **Next**, the wizard will offer you to use region settings associated with this instance.

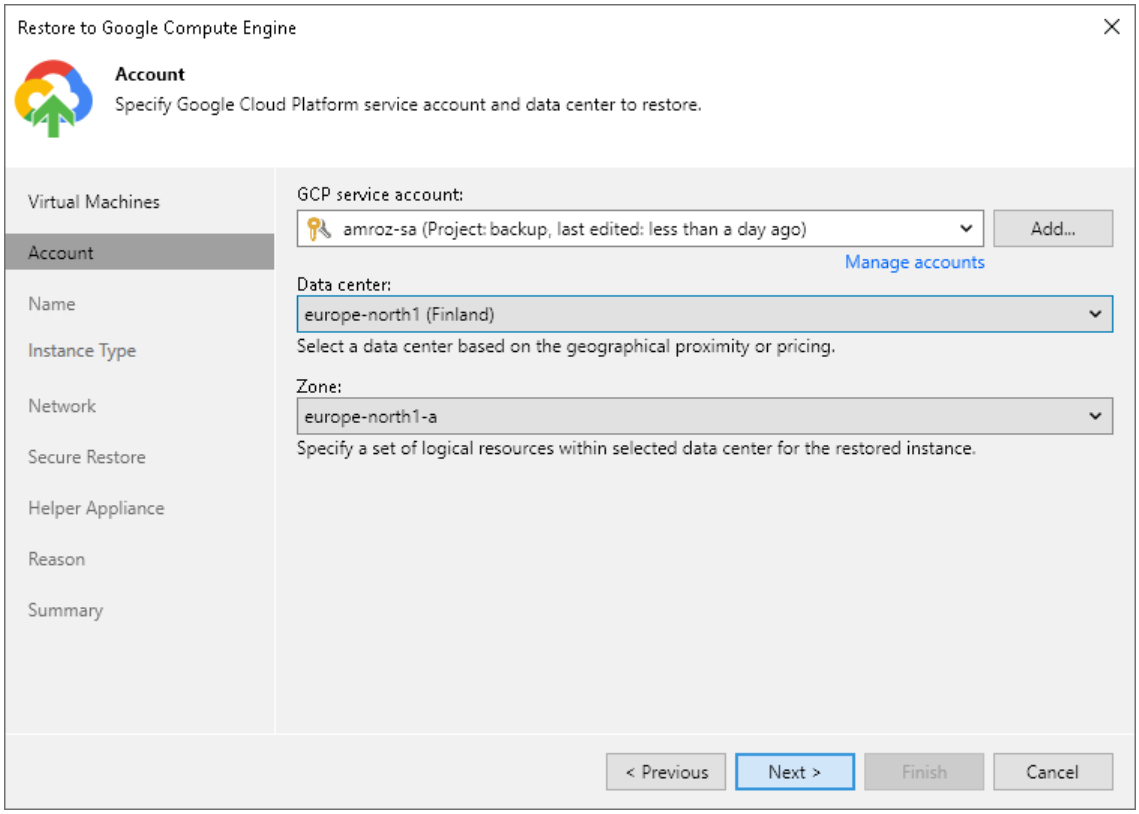

## Step 4. Specify Name and Labels

At the **Name** step of the wizard, you can specify names and manage Google labels for the restored workloads. By default, Veeam Backup & Replication uses original workload names.

# Specifying New Name

To define a new name for a restored workload:

1. In the **Instances** list, select a workload and click **VM Name**.
2. In the **Set name to** field of the **Change Name** window, enter a new name for the restored workload.

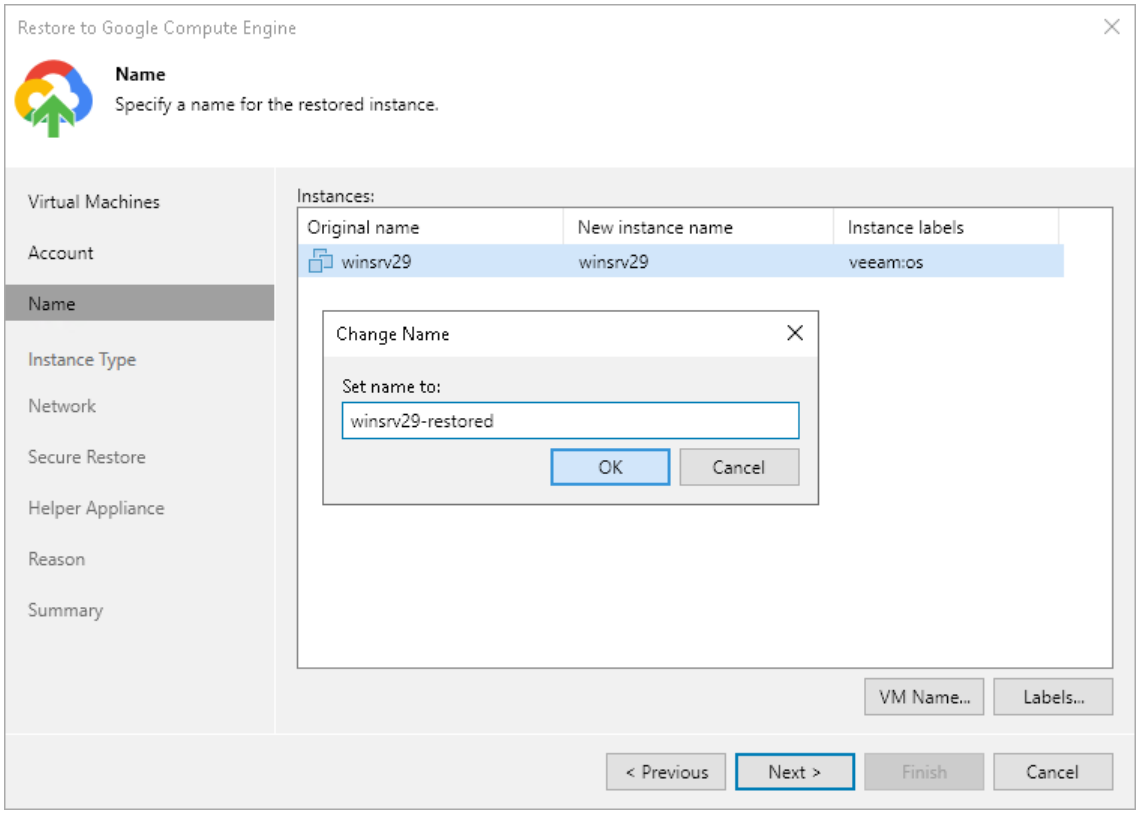


# Managing Google Cloud Labels

You can use Google Cloud labels to categorize instances in Google Compute Engine. A label is a tag with metadata that includes two properties: a key and a value. For more information on Google Cloud labels their format and limitations, see the Google Cloud documentation.

## Adding Label

To add a new label:

1. In the **Instances** list, select a workload and click **Labels**.
2. In the **Labels** window, click **Add**.
3. In the **GCE Instance Label** window, specify the **Key** and **Value** properties.

## Modifying Label

To modify a label:

1. In the **Instances** list, select a workload and click **Labels**.
2. In the **Labels** window, select the required label and click **Edit**.
3. In the **GCE Instance Label** window, edit the **Key** or **Value** properties.

## Deleting Label

To delete a label:

1. In the **Instances** list, select a workload and click **Labels**.
2. In the **Labels** window, select the required label and click **Remove**.

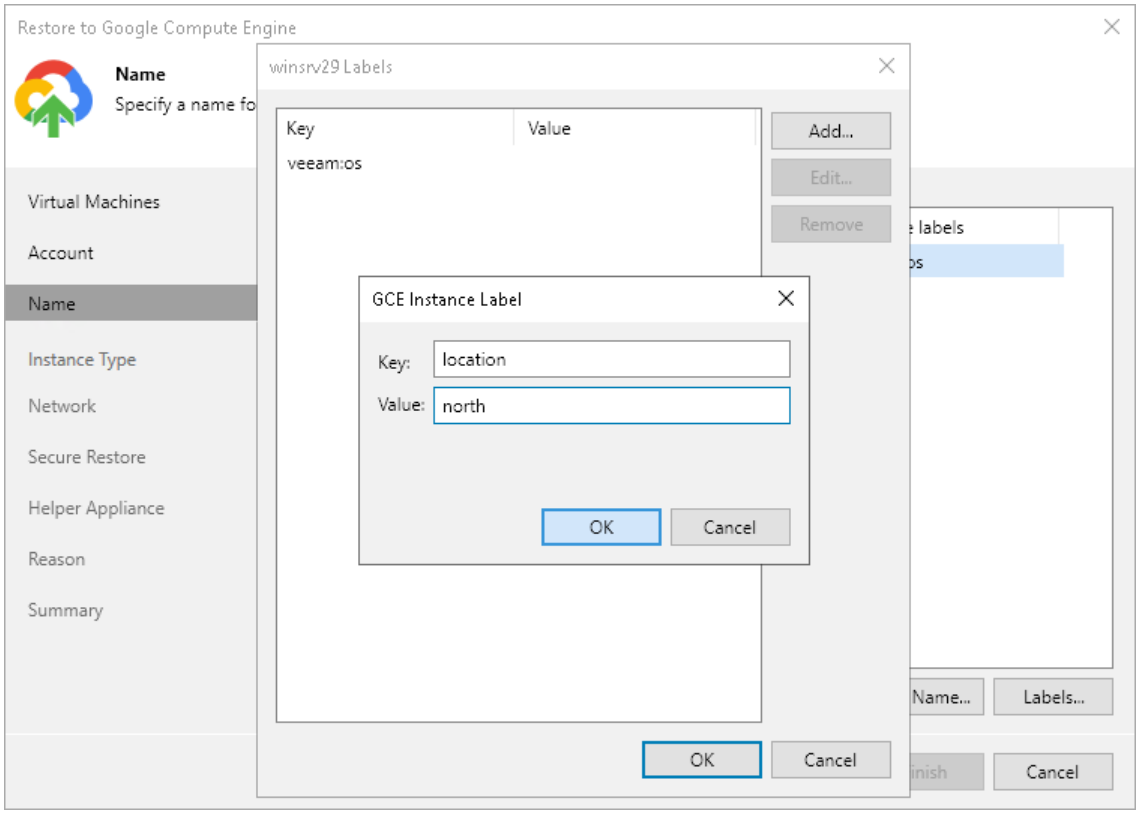

## Step 5. Specify Instance Type and Disks

At the **Instance Type** step of the wizard, select instance types and disk types for the restored workloads. By default, Veeam Backup & Replication restores all disks as Google Compute Engine disks of the Balanced persistent disk type. For information on types of Google Compute Engine disks, see the Google Cloud documentation.

# Selecting Instance Type

You can select the amount of computing resources that Google Compute Engine will provision for your restored workload — a Google Compute Engine instance type. Each instance type offers a unique combination of CPU and memory resources.

To select an instance type for a workload:

1. In the **Instances** list, select a workload and click **Edit**.
2. From the **Machine type** list, select the instance type for the restored workload.

   Make sure that you select the right instance type that corresponds to the initial workload configuration. For the information on instance types, see Google Cloud Documentation.

   Note that if you restore a Google Compute Engine instance from the backup created by Veeam Backup for Google Cloud, Veeam Backup & Replication will identify the type of a backed-up instance and select it by default.
3. From the **OS license** list, select an option that will define what license Google Compute Engine will use for the OS on the restored workload:
   - **Provided by GCE** — the OS license will be provided by Google Compute Engine.
   - **Bring Your Own License (BYOL)** — the OS license will be restored from the backup. For more information, see Google Cloud documentation.

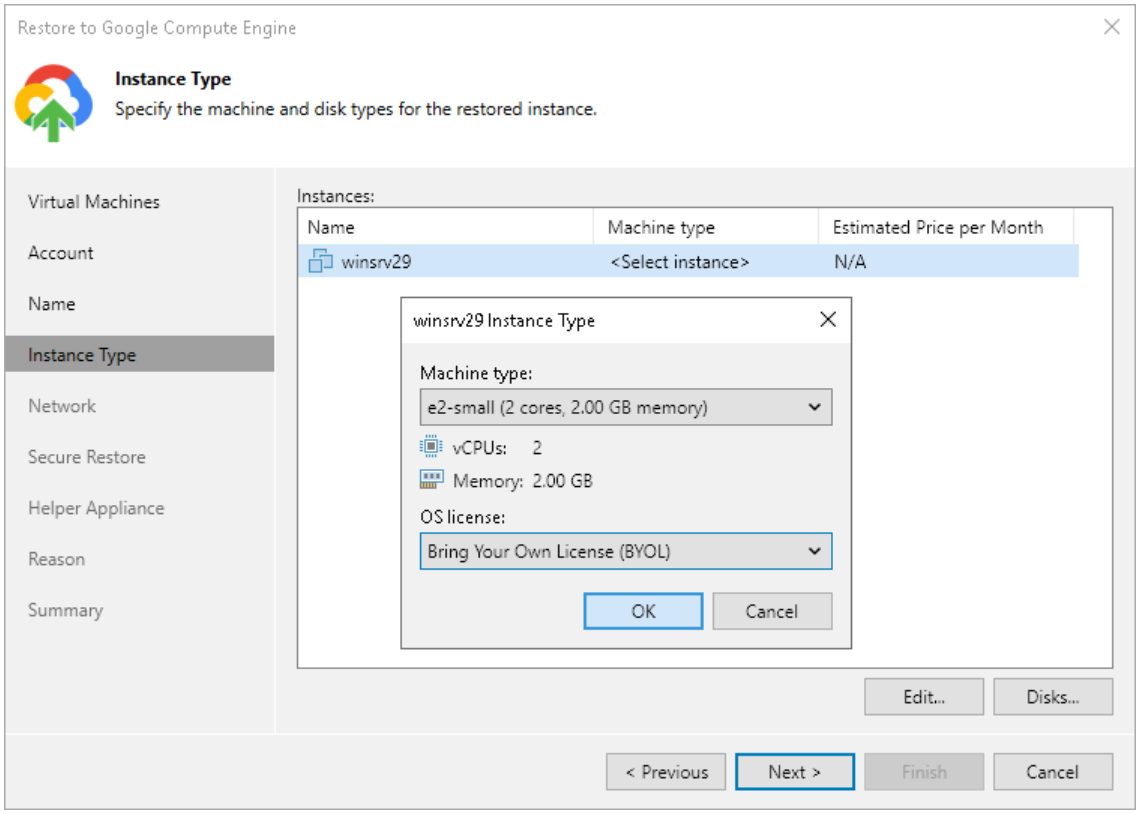

# Selecting Disk Type

You can restore all disks or specific disks of a workload. You can also change disk types of the restored disks.

To select workload disks for restore:

1. In the **Virtual machines** list, select a workload and click **Disks**.
2. In the **Disks To Restore** window, make sure that check boxes next to disks that you want to restore are selected. Clear check boxes next to disks that you do not want to restore.
3. Select a disk whose type you want to change and click **Disk type**.
4. In the **Select EC2 Disk Type** window, choose the disk type.

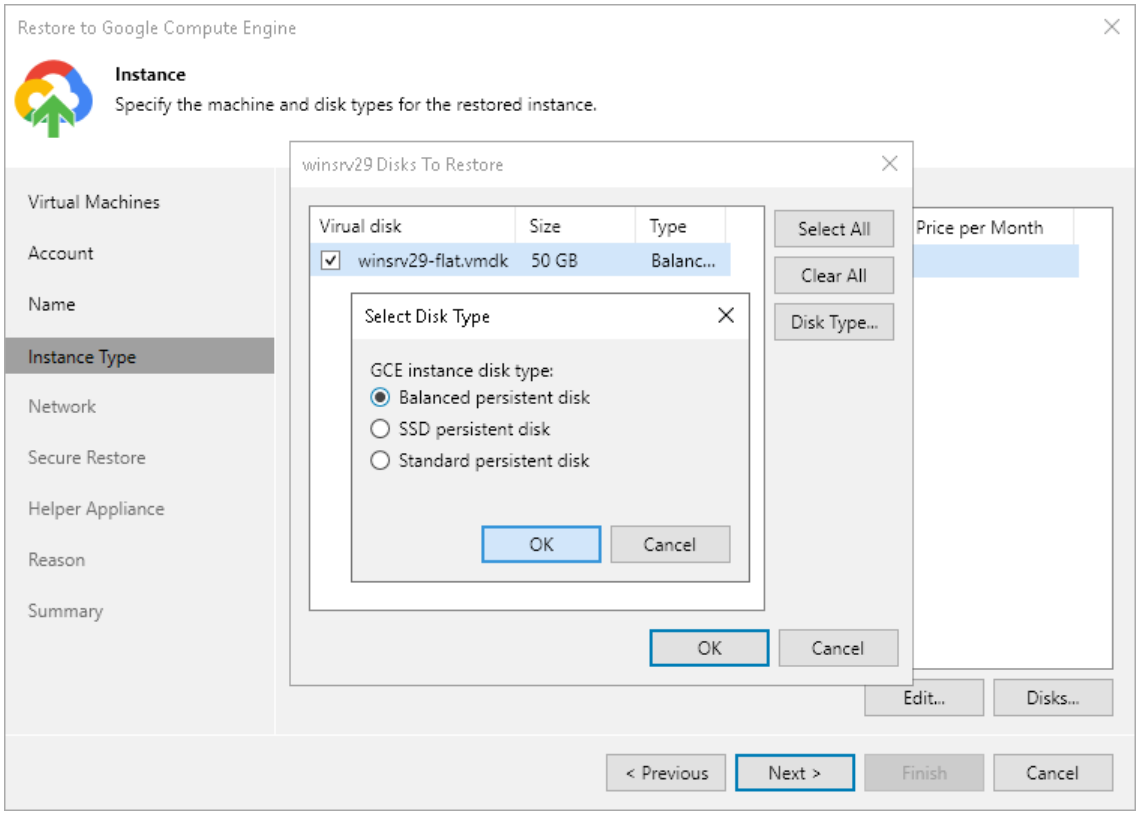

# Step 6. Select Google VPC

At the **Network** step of the wizard, you can select to which Google Virtual Private Cloud (Google VPC) network the workload must be connected after restore. You can also specify a subnet. For the information on Google VPC, see the Google Cloud documentation.

To define network settings for the restored workload, do the following:

1. From the **VPC** list, select the VPC where the restored workload will be launched.
2. From the **Subnet** list, select an IP address range for the selected VPC.
3. From the **Public IP address** list, select one of the following:
   - **Assign (restored VM will be accessible from the Internet)** — if you want to assign a public IP to the restored workload. For security reasons, make sure firewall rules are properly configured in the target VPC.
   - **Do not assign (more secure)** — if you do not want to assign a public IP.

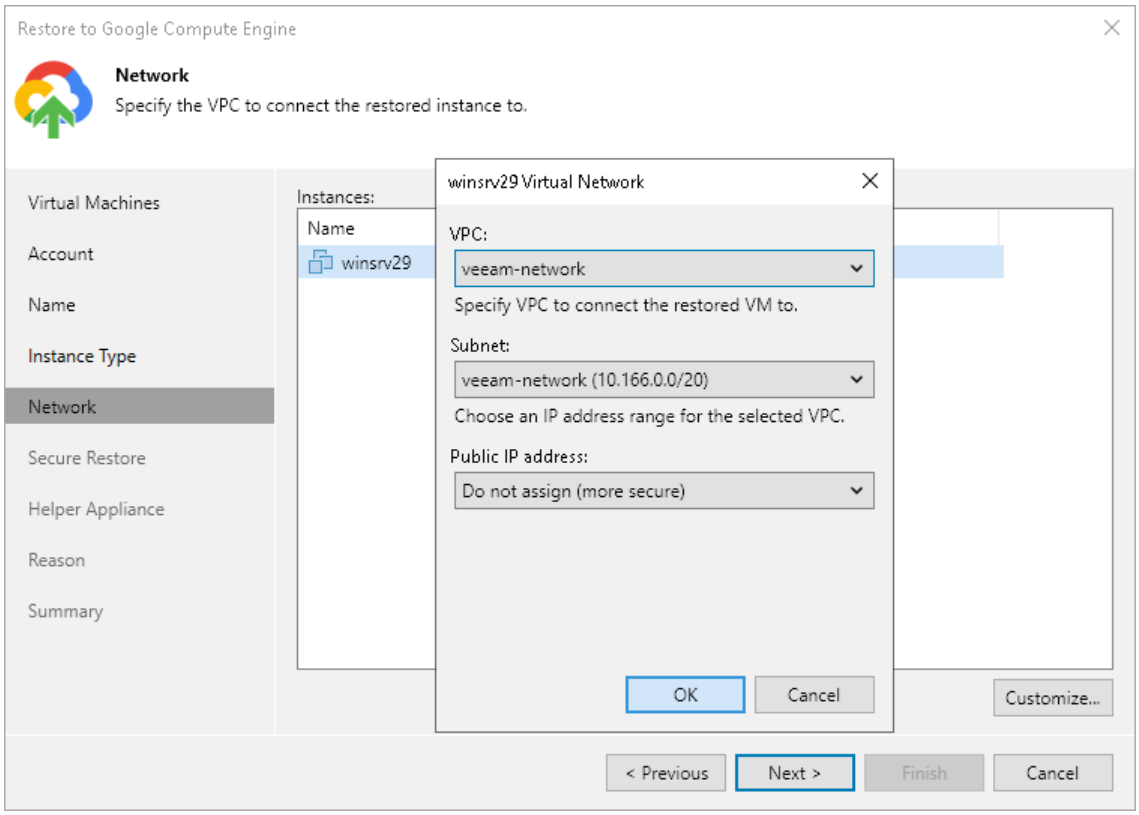

# Step 7. Specify Secure Restore Settings

This step is available if you restore Microsoft Windows workloads.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Select which action Veeam Backup & Replication will take if scan finds a virus threat:
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.
   - **Proceed with recovery but connect the VM to a different network**. Select this action if you want to restore the workload to a different Google Cloud network.

     Click the **Click to change** link to specify the VPC and subnet for this network.

3. Select the **Continue scanning all remaining files after the first occurrence** check box if you want to continue workload scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Malware Scan Results.

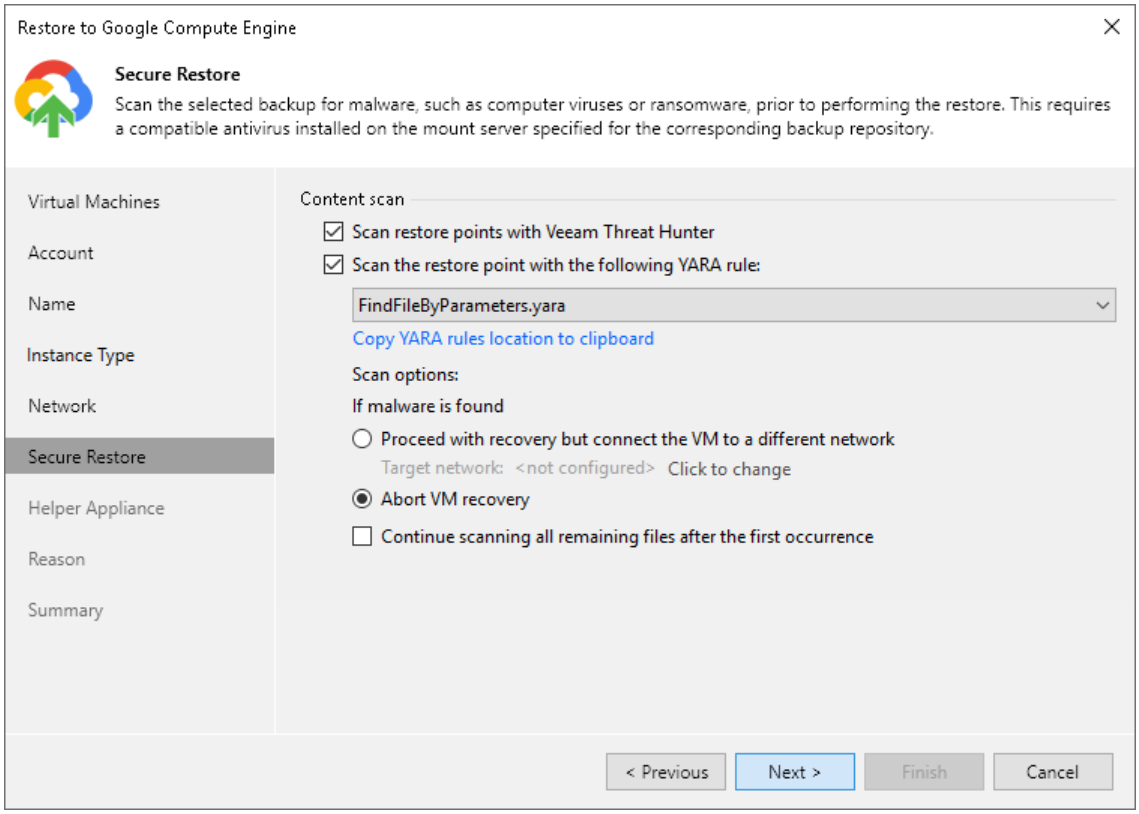

# Step 8. Configure Helper Appliance

At the **Helper Appliance Settings** step of the wizard, you can specify helper appliance settings. A helper appliance is an auxiliary Linux-based instance used to upload disks of a backed-up workload to Google Compute Engine. For more information on the helper appliance and requirements for it, see Helper Appliances.

To specify helper appliance settings, do the following:

1. Select the **Use the helper appliance** check box.
2. Click **Customize**.
3. From the **Machine type** list, select the instance type for the helper appliance.
4. From the **VPC** list, select the VPC network for the helper appliance.
5. From the **Subnet** list, select the subnet for the helper appliance.
6. In the **Redirector port** field, specify the port that Veeam Backup & Replication will use to route requests between the helper appliance and backup infrastructure components.

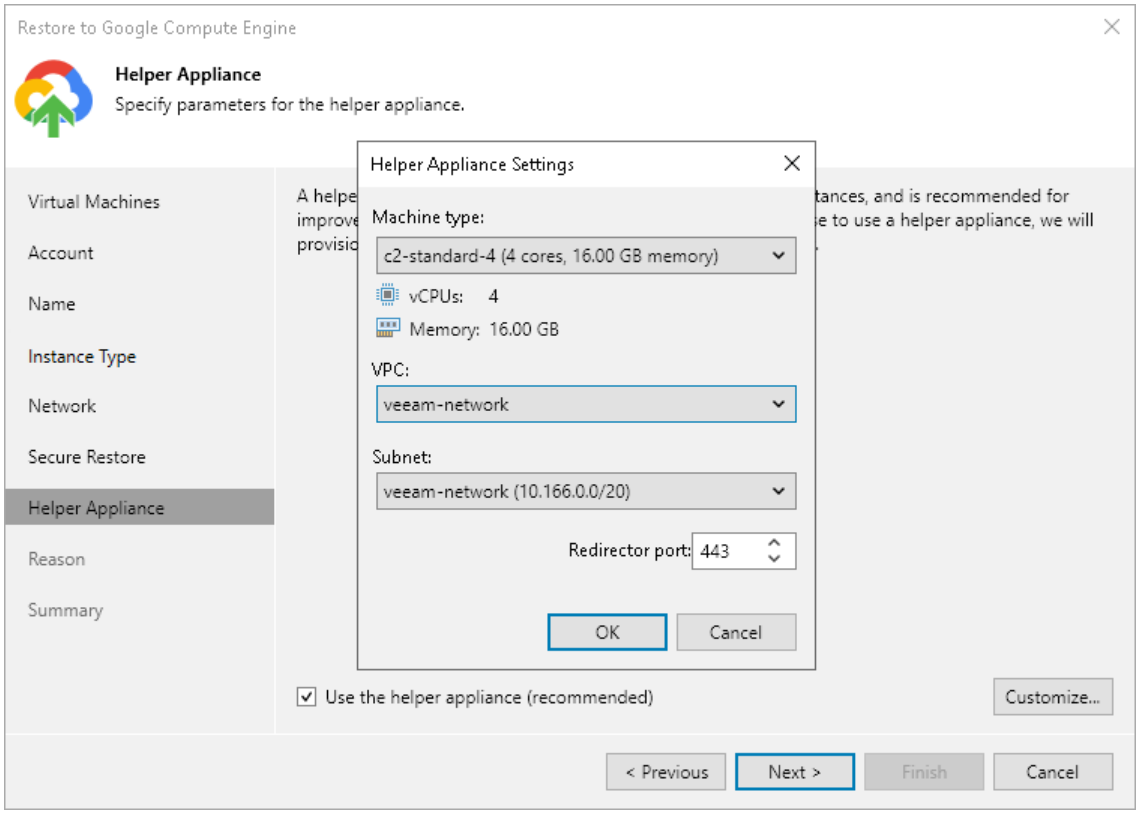

# Step 9. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the workload. The information you provide will be saved in the session history in Veeam Backup & Replication, and you can view it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

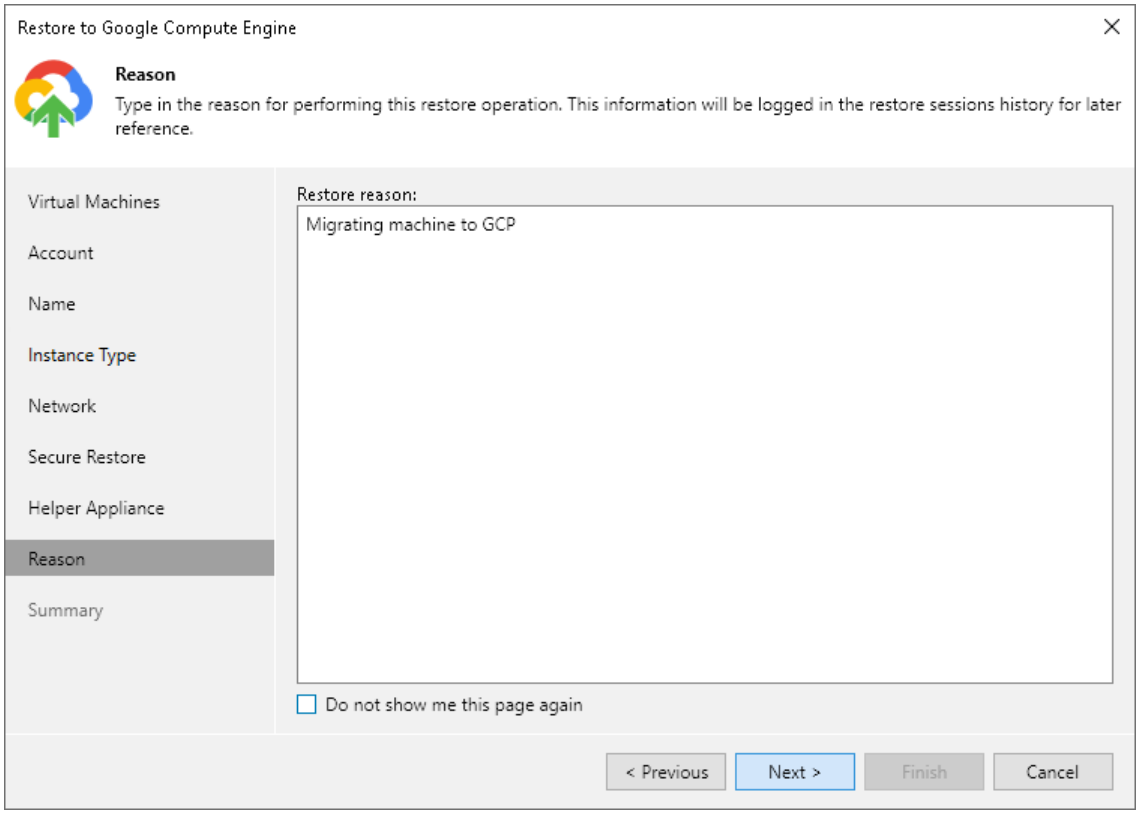

# Step 10. Verify Restore Settings

At the **Ready to Restore** step of the wizard, check the specified settings and click **Finish**. If you want to start the VM instance right after restore, select the **Power on VM after restoring** check box.

You can track the restore process in the **Restore Session** window. If you need to cancel the workload restore, click the **Cancel** restore task link.

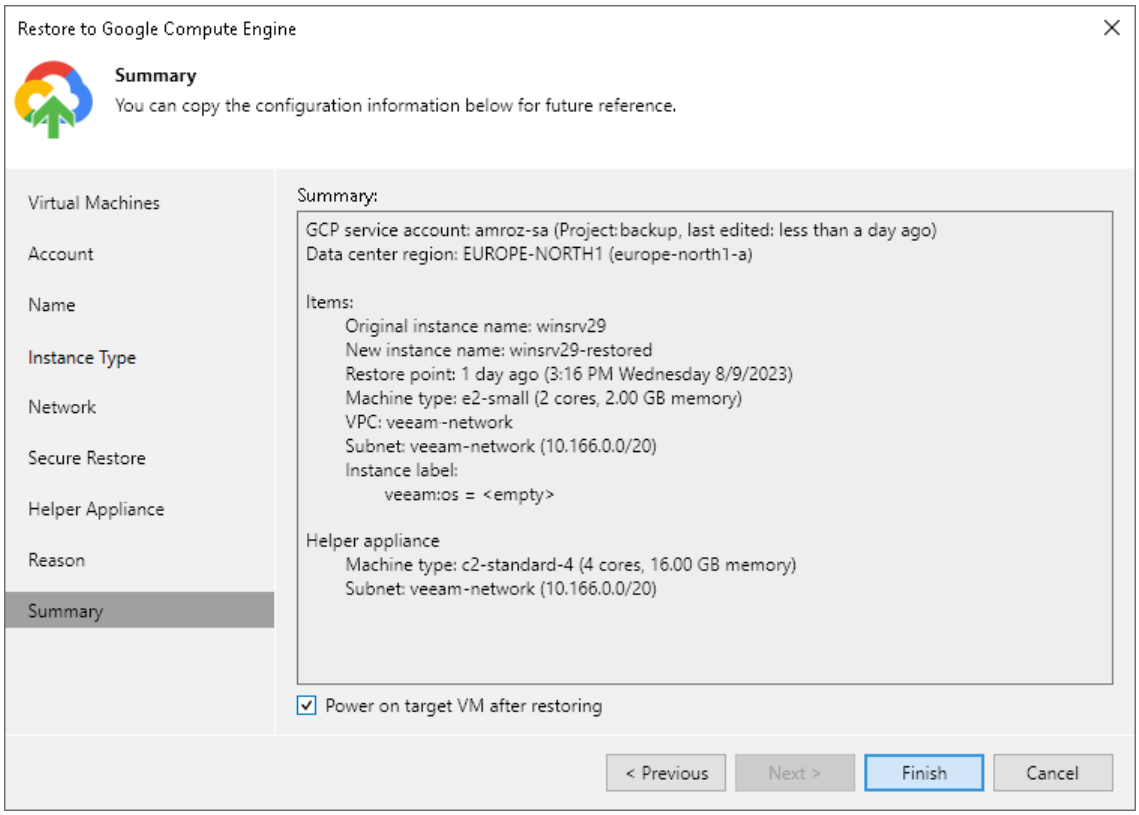

# Restore to Nutanix AHV

Veeam Backup & Replication allows you to recover different workloads as Nutanix AHV VMs.

> **IMPORTANT**
>
> To restore to Nutanix AHV, you must install Nutanix AHV Plug-in on the backup server. To learn more, see the Installation section in the Veeam Backup for Nutanix AHV User Guide.

The restore procedure of entire workloads to Nutanix AHV practically does not differ from the procedure described in the Performing VM Restore section in the Veeam Backup for Nutanix AHV User Guide.

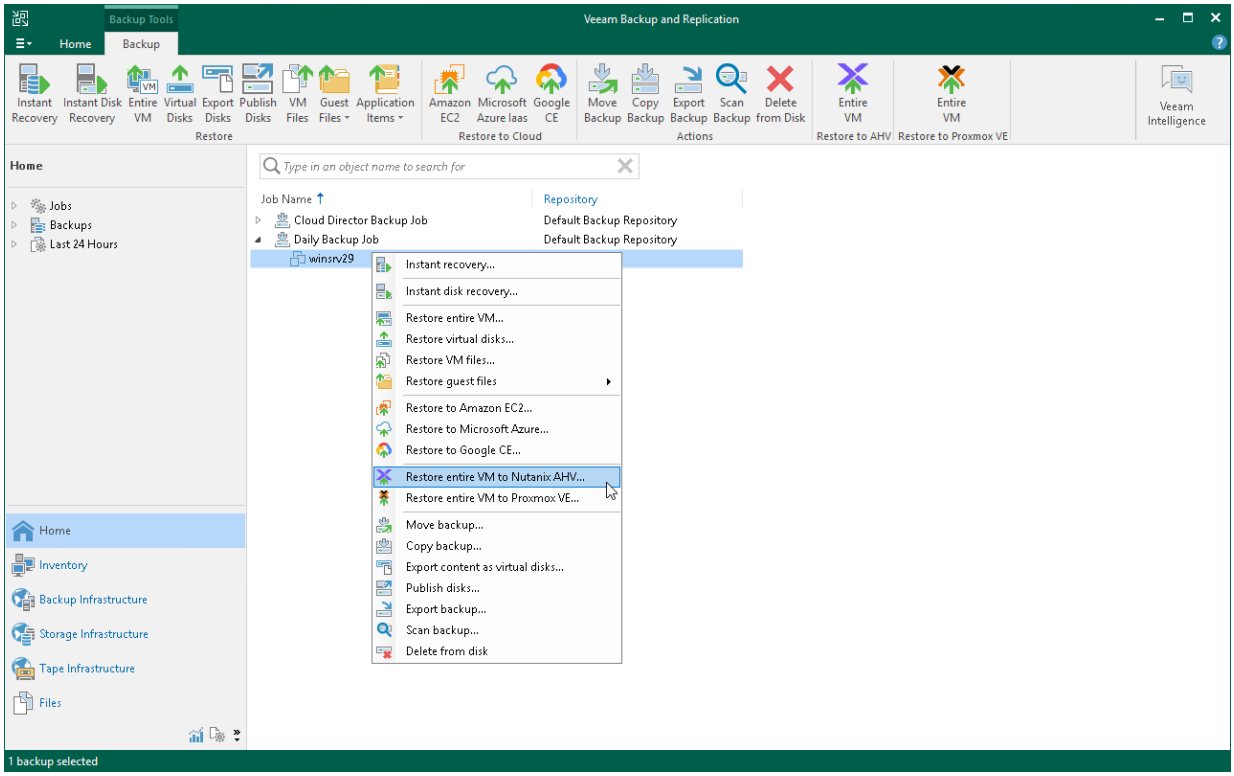

# Restore to Proxmox VE

Veeam Backup & Replication allows you to recover different workloads as Proxmox VE VMs.

The restore procedure of entire workloads to Proxmox VE practically does not differ from the procedure described in the Performing VM Restore section in the Veeam Backup for Proxmox VE User Guide.

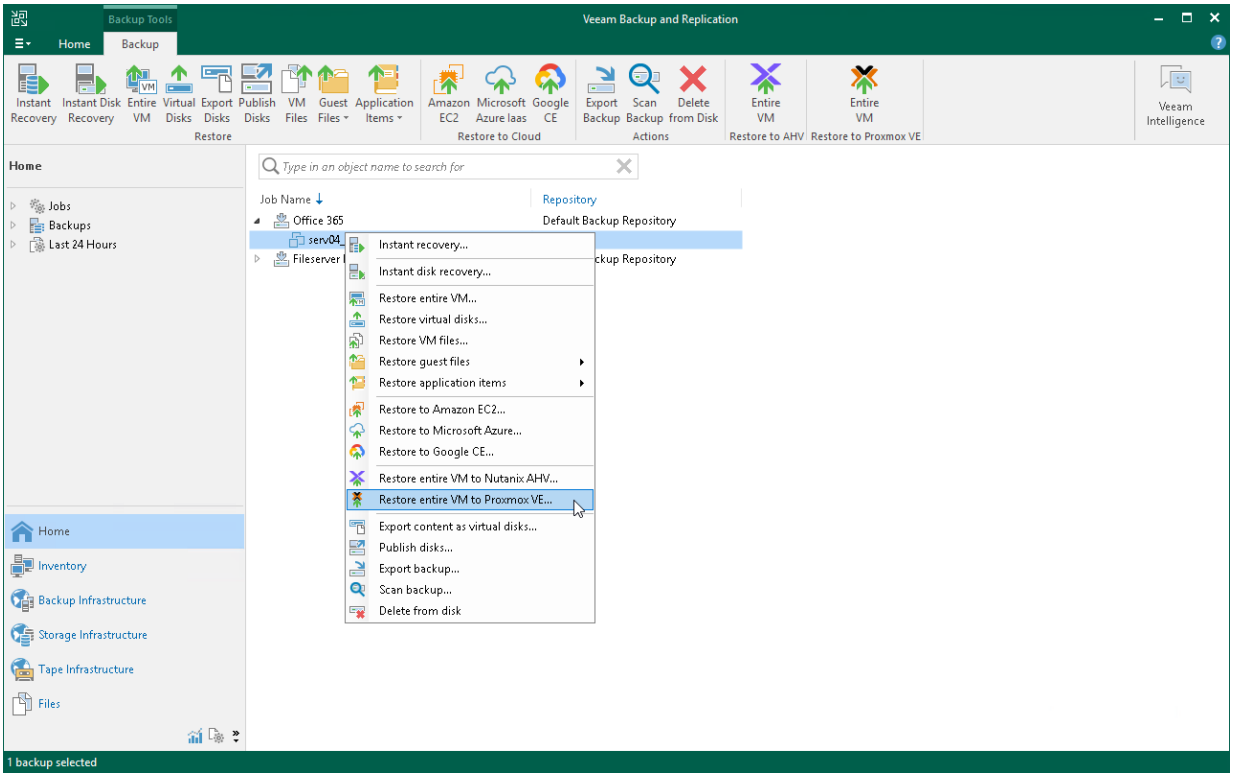

# Restore to OLVM and RHV

Veeam Backup & Replication allows you to recover different workloads as OLVM and RHV VMs.

> **IMPORTANT**
>
> To restore to OLVM and RHV, you must install oVirt KVM Plug-in for Veeam Backup & Replication on the backup server. To learn more, see the Installation section in the Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization User Guide.

The restore procedure of entire workloads to oVirt KVM practically does not differ from the procedure described in the Performing VM Restore section in the Veeam Backup for OLVM and RHV User Guide.

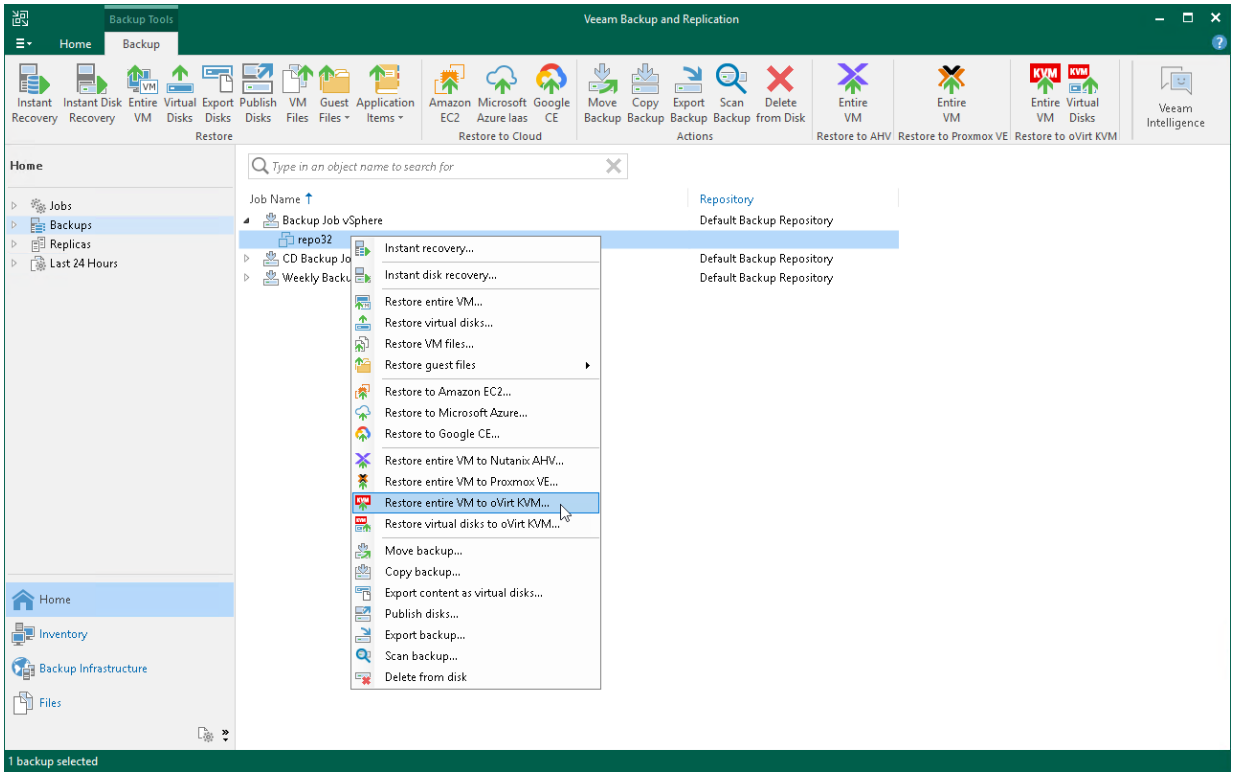

# Disk Recovery

Disk recovery includes the following methods:

- Instant Disk Recovery — to instantly recover VM disks from backup files. The disks are recovered in their initial format.

  Instant Disk Recovery helps improve recovery time objectives (RTO) but publishes disks with the reduced I/O performance, that is, provides "temporary spares". To provide the full I/O performance, you must finalize Instant Disk Recovery — migrate the recovered disks to the production environment. If you do not want to migrate the recovered disks, you can stop publishing them. This removes the recovered disks.

  Use Instant Disk Recovery if you want to keep the target VM (VM to which you want to attach recovered disks) turned on. If the VM can be turned off, you can use virtual disk restore.

- Instant First Class Disk Recovery — to instantly recover VM disks from backup files and register them as First Class Disks (FCDs). FCDs are a type of improved VMDK that allow you to perform different tasks: create snapshots, delete and restore data stored on VMDK without attaching these disks to a VM.

  This method is similar to Instant Disk Recovery except for the format in which disks are recovered.

  Use Instant FCD Recovery if you want to recover recovered disks as FCDs.

- Virtual disk restore — to restore VM disks. When you restore disks, you extract them from backups to the production storage. Virtual disk restore takes more resources to complete than Instant Disk Recovery but restores disks with full I/O performance. You also do not need to perform additional steps to finalize the restore process.

  Use virtual disk restore if the target VM can be turned off. If the VM must stay turned on, perform Instant Disk Recovery.

- Disk export — to convert disks of different workloads (EC2 instances, Microsoft Azure VMs and so on) in the VMDK, VHD or VHDX formats. Then you can export these disks to the backup infrastructure or place them on a datastore connected to an ESXi host (for VMDK disk format only).

  Use disk export to convert disks to the VMDK, VHD or VHDX formats.

- Disk publishing — to get the backup content without restoring all disks from a backup. After you publish a disk, you can browse its content, perform antivirus scan and other testing.

- Disk Restore to oVirt KVM — to recover different workloads as OLVM and RHV VM disks.

# Instant Disk Recovery

With Instant Disk Recovery, you can immediately restore VM disks from a backup file and publish them in the initial format. If you want to register disks as First Class Disks (FCD), see Instant First Class Disk (FCD) Recovery. For the list of all disk recovery methods and their brief descriptions, see Disk Recovery.

> **NOTE**
>
> You can also recover VMware vSphere VM disks directly from storage snapshots.

Use Instant Disk Recovery if you need to:

- Recover VM disks, not an entire VM. Otherwise, use Instant Recovery to VMware vSphere.
- Recover VM disks and keep the target VM (to which you want to attach recovered disks) turned on. If the VM can be turned off, use Virtual Disk Restore.

The Instant Disk Recovery is a process that must be finalized. First, test the recovered disks and then decide whether to migrate them to the production environment or stop publishing the disks. To migrate the disks, Veeam Backup & Replication uses the Quick Migration mechanism. For more information on how to finalize Instant Disk Recovery, see Finalizing Instant Disk Recovery.

## How Instant Disk Recovery Works

Veeam Backup & Replication performs Instant Disk Recovery using vPower technology.
Veeam Backup & Replication performs the following steps to recover disks:

1. Checks whether the vPower NFS datastore is mounted to the ESXi host and contains VMDK files of the recovered disks.
2. [If you replace an existing disk] Unmounts and deletes the existing disk.
3. Connects the recovered disks to the target VM.
4. Initiates a creation of a protective snapshot for the target VM. If Instant Disk Recovery process fails, the protective snapshot guarantees no data loss.
5. Writes changes made to the recovered disks to redo logs on the vPower NFS server.

## Performing Instant Disk Recovery

To recover disks using Instant Disk Recovery, use the **Instant Disk Recovery** wizard.

### Before You Begin

Before you perform Instant Disk Recovery, consider the following:

- Prerequisites for Instant Disk Recovery from storage snapshots are listed in the Data Recovery from Storage Snapshots section in the Storage System Snapshot Integration Guide.
- You can restore a virtual disk from a backup that has at least one successfully created restore point.
- If you want to scan disk data for viruses, check the secure restore requirements and limitations.
- Instant Disk Recovery to VMware Cloud Director is not supported.

- You must have at least 10 GB of free disk space on the datastore where write cache folder is located. This disk space is required to store virtual disk updates for the restored VM.

  By default, Veeam Backup & Replication writes virtual disk updates to the `IRCache` folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`.

## Step 1. Launch Disk Recovery Wizard

To launch the **Instant Disk Recovery** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Disk restore** > **Instant disk recovery**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand a backup or storage snapshot, select a VM whose disks you want to restore and click **Instant Disk Recovery** on the ribbon. Alternatively, you can right-click a VM whose disks you want to restore and select **Instant disk recovery**.

Alternatively, to recover VMware vSphere VMs from storage snapshots, you can open the **Storage Infrastructure** view. In the inventory pane, expand the storage system tree and select the necessary volume snapshot. In the working area, select a VM whose disks you want to restore and click **Instant Disk Recovery** on the ribbon. You can also right-click a VM and select **Instant disk recovery**.

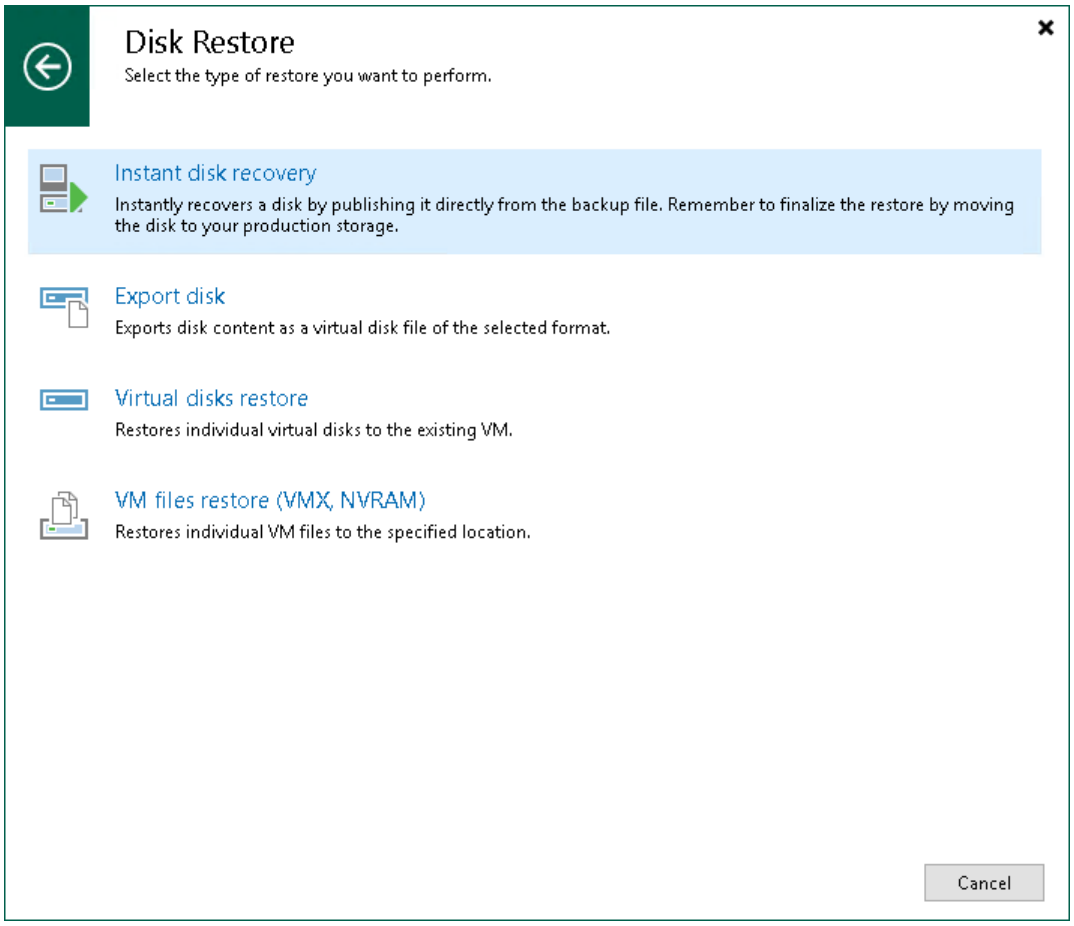

# Step 2. Select Source VM

At the **Virtual Machine** step of the wizard, select a VM whose disks you want to recover.

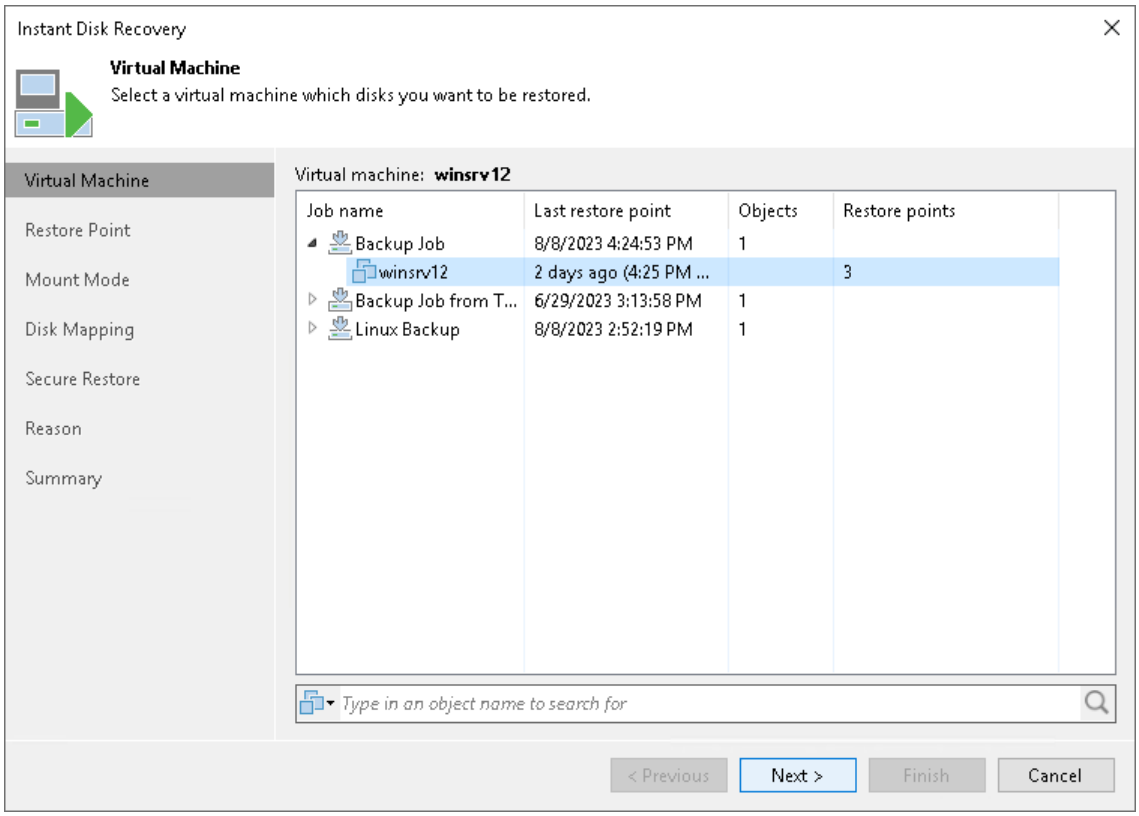

## Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point from which you want to recover VM disks.

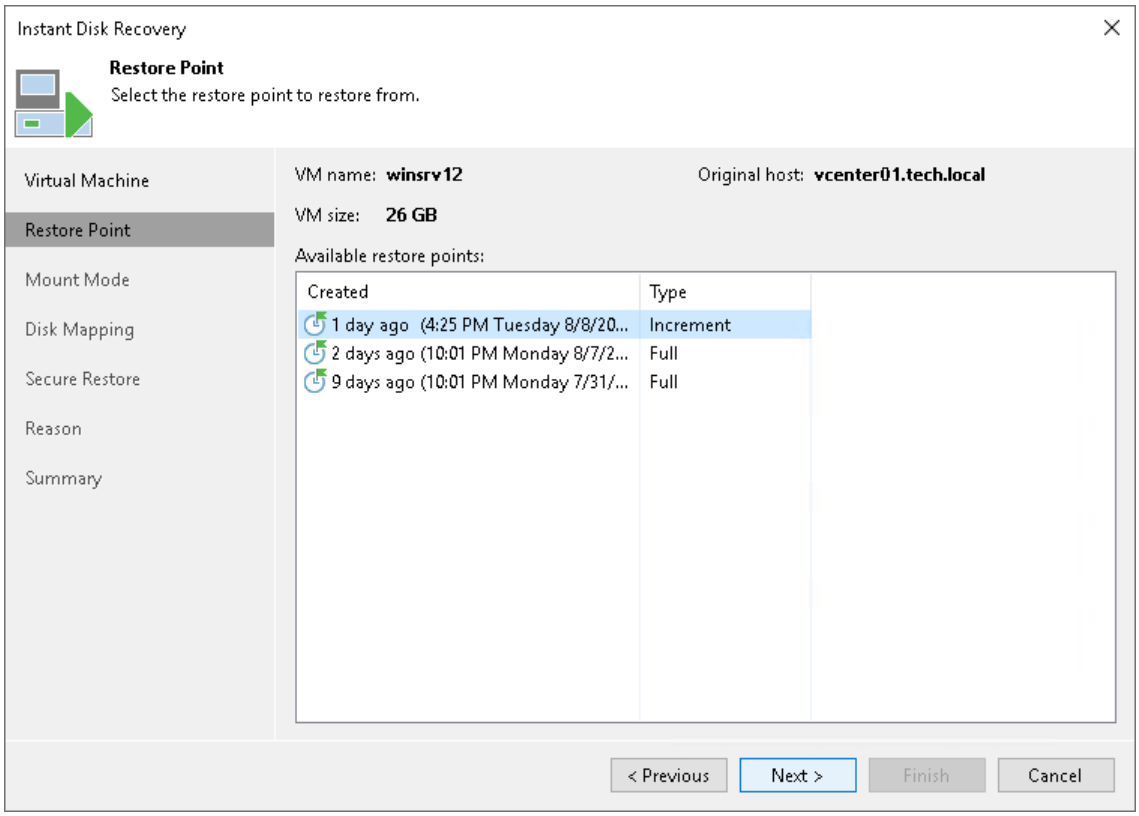

# Step 4. Select Mount Mode

This step is not available if you restore VMware vSphere VM disks from a storage snapshot.

At the **Mount Mode** step of the wizard, select the **VM disk** option to register virtual disks on a VM added to an ESXi host.

If you want to register virtual disks on a cluster as FCDs, select the First class disk (FCD) option. In this case, steps of the wizard differ and Veeam Backup & Replication performs Instant FCD Recovery as described in section Performing Instant First Class Disk (FCD) Recovery.

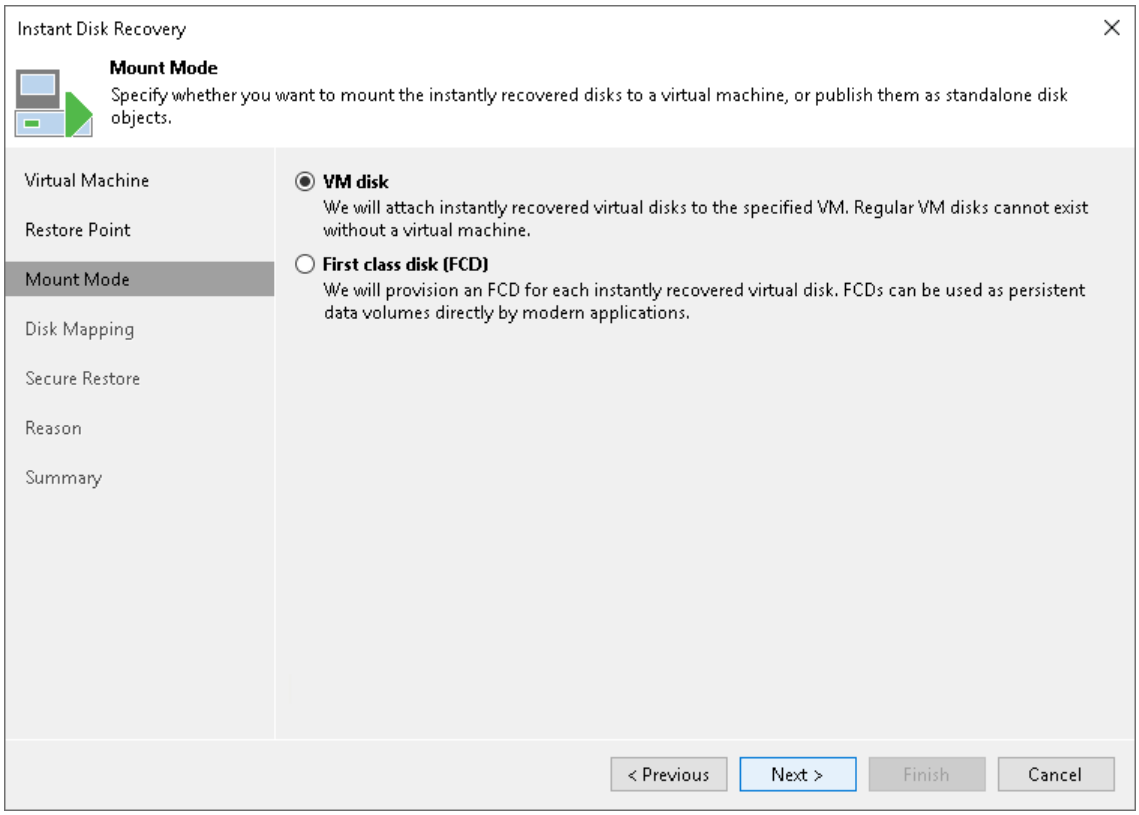

# Step 5. Select Virtual Disks to Restore

At the **Disk Mapping** step, select virtual disks that you want to restore and choose a VM to which the disks must be attached:

1. By default, Veeam Backup & Replication maps the recovered disks to the original VM. If the original VM was relocated or if you want to attach disks to another VM, select a target VM. For this, click **Choose** and select a VM from the virtual environment.

   > **IMPORTANT**
   >
   > Do not use the backup server or a VM where the vPower NFS Service is installed as a target VM.

2. In the **Disk mapping** section, select virtual disks that you want to restore.
3. To change a virtual device node where a virtual disk will be restored, select a disk in the **Disk mapping list** and click **Change**. In the **Virtual Disk Properties** window, select the node:
   - If you want to replace an existing virtual disk, select an occupied virtual node. The original disk will be deleted.
   - If you want to attach the restored disk to the VM as a new drive, select a node that is not occupied yet.

     > **NOTE**
     >
     > If you restore a virtual disk to an unoccupied node, Veeam Backup & Replication restores the disk in the offline state. To work with it, you need to bring the disk online.

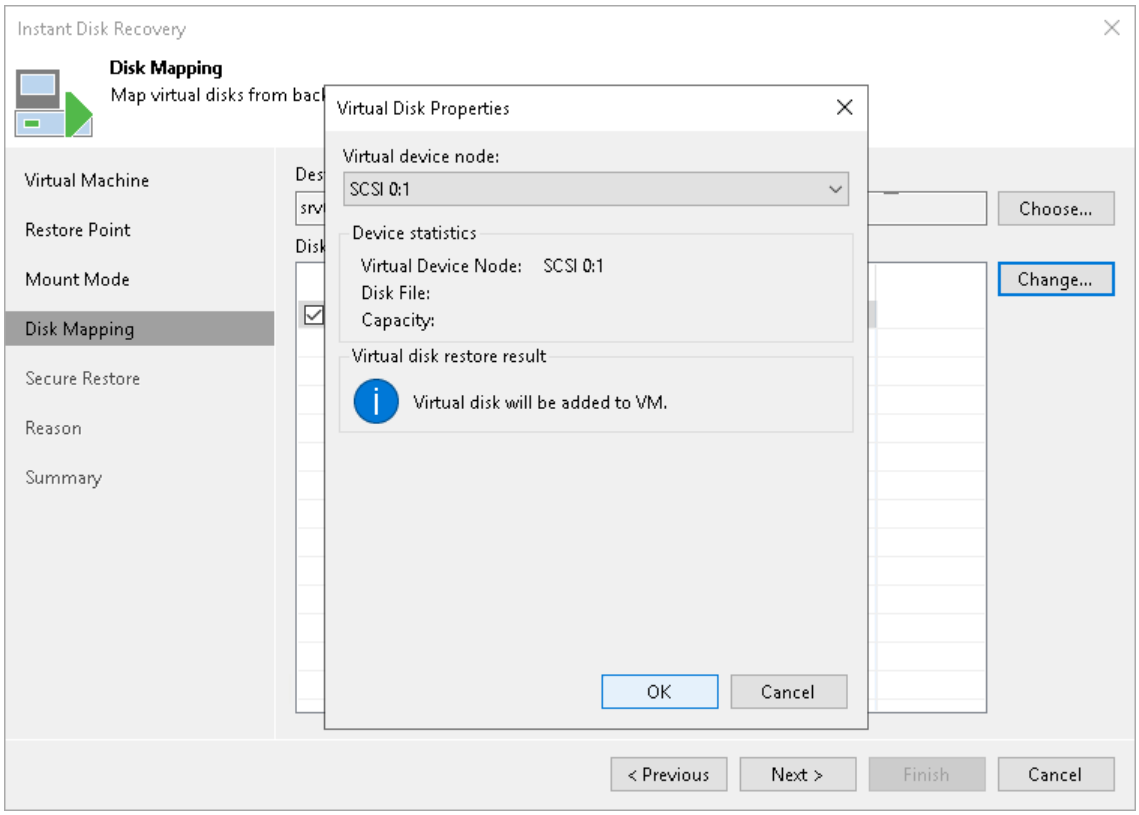

# Step 6. Specify Secure Restore Settings

This step is available if you restore disks of Microsoft Windows VMs. However, this step is not available if you restore disks from storage snapshots.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.

2. Select the **Continue scanning all remaining files after the first occurrence** check box if you want to continue the virtual disk scan after a virus threat is detected. For information on how to view results of the malware scan, see Viewing Malware Scan Results.

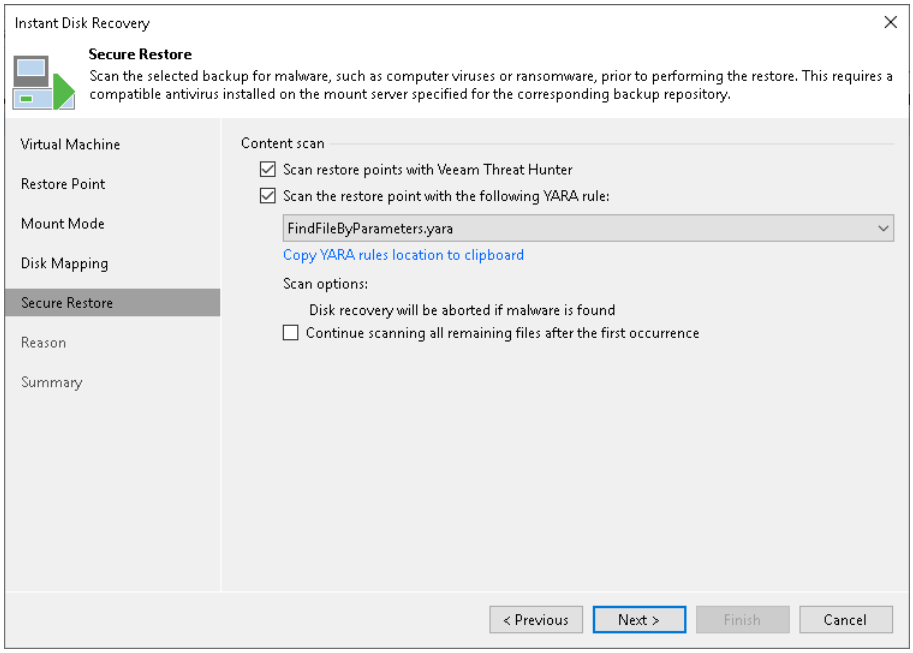

# Step 7. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring VM disks. This information will be saved in the session history so that you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

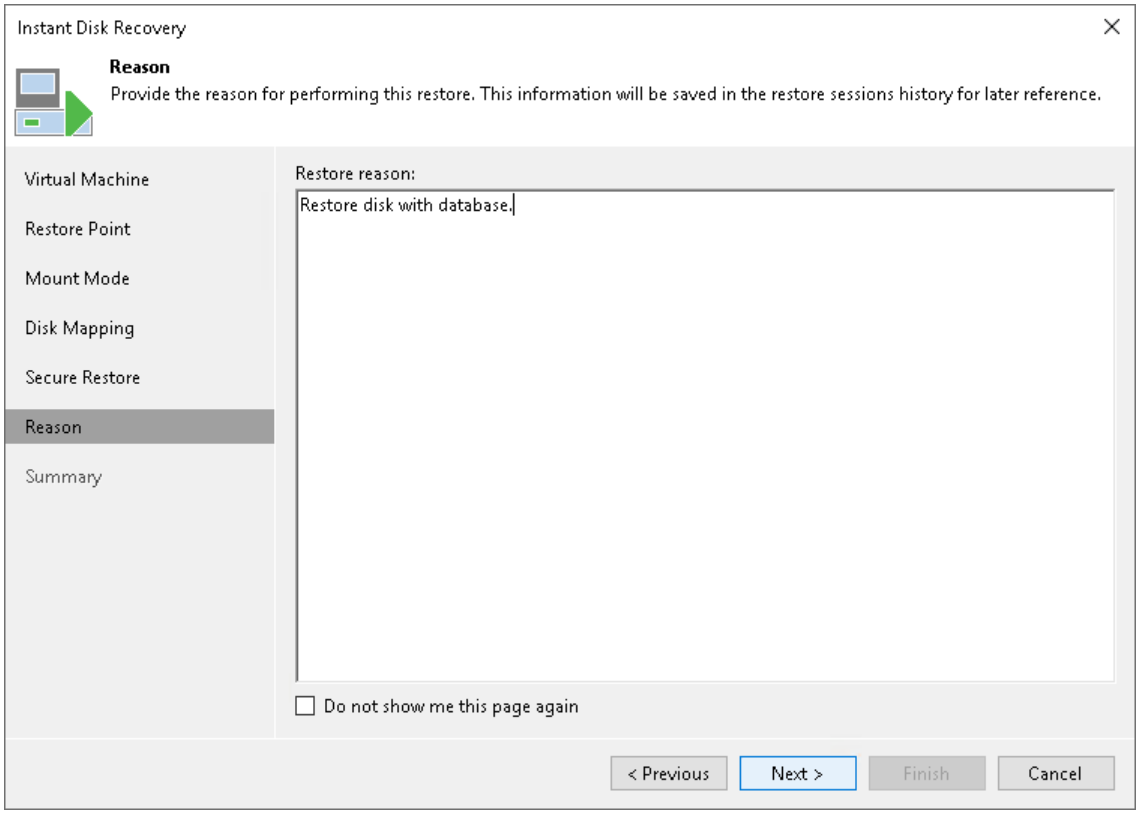

## Step 8. Verify Recovery Settings

At the **Summary** step of the wizard, check settings of Instant Disk Recovery and click **Finish**.

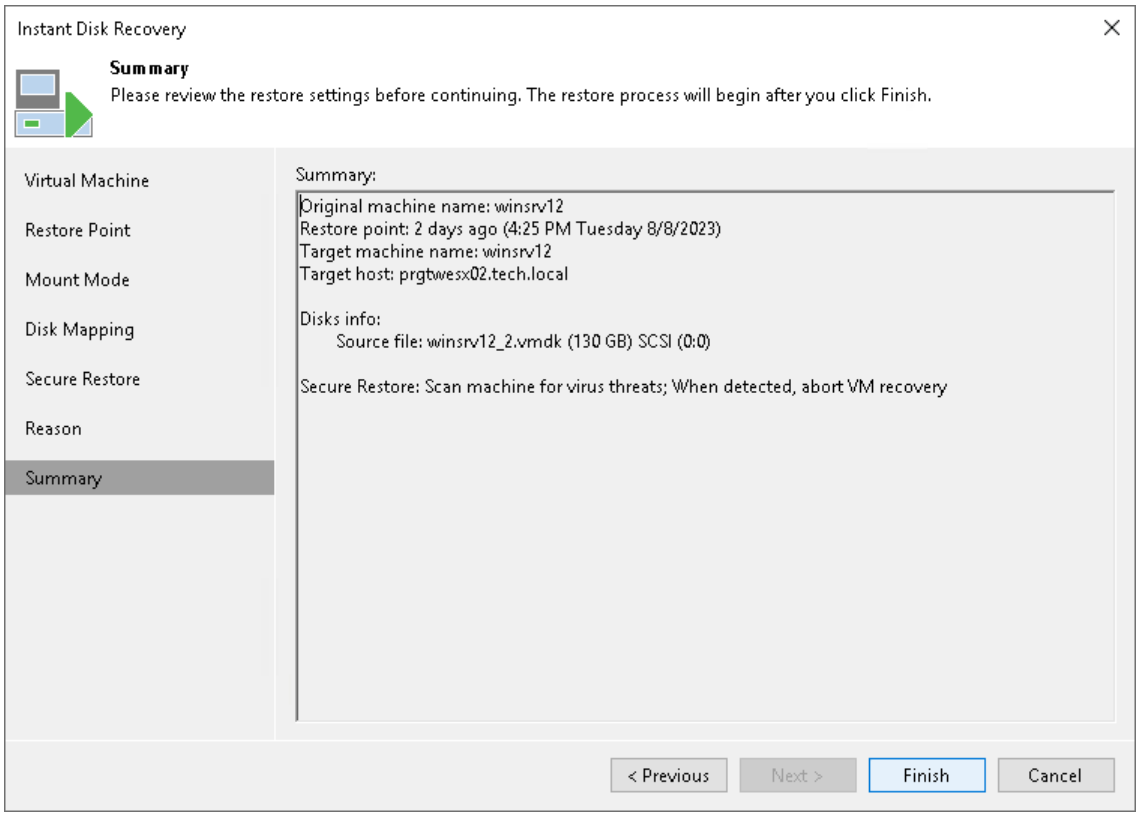


# Finalizing Instant Disk Recovery

After the disks have been successfully recovered, you must finalize the process. For this, test the recovered disks and decide whether to migrate them to production environment or stop publishing them.

# Testing Recovered VM Disks

To test the recovered disks before you migrate them to production, you can launch from the Veeam Backup & Replication console the VMware Remote Console software of the VM to which the disks were recovered.

> **IMPORTANT**
>
> Before you launch VMware Remote Console, make sure that this software is installed on the machine where the Veeam Backup & Replication console runs.

To open a VM console in Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.

3. In the working area, right-click a VM and select **Open VM console**.

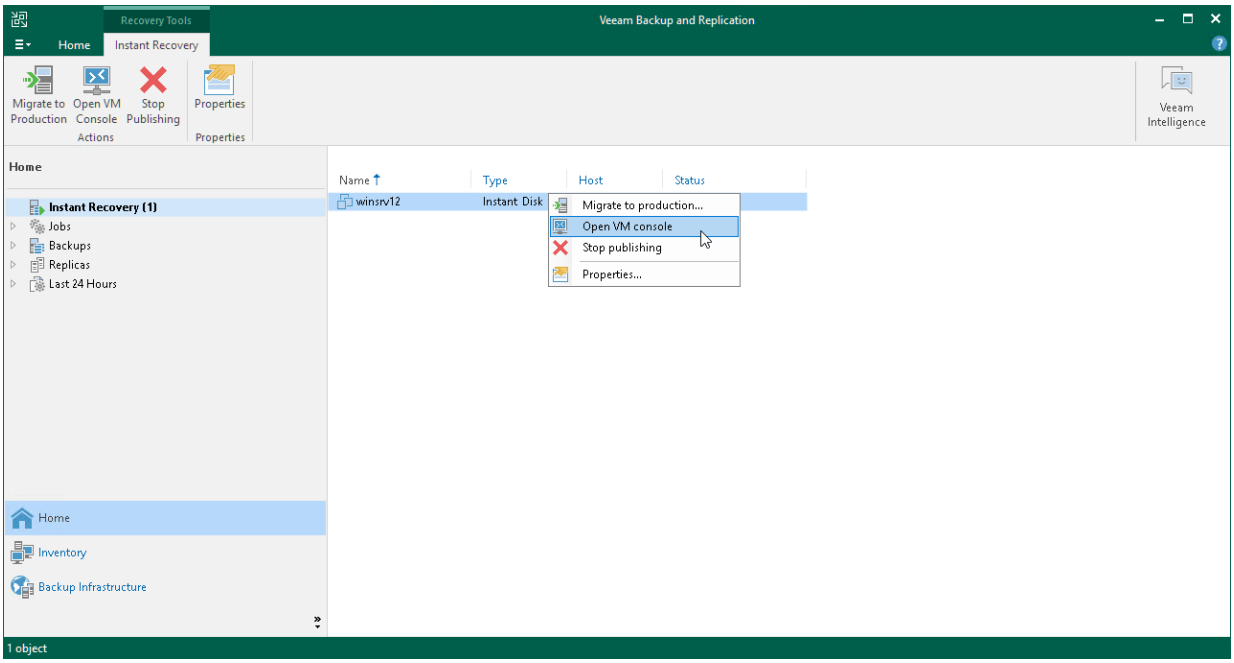


# Migrating Recovered VM Disks

To migrate recovered disks to the production environment:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click the VM to which disks were recovered and select **Migrate to production**. Veeam Backup & Replication will launch the Quick Migration wizard.

   The wizard will be opened on the **Destination** step. At this step, you can change the datastore where the virtual disks will be placed. By default, Veeam Backup & Replication places the disks in the datastore where the VM configuration file is stored.

   Other fields (host, resource pool and VM folder) are populated with data of the VM where you restore the disks. You cannot change this data.

After you finish working with the wizard, Veeam Backup & Replication migrates the disks with all changes made after the disk recovery and before its migration.

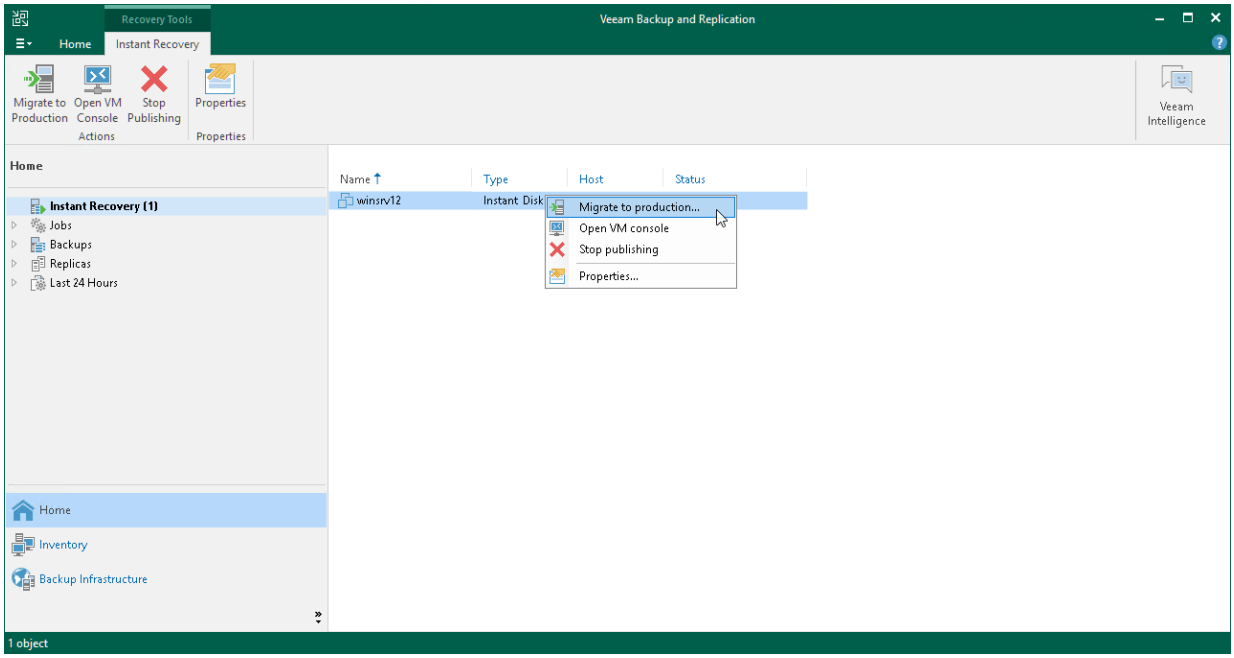


# Stop Publishing Recovered VM Disks

If your tests have failed, you can stop publishing the recovered disks. This will remove the recovered disks from the VM that you selected as the destination for recovery. Note that all changes made to the recovered disks will be lost.

To remove the recovered disks:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click the VM to which the disks were recovered and select **Stop publishing**.

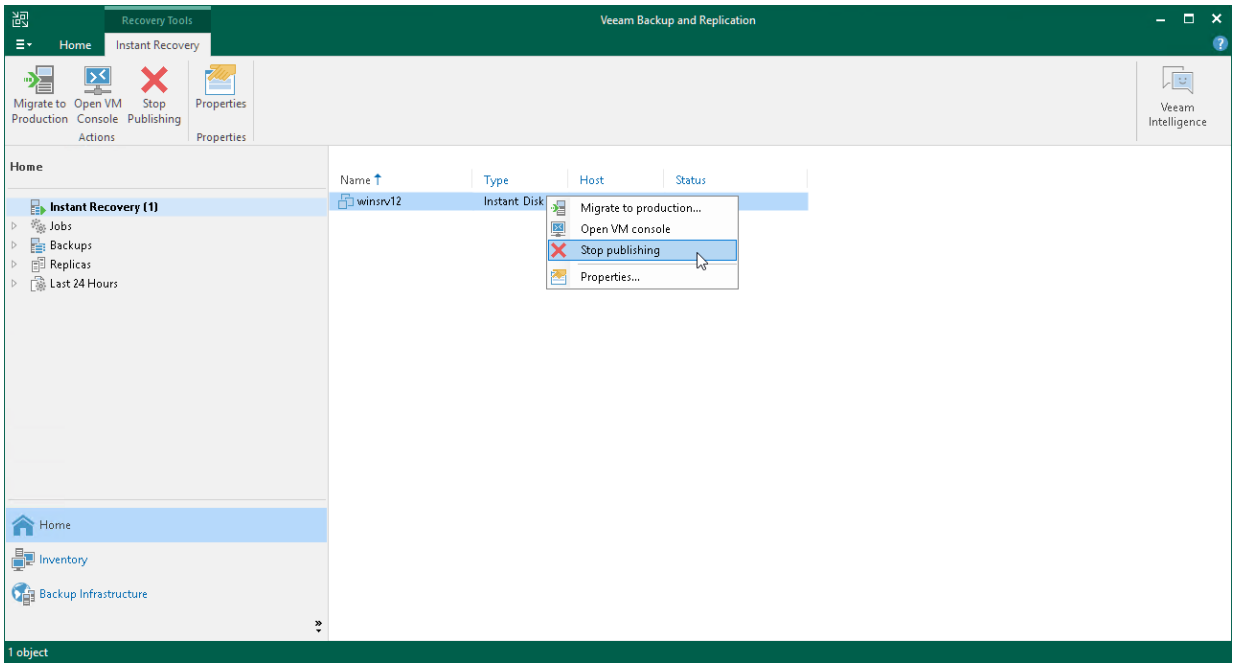

# Instant First Class Disk (FCD) Recovery

With Instant First Class Disk (FCD) Recovery, you can immediately restore VM disks from a backup file and register them as FCDs. If you want to recover disks in their initial format, see Instant Disk Recovery. For the list of all disk recovery methods and their brief descriptions, see Disk Recovery.

The Instant FCD Recovery is a process that must be finalized. You must migrate FCDs to the production environment or stop publishing them. To migrate the disks, Veeam Backup & Replication uses the Quick Migration mechanism. For more information on how to finalize Instant FCD Recovery, see Finalizing Instant FCD Recovery.

## How Instant FCD Recovery Works

Veeam Backup & Replication performs the following steps to recover disks:

1. Initiates a recover session.
2. Connects to a target VMware cluster and mounts vPower NFS datastore to all ESXi hosts added to the target VMware cluster.
3. Mounts recovered disks and publishes them to the target vPower NFS datastore.
4. Registers the published disks as FCDs to the target vPower NFS datastore.

Veeam Backup & Replication performs Instant FCD Recovery using vPower technology.

> **NOTE**
>
> Consider the following:
>
> - If you specify the custom datastore to keep the redo logs, information on the FCD snapshots will be stored in the instant recovery session logs.
> - Veeam Backup & Replication does not remove the mounted vPower NFS datastore from a VMware cluster.

## Performing Instant FCD Recovery

To register VM disks as FDCs, use the **Instant Disk Recovery** wizard.

### Before You Begin

Before you perform Instant FCD Recovery, consider the following:

- The vCenter Server where the target cluster (cluster to which you plan to recover disks) is located must be version 6.7.U3 or later.
- The target cluster must be configured beforehand and added to the backup infrastructure. The hosts located in this cluster must meet the following requirements:
  - ESXi hosts must be powered on.
  - Veeam Backup & Replication must be able to connect to ESXi hosts.
  - API version of ESXi hosts must be 6.7.U3 or later.
  - ESXi must have the VMkernel network interface enabled.

- Veeam Backup & Replication does not support restore to a single ESXi host.
- You can restore a virtual disk from a backup that has at least one successfully created restore point.
- Instant FCD Recovery to VMware Cloud Director is not supported.
- The datastore which you specify to store redirect logs must meet the following requirements:
  - The datastore must be added to the vCenter infrastructure and must be available for Veeam Backup & Replication to set up a connection to it.
  - The datastore must be located in the same vCenter, where a cluster to which Veeam Backup & Replication will register FCDs, is located.
  - The datastore must be mounted to all ESXi hosts of a cluster to which Veeam Backup & Replication will register FCDs.
  - You must have at least one more additional datastore in your vCenter infrastructure, otherwise restore will fail.
  - The datastore must not be the Veeam NFS datastore.
- You can assign storage policies to FCDs only if you redirect redo logs to another datastore.
- The encryption storage policies are not supported. If you select this type of policy, Veeam Backup & Replication will not be able to apply this policy to registered FCDs.
- You must have at least 10 GB of free disk space on the datastore where write cache folder is located. This disk space is required to store virtual disk updates for the restored VM.

  By default, Veeam Backup & Replication writes virtual disk updates to the `IRCache` folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`.

## Step 1. Launch Instant Disk Recovery Wizard

To launch the **Instant Disk Recovery** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Disk restore** > **Instant disk recovery**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand a backup, select a VM whose disks you want to restore and click **Instant Disk Recovery** on the ribbon. Alternatively, you can right-click a VM whose disks you want to restore and select **Instant disk recovery**.

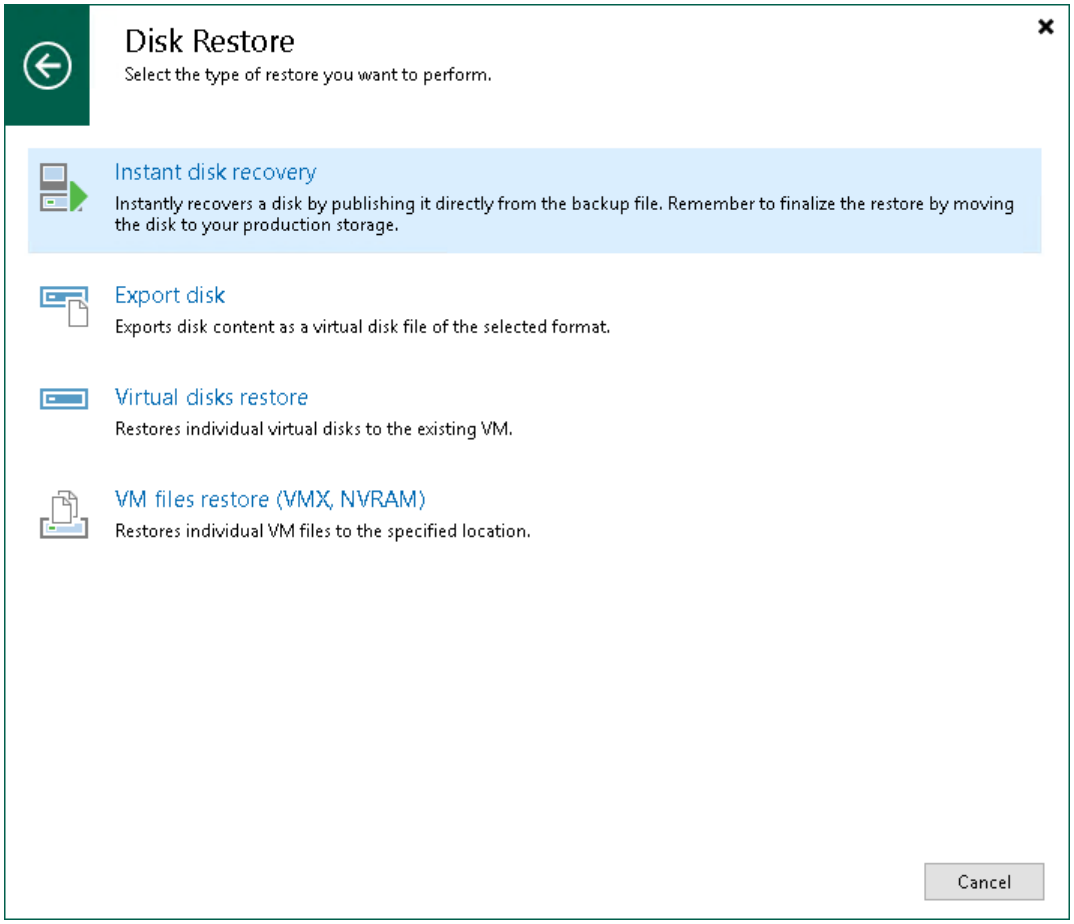

## Step 2. Select Source VM

At the **Virtual Machine** step of the wizard, select a VM whose disks you want to register as FCDs.

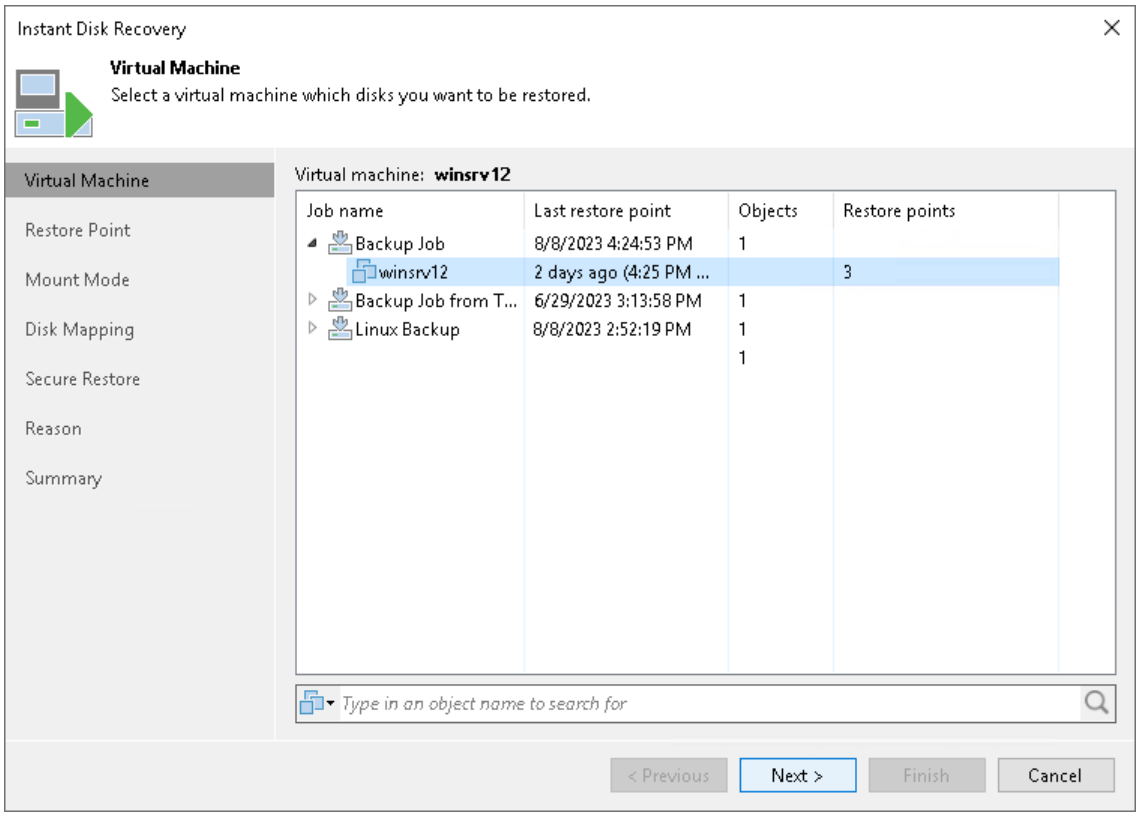

# Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point for the VM whose virtual disks you want to register as FCDs.

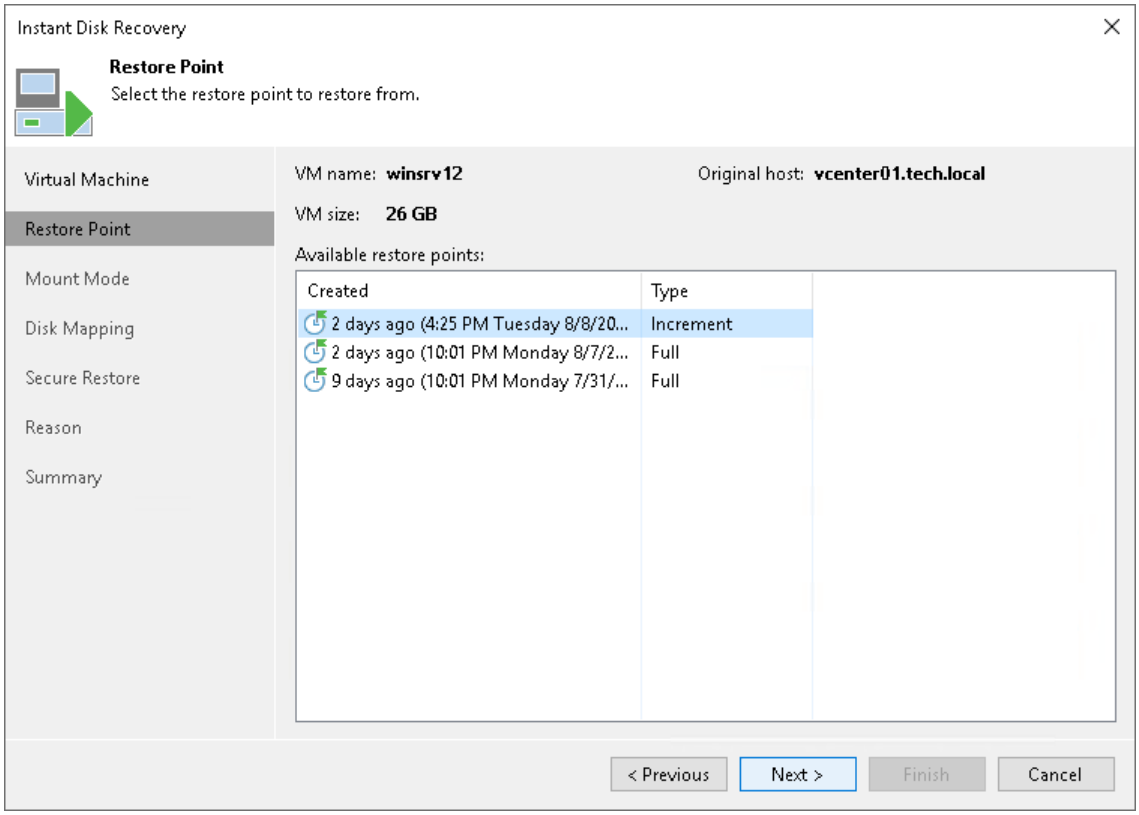

# Step 4. Select Mount Mode

At the **Mount Mode** step of the wizard, select the **First class disk (FCD)** option to register virtual disks on a cluster as FCDs.

If you want to register virtual disks on a VM added to an ESXi host, select the **VM disk** option. In this case, steps of the wizard differ and Veeam Backup & Replication performs the instant disks recovery as described in section Performing Instant Disk Recovery.

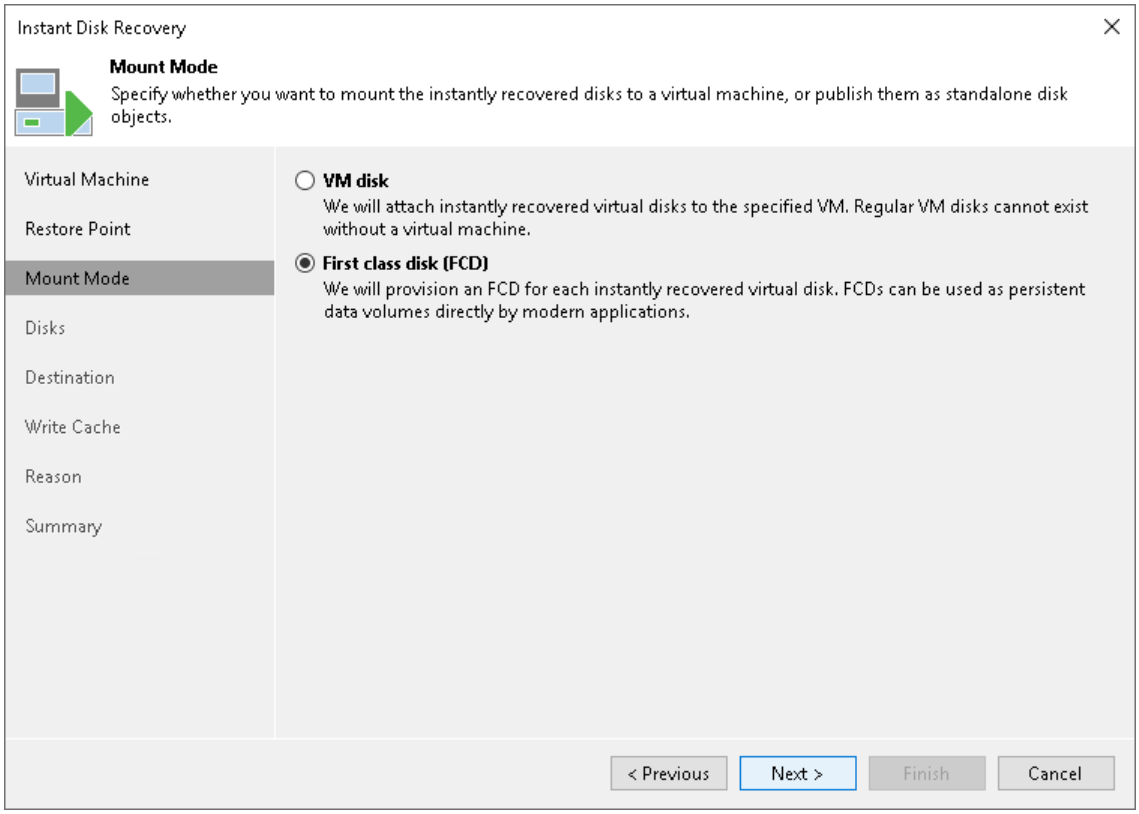

## Step 5. Select Virtual Disks to Restore

At the **Disks** step of the wizard, select virtual disks that you register as FCDs and define names for the recovered disks and FCDs.

1. Select check boxes next to virtual disks that you want to register as FCDs.
2. By default, Veeam Backup & Replication recovers disks and registers FCDs under the original disk names. To specify new names, select a disk and click **Change**. In the **Target Disk Name** window, do the following:
   a. In the **Disk name** field, specify a name under which the disk will be recovered.
   b. In the **Display name** field, specify a name under which FCD will be registered on the cluster.

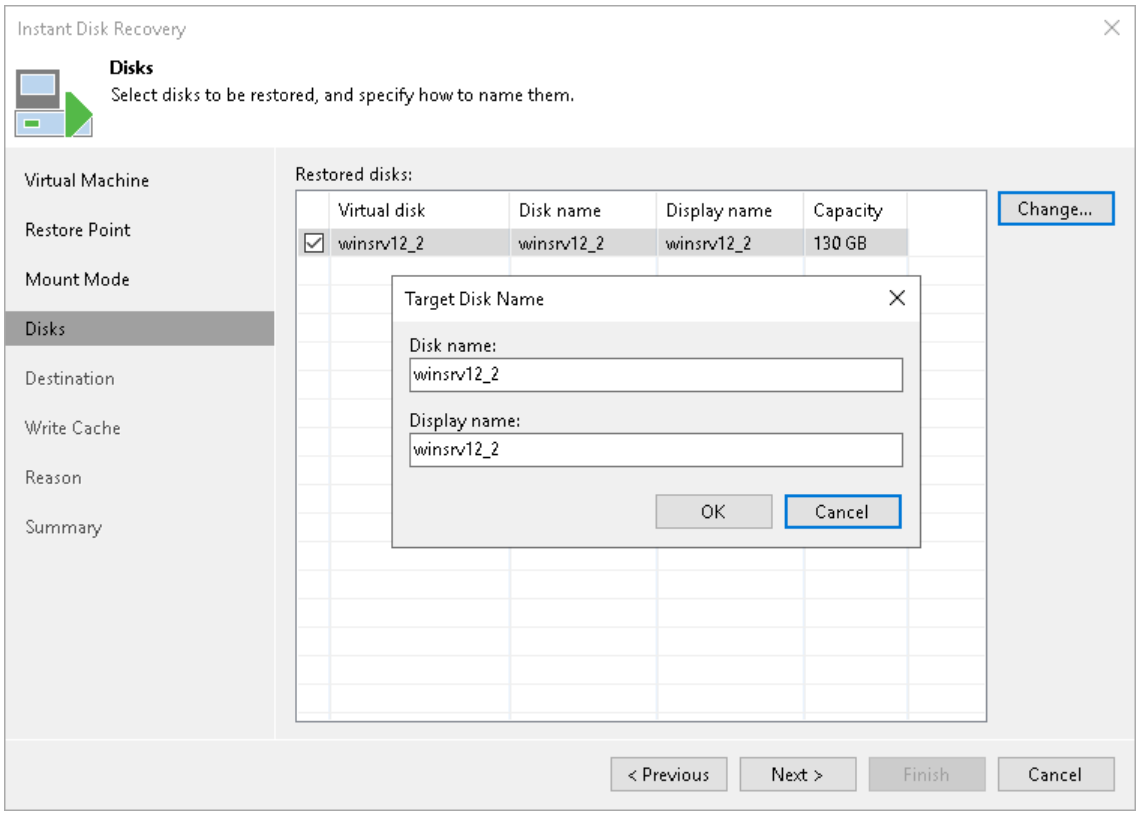

## Step 6. Specify Target Cluster

At the **Destination** step of the wizard, specify a cluster where Veeam Backup & Replication will register FCDs.

> **IMPORTANT**
>
> To register FCDs on a cluster, Veeam Backup & Replication mounts to a cluster the vPower NFS datastore that contains virtual disks. Veeam Backup & Replication checks the cluster hosts to verify that they are powered on and Veeam Backup & Replication is able to connect to these hosts. If one of these hosts is not available, Veeam Backup & Replication will not include them to the FCD restore scenario.
> Veeam Backup & Replication will not start the Instant FCD Recovery process if there are no powered on hosts or Veeam Backup & Replication is not able to connect to them.

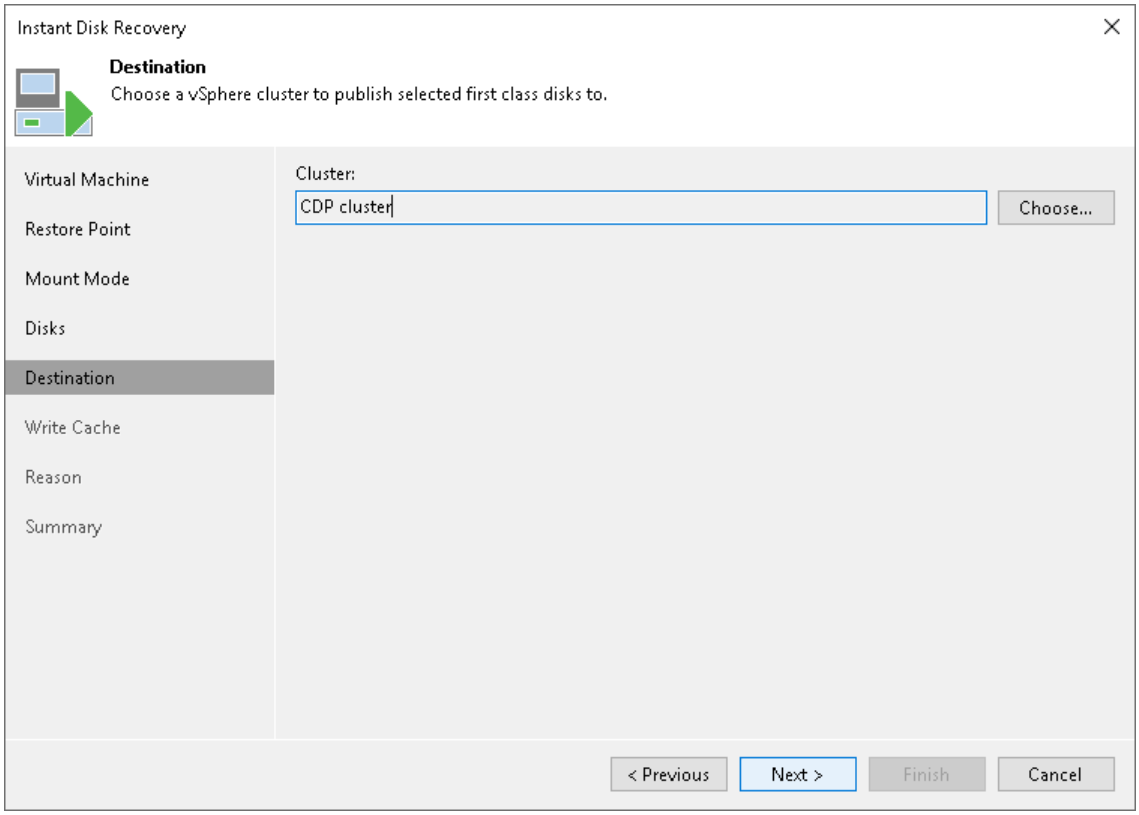

# Step 7. Select Destination for FCD Updates

At the **Write Cache** step of the wizard, specify where to store the redo logs. Veeam Backup & Replication uses redo logs to write all changes of the virtual disks that take place while performing Instant FCD Recovery.

By default, redo logs are stored on the vPower NFS server. You can store redo logs on any datastore in the virtual environment. As soon as a recovery verification job completes, Veeam Backup & Replication deletes redo logs.

To redirect redo logs:

1. Select the **Redirect write cache** check box.
2. Click **Choose** and select a datastore from the list.

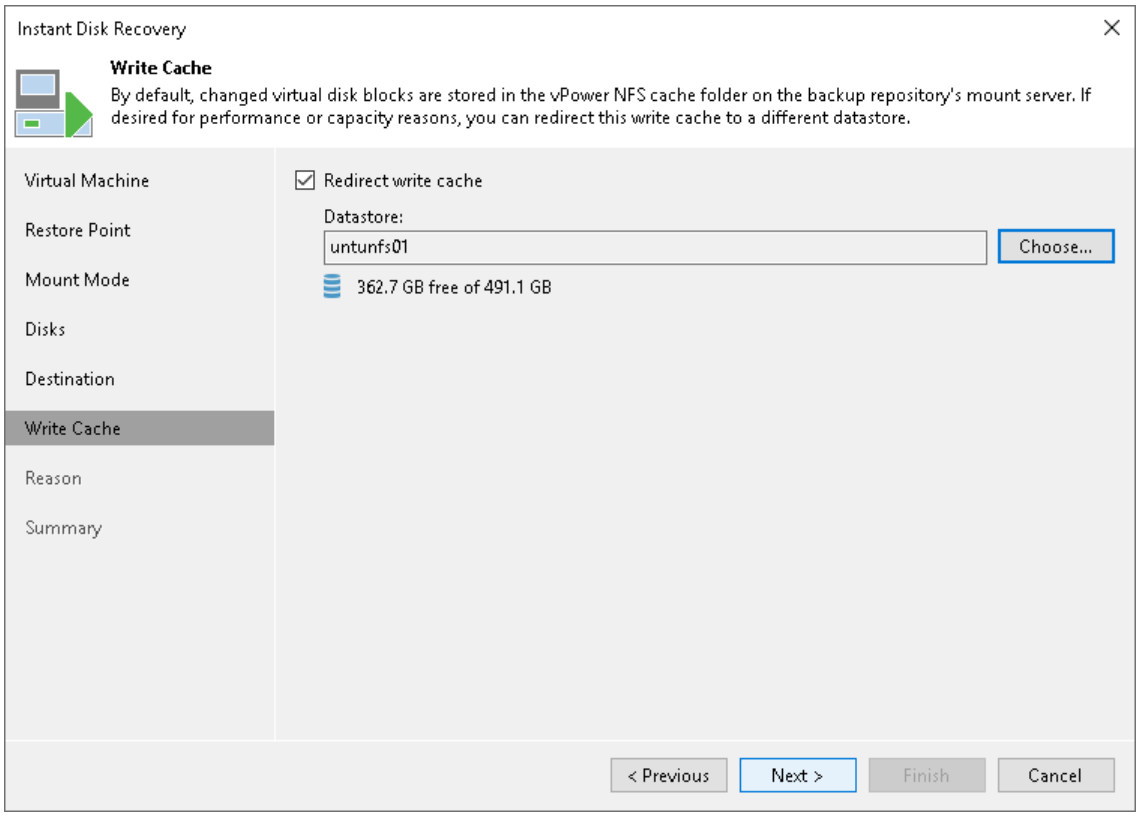

# Step 8. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring VM disks. This information will be saved in the session history so that you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

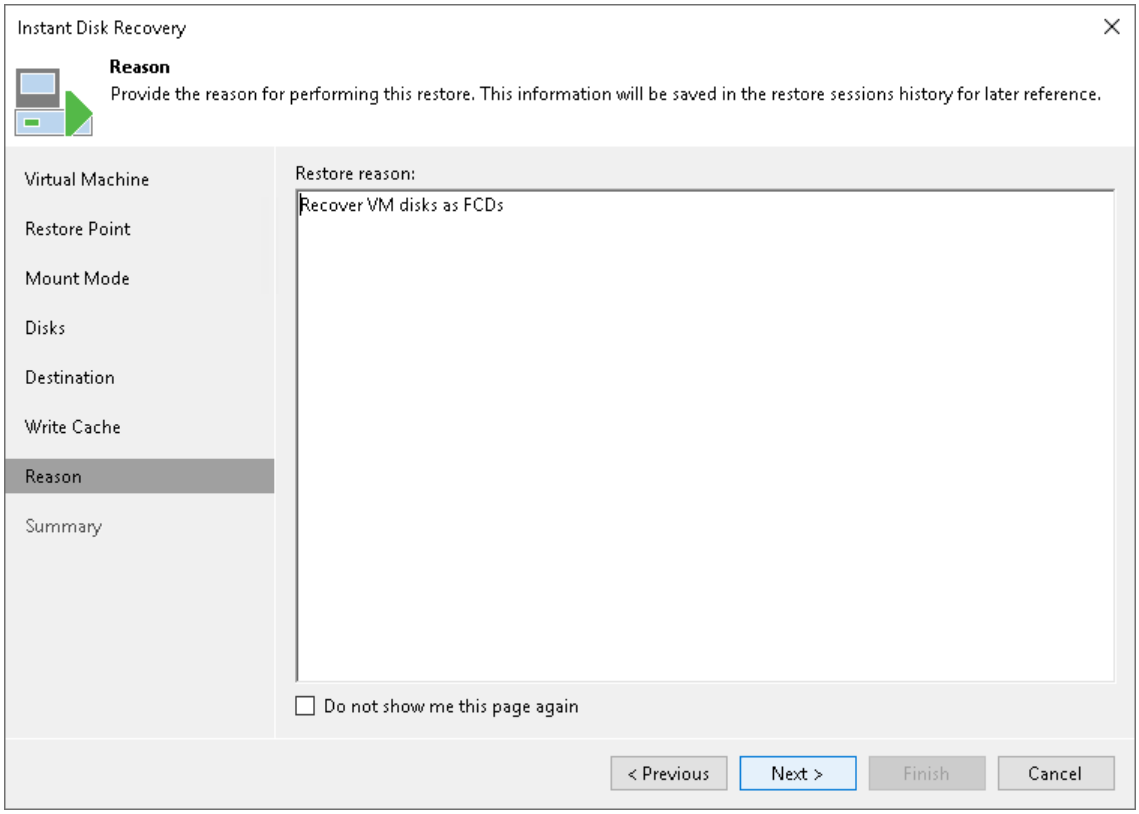

## Step 9. Verify Instant FCD Recovery Settings

At the **Summary** step of the wizard, check settings of Instant FCD Recovery and click **Finish**.

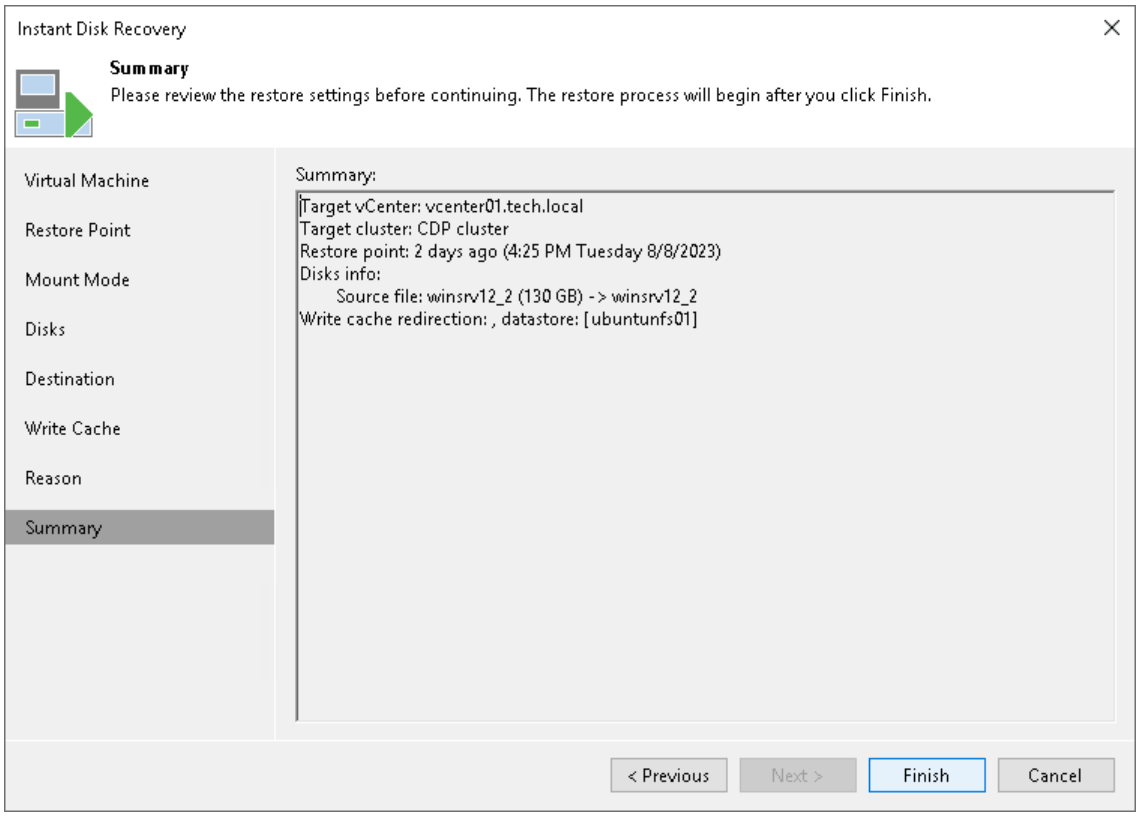


# Finalizing Instant FCD Recovery

After FCDs have been successfully registered, you must finalize the process and decide whether to migrate the disks to the production environment or stop publishing them.

# Migrating FCDs

To migrate FCDs to the production environment:

1. Check requirements for FCD migration.
2. Open the **Home** view.
3. In the inventory pane, select the **Instant Recovery** node.
4. In the working area, right-click the VM whose disks you recovered and select **Migrate to production**. Veeam Backup & Replication will launch the FCD Quick Migration wizard.

After you finish working with the wizard, Veeam Backup & Replication migrates the disks with all changes made after the disk recovery and before its migration.

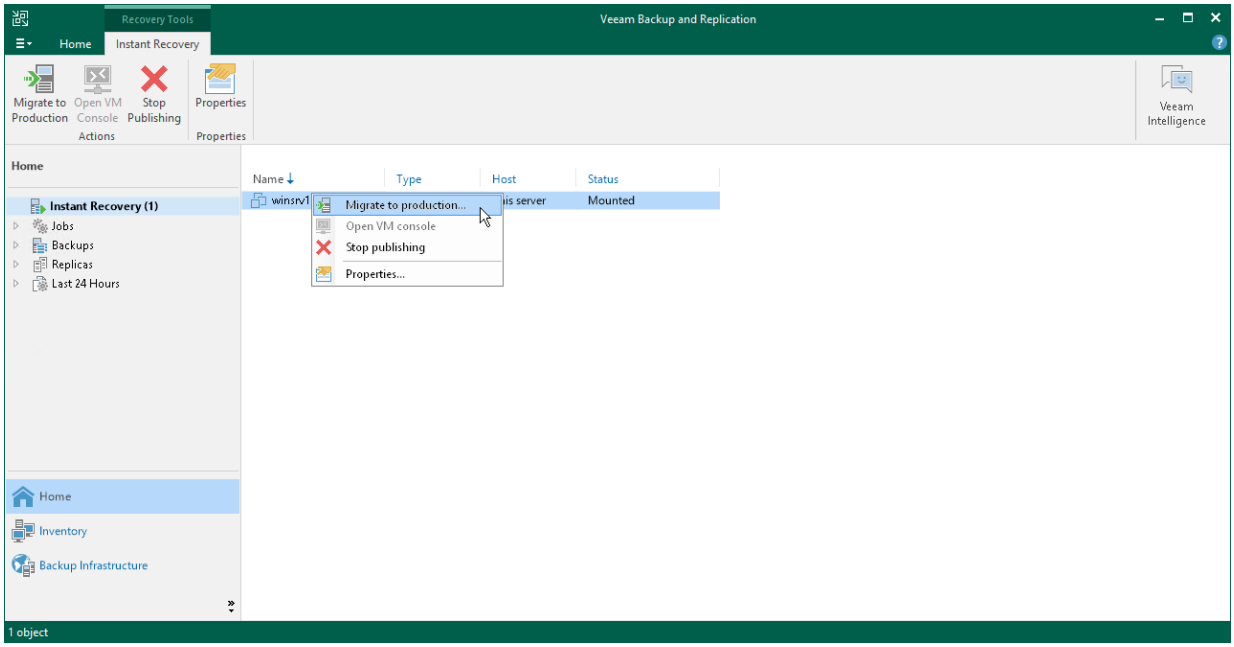


# Stop Publishing FCDs

You can stop publishing FCDs. This will remove the registered FCDs from the cluster that you selected as the destination for recovery. Note that all changes made to the FCDs will be lost.

To stop publishing FCDs:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click a VM whose disks you recovered and select **Stop publishing**.

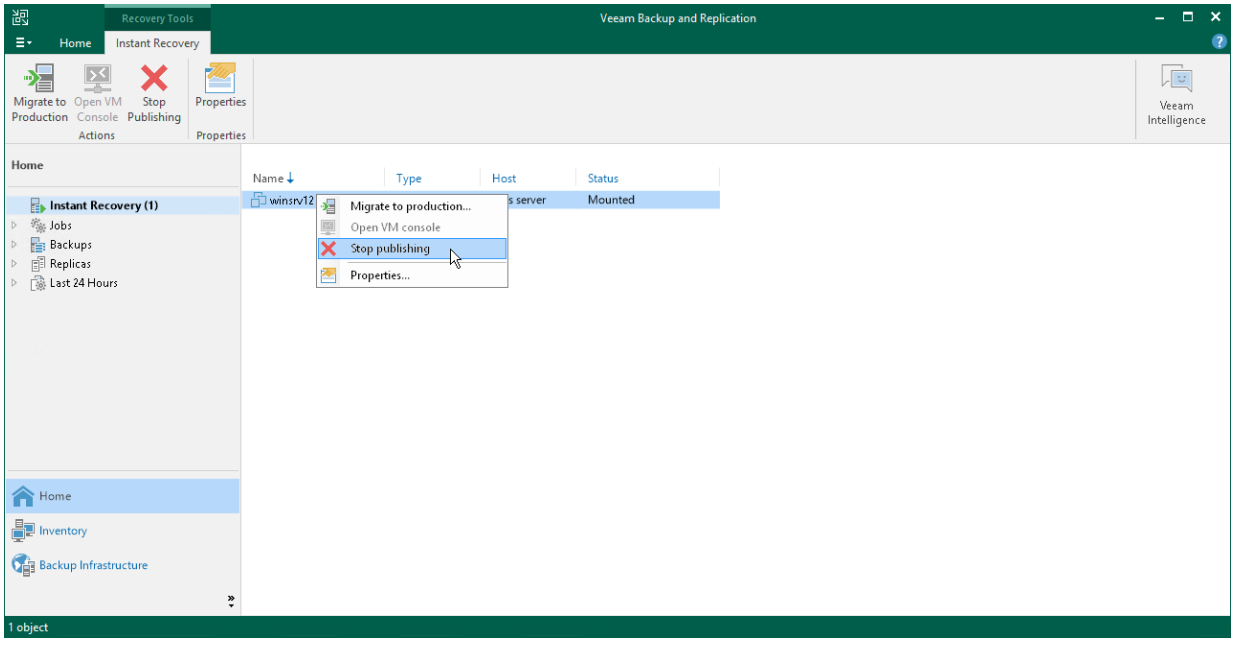

# Virtual Disk Restore

Veeam Backup & Replication allows you to recover individual virtual disks of a VM from a backup. Recovered virtual disks can be attached to the original VM (for example, if you want to replace a corrupted disk) or mapped to any other VM in the virtual infrastructure. This recovery option can be helpful if a VM disk becomes corrupted.

You can restore VM virtual disk to the latest state or any valid restore point. You can preserve the format of a recovered virtual disk or convert it to the thin or thick provisioned disk format.

> **NOTE**
>
> If a VM has several VM disks, Veeam Backup & Replication restores VM disks in parallel.

## Restoring Virtual Disks

To restore virtual disks, use the **Virtual Disk Restore** wizard.

### Before You Begin

Before you restore virtual disks, check the following prerequisites:

- You can restore virtual disks from a backup that has at least one successfully created restore point.
- During the virtual disk restore, Veeam Backup & Replication turns off the target VM to reconfigure its settings and attach restored disks. It is recommended that you stop all activities on the target VM for the restore period.
- If you want to scan virtual disk data for viruses, check the secure restore requirements and limitations.

> **NOTE**
>
> If you back up a VM with vRDM disks, Veeam Backup & Replication converts them into VMDK files. Thus, when you restore a vRDM disk, Veeam Backup & Replication will restore it as a VMDK file. If you want to preserve the vRDM format for restored disks, use Quick Rollback. For more information, see Quick Rollback.

## Step 1. Launch Virtual Disk Restore Wizard

To launch the **Virtual Disk Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Disk restore** > **Virtual disks restore**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, click the VM whose files you want to restore and click **Virtual Disks** on the ribbon. Alternatively, you can right-click the VM whose files you want to restore and select **Restore virtual disks**.
- Double-click the VBK or VBM file (for example, in Microsoft Windows Explorer). In the displayed window, select the VM and click **Restore** > **Virtual disks**.

  You can use this option if you perform restore on the backup server. You cannot use this option if you perform restore remotely over the Veeam Backup & Replication console.

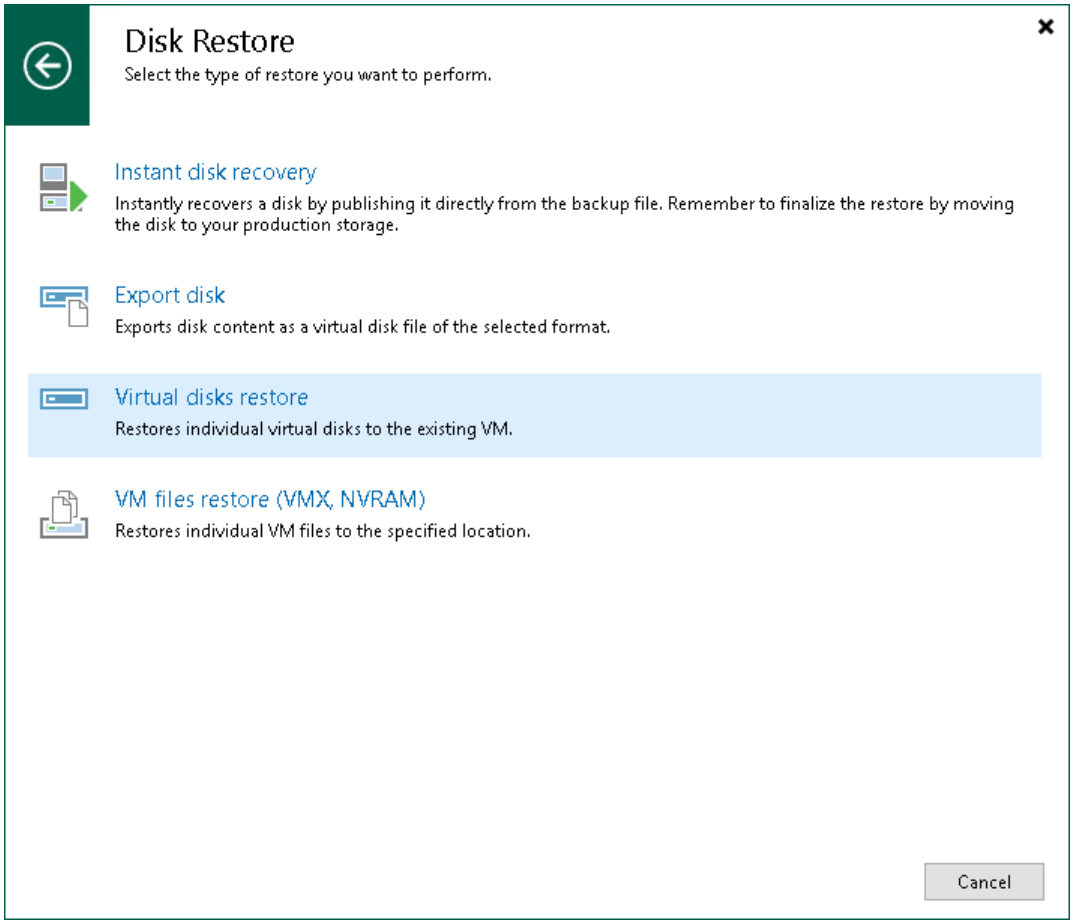

## Step 2. Select VM

At the **Virtual Machine** step of the wizard, expand a backup and select a VM whose disks you want to restore.

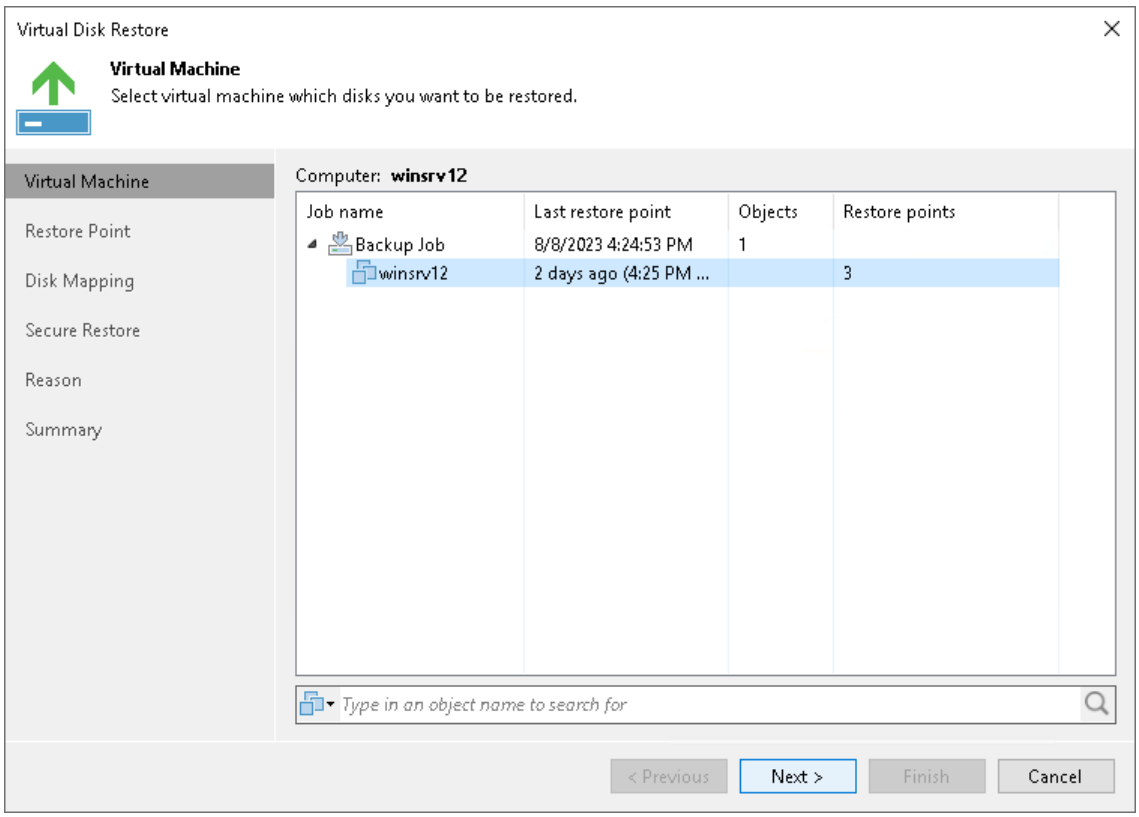

# Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point from which you want to restore VM disks.

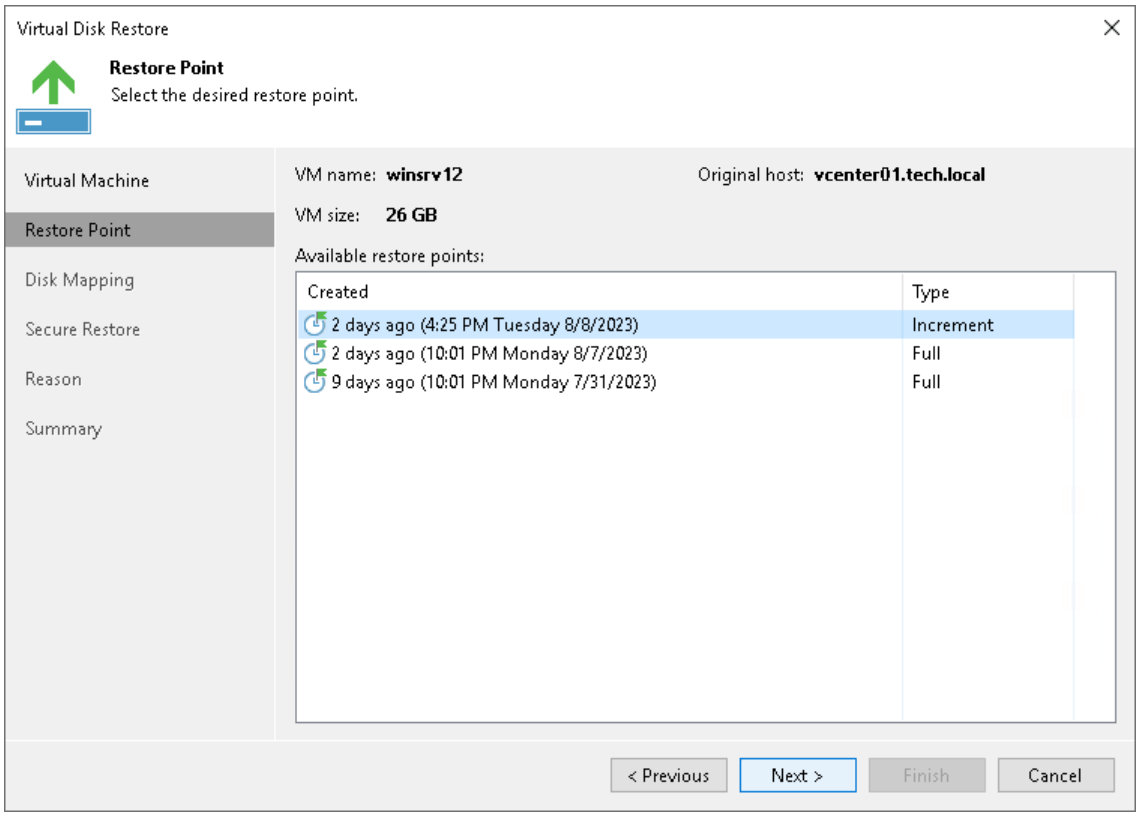

## Step 4. Select Virtual Hard Disks to Restore

At the **Disk Mapping** step, select virtual hard disks to restore, choose a VM to which the disks must be attached and define additional restore settings.

1. By default, Veeam Backup & Replication maps restored disks to the original VM. If the original VM was relocated or if you want to attach disks to another VM, you need to select the target VM manually. Click **Choose** and select the necessary VM from the virtual environment.

   > **IMPORTANT**
   >
   > You cannot attach restored disks to a VM that has one or more snapshots.

2. In the **Disk mapping** list, select virtual disks that you want to restore.
3. To configure virtual disk properties, select a disk in the **Disk mapping** list and click **Change**. In the **Virtual Disk Properties** window, do the following:

   a. Click **Choose** to pick a datastore where the restored hard disk must be placed.

      If you use storage policies in the virtual environment, Veeam Backup & Replication displays information about storage policies in the **Select Datastore** window. You can select a datastore associated with the necessary storage policy.

   b. From the **Virtual device node** drop-down list, select a virtual device node. If you want to replace an existing virtual disk, select an occupied virtual node. If you want to attach the restored disk to the VM as a new disk, select a node that is not occupied yet.

   c. Click **OK**.

4. Veeam Backup & Replication preserves the format of the restored virtual hard disks. To change disk format, select the required option from the **Restore disk type** drop-down list. For more information about disk types, see VMware Docs.

   > **NOTE**
   >
   > Disk format change is supported only for VMs with Virtual Hardware version 7 or later.

5. [For disk restore to the original location and with original format] Select the **Quick rollback** check box if you want to use incremental restore for the VM disk. For more information on quick rollback, its requirements and limitations, see Quick Rollback.

   It is recommended that you enable this option if you restore a VM disk after a problem that occurred at the level of the VM guest OS: for example, there has been an application error or a user has accidentally deleted a file on the VM guest OS. Do not enable this option if the problem has occurred at the VM hardware level, storage level or due to a power loss.

6. Click the **Pick proxy to use** link to select backup proxies over which VM data must be transported to the target datastore. By default, Veeam Backup & Replication assigns proxies automatically.

   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that have access to the source datastore and will automatically assign optimal proxy resources for processing VM data.

     During the restore process, VM hard disks are processed simultaneously. Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use for writing data to target, current workload on these backup proxies, and selects the most appropriate resources for VM hard disk processing.

- If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that must be used for restore. It is recommended that you select at least two backup proxies to ensure that VM hard disks are recovered if one of backup proxies fails or loses its connectivity to the target datastore during restore.

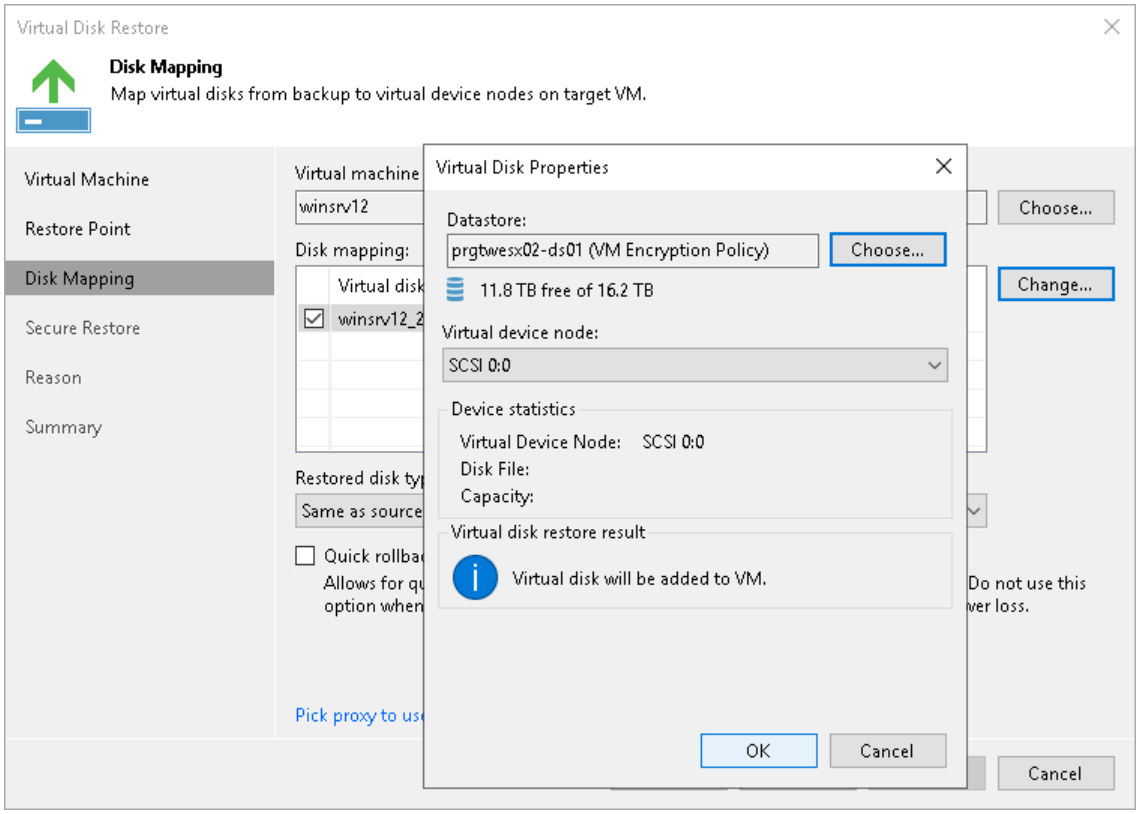

# Step 5. Specify Secure Restore Settings

This step is available if you restore disks of Microsoft Windows VMs.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Select which action Veeam Backup & Replication will take if scan detects a virus threat:
   - **Proceed with recovery but do not attach infected disks to the target VM**. Select this action if you want to continue the virtual disk restore. In this case, the restored disk will not be attached to the target VM.
   - **Abort disk recovery**. Select this action if you want to cancel the restore session.

3. Select the **Continue scanning all remaining files after the first occurrence** check box if you want to continue the virtual disk scan after the first virus threat is detected. For information on how to view results of the malware scan, see Viewing Malware Scan Results.

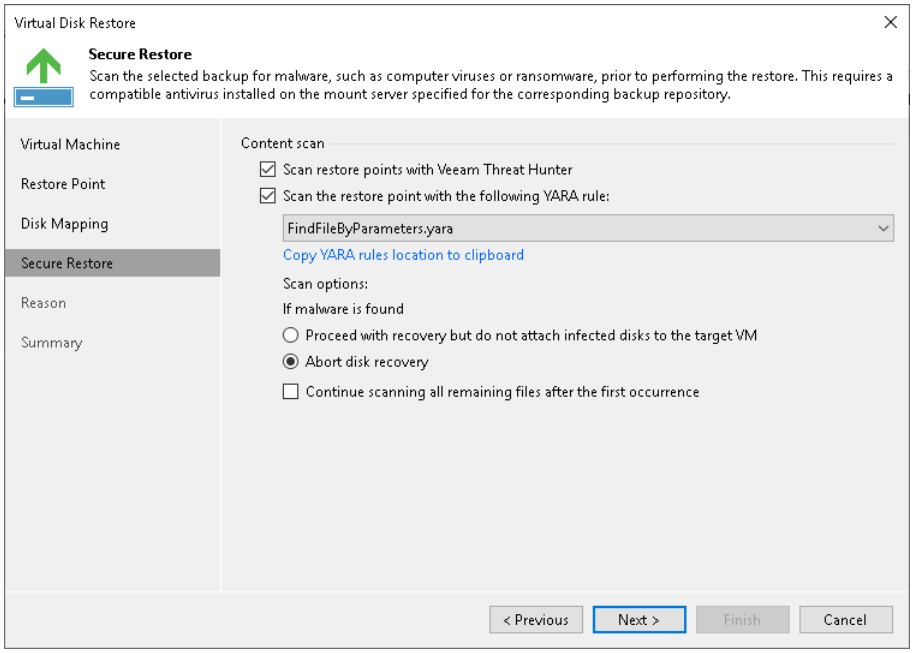

# Step 6. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring VM disks. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

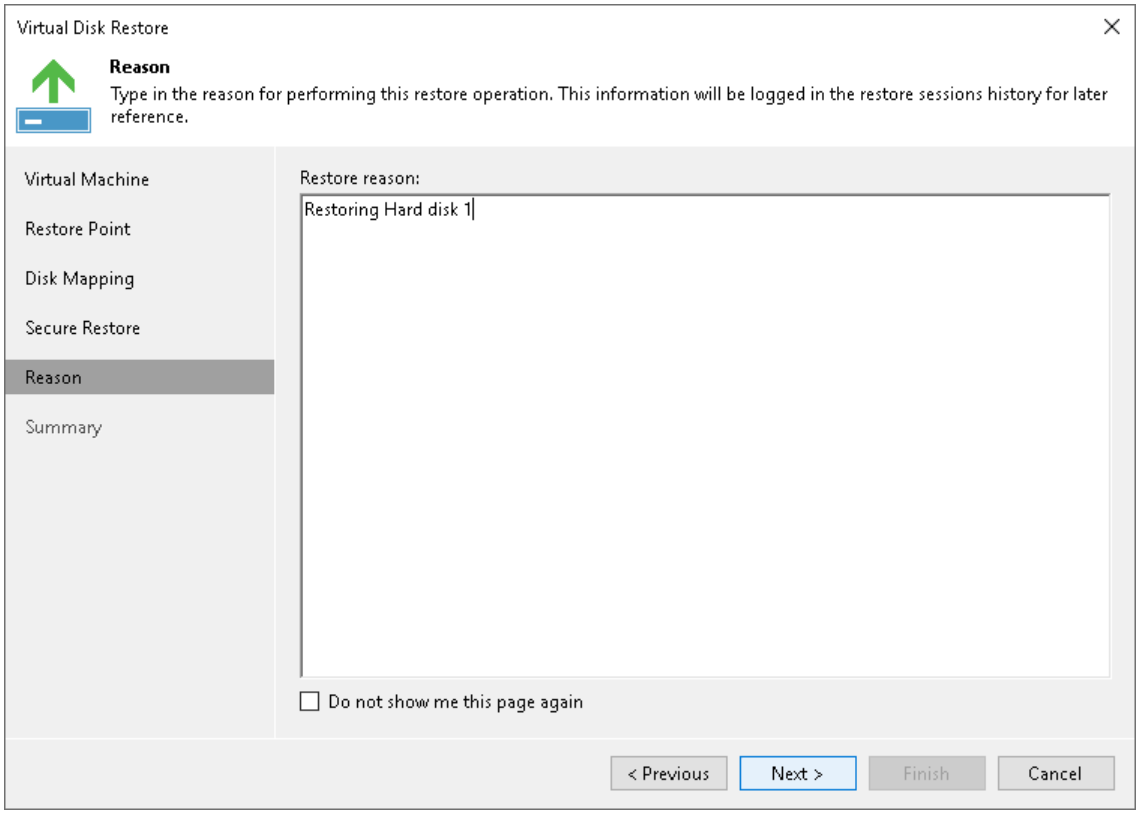

## Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured settings and click **Finish**.

If you want to start the VM to which the disk will be attached right after restore, select the **Power on target VM after restoring** check box.

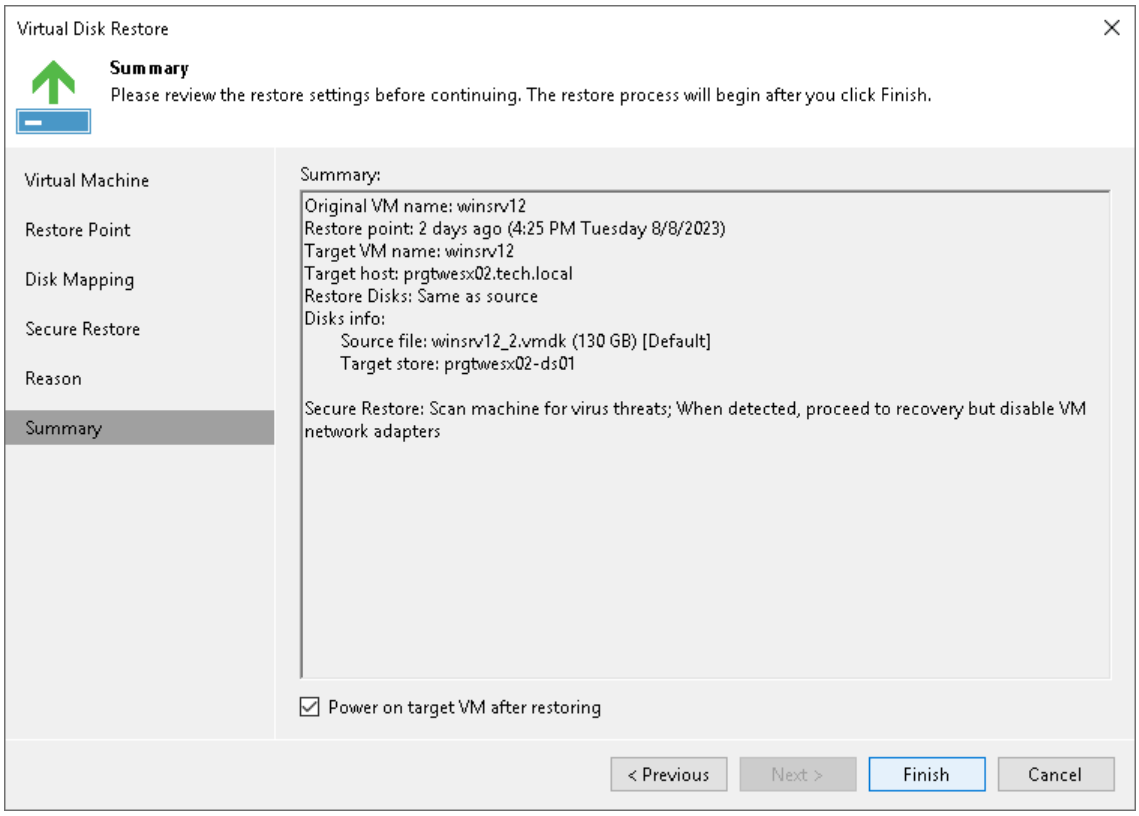

# Disk Export

Veeam Backup & Replication allows you to export disks, that is, restore disks from backups of physical or virtual machines and convert them to the VMDK, VHD or VHDX formats.

During disk export, Veeam Backup & Replication creates the following files that can be used by VMware vSphere and Microsoft Hyper-V VMs:

- When you export a disk in the VMDK format, Veeam Backup & Replication creates a pair of files that make up the VM virtual disk: a descriptor file and file with the virtual disk content.
- When you export a disk in the VHD/VHDX format, Veeam Backup & Replication creates a file of the VHD or VHDX format.

You can save the exported disks to any server added to the backup infrastructure or place disks on a datastore connected to a host.

Veeam Backup & Replication supports batch disk export. For example, if you choose to export 2 disks, Veeam Backup & Replication will convert them to 2 virtual disks and store these disks in the specified location.

## Supported Backup Types

You can restore disks from the following backups:

- Backups of VMware vSphere virtual machines created by Veeam Backup & Replication
- Backups of VMware Cloud Director virtual machines created by Veeam Backup & Replication
- Backups of Microsoft Hyper-V virtual machines created by Veeam Backup & Replication
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for Mac, Veeam Agent for Oracle Solaris or Veeam Agent for IBM AIX
- Backups of EC2 instances created by Veeam Backup for AWS
- Backups of Azure VMs created by Veeam Backup for Microsoft Azure
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV
- Backups of Google VM instances created by Veeam Backup for Google Cloud
- Backups of oVirt VMs created by Veeam Backup for OLVM and RHV
- Backups exported by Kasten policies
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE

## Exporting Disks

You can export disks of different workloads from backups and convert disks to the VMDK, VHD or VHDX format. For the list of backups that you can use for export, see Supported Backup Types.

To export disks, use the **Export Disk** wizard.

## Step 1. Launch Export Disk Wizard

To launch the **Export Disk** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select one of the following:
  - **VMware vSphere** > **Restore from backup** > **Disk Restore** > **Export disk** — to export disks of VMware vSphere VMs from a VM backup created by Veeam Backup & Replication.
  - **VMware Cloud Director** > **Restore from backup** > **VM restore** > **Disk Restore** > **Export disk** — to export disks of VMware Cloud Director VMs from a VM backup created by Veeam Backup & Replication.
  - **Microsoft Hyper-V** > **Restore from backup** > **Entire VM restore** > **Export disk** — to export disks of Hyper-V VMs from a VM backup created by Veeam Backup & Replication.
  - **Agent** > **Disk Restore** > **Export disk** — to export disks of physical machines and virtual machines from backups created by Veeam Agent for Microsoft Windows, Veeam Agent for Linux or Veeam Agent for Mac.
  - **AWS** > **Amazon EC2** > **Entire machine restore** > **Export disk** — to export disks of EC2 instances from backups created by Veeam Backup for AWS.
  - **Azure IaaS backup** > **Entire machine restore** > **Export disk** — to export disks of Azure VMs from backups created by Veeam Backup for Microsoft Azure.
  - **GCE backup** > **Entire machine restore** > **Export disk** — to export disks of VM instances from backups created by Veeam Backup for Google Cloud.
  - **Nutanix backup** > **Entire machine restore** > **Export disk** — to export disks of VMs from backups created by Veeam Backup for Nutanix AHV.
  - **oVirt KVM** > **Entire machine restore** > **Export disk** — to export disks of VMs from backups created by Veeam Backup for OLVM and RHV.
  - **Kasten backup** > **Export disk** — to export disks of VMs whose backups were exported by Kasten policies.
  - **Proxmox VE** > **Export disk** — to export disks of Proxmox VE VMs from backups created by Veeam Backup for Proxmox VE.

- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary backup, select a workload whose disks you want to export and click **Export Disks** on the ribbon. Alternatively, you can right-click the workload and select **Export content as virtual disks**.

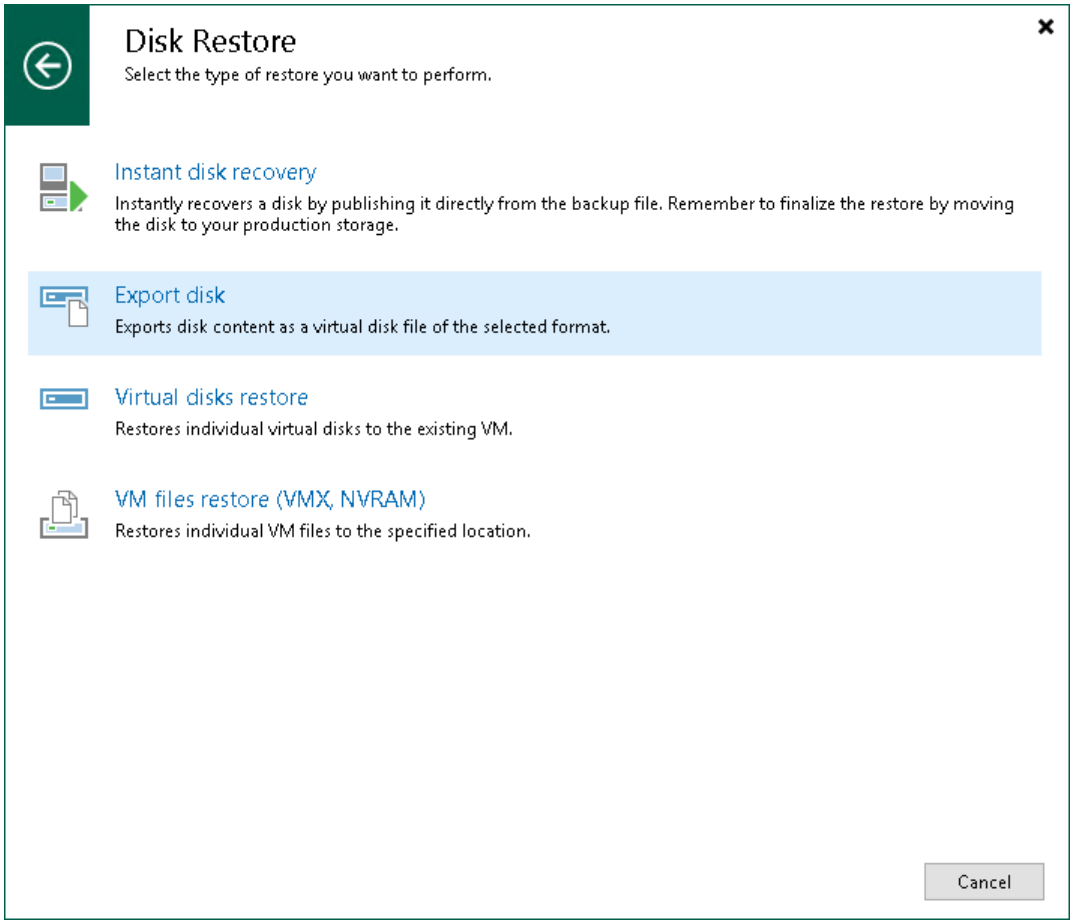

## Step 2. Select Backup

At the **Backup** step of the wizard, expand a backup and select the workload whose disks you want to export.

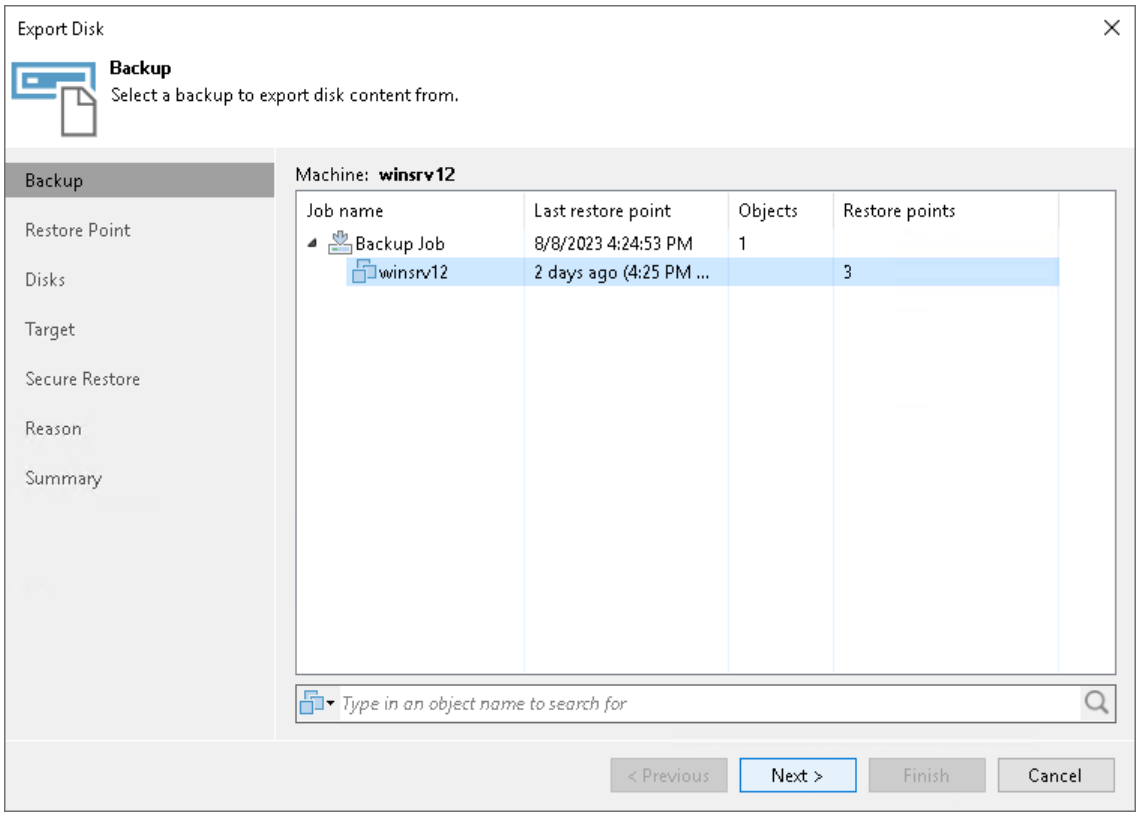

# Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select the restore point from which you want to restore disks. By default, Veeam Backup & Replication restores disks to the latest state. However, you can restore disks to an earlier state.

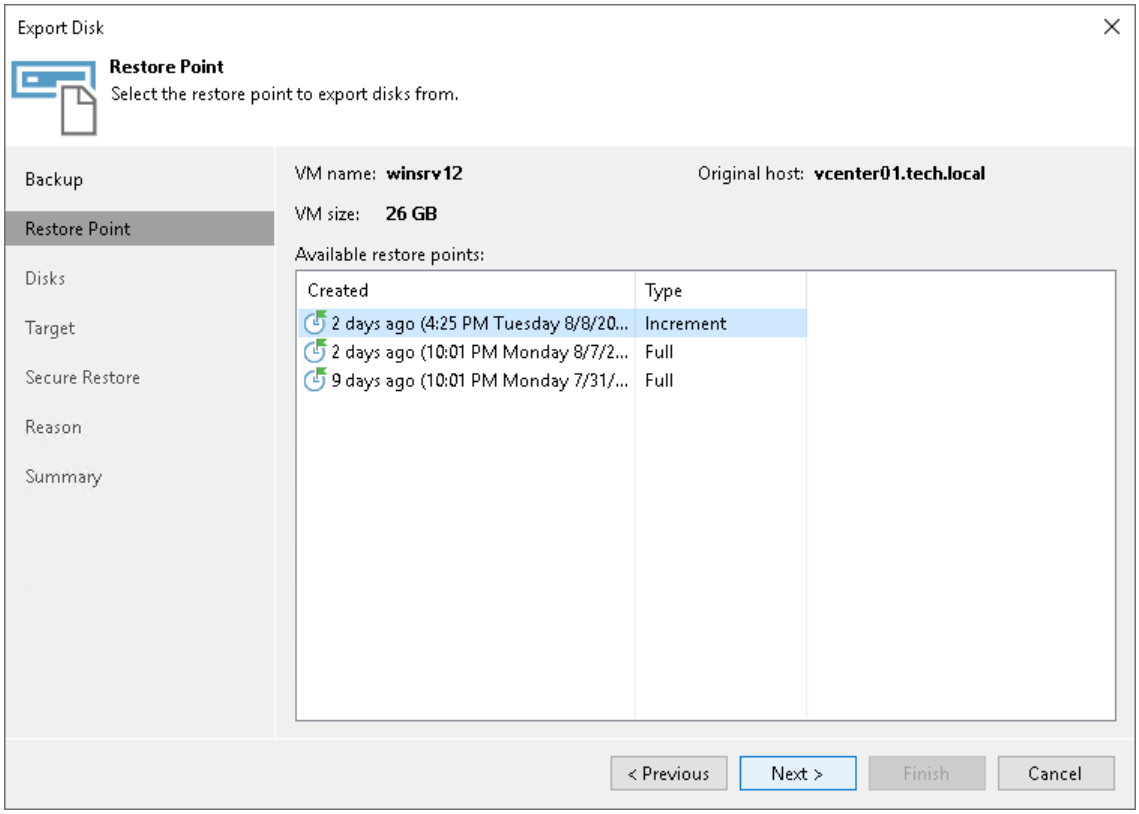

## Step 4. Select Disks

At the **Disks** step of the wizard, select check boxes next to the disks that you want to export.

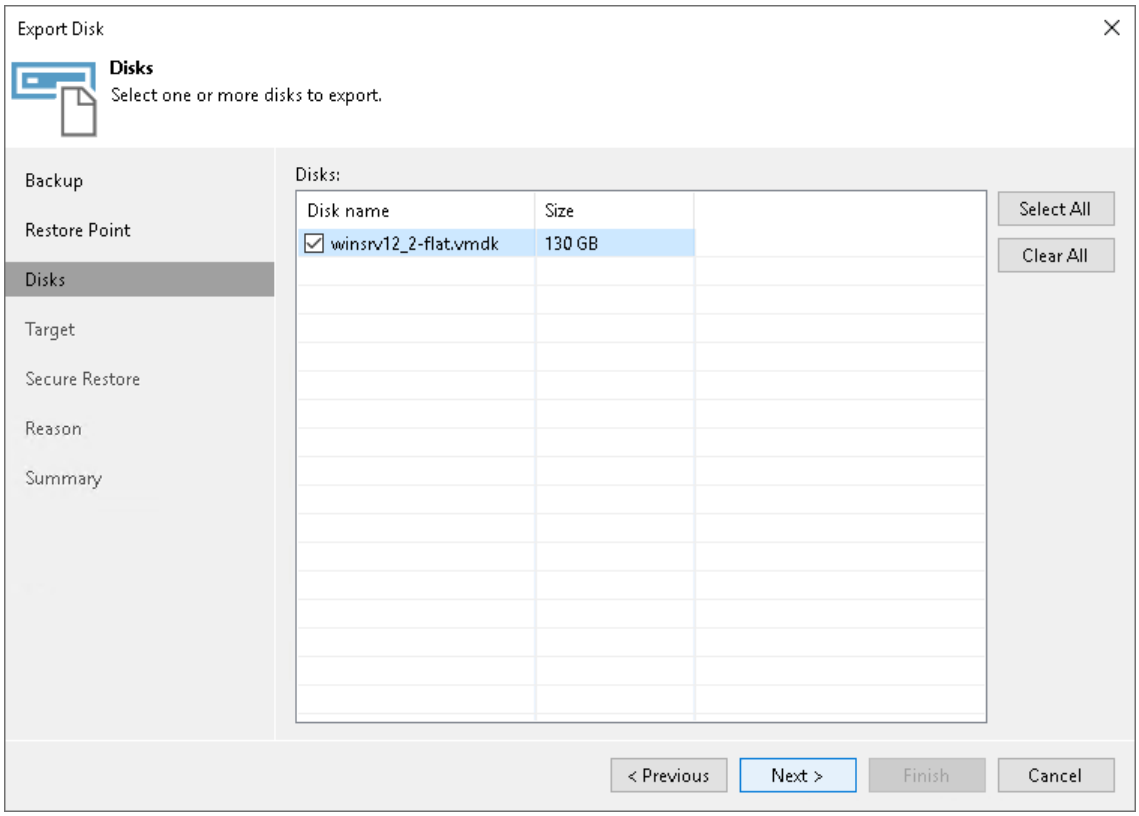

## Step 5. Select Destination and Disk Format

At the **Target** step of the wizard, select the destination for disk export and format in which you want to save the resulting virtual disks:

1. From the **Server** list, select a server on which the resulting virtual disks must be saved. If you plan to save the disks to a datastore, select a host to which this datastore is connected.

   > **NOTE**
   >
   > Consider the following:
   >
   > - If you select to export the resulting virtual disk to an ESXi datastore, you can save the virtual disk in the VMDK format only. Other options are disabled.
   > - If you export disks as VMDK to a vSAN datastore, disk are exported according to the storage policy set on the datastore.
   > - If you export disks as VMDK to a place other than an ESXi datastore, disks are exported as thick provision lazy zeroed.

2. In the **Path** to folder field, specify a datastore or folder on the server where the virtual disks must be placed.
3. Select the export format for the disks:
   - **VMDK** – select this option if you want to save the resulting virtual disk in the VMware VMDK format. This is the only available option if you export disks to an ESXi datastore.
   - **VHD** – select this option if you want to save the resulting virtual disk in the Microsoft Hyper-V VHD format.
   - **VHDX** – select this option if you want to save the resulting virtual disk in the Microsoft Hyper-V VHDX format (supported by Microsoft Windows Server 2012 or later).
4. [For disks exported as VMDK to an ESXi datastore] Click the **Pick proxy to use** link to select backup proxies over which disk data must be transported to the target datastore. You can assign backup proxies explicitly or instruct Veeam Backup & Replication to automatically select backup proxies.

5. From the **Disk Type** drop-down list, select a disk type for the exported disks.

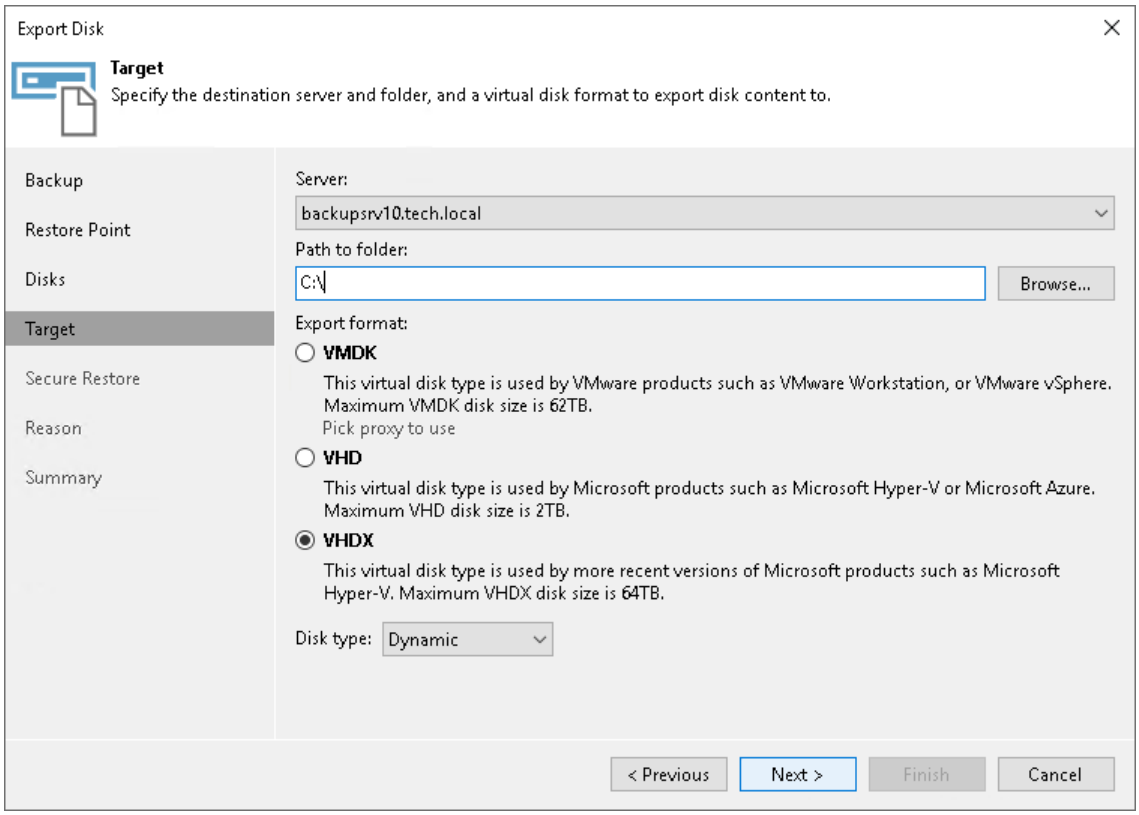

# Step 6. Specify Secure Restore Settings

This step is available if you export disks of Microsoft Windows workloads.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Specify how Veeam Backup & Replication must behave after the first malware threat is found: continue or abort disk recovery.

4. Select the **Continue scanning all remaining files after the first occurrence** check box if you want to continue the disk scan after the first virus threat is detected. For information on how to view results of the malware scan, see Viewing Malware Scan Results.

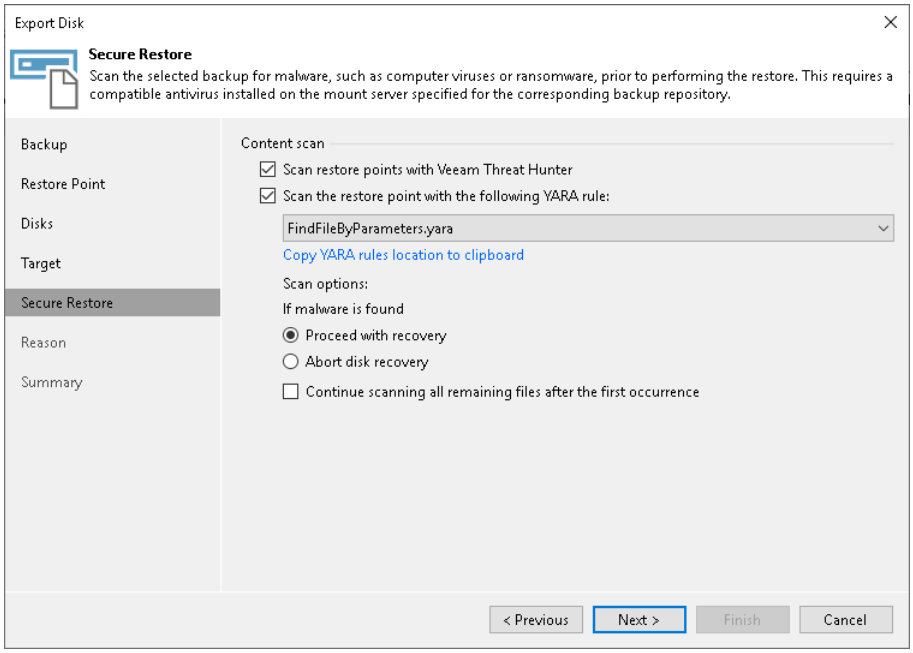

## Step 7. Specify Export Reason

At the **Reason** step of the wizard, enter a reason for disk export.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

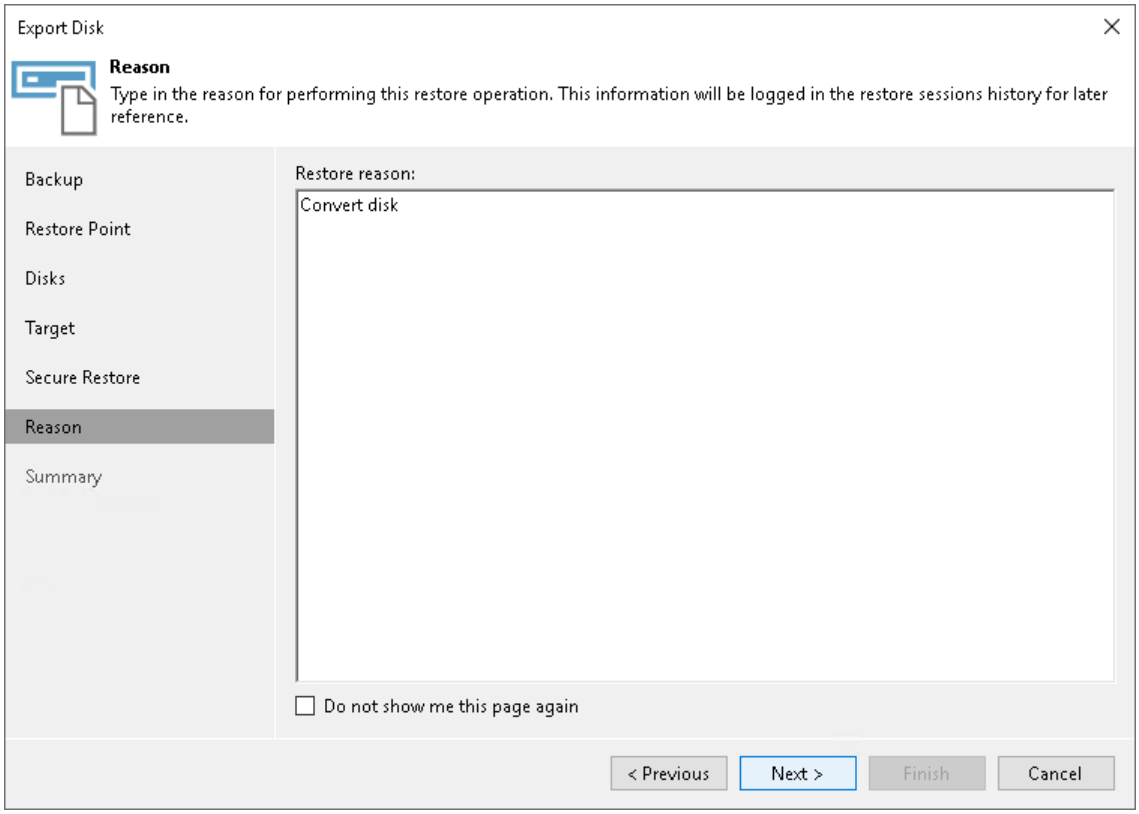

# Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured settings and click **Finish**.

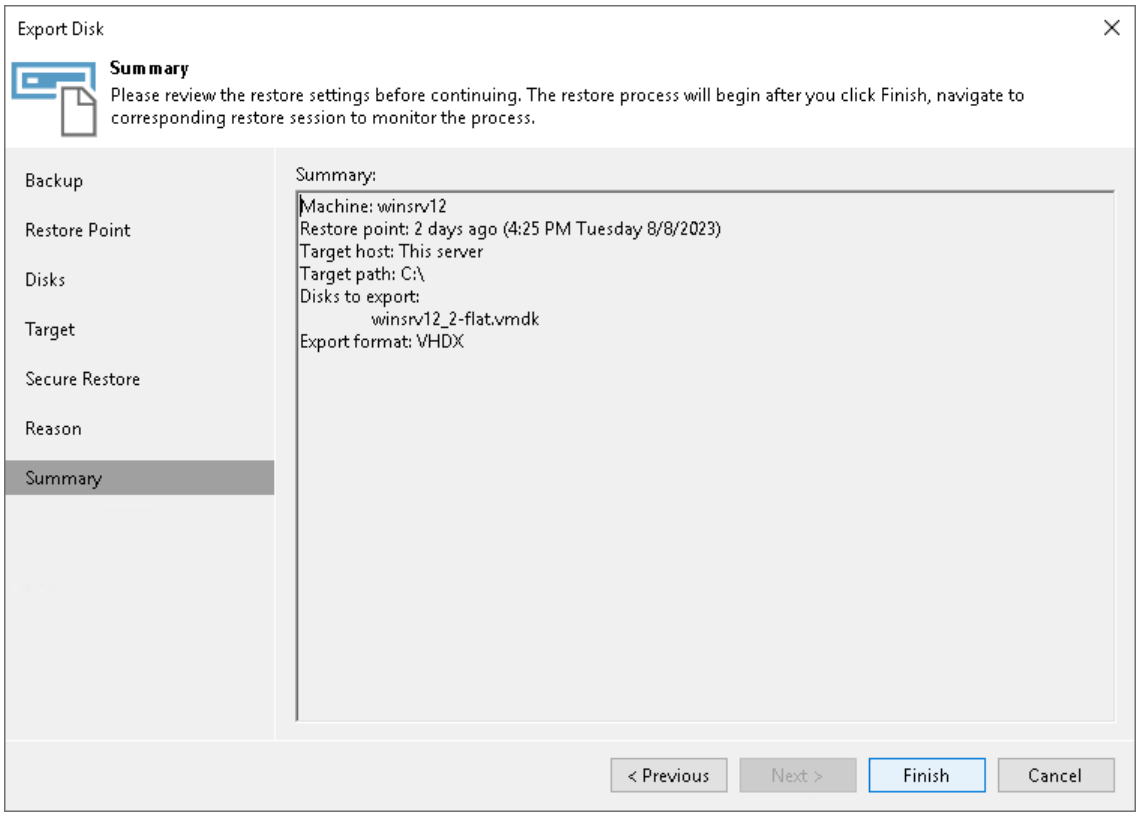

# Disk Publishing (Data Integration API)

Disk publishing allows you to save time by getting backup content of one or multiple disks instead of all disks from a backup. This technology gives read-only access to data and helps if you want to analyze data of your backup. For example, look for specific documents or usage patterns, or perform antivirus scan of backed-up data.

You can publish a disk from different types of backups. The disk can have a Microsoft Windows file system or Linux, Unix or other file system. For the full list of supported file systems, see Supported Platforms and Applications. To present the backed-up disk to the target server, disk publishing uses the iSCSI protocol for Windows file systems or the FUSE protocol for other OSes. After the publishing, the target server can access the backup content using the iSCSI initiator or FUSE protocol, and read the necessary data from the disk.

## Supported Backup Types

You can publish disks from the following types of backups:

- Backups of VMware vSphere virtual machines created by Veeam Backup & Replication
- Backups of Microsoft Hyper-V virtual machines created by Veeam Backup & Replication
- Backups of VMware Cloud Director virtual machines created by Veeam Backup & Replication
- Backups of oVirt VMs created by Veeam Backup for OLVM and RHV
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for Mac, Veeam Agent for Oracle Solaris or Veeam Agent for IBM AIX
- Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV
- Backups of Amazon EC2 instances created by Veeam Backup for AWS
- Backups of Microsoft Azure virtual machines created by Veeam Backup for Microsoft Azure
- Backups of Google instances created by Veeam Backup for Google Cloud
- Backups exported by Kasten policies
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE

> **NOTE**
>
> You cannot publish disks from replicas using UI. To do that, use the Publish-VBRBackupContent cmdlet.

## Considerations and Limitations

Disk publishing has the following requirements and limitations:

- The necessary ports must be opened on the target server. For more information, see Ports.
- The file system of a workload whose disks you plan to restore must be supported. For details, see the File-Level Restore.
- To restore disks that have file system other than Microsoft Windows, you can use only a Linux-based server as the target server.
- The 32-bit version of a Linux-based server is not supported as the target server.

- The target server must support the file system of the disk that you plan to publish.
- A server on which a hardened repository is configured cannot be used as the target server.
- Disk publishing is not available for Linux, Unix, macOS (and other) workload backups stored on Veeam Cloud Connect Repositories. For more information on Veeam Cloud Connect repositories, see the Cloud Repository section in the Veeam Cloud Connect Guide.
- Basic disks, Linux LVM (Logical Volume Manager) and ZFS pools can be published. Encrypted, RAID1 and mirrored LVM volumes are not supported.
- RAID disks publish for Linux workloads is not supported.
- The target server must support the same ReFS version or later than the version used on the workload from which you plan to restore files. For more information on which OSes support which ReFS, see ReFS versions and compatibility matrix.
- If data deduplication is enabled for some disks in a backup, data deduplication must be enabled on the target server.
- You cannot publish disks from a backup created in the reverse incremental mode if the backup job is being performed. If the backup is created in the incremental backup mode and the backup job is being performed, you can publish disks from any available restore point.
- You can publish disks from Novell Storage Services (NSS) file system. The target host must differ from the backed-up workload.
- Mounting of LVM snapshots is not supported. LVM snapshots are skipped from processing.
- You can publish disks with ZFS FS if the `zfsutils-linux` package is installed on the specified target host. The `zfs-fuse` package is not supported.
- If you want to publish a Btrfs disk and select the original server as the target server, the mount of the Btrfs disk will fail. The issue occurs due to restriction of mounting two Btrfs disks with identical IDs to the same machine. To avoid this issue, use another target server. However, note that Btrfs disks can be restored for backups created by Veeam Agent for Linux. During the backup process, Veeam Agent for Linux changes disk IDs in the backup file.

# How Publishing Works for Microsoft Windows File Systems

After you start the publishing session for disks with Microsoft Windows file system, the following applies:

1. Veeam Backup & Replication connects to the backup repository where the backup locates and to the mount server associated with this repository.
2. Veeam Backup & Replication starts the following agents and services in the infrastructure:
   - The repository agent. This agent runs on the backup repository or on the gateway server associated with this repository.
   - The mount agent. This agent runs on the mount server. The agent configures the iSCSI Target Server on the mount server and triggers the repository agent to read the backup data.
   - Veeam Installer Service. This service is installed and runs on the target server and remains on it after the publishing finishes.
3. Veeam Backup & Replication accesses the target server and configures the iSCSI initiator on this server.
4. After Veeam Backup & Replication adds the target server to the iSCSI Target Server settings, the target server can access disk content. The disk content is available in the `C:\VeeamFLR\` folder on the target server. Note that disk data is read-only.

After that, you can browse the disks and perform data analysis operations with them. The published disks are available while the mount server is up and running. Even if you turn off the backup server, you will be able to access the published disks.

After you finish working with the disks, you should stop publishing them as described in Managing Published Disks.

# How Publishing Works for Linux, Unix and Other File Systems

After you start the publishing session for disks with Linux, Unix or other file system, the following applies:

1. Veeam Backup & Replication connects to the backup repository where the backup locates and to the mount server associated with this repository.
2. Veeam Backup & Replication starts the following agents in the infrastructure:
   - The repository agent. This agent runs on the backup repository or on the gateway server associated with this repository.
   - A temporary agent. This agent runs on the target server and is removed after the publishing session finishes.
3. Veeam Backup & Replication uses the FUSE protocol to publish the content of the backup automatically.

   The published disk images are available in the `/tmp/Veeam.Mount.Disks` location. The disk content is available in the `/tmp/Veeam.Mount.FS` location. Note that disk data is read-only.

After that, you can browse the disks and perform data analysis operations with them.

After you finish working with the disks, you should stop mounting them as described in Managing Published Disks.

# Mount Modes

When you publish disks with Microsoft Windows file system from the UI, Veeam Backup & Replication automatically configures the iSCSI session and gives the target server (iSCSI initiator) access to the published disks. To manually start the iSCSI session from any server that has access to the iSCSI Target Server, you can use the Publish-VBRBackupContent cmdlet.

# Publishing Disks

To publish a disk, use the **Publish Disks** wizard.

## Step 1. Launch Wizard

To launch the **Publish Disks** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select one of the following:
  - **VMware vSphere** > **Restore from backup** > **Disk Restore** > **Publish disk** — to publish a disk of a VMware vSphere VM from a backup created by Veeam Backup & Replication.
  - **VMware Cloud Director** > **Restore from backup** > **VM restore** > **Disk Restore** > **Publish disk** — to publish a disk of a VMware Cloud Director VM from a backup created by Veeam Backup & Replication.
  - **Microsoft Hyper-V** > **Restore from backup** > **Entire VM restore** > **Publish disk** — to publish a disk of a Hyper-V VM from a backup created by Veeam Backup & Replication.
  - **Agent** > **Disk Restore** > **Publish disk** — to publish a disk of a physical machine or a virtual machine from a backup created by Veeam Agents.
  - **AWS** > **Amazon EC2** > **Entire machine restore** > **Publish disk** — to publish a disk of an EC2 instance from a backup created by Veeam Backup for AWS.
  - **Azure IaaS backup**> **Entire machine restore** > **Publish disk** — to publish a disk of an Azure VM from a backup created by Veeam Backup for Microsoft Azure.
  - **GCE backup** > **Entire machine restore** > **Publish disk** — to publish a disk of a VM instance from a backup created by Veeam Backup for Google Cloud.
  - **Nutanix backup** > **Entire machine restore** > **Publish disk** — to publish a disk of a VM from a backup created by Veeam Backup for Nutanix AHV.
  - **oVirt KVM** > **Entire machine restore** > **Publish disk** — to publish a disk of a VM from backups created by Veeam Backup for OLVM and RHV.
  - **Kasten backup** > **Publish disk** — to publish a disk of a VM whose backups were exported by Kasten policies.
  - **Proxmox VE** > **Publish disk** — to publish a disk of a Proxmox VE VM from a backup created by Veeam Backup for Proxmox VE.

- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary backup, select a workload whose disks you want to publish and click **Publish Disks** on the ribbon. Alternatively, you can right-click the workload and select **Publish disks**.

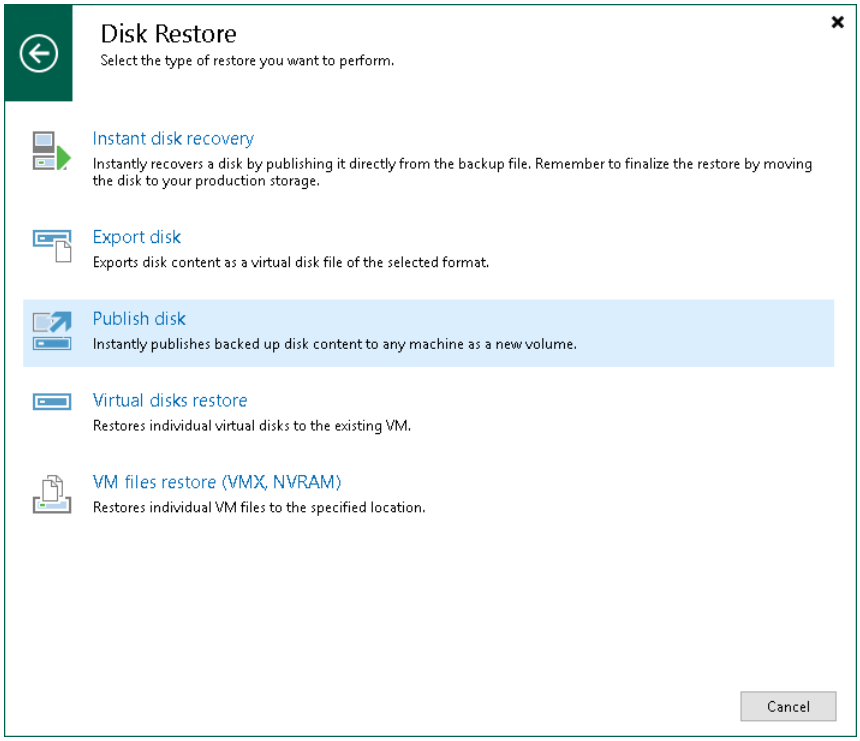

# Step 2. Select Workload

At the **Machine** step of the wizard, expand a backup and select a workload whose disks you want to publish.

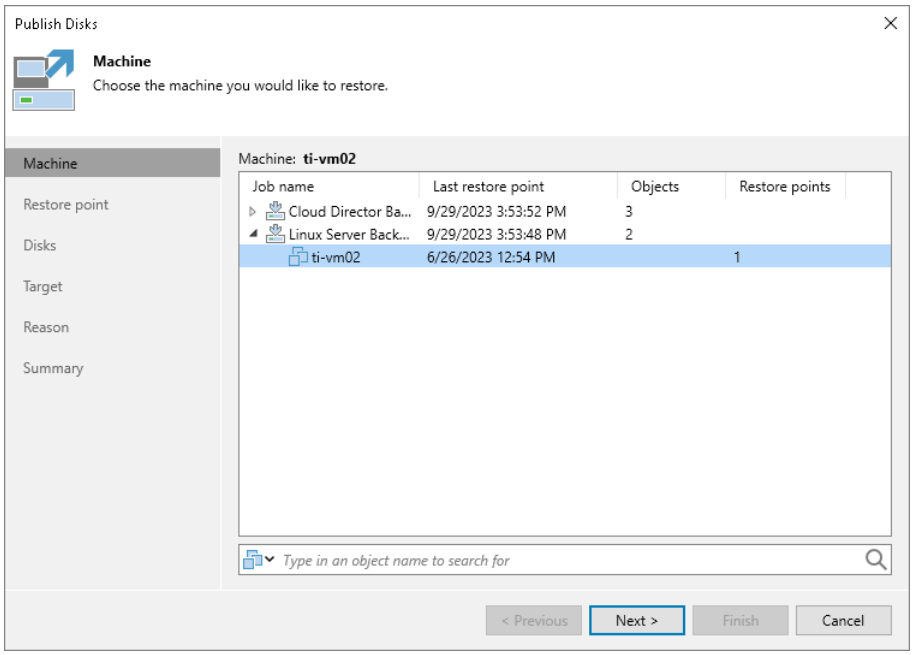

# Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point from which you want to publish disks.

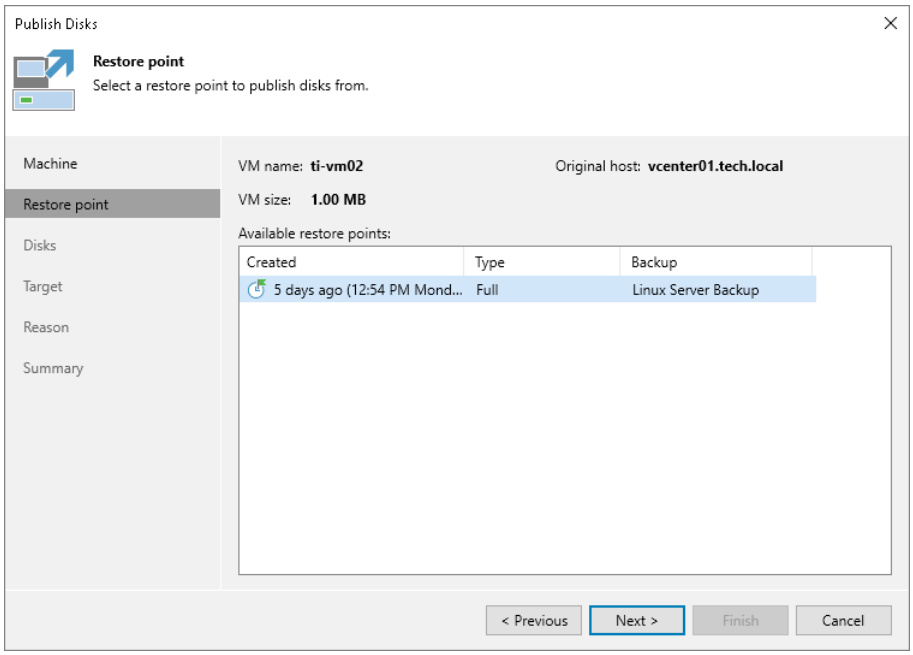

# Step 4. Select Disks

At the **Disks** step of the wizard, select a check box next to the disks that you want to publish.

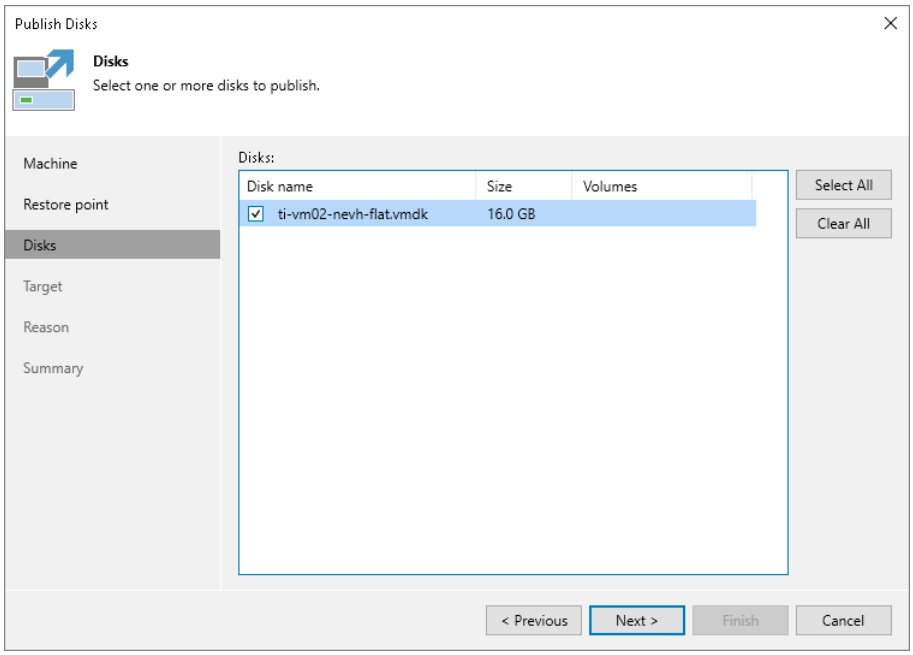

# Step 5. Specify Target Server

At the **Target** step of the wizard, select a server that will have access to disk content — for Microsoft Windows file system a Microsoft Windows server; for Linux, Unix or other file system a Linux-based server. You can select one of the following types of servers:

- A server added to the backup infrastructure.
- A temporary server. In this case, select *Specify a different host* from the drop-down list. In the **Target Server** window, specify a server name or IP and credentials to the server.
- The original server if you publish disks from a backup created by Veeam Agent for Microsoft Windows or Veeam Agent for Linux. In this case, select *Original server* from the drop-down list.

If prompted, specify credentials for the target server.

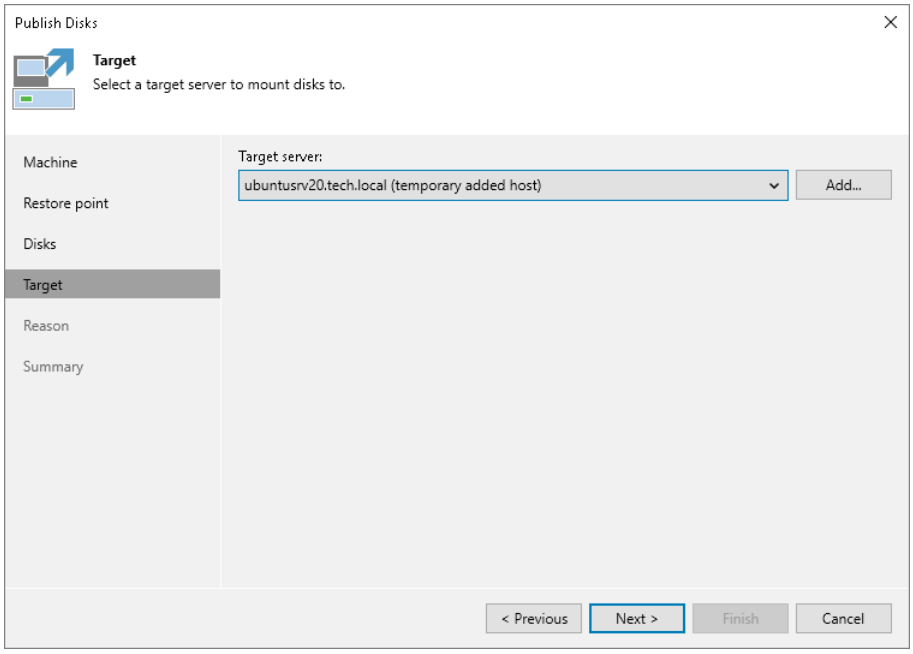

# Step 6. Specify Reason

At the **Reason** step of the wizard, enter a reason for publishing disks.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

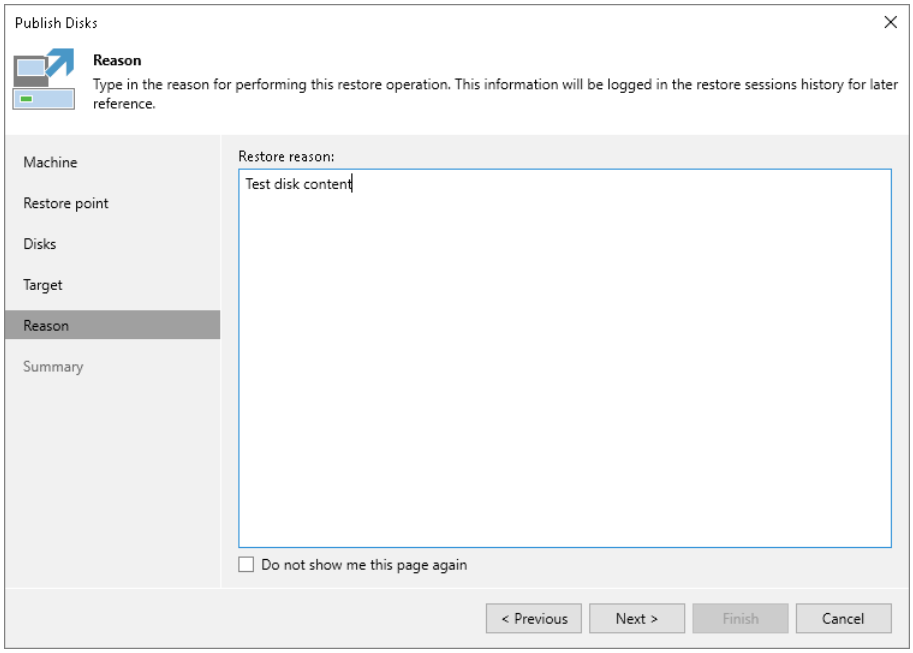

## Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured settings and click **Finish**.

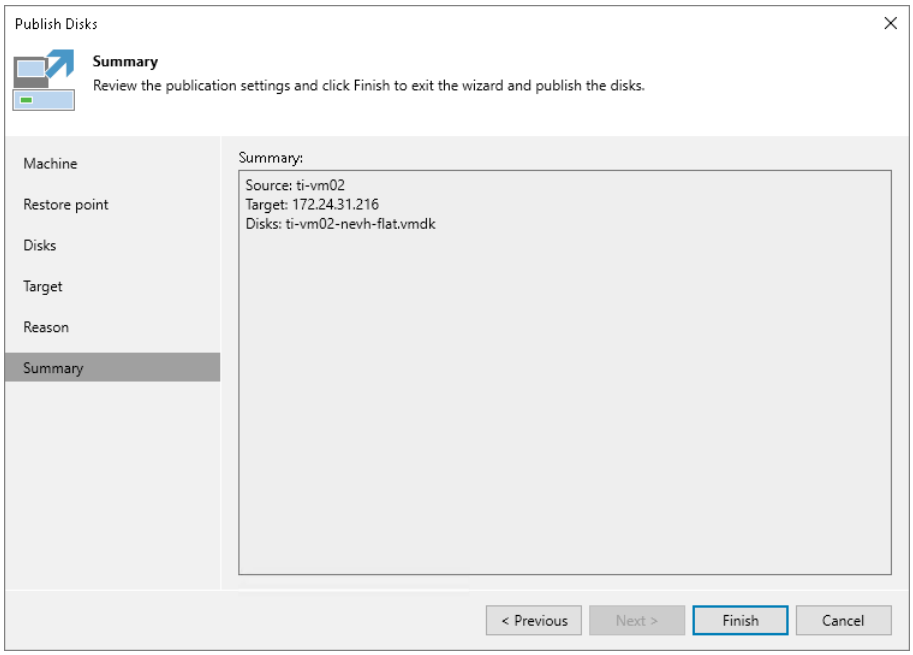


# What You Do Next

Depending on the file system of the disk, go to the following locations on the target server to browse disk content:

- [Microsoft Windows file systems] Go to `C:\VeeamFLR\` folder.
- [Linux, Unix and other file systems] Go to the `/tmp/Veeam.Mount.Disks` location to browse disk images. Go to the `/tmp/Veeam.Mount.FS` location to browse disk content.

# Managing Published Disks

After you started a publishing session, you can check details on this session or stop it.

## Viewing Statistics on Publishing Session

To view publishing session statistics, do one of the following:

- Open the **Home** view. In the inventory pane, select **Instant Recovery.** In the working area, select the necessary publishing session and click **Properties** on the ribbon. Alternatively, right-click the session and **Properties**.
- Open the **Home** view. In the inventory pane select **Last 24 hours**. In the working area, double-click the necessary publishing session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

- Open the History view. In the inventory pane select **Restore**. In the working area, double-click the necessary mount session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

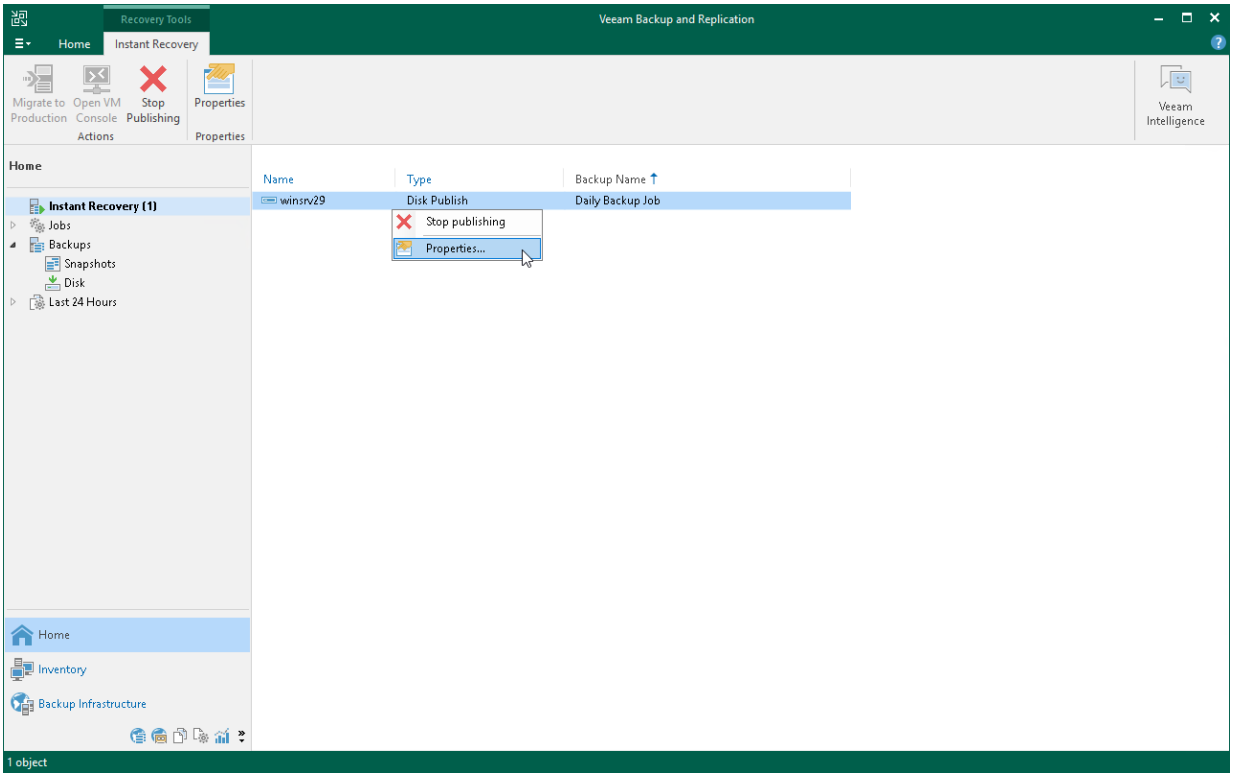


The publishing statistics provides the following data:

- At the top of the **Restore Session** window, Veeam Backup & Replication shows general session statistics. It includes a name of the workload whose disk you want to publish, a name of the backup server which initiated the publishing session, a user name of the account under which the session was started, session status, and duration details.
- The **Reason** tab shows the reason for the publishing session.
- The **Parameters** tab shows information about the target server, the machine whose disks you publish and the restore point selected for publishing.
- The **Log** tab shows the list of operations performed during the session.

## Stopping Publishing Session

To stop a publishing session, do one of the following:

- Open the **Home** view. In the inventory pane select **Instant Recovery**. In the working area, select the necessary publishing session and click **Stop Publishing** on the ribbon. Alternatively, right-click the session and **Stop Publishing**.
- Open the **Home** view. In the inventory pane select **Last 24 hours**. In the working area, select the necessary publishing session and click **Stop Publishing** on the ribbon. Alternatively, right-click the session and **Stop Publishing**.

- Open the History view. In the inventory pane select **Restore**. In the working area, select the necessary publishing session, double-click it. In the **Restore Session** window, click **Cancel restore task**.

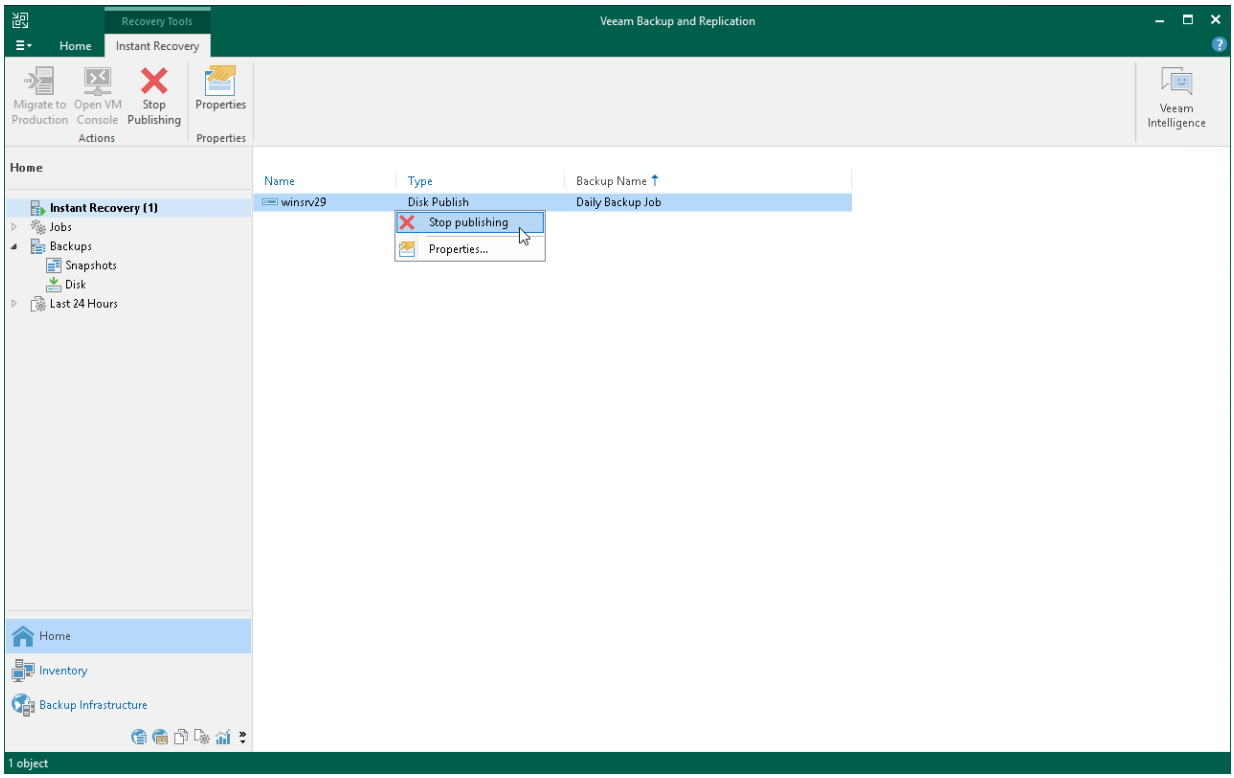

# Disk Restore to oVirt KVM

Veeam Backup & Replication allows you to recover different workloads as OLVM and RHV VM disks.

> **IMPORTANT**
>
> To restore to OLVM and RHV, you must install oVirt KVM Plug-in for Veeam Backup & Replication on the backup server. To learn more, see the Installation section in the Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization User Guide.

The restore procedure of virtual disks to oVirt KVM practically does not differ from the procedure described in the Performing Disk Restore section in the Veeam Backup for OLVM and RHV User Guide.

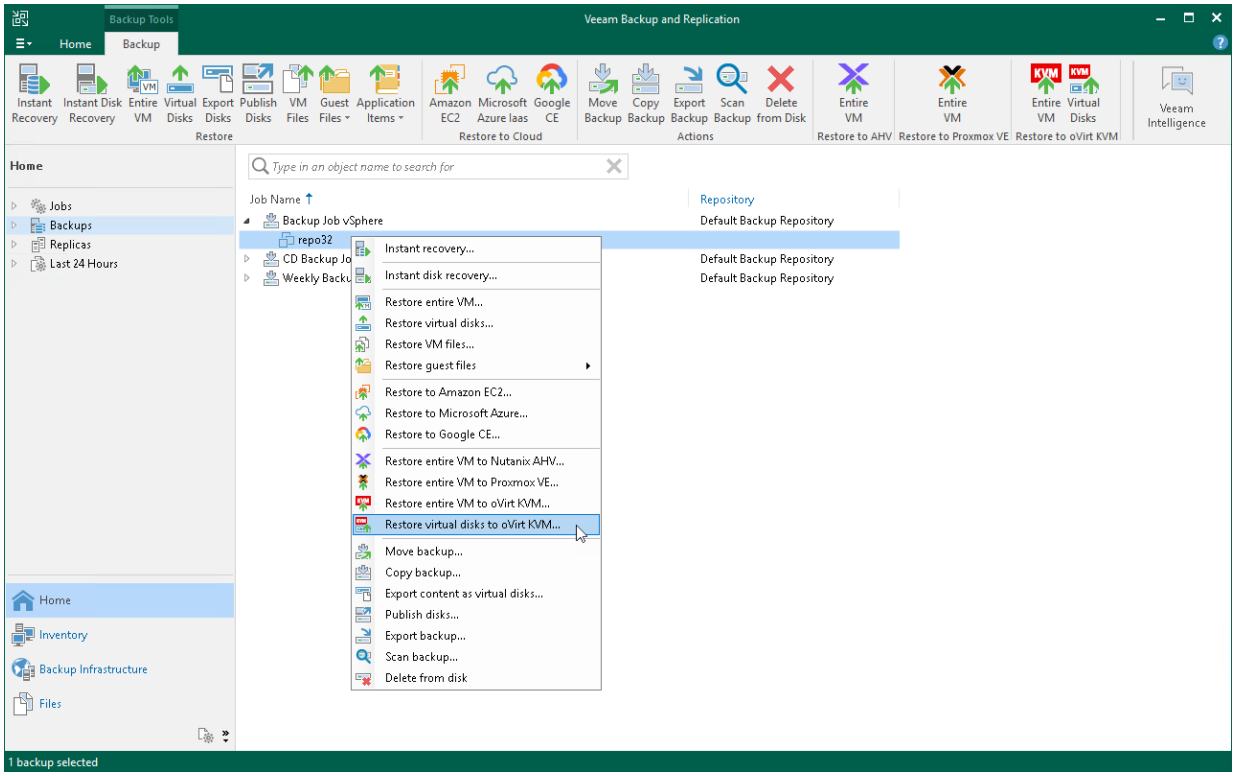

# Item Recovery

Item recovery includes the following methods:

- VM files restore — to restore VM files (.VMX, .VMXF and so on) without restoring the entire VM.
- Guest OS file restore — to restore individual guest OS files from Windows, Linux, Mac and other guest OS file systems. You can restore files and folders directly from a regular image-level backup, storage snapshot or replica.
- Application items restore — to restore items from different applications such as Microsoft Active Directory, Microsoft SQL Server and so on. Application items are recovered directly from VM backups and replicas. To recover application items, Veeam Backup & Replication uses the capabilities of Veeam Backup Explorers.

# VM Files Restore

You can restore specific VM files (.VMX, .VMXF, .NVRAM, .VMDK including flat files) if any of these files are deleted or the datastore is corrupted. This option provides a great alternative to entire VM restore, for example, when your VM configuration file is missing and you need to restore it. Instead of restoring the whole VM image to the production storage, you can restore a specific VM file only.

When you perform VM file restore, VM files are restored directly from regular image-level backups, without prior de-staging of VM images from backups. VM files can be restored to the original VM location or to a new location.

## Restoring VM Files

To restore VM files from a backup, use the **Virtual Machine Files Restore** wizard.

### Before You Begin

Before you restore VM files, check the following prerequisites:

- You can restore VM files from a backup that has at least one successfully created restore point.
- The server on which you plan to save restored VM files must be added to the backup infrastructure.

## Step 1. Launch Restore Wizard

To launch the **Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Disk restore** > **VM files restore (VMX, NVRAM)**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, click the VM whose files you want to restore and click **VM Files** on the ribbon. Alternatively, you can right-click the VM whose files you want to restore and select **Restore VM files**.

  In this case, you will pass to the Restore Point step of the wizard.
- Double-click the VBK or VBM file (for example, in Microsoft Windows Explorer). In the displayed window, select the VM and click **Restore** > **VM files**. In this case, you will pass to the Restore Point step of the wizard.

  You can use this option if you perform restore on the backup server. You cannot use this option if you perform restore remotely over the Veeam Backup & Replication console.

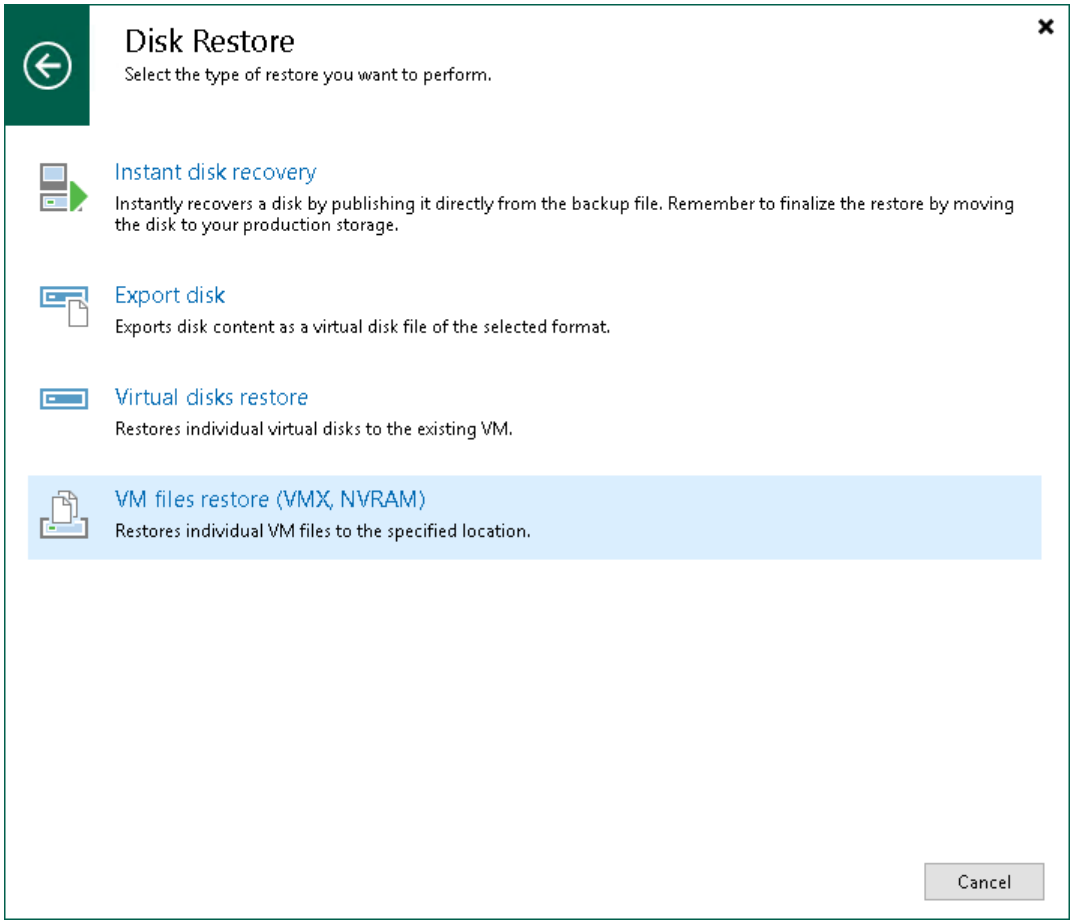

## Step 2. Select VM

At the **Virtual Machine** step of the wizard, select the VM whose files you want to restore:

1. In the **Virtual machine** list, expand the necessary backup.
2. Select the VM.

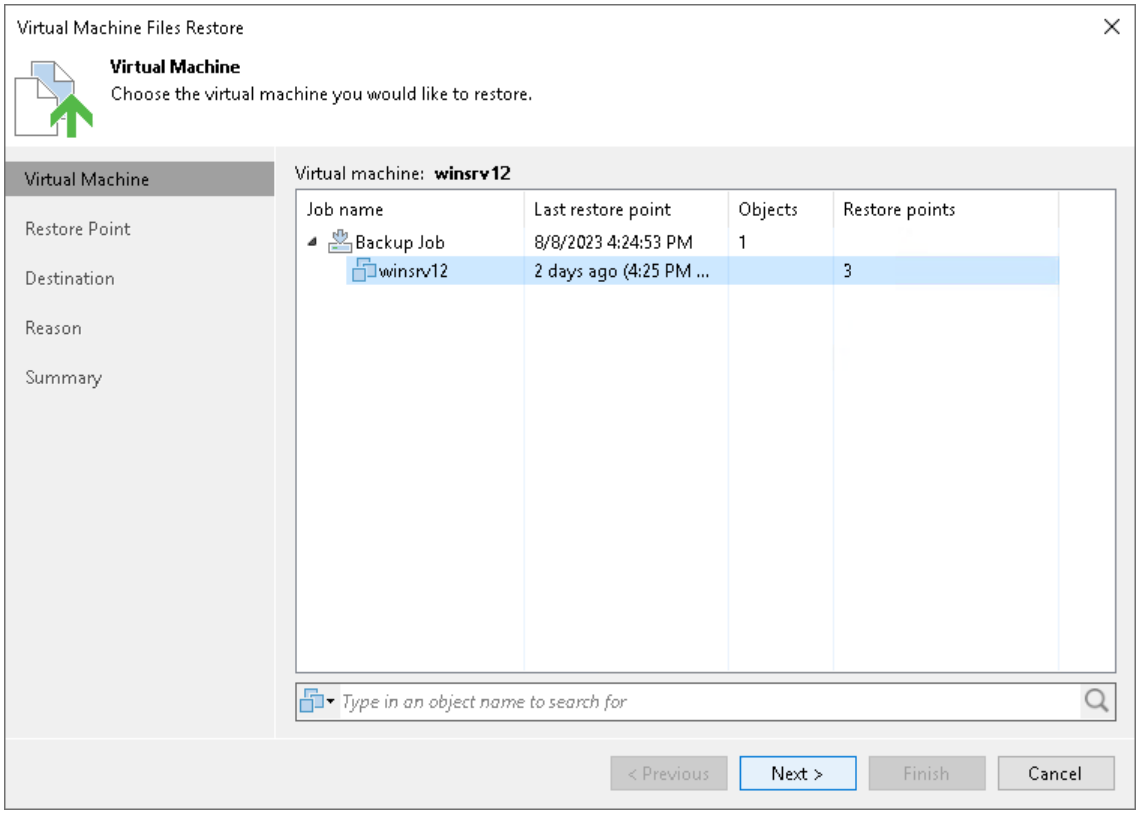

# Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point from which you want to restore VM files.

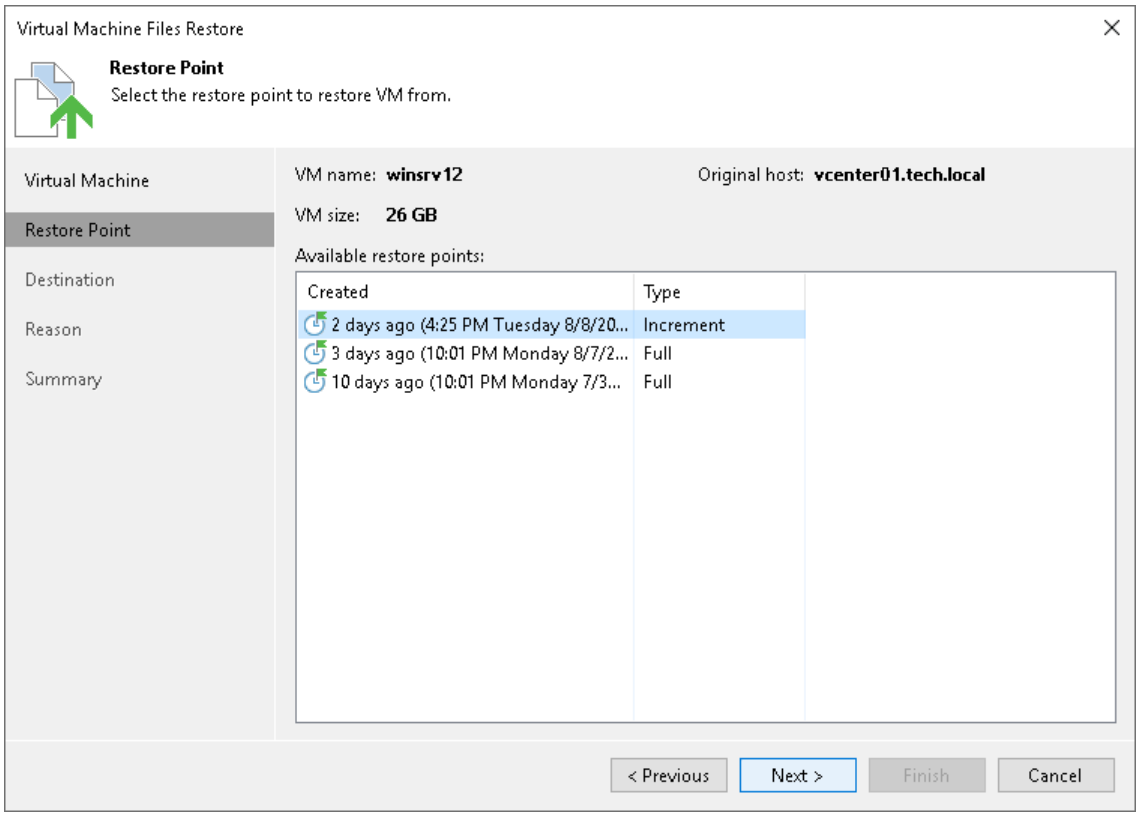

# Step 4. Select VM Files and Destination

At the **Destination** step of the wizard, select VM files that you want to restore and destination where the restored files must be stored:

1. From the **Server** list, select where to store VM files: on an ESXi host, on the backup server or on a Microsoft Windows or Linux server added to the backup infrastructure. Use the **Details** button to view or change connection settings of the target host or server.
2. In the **Path to folder** section, specify a path to the folder on the selected host where files must be restored.

   To create a new folder for restored files, click **Browse**. In the **Select Folder** window, select the target location for VM files and click **New Folder** at the bottom of the window.
3. In the **VM files to restore** section, select check boxes next to files that you want to restore.

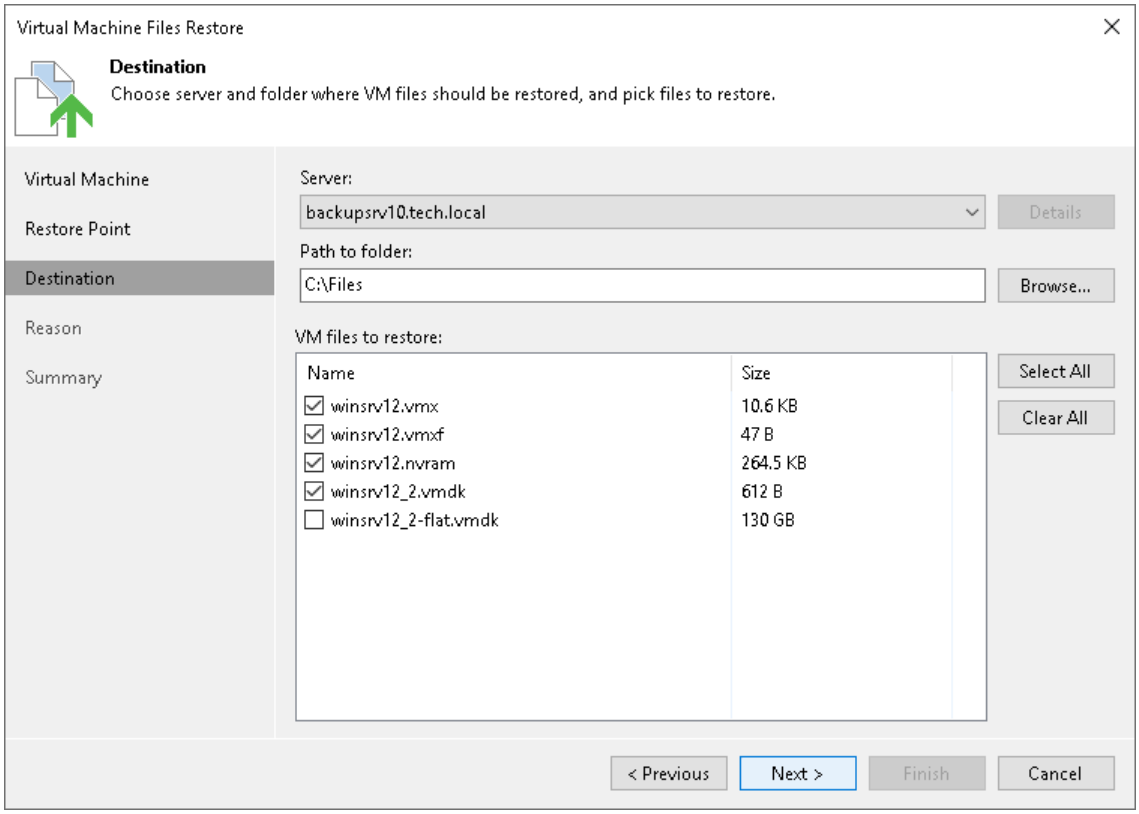

## Step 5. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring VM files. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

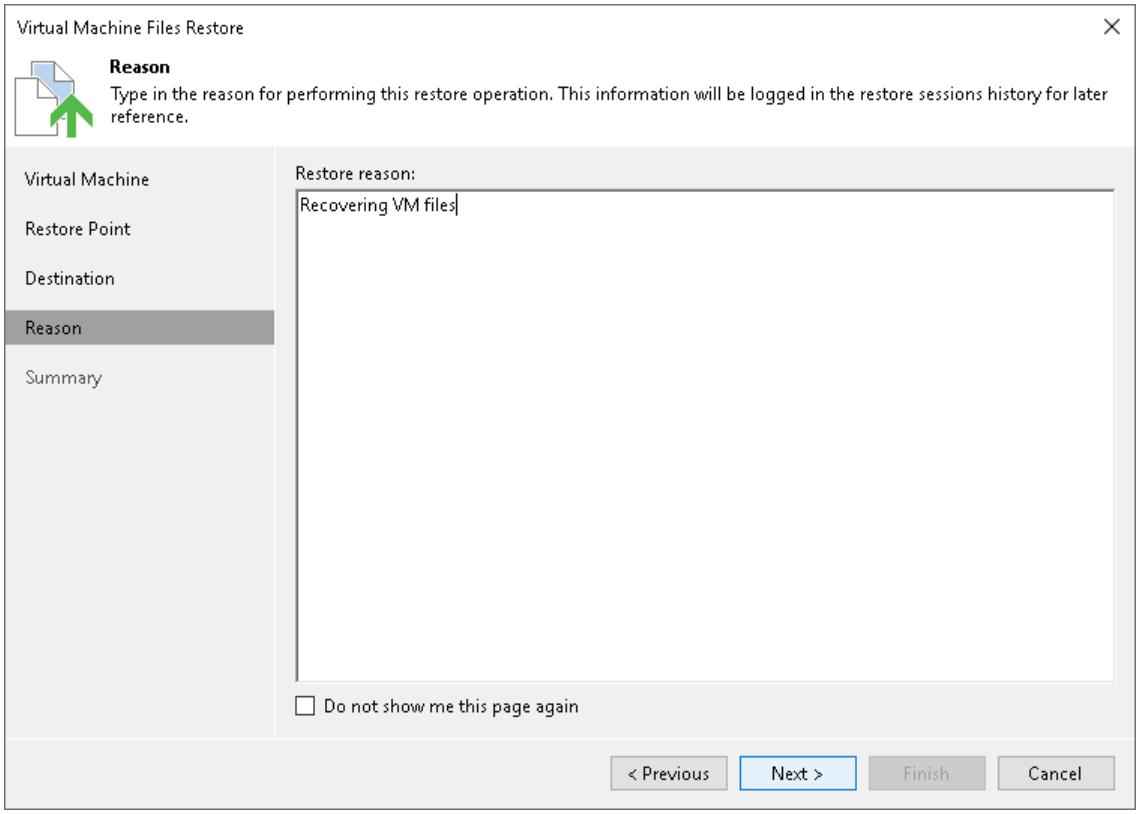

# Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, check details for the restore task and click **Finish** to start VM files restore.

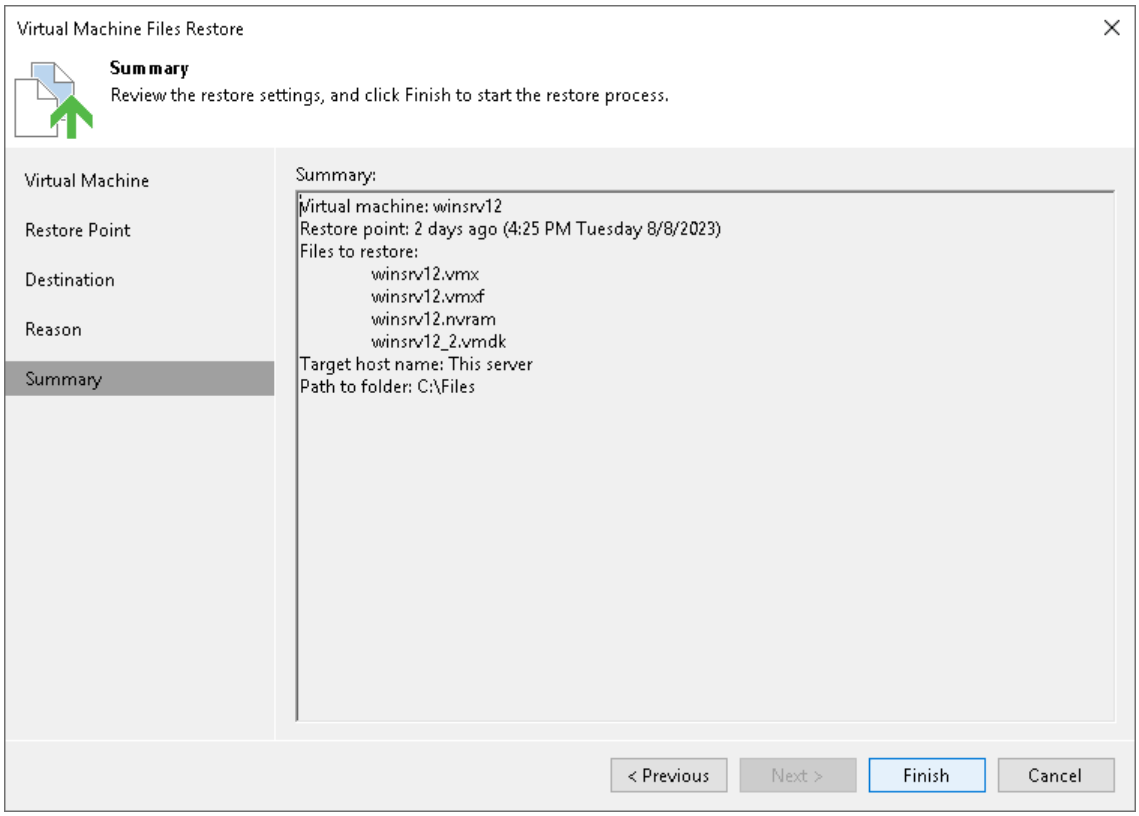

# Guest OS File Restore

With guest OS file restore, you can recover individual guest OS files and folders from backups, replicas storage snapshots, Nutanix AHV snapshots and so on. When restoring files or folders, you do not need to extract the VM image to a staging location or start the VM prior to restore. You can restore files and folders directly from a regular image-level backup or replica to the necessary point in time.

Using the following methods, you can restore files from different guest OS file systems:

- Restore from Microsoft Windows File Systems (FAT, NTFS or ReFS)

  This method helps you restore files from Microsoft Windows VMs with NTFS, FAT and ReFS file systems.

- Restore from Linux, Unix and Other File Systems (multi-OS guest OS file restore)

  This method helps you restore files from Linux, Solaris, BSD, Novell Storage Services, Unix and Mac machines.

  Note that this method supports recovery of files and folders only. Recovery of other file system objects such as pipes is not supported.

- Restore from Other File Systems

  This method helps you restore files from file systems that other methods from the list do not support.

## Restore from Microsoft Windows File Systems (FAT, NTFS or ReFS)

The restore from FAT, NTFS and ReFS method helps you restore files of Microsoft Windows workloads with NTFS, FAT and ReFS file systems. For the full list of supported file systems, see Supported Platforms and Applications.

You can restore files from the following types of data:

- Backups
- Replicas
- Storage snapshots

You can restore files to the original or new location, restore changes and permissions, work with the restored files using Microsoft Windows File Explorer or launch application item restore for the files. For more information, see Finalize Restore.

## How Restore Works

During the guest OS file restore, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication creates a mount point under the `C:\VeeamFLR\<vmname>` folder and mounts workload disks from the backup or replica to it. For more information on which machines Veeam Backup & Replication creates mount points, see Mount Points and Restore Scenarios.

   If a workload has ReFS disks, Veeam Backup & Replication uses virtual hard disk (VHD) mount to access the content of all disks. If a workload does not have ReFS disks, Veeam Backup & Replication uses a separate program — Virtual Disk Driver (VDK) that is provided with the product. Workload disks are not physically extracted from the backup file or workload replica. Veeam Backup & Replication emulates their presence on the backup server or Veeam Backup & Replication console. The backup file or workload replica itself remains in the read-only state.

2. Veeam Backup & Replication launches the Veeam Backup browser on the Veeam Backup & Replication console. The Veeam Backup browser shows the content of disks mounted to the machine where the mount point is created.

   You can browse the workload guest file system in the Veeam Backup browser and restore files or folders to the original or new location.

3. Depending on which restore command you use, the operations differ:

   - If you restore to the original location, select the **Restore** > **Keep** or **Restore** > **Overwrite** command.

     The machine where the mount point is created connects to the workload over network or VIX API/vSphere Web Services if a connection over the network cannot be established.

   - If you restore permissions, select the **Restore** > **Permissions only** command.

     The machine where the mount point is created connects to the workload over network.

   - If you restore files to a new workload, select the **Restore to** > **Keep** or **Restore to** > **Overwrite** command.

     The machine where the mount point is created connects to the workload over network or VIX API/vSphere Web Services if a connection over the network cannot be established.

   - If you restore files to the Veeam Backup & Replication console or a shared folder, select the **Copy to** command.

     The machine where the mount point is created connects to the Veeam Backup & Replication console over network. If you recover files to a shared folder, the Veeam Backup & Replication console also connects to this folder over network.

   - If you restore changes, select the **Restore changed only** command.

     During the comparison, Veeam Backup & Replication connects to the original workload over network and installs Veeam Installer Service. If the service is already installed on the workload, Veeam Backup & Replication upgrades it to the required version. Note that after the comparison or restore is finished, the service is still installed on the workload.

     When restore is launched, the machine where the mount point is created connects to the workload over network.

4. When the restore process is finished or the Veeam Backup browser is closed by timeout, Veeam Backup & Replication removes all the created mount points.

When you perform guest OS file restore, you can also open files in Microsoft Windows explorer to view file content. In this case, Veeam Backup & Replication creates an additional mount point on the Veeam Backup & Replication console. For more information on how to view file content, see Open Files in Microsoft Windows Explorer.

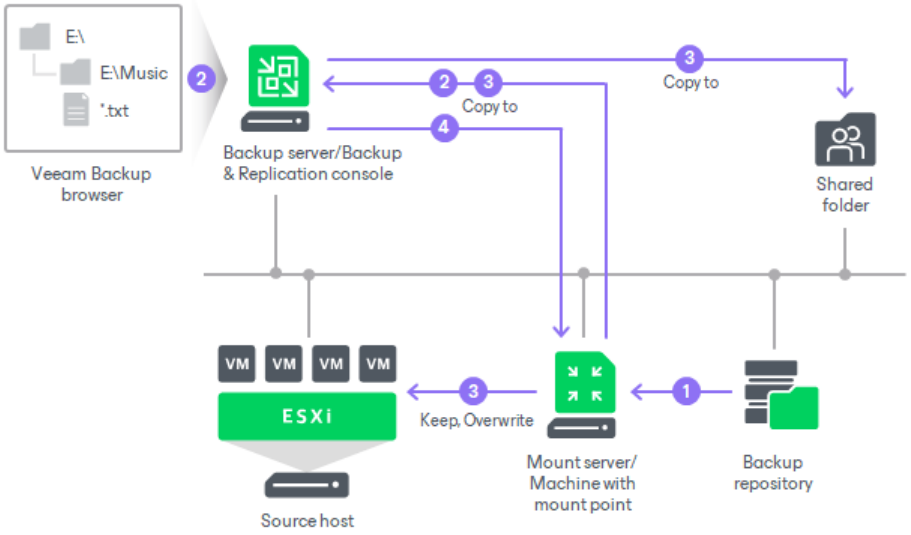


## Mount Points and Restore Scenarios

Veeam Backup & Replication creates mount points on different machines depending on a restore scenario you use.

# Restoring Files from Backups

When you restore files from backups that reside in the backup repository, Veeam Backup & Replication creates a mount point on the following machines:

- Mount server associated with the backup repository on which the backup file resides. Veeam Backup & Replication uses this mount point when the restore process starts and allows you to browse the VM file system and restore files.
- Veeam Backup & Replication console. Veeam Backup & Replication uses this mount point only for the mount to console functionality, that is, while you use Microsoft Windows File Explorer to view files and folders of the backed-up VM. For more information, see Open Files in Microsoft Windows Explorer.

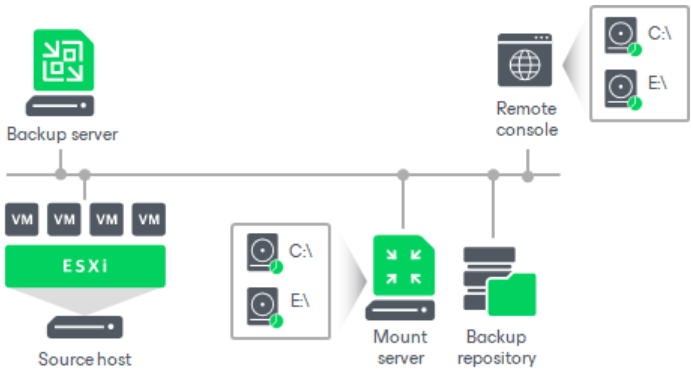

NOTE

Backup files on HPE StoreOnce are locked exclusively by a restore task. For this reason, Veeam Backup & Replication uses only one mount point on the backup server or Veeam Backup & Replication console machine for backups on HPE StoreOnce.

# Restoring Files from Replicas

When you restore files from a VM or CDP replica, Veeam Backup & Replication creates a mount point on the following machines:

- Backup server. Veeam Backup & Replication uses this mount point while the restore process and allows you to browse the VM file system and restore files.
- Veeam Backup & Replication console. Veeam Backup & Replication uses this mount point only for the mount to console functionality, that is, while you use Microsoft Windows File Explorer to view files and folders of the VM replica. For more information, see Open Files in Microsoft Windows Explorer. If you also restore files to the console or a shared folder (use the **Copy to** command) while the mount to console functionality is enabled, Veeam Backup & Replication uses the same mount point for restore. If you launch other restores, Veeam Backup & Replication uses the backup server as a mount point.

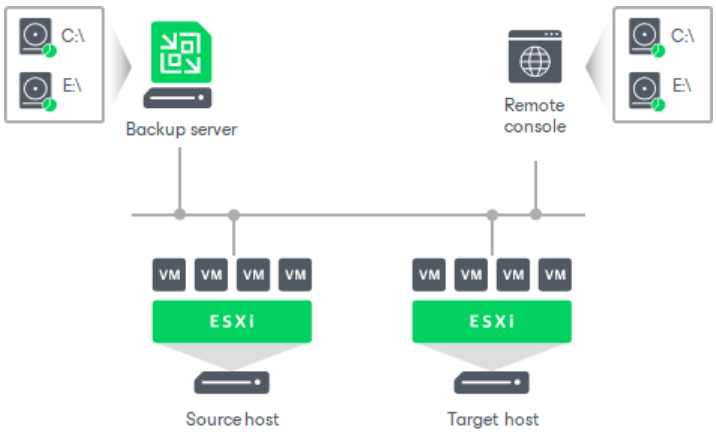


# Restoring Files from Storage Snapshots

When you restore files from storage snapshots, Veeam Backup & Replication uses one mount point on the backup server or Veeam Backup & Replication console machine.

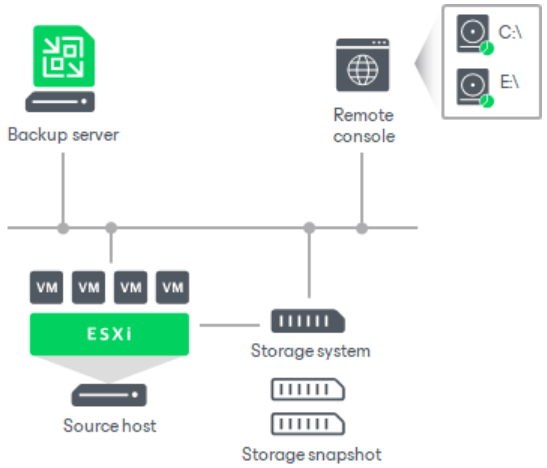

# Restoring Files for Veeam Explorers

If you have launched guest OS file restore and then launch application items restore from the Veeam Backup browser, Veeam Backup & Replication uses additional restore points apart from the the mount points described in the subsections above:

- If you restore Microsoft Active Directory or Microsoft Exchange items, Veeam Backup & Replication creates an additional mount point on the Veeam Backup & Replication console. However, if during guest OS file restore the mount point was already created on a machine where the Veeam Backup & Replication console is installed, an additional point is not created.
- If you restore Microsoft SharePoint items, Veeam Backup & Replication creates an additional mount point on a staging Microsoft SQL Server.
- If you restore Microsoft SQL Server items, Veeam Backup & Replication creates two additional mount points: one on a staging Microsoft SQL Server and the other one on the target VM to which you restore the application items.
- If you restore Oracle server items, Veeam Backup & Replication creates two additional mount points: one on a staging Oracle server and the other one on the target VM to which you restore the application items.
- If you restore PostgreSQL server items, Veeam Backup & Replication creates two additional mount points: one on a staging PostgreSQL server and the other one on the target VM to which you restore the application items.

To create mount points on Linux and Windows machines, Veeam Backup & Replication uses different methods:

- To create a mount point on Microsoft Windows machines, Veeam Backup & Replication uses the iSCSI protocol. The remote machine or staging server acts as an iSCSI initiator. The machine on which the Veeam Explorer runs acts as an iSCSI target. The iSCSI mount point is non-persistent — it is created only for duration of the restore process.
- To create a mount point on Linux VMs (for Oracle and PostgreSQL running on Linux), Veeam Backup & Replication uses FUSE.

# Restoring Files from Veeam Backup Enterprise Manager

When you restore files from Veeam Backup Enterprise Manager, Veeam Backup & Replication creates a mount point on the mount server associated with the backup repository on which the backup file resides. Veeam Backup & Replication uses this mount point when the restore process starts and allows you to browse the VM file system and then restore files.

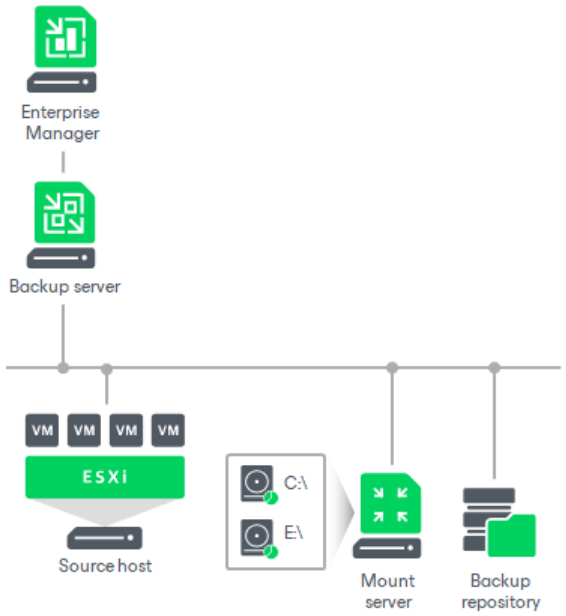


# Restoring Files from Nutanix AHV Snapshots

To learn how Veeam Backup & Replication performs guest OS file restore from backup snapshots, user snapshots and persistent disk snapshots created by AHV Backup Proxy and which mount points are used, see the Performing File-Level Restore section of the Veeam Backup for Nutanix AHV User Guide.

# Restoring Files from Amazon EC2 Instances

To perform guest OS file restore for EC2 instances, Veeam Backup & Replication utilizes the functionality of AWS Plug-in for Veeam Backup & Replication. To learn more how restore is performed, see the Performing File-Level Restore section of the Veeam Backup for AWS User Guide.

# Restoring Files from Microsoft Azure VMs

To perform guest OS file restore for Microsoft Azure VMs, Veeam Backup & Replication utilizes the functionality of Microsoft Azure Plug-in for Veeam Backup & Replication. To learn how restore is performed, see the Performing File-Level Recovery section of the Veeam Backup for Microsoft Azure User Guide.

# Restoring Files from Google Cloud VMs

To perform guest OS file restore for Google Cloud VMs, Veeam Backup & Replication utilizes the functionality of Google Cloud Plug-in for Veeam Backup & Replication. To learn how restore is performed, see the Performing File-Level Recovery section of the Veeam Backup for Google Cloud User Guide.

## Considerations and Limitations

Before you restore VM guest OS files, check the following considerations and limitations.

# Licensing

Restore changes functionality is included in the Veeam Universal License. When using a legacy socket-based license, the Enterprise or Enterprise Plus editions of Veeam Backup & Replication are required.

# Infrastructure Components

- You can restore files from the file systems listed in section Supported Platforms and Applications.
- The account that you use to start the Veeam Backup & Replication console and to connect to the backup server must have permissions and privileges described in section Veeam Backup & Replication Console Permissions.
- You can restore files from basic disks and dynamic disks (including simple, mirrored, striped, spanned and RAID5 volumes).
- [For restore to another VM, to original location, or permissions only] If the target VM uses the gMSA account and you restore files from a backup, you must also install this account on the mount server associated with the backup repository on which the backup resides. If you restore from a replica, you must install the gMSA account on the backup server.
- [For restore to original location] The mount server must have access to the VM guest OS (if restore is performed over the network) or vCenter Server and ESXi host where the target VM runs (if restore is performed over VIX API/vSphere Web Services).

# Restore Scenarios

If you restore files from backups or replicas using Veeam Backup & Replication or Veeam Backup Enterprise Manager and create a mount point on the mount server, consider the following:

- If the antivirus software installed on the mount server blocks or deletes from the mounted disk the objects you are trying to restore, Veeam Backup & Replication will not be able to access these objects. In this case, the restore session will be finished with the *Failed* state.
- To avoid such issues, you can add the `C:\VeeamFLR` folder to the antivirus exclusions.

# Source for Data Recovery

- You can restore VM guest OS files from a backup, VM replica, Cloud Director replica or CDP replica that has at least one successfully created restore point.
- You cannot restore files from a backup created in the reverse incremental mode if the backup job is being performed. If the backup is created in the incremental backup mode and the backup job is being performed, you can restore files from any available restore point.
- You can restore guest OS files from disks that use either the GPT or MBR partitioning scheme. Restore from disks without a partitioning scheme is not supported.
- You cannot restore pipes and other file system objects. Guest OS file restore supports recovery of files and folders only.

- You cannot restore guest OS files from a running VM replica or if the replication job with the necessary VM is being performed. However, restore is possible for CDP replicas if the CDP policy is running.
- You cannot restore guest OS files encrypted with Windows EFS.
- You cannot restore and browse guest OS files on disks encrypted by BitLocker.
- Processing of reparse points is supported only for NTFS. Note that reparse points with reparse tag values other than IO_REPARSE_TAG_MOUNT_POINT, IO_REPARSE_TAG_SYMLINK and IO_REPARSE_TAG_DEDUP may be processed and restored incorrectly.
- [For restore to original location] You cannot restore VM guest OS files if you have excluded the system disk from the VM backup used for restore and the volume GUID of the system disk was changed after the VM backup creation.
- [For comparison functionality and restore of permissions only] Check that VMware Tools are installed on the original machine and the machine is accessible over the network.
- You cannot use the comparison functionality if the Group Managed Service Account (gMSA) account is used for the VM whose files you plan to restore.
- The comparison functionality is not available for backups created by Veeam Backup for OLVM and RHV, for backups exported with Kasten policies and for backups stored in external repositories (for example, backups created by Veeam Backup for AWS, Veeam Backup for Microsoft Azure and so on).
- [For permission restore] Permissions can be restored only for files and folders that are still present on the original VMs. If files and folders are missing, restore fails.
- The comparison functionality uses Veeam Deployer Service. This service is a 32-bit service. During the comparison, the service converts some 64-bit objects in 32-bit objects. That is why such objects are shown as deleted in the Veeam Backup browser, for example, some objects in the `Windows` folder.

## CDP Replicas

- You can launch restore for a CDP replica if its CDP policy is currently running. CDP will continue working.
- During restore, the CDP policy does not create new long-term restore points and does not delete the existing ones. Short-term restore points are still created.
- You cannot launch guest OS file restore, SureReplica, application item restore and failover in parallel for one VM or replica.

## Target for Data Recovery

- [For restore to original location] VMware Tools must be installed on the target VM. Application-aware processing must be supported for the Microsoft Windows OS of the original machine. If this is not possible, you can use 1-click file-level restore or copy files to the selected folder and then move them to their original location.
- [For restore to another VM] You can restore items only to Microsoft Windows-based VMs. You can select a VM only within the same virtual infrastructure where the original VM resides. For example, if the original VM resides in VMware vSphere, you can select a VM that resides in VMware vSphere only.
- [For permission restore] Veeam Backup & Replication restores only permissions. If you restore a folder, permissions are restored for all subfolders and nested files. Attributes such as Read-only, Encrypted and so on are not restored.

## ReFS

- The machine on which a mount point is created (for example, the mount server) must run Microsoft Windows Server 2012 or later.
- The machine on which a mount point is created must support the same ReFS version or later than the version used on the VM from which you plan to restore files. For more information on which OSes support which ReFS, see ReFS versions and compatibility matrix.

To learn in which scenarios on which machines mount points are created, see Mount Points and Restore Scenarios.

## Data Deduplication

If you plan to restore files from a VM running Microsoft Windows Server 2012 or later and data deduplication is enabled for some VM volumes, consider the following:

- The machine on which a mount point is created (for example, the mount server) must run Microsoft Windows Server 2012 or later.
- The machine on which a mount point is created must run Microsoft Windows Server of the same version or later than the guest OS of a VM from which you plan to restore files.
- Data deduplication must be enabled on the machine on which a mount point is created.

To learn in which scenarios on which machines mount points can be created, see Mount Points and Restore Scenarios.

## Storage Snapshots

Requirements for guest OS file restore from storage snapshots are listed in the Data Recovery from Storage Snapshots section in the Storage System Snapshot Integration Guide.

### Restoring VM Guest OS Files (FAT, NTFS or ReFS)

To restore VM guest OS files and folders, use the **File Level Restore** wizard.

## Step 1. Launch File Level Restore Wizard

To launch the **File Level Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Guest files restore** > **Microsoft Windows** or **Restore from replica** > **Guest files restore** > **Microsoft Windows**.
- Open the **Home** view. In the inventory pane, select **Backups** or **Replicas**. In the working area, expand the necessary backup, click the VM whose files you want to restore and click **Guest files** > **Microsoft Windows** on the ribbon. Alternatively, right-click the VM whose files you want to restore and select **Restore guest files** > **Microsoft Windows**.

Alternatively for restore from regular backups, you can double-click the VBK or VBM file (for example, in Microsoft Windows Explorer). In the displayed window, select the VM and click **Restore** > **Guest files (Microsoft Windows)**. You can use this option if you perform restore on the backup server. You cannot use this option if you perform restore remotely over the Veeam Backup & Replication console.

Alternatively for restore from storage snapshots, you can open the **Storage Infrastructure** view. In the inventory pane, expand the storage system tree and select the necessary volume snapshot. In the working area, select a VM whose files you want to restore and click **Guest Files > Microsoft Windows** on the ribbon. You can also right-click a VM and select **Restore guest files > Microsoft Windows**.

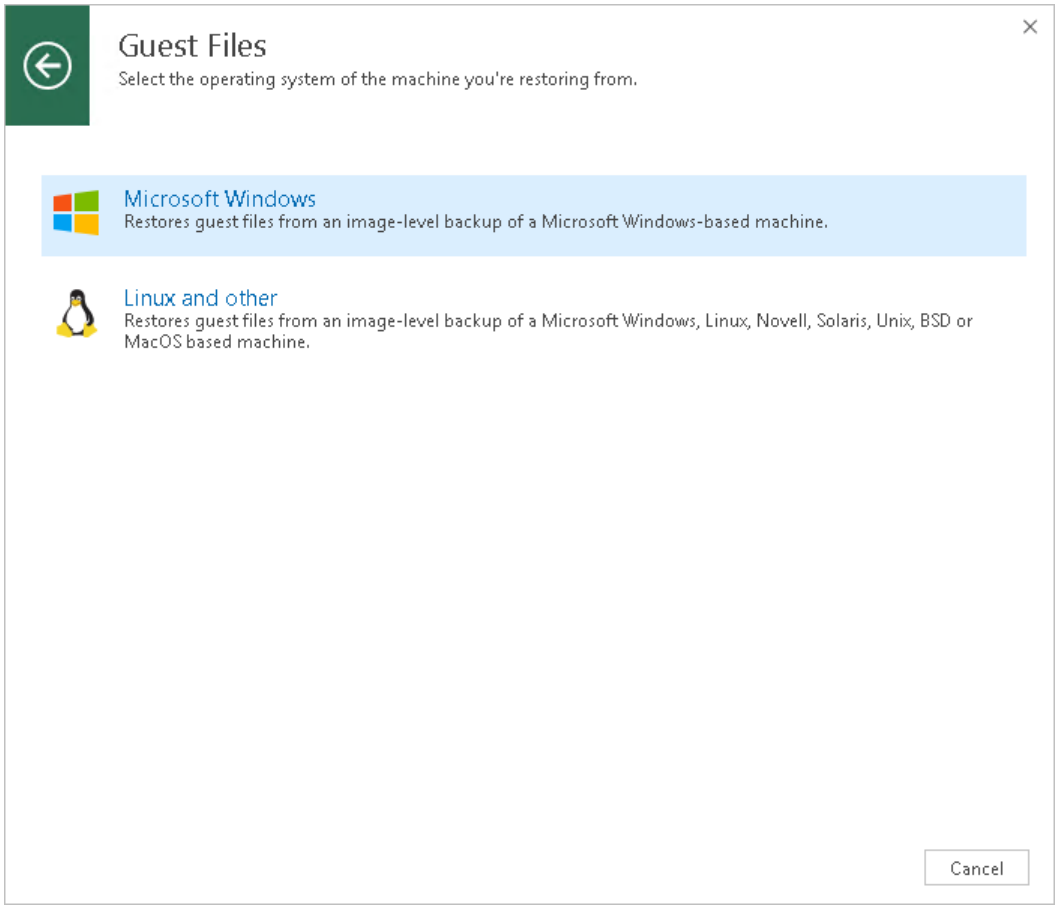

## Step 2. Select VM

At the **Machine** step of the wizard, select the VM whose guest OS files you want to restore:

1. In the **Machine** list, expand the necessary backup.
2. Select the VM.

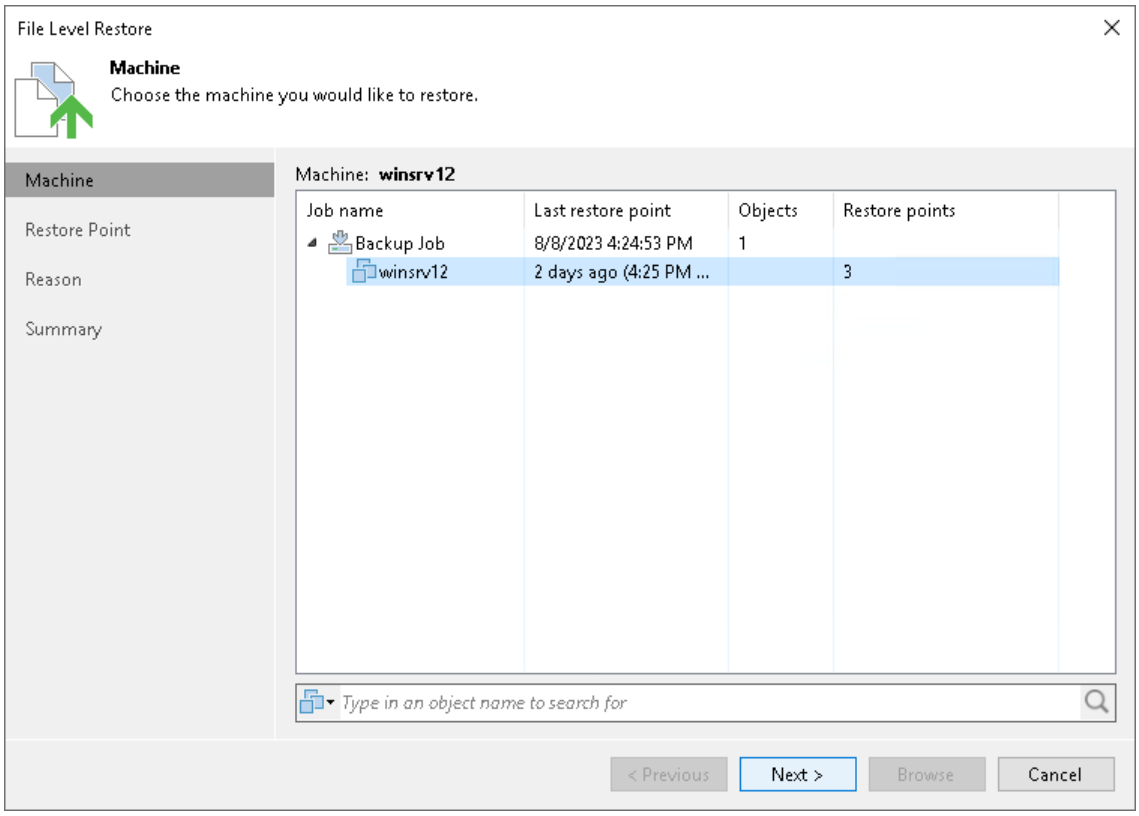

### Step 3. Select Restore Point

The UI of this step differs depending on the source for restore.

## Restore from Backups and VM Replicas

At the **Restore Point** step of the wizard, select a restore point from which you want to restore VM guest OS files.

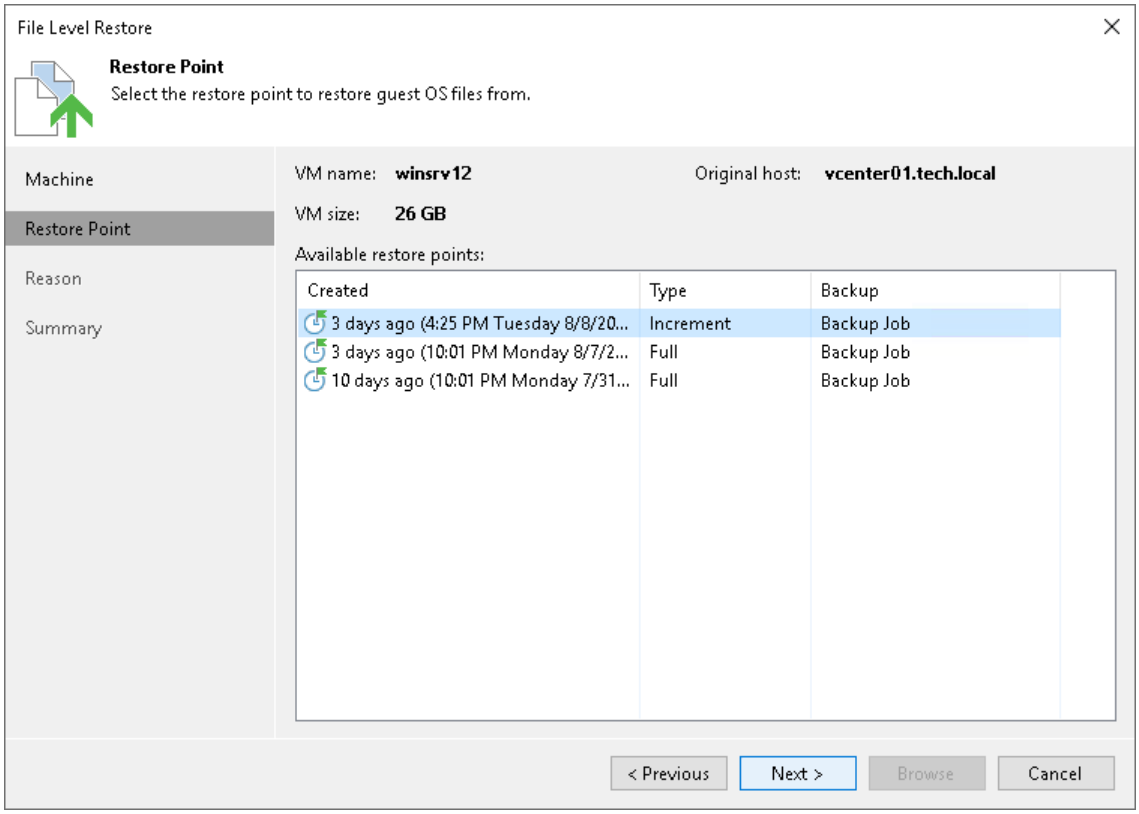


## Restore from CDP Replicas

At the **Restore Point** step of the wizard, select a short-term or long-term restore point from which you want to restore VM guest OS files. Use the right and left arrows on the keyboard to select the required restore point.

To restore to a short-term restore point, select a point in the green area. The darker the green, the more I/O load was produced on the source VM. To restore to a crash-consistent or application consistent long-term point, select a violet or turquoise vertical bar with a circle at the top.

To quickly find a long-term restore point, click a link that shows a date. In the opened window, you will see a calendar where you can select the necessary day. In the **Timestamp** section, you will see long-term restore points created during the selected day.

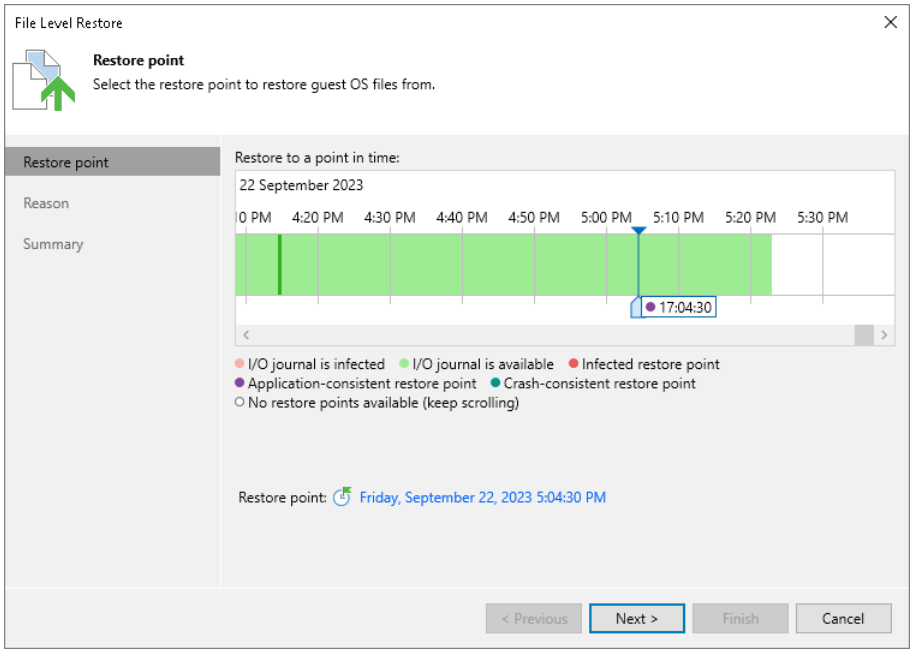

## Step 4. Select Host for Snapshot Mounting

This step is available if you restore files from storage snapshots.

At the **Staging Host** step of the wizard, select an ESXi host to which the clone/virtual copy of the storage snapshot will be mounted. On the selected ESXi host, Veeam Backup & Replication will create a temporary VM and mount disks of the restored VM to this temporary VM.

To specify destination for a snapshot clone/virtual copy and temporary VM:

1. At the **Staging Host** step of the wizard, click **Customize**.
2. Next to the **Host** field, click **Choose** and select an ESXi host to which the snapshot clone/virtual copy must be mounted and on which the temporary VM must be created.
3. Next to the **Resource pool** field, click **Choose** and select a resource pool to which you want to place the temporary VM.
4. Next to the **VM folder** field, click **Choose** and select a folder to which you want to place the temporary VM.

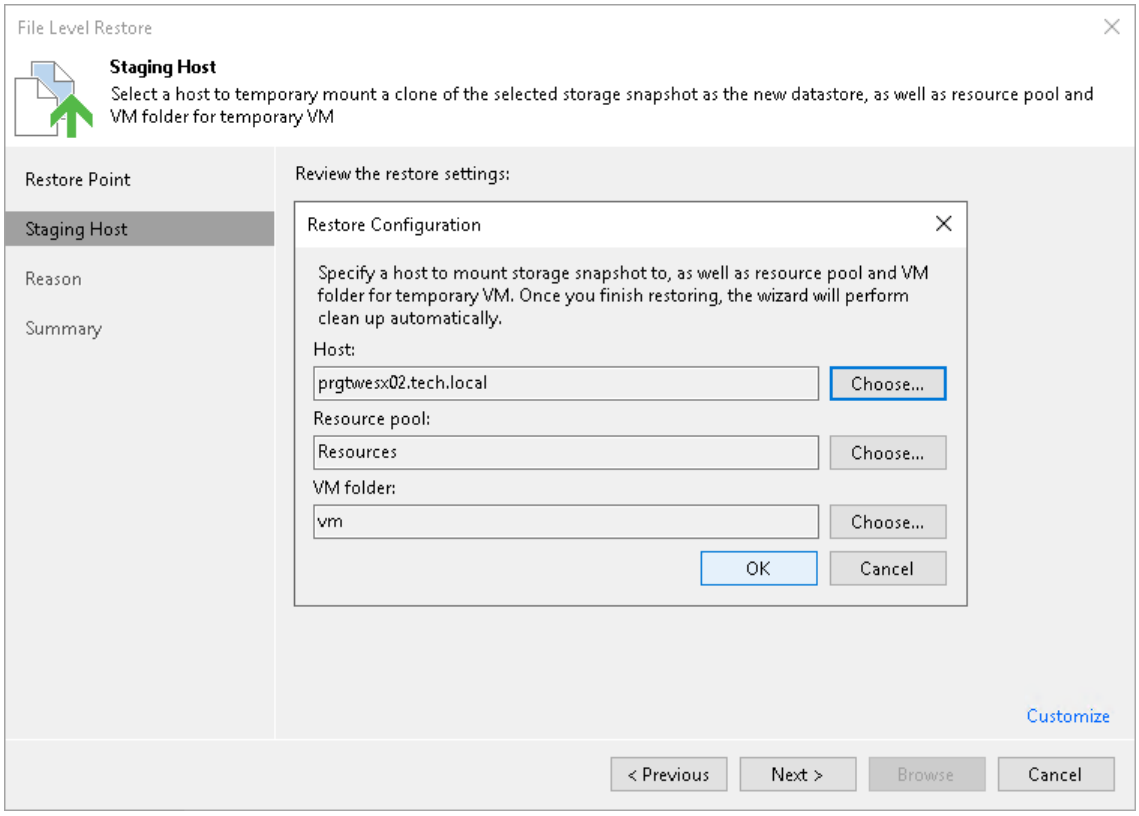

## Step 5. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring VM guest OS files. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

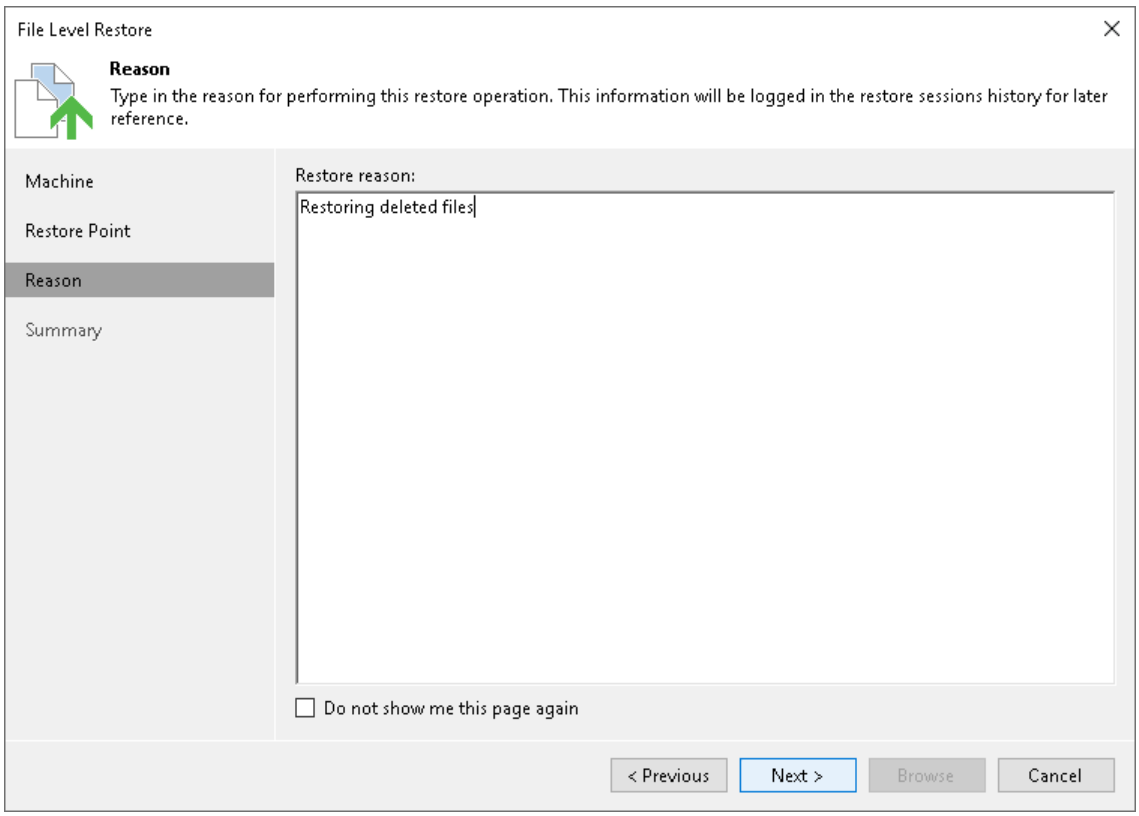

## Step 6. Verify Restore Settings

At the **Summary** step of the wizard, check details of the restore task and click **Browse** to close the wizard and open the Veeam Backup browser.

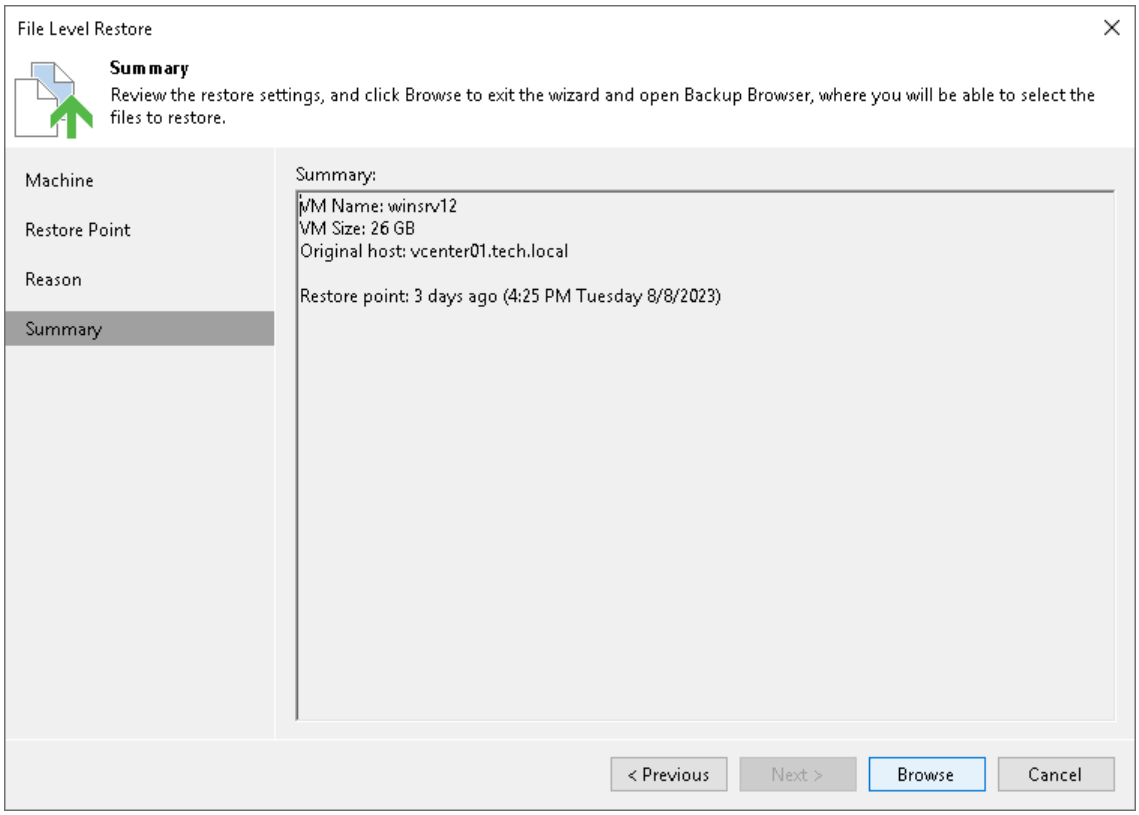

### Step 7. Finalize Restore

After you close the wizard, Veeam Backup & Replication opens the Veeam Backup browser with the file system tree of the restored VM.

You can perform the following operations in the Veeam Backup browser:

- Compare files and folders from a backup with the original files and folders (the **Compare** command).
- Restore only changed files and folders to the original location (the **Restore changed only** command).
- Copy files and folders to the Veeam Backup & Replication console or network shared folder (the **Copy to** command).
- Restore files and folders to another machine in the virtual infrastructure (the **Restore to** command).
- Restore files and folders to the original location (the **Restore** command).
- Restore permissions only (the **Permissions only** command).
- Launch application item restore (the **Application items** command).
- Open files in Microsoft Windows File Explorer.

After you finish restoring files, close the Veeam Backup browser.

> **NOTE**
>
> Consider the following:
>
> - To use the restore and comparison functionality, check Considerations and Limitations.
> - In the Veeam Backup browser, names of the restored VM hard disks may differ from the original ones.
> - Folder symbolic links are displayed under the icon.
> - When restoring symbolic links to the original location or to another machine, Veeam Backup & Replication restores only links, not the content they point to.
> - Veeam Backup & Replication restores junction points as folders with the content they refer to. If a junction point refers to the folder in which it lies or to the parent folder, restore will fail. For more information, see this Veeam KB article.
> - Hard links are displayed and restored as files.

# Comparing Files and Folders

When Veeam Backup & Replication compares files, it compares their attributes. When Veeam Backup & Replication compares folders, it compares folder attributes, and also folder and file attributes inside the folder recursively. Once a file or folder with changed attributes is found, Veeam Backup & Replication stops comparing items and marks the folder as changed. If you further browse the folder in the compared state, Veeam Backup & Replication continues comparing non-compared files and folders. Veeam Backup & Replication compares the following attributes: Date Created, Date Modified, Size (only for files), Read-Only, Hidden, Archive, NTFS Encryption.

To compare files and folders from a backup with the files and folders stored in the original location:

1. Select the necessary files and folders in the file system tree or in the details pane on the right. You can also select disks. In this case, Veeam Backup & Replication will compare files and folders stored on the disks.
2. Right-click one of the selected items and select **Compare** > **Compare**. Alternatively, click **Compare with Production** on the ribbon.

3. If prompted, in the **Credentials** window, specify user credentials to access the original location.

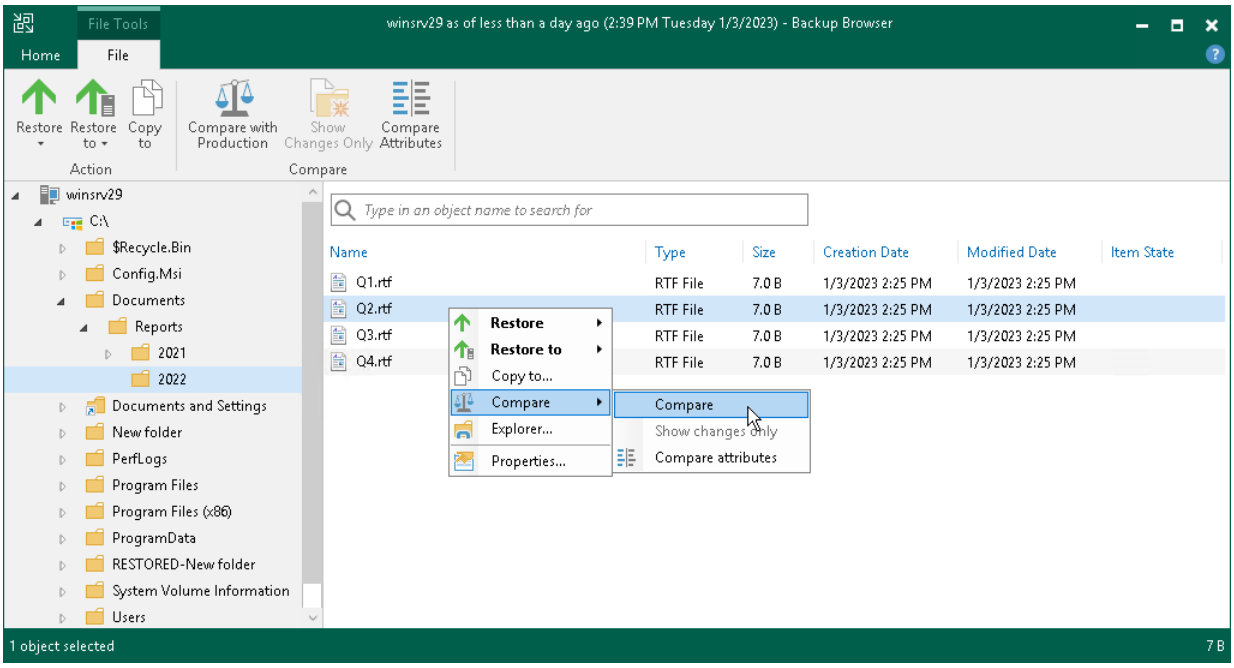


After the comparison, files and folders will have the following comparison states in the **Item State** column: *changed*, *unchanged*, *deleted*, *pending*, or *failed to compare*. The states are updated when you turn off and then turn on the comparison mode, and when you start restoring changes of files and folders. Note that when comparing symbolic links, Veeam Backup & Replication compares attributes of the links, not the attributes of files and folders which the symbolic link points to.

For files and folders in the comparison states, Veeam Backup & Replication provides other restore operations than for files and folders in the non-comparison state. For example, you can restore only changed files and folders. For more information, see Restoring Changed Files and Folders.

> **TIP**
>
> To show only changed files and folders (in the *changed* and *deleted* states), right-click any area in the Veeam Backup browser and select **Compare** > **Show changes only**. Alternatively, click **Show Changes Only** on the ribbon. Note that in this case the search field becomes unavailable. To show all files and folders, click the **Show changes only** or **Show Changes Only** option once again.
>
> To switch off the comparison states, select an item in the comparison state, click **Compare** > **Compare**. Alternatively, click **Compare with Production** on the ribbon. Note that if you switch off comparison for child files and folders, comparison for parent folders will also be switched off.

You can view which attributes were changed for files and folders:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected items and select **Compare > Compare attributes**. Alternatively, click **Compare Attributes** on the ribbon.

In the **Compare Attributes** window, Veeam Backup & Replication shows changed attributes. If you want to show all attributes, click the **Show unchanged attributes** check box at the top right corner. Note that Veeam Backup & Replication shows attributes maximum for 500 files and folders and shows attributes for the selected files and folders, not for the nested files.

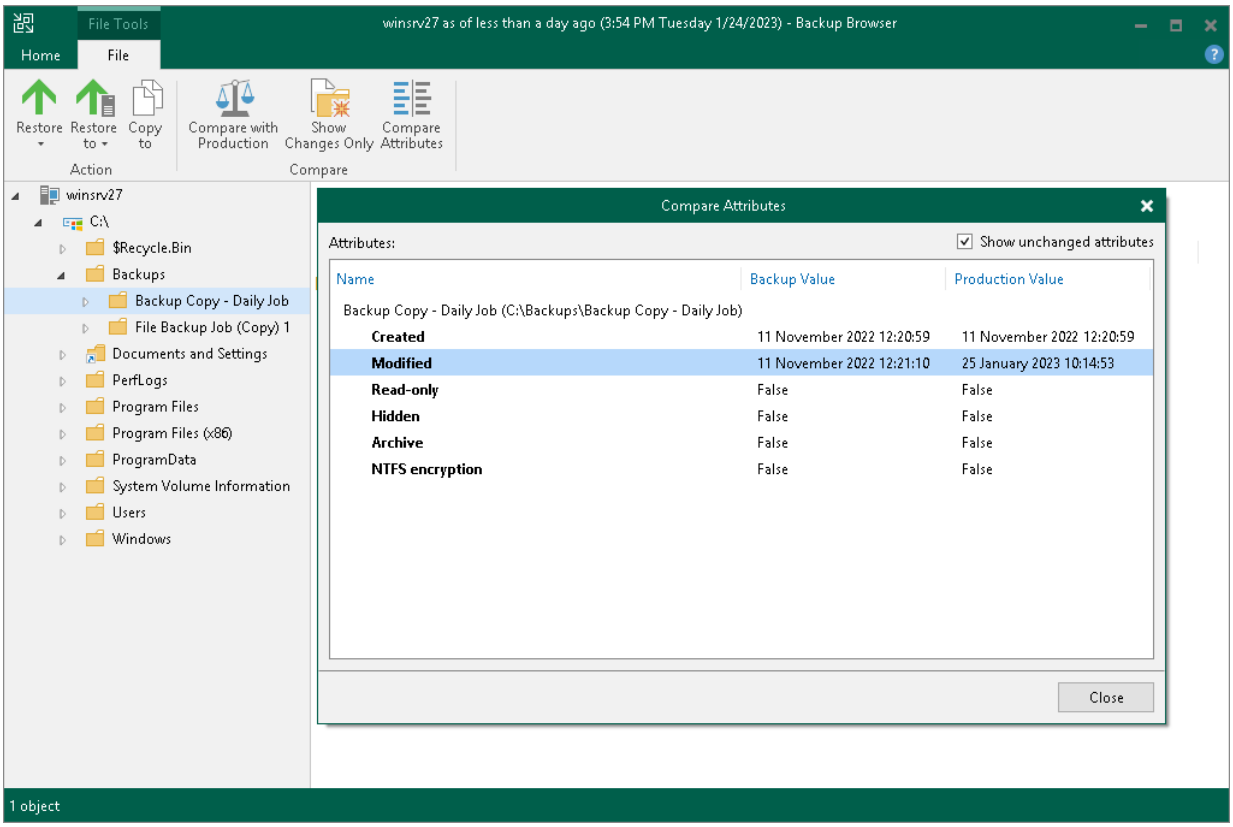


# Restoring Changed Files and Folders

To restore only changed files and folders to the original location, do the following:

1. Select the necessary files and folders in the file system tree or in the details pane on the right. Note that at least one file or folder must be in a comparison state. Files and folders in the non-comparison state, Veeam Backup & Replication will compare automatically.
2. Right-click one of the selected items and select one of the following:
   - To overwrite the original files and folders with the ones restored from the backup, select **Restore changed only** > **Overwrite**.
   - To save the files and folders restored from the backup next to the original ones, select **Restore changed only** > **Keep**.

     Veeam Backup & Replication will add the *RESTORED_YYYYMMDD_HHMMSS* postfix to the original names and store the restored items in the same folder where the original items reside.

   Alternatively, you can select the same commands on the ribbon.

If you want to restore entire files and folders to the original location, see Restoring to Original Location.

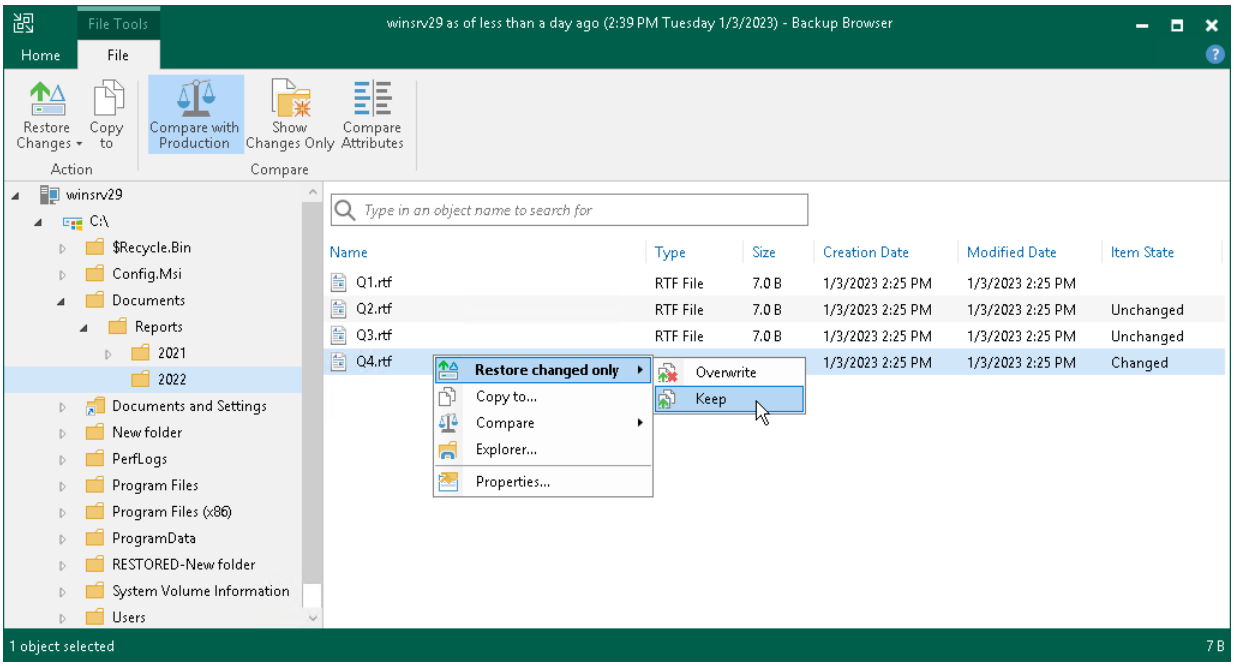


# Copying Files and Folders to Console or Shared Folder

To copy files and folders to the VM where the Veeam Backup & Replication console is installed or to a network shared folder:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected items and click **Copy to**. Alternatively, click **Copy to** on the ribbon.
3. In the **Choose Folder** window, select the necessary destination:
   - To recover files and folders to a folder on the VM where the Veeam Backup & Replication console is installed, click **Browse** to find the necessary folder.
   - To recover files and folders to a network shared folder, enter a path to the destination folder in the **Choose folder** field.
4. In the **Choose Folder** window, choose whether to preserve original NTFS permissions or not:
   - To keep the original ownership and security permissions for the restored items, select the **Preserve permissions and ownership** check box.

     Veeam Backup & Replication will copy selected files and folders along with associated Access Control Lists, preserving granular access settings.
   - If you do not want to preserve the original ownership and access settings for the restored items, leave the **Preserve permissions and ownership** check box not selected.

     Veeam Backup & Replication will change security settings: the user who launched the Veeam Backup & Replication console will be set as the owner of the restored item, while access permissions will be inherited from the folder to which the restored item is copied.
5. If prompted, in the **Credentials** window, specify user credentials to access the destination location.

NOTE

Consider the following:

- The copy to operation does not use the comparison states and copies all selected files and folders.
- When copying symbolic links, Veeam Backup & Replication copies the content which the links point to.

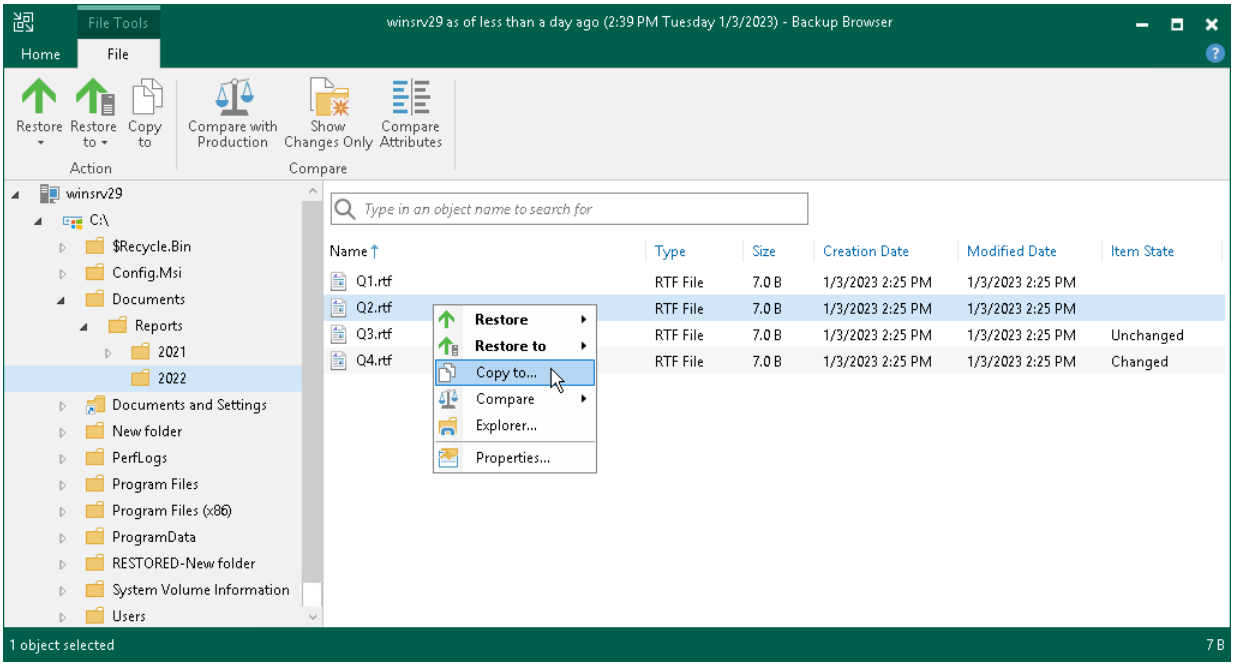


# Restoring Files and Folders to Another VM

This functionality is available for backups of VMware vSphere VMs and backups of VMware Cloud Director VMs created by Veeam Backup & Replication. Note that files and folders that you plan to restore must be in the non-comparison state.

To restore files and folders to a new location over the network or without the network, do the following:

3. Select the necessary files and folders in the file system tree or in the details pane on the right.
4. Right-click one of the selected items and select one of the following:
   - If you want to overwrite files and folders with identical names in the target location, select **Restore to** > **Overwrite**.
   - If you want to keep files and folders with identical names in the target location, select **Restore to** > **Keep**.

     If there are items with identical names, Veeam Backup & Replication will add the *RESTORED_YYYYMMDD_HHMMSS* postfix to the original names and store the restored items in the target location.

   Alternatively, you can select the same commands on the ribbon.
5. In the **Select Virtual Machine** window, select the target workload.
6. In the **Credentials** window, provide credentials to connect to the target workload.

7. In the **Choose Target Folder** window, click **Browse** and select a folder where items will be restored.

   Note that the browse option is not available for backups created by Veeam Agents.

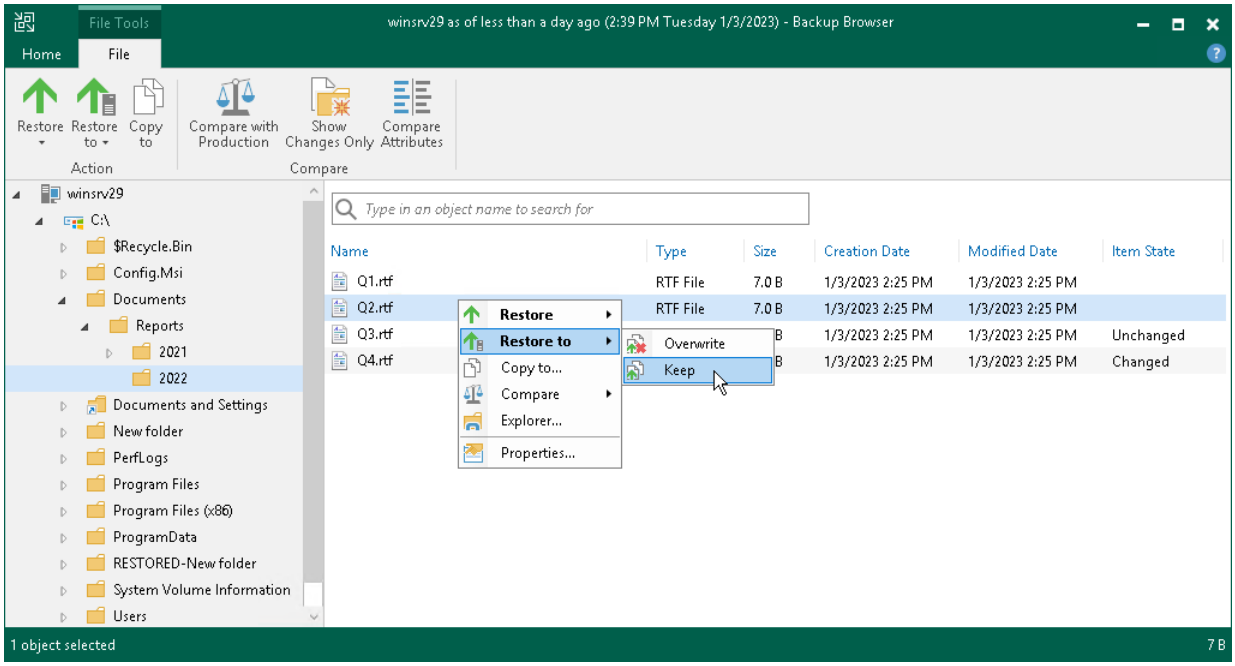


# Restoring Files and Folders to Original Location

This functionality is available for files and folders in the non-comparison state.

To restore files and folders to the original location, do the following:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected items and select one of the following:
   - To overwrite the original files and folders with the ones restored from the backup, select **Restore** > **Overwrite**.
   - To save the restored files and folders next to the original ones, select **Restore** > **Keep**.

     Veeam Backup & Replication will add the *RESTORED_YYYYMMDD_HHMMSS* postfix to the original names and store the restored items in the same folder where the original items reside.

   Alternatively, you can select the same commands on the ribbon.

Veeam Backup & Replication will restore all files and folders. If you want to restore changed files and folders only, see Restoring Changed Files and Folders.

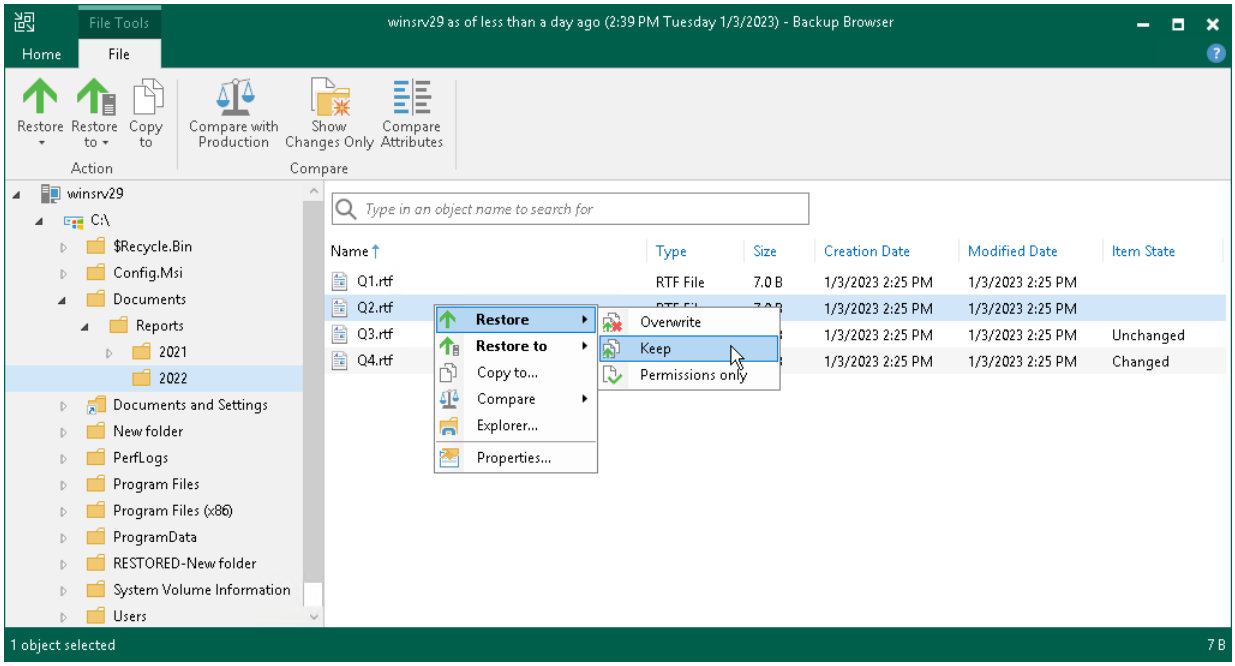


# Restoring Permissions

This functionality is available for files and folders in the non-comparison state.

To restore permissions for files and folders, do the following:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected items and select **Restore** > **Permissions only**. Alternatively, you can select **Restore** > **Permissions only** on the ribbon.

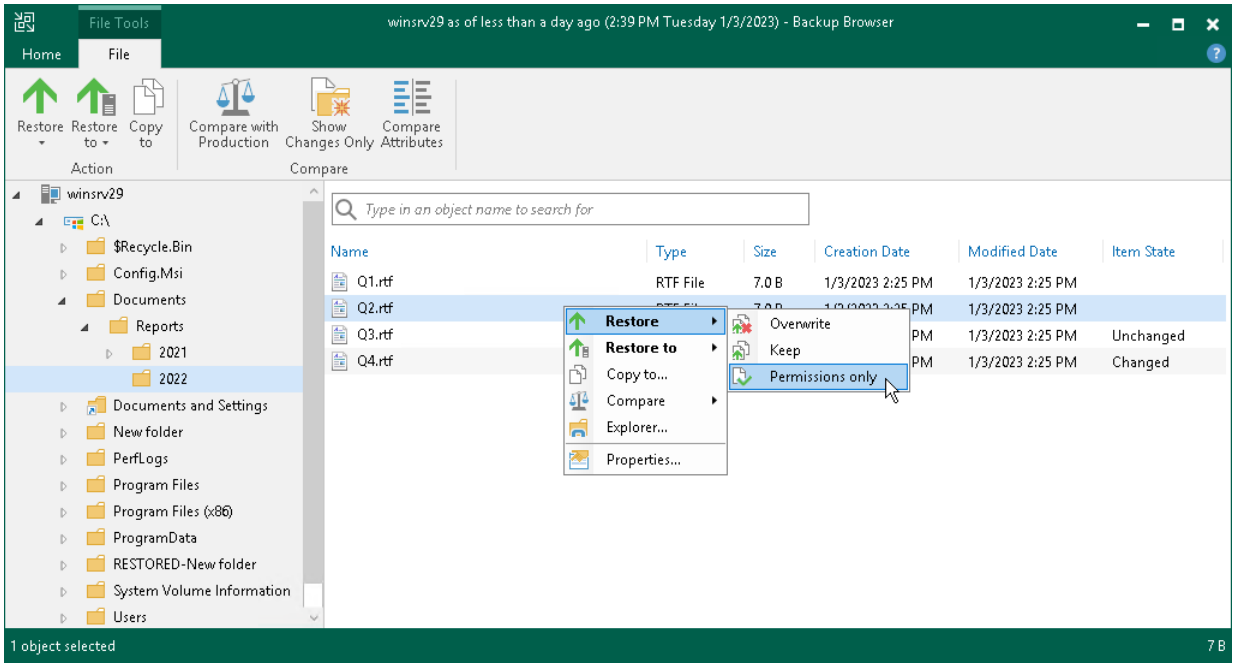

# Launching Application Item Restore

If you are restoring guest OS files from VMs where the supported applications (Microsoft Active Directory, Oracle and so on) are installed, you can also launch application item restore directly from the Veeam Backup browser. To restore application items, Veeam Backup & Replication uses special tools called Veeam Explorers.

To launch application item restore, do the following:

1. On the ribbon, switch to the **Home** tab.
2. Click **Application Items** and select the required application.
3. In the opened Veeam Explorer, perform the necessary operations. For more information on Veeam Explorers, see the Veeam Explorers User Guide.

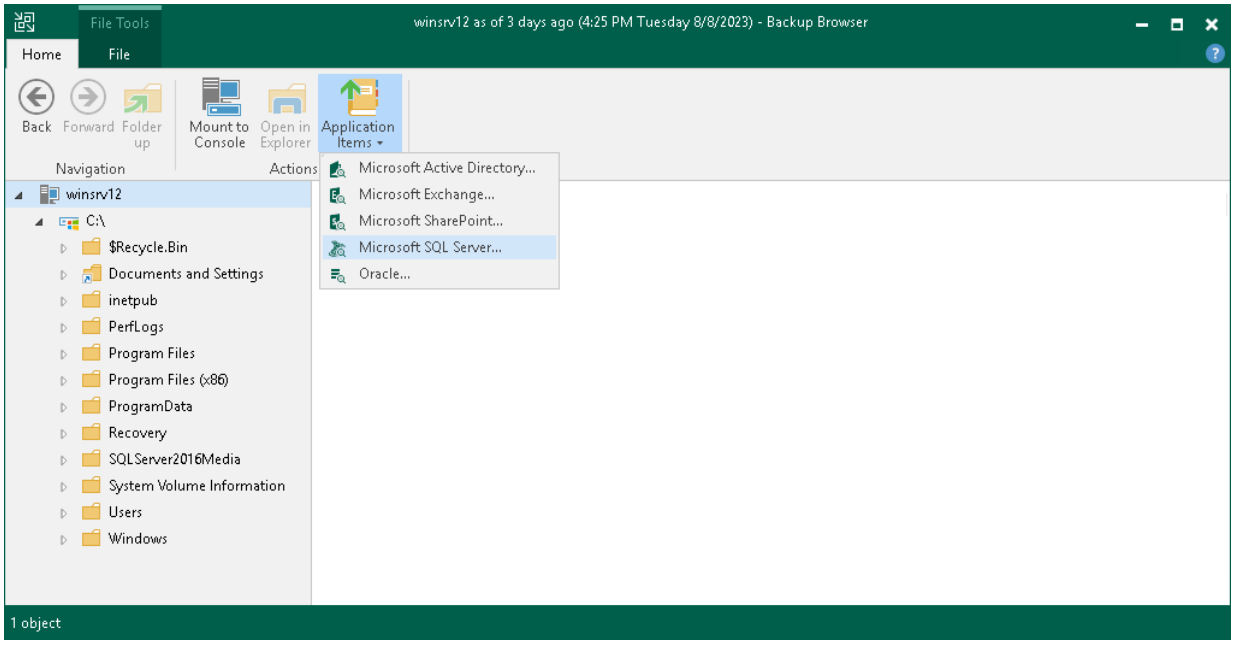


# Working with Microsoft Windows File Explorer

You can use Microsoft Windows File Explorer to work with restored files and folders:

1. On the ribbon of the Veeam Backup browser, switch to the **Home** tab and click **Mount to Console** to mount the VM disks to the Veeam Backup & Replication console.
2. To open Microsoft Windows File Explorer, do the following:
   - Click **Open in Explorer** on the Veeam Backup browser ribbon or right-click the necessary folder and select **Explorer**.
   - Click **File Explorer** in the **Start** menu of the VM where Veeam Backup & Replication console is installed. Browse to the `C:\VeeamFLR\<vmname>\<volume n>` folder where the VM disks are mounted and find the necessary files.

> **NOTE**
>
> The **Mount to Console** button is not available if the mount point is already created on the Veeam Backup & Replication console.

It is recommended that you use Microsoft Windows File Explorer only to view file content, not to restore files. For guest OS file restore, use Veeam Backup browser. This browser has the following advantages:

1. You can browse the VM guest OS file system ignoring the file system ACL settings.
2. You can preserve permissions and ownership during file-level restore.

If you open the VM file system in Microsoft Windows Explorer, these capabilities are not available. For more information, see Microsoft Docs.

# Closing Veeam Backup Browser

You can browse VM guest OS files only while the Veeam Backup browser is open. After the Veeam Backup browser is closed, Veeam Backup & Replication unmounts VM disks from the VM where the Veeam Backup & Replication console is installed and from the mount server (if you have restored VM guest OS files to the original location).

It is recommended that you close the Veeam Backup browser after you have finished restoring VM guest OS files. When the Veeam Backup browser is open, the backup file whose VM guest OS file system is displayed in the browser is locked in the backup repository. As a result, some scheduled operations that use this backup file may fail.

Veeam Backup & Replication checks if there is any activity in the Veeam Backup browser with an interval of 5 minutes. If the user or Veeam Backup & Replication components and services do not perform any actions for 30 minutes, Veeam Backup & Replication displays a warning that the Veeam Backup browser is to be closed in 5 minutes.

After the warning is displayed, you can perform one of the following actions:

- You can close the Veeam Backup browser manually.
- You can click **Cancel** to postpone the close operation. In this case, the Veeam Backup browser will remain open for 5 minutes. After this period expires, Veeam Backup & Replication will display the warning again.
- You can perform no action at all. In this case, the Veeam Backup browser will close automatically in 5 minutes.

# Restore from Linux, Unix and Other File Systems

The method of restore from Linux, Unix and other file systems helps you restore files of Linux, Solaris, BSD, Novell Storage Services, Unix and Mac workloads. For the list of supported OSes and file systems, see Supported Platforms and Applications.

You can restore files from the following types of data:

- Backups
- Replicas
- Storage snapshots

  For more information on how restore from storage snapshots works, see the Linux, Unix and Other File System Restore from Storage Snapshots section in the Storage System Snapshot Integration Guide.

You can restore files to the original or new location, access restored files using FTP or launch application item restore for the files. For more information, see Finalize Restore.

# How Restore Works

When you perform guest OS file restore, Veeam Backup & Replication provides the following options for mounting VM disks from the backup or replica:

- Mounting disks to a helper host. As a helper host, you can select the target host where you want to restore files from the backup or any other Linux server. We recommend you to specify the same server to which you want to restore the files. This will improve the performance.
- Mounting disks to a helper appliance. The helper appliance is a helper VM running a stripped down Linux kernel that has minimal set of components. The appliance is quite small — around 50 MB. It requires 2048 MB RAM and 2 CPU.

When you perform guest OS file restore, Veeam Backup & Replication performs the following operations:

1. [If you have selected to mount disks to helper appliance] Veeam Backup & Replication deploys the helper appliance on the ESXi host in the virtual infrastructure.
2. Veeam Backup & Replication mounts disks of a VM from the backup or replica to the host selected as a helper host or helper appliance. The backup file or VM replica itself remains in the read-only state in the backup repository or datastore.
3. Veeam Backup & Replication launches the Veeam Backup browser where mounted VM disks are displayed. You can browse the VM guest file system in the Veeam Backup browser and restore files or folders to the original VM or to another location. Also, you can enable an FTP server on the virtual appliance and allow VM owners to restore files themselves.
4. The operations differ depending on which restore command you use:
   - The **Restore** or **Restore to** command to restore files to the original location or to another VMware vSphere VM.

     If you have installed the Linux Management Agent, the helper host or helper appliance connects to the VM to which you restore files (target VM) using this agent and restores files. This is the preferred way of connection.

     If the backup server fails to connect to the Management Agent, the backup server connects to the target VM over SSH. If the SSH connection also fails, Veeam Backup & Replication uses networkless processing over VIX API/vSphere Web Services. Then Veeam Backup & Replication deploys on the VM a temporary agent which performs restore.

     > **NOTE**
     >
     > The backup server uses SSH to connect to the target machine and perform the restore if the following conditions are met:
     >
     > - You use the **Restore to** command.
     > - The target machine is a member of a protection group for pre-installed Veeam Agents.

   - The **Copy to** command to restore files to a new location.

     The helper host or helper appliance connects to the VM to which you restore files (target VM) over SSH. Then Veeam Backup & Replication deploys on the VM the agent which performs restore.
5. When you close the Veeam Backup browser or it is closed by timeout, Veeam Backup & Replication unmounts the content of the backup file or replica from the helper appliance or helper host.
6. [If you have selected to mount disks to helper appliance] Veeam Backup & Replication unregisters the helper appliance on the ESXi host.

## Considerations and Limitations

Before you restore VM guest OS files, check the following considerations and limitations.

# Infrastructure Components

- Check the supported file systems. For details, see the File-Level Restore section of Supported Platforms and Applications.
- [For source and target VM] Veeam Backup & Replication uses the ICMP `ping` command to define whether a VM is available over network. If the VM must be available over the network, check that ICMP protocol is enabled on the VM.
- Veeam Backup & Replication must have access over the network to the guest OS of the target VM or direct access to the vCenter or ESXi host where the target VM resides to deploy a coordination process. The coordination process performs a number of administrative actions on the target VM guest OS, for example, collects information about mount points.
- The mount server, helper host and helper appliance, must have access over the network to a VM whose files you restore or direct access to vCenter or ESXi host where the VM resides. If a connection over the network cannot be established, the mount server connects to the VM over VIX API/vSphere Web Services. In this case, you must use a root account for the target VM and check that `/tmp` directory on the target VM is mounted with the `exec` option. Otherwise the restore process will fail.

  If you use the FQDN of the ESXi host in the helper appliance configuration window, the helper appliance must be able to resolve the FQDN of the ESXi host.

  If your DHCP does not provide the DNS configuration, assign the DNS server address manually using the `[HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication\DnsServersList]` registry value. This value accepts data of the REG_SZ type, for example, 10.11.12.13. If you want to specify multiple addresses, separate them using a semicolon. After you specify the registry value, restart the restore process.
- [For helper appliance] If you recover data from a CDP replica, the helper appliance must be deployed on an ESXi host or cluster where the I/O filter is installed.
- 
- You can restore from Novell Storage Services (NSS) file system if you use FUSE or a helper appliance. When using FUSE, the helper host must differ from the backed-up VM. When using the helper appliance, you can perform restore only in the IPv4 environment.

# Helper Host

If you mount disks to a helper host, consider the following:

- Check the supported OSes listed in System Requirements.
- The backup server and the mount server, associated with the repository where the backup is stored, must be able to resolve the FQDN of the helper host.
- You can restore from ZFS using the helper host option if the `zfsutils-linux` package is installed on the specified helper host. The package `zfs-fuse` package is not supported.
- The helper host OS kernel must support the file system that you plan to mount on this host. Otherwise, the mount will be refused and in rare cases it may cause kernel panic.

- Make sure that the `/tmp` directory is mounted with the `exec` option. Otherwise, you will get an error with the permission denial.
- For the helper host option, mounting of LVM snapshots is not supported. Thus, LVM snapshots are skipped from processing.
- [For backups created by Veeam Agent for Linux] The original machine cannot be used as the helper host for the 32-bit versions of the Linux-based OSes. However, for VM backups, you can still use the original VM as the helper host for the 32-bit versions of the Linux-based OSes.
- Restore from a Btrfs disk using the original host as the helper host is not possible. The issue occurs due to restriction of mounting two Btrfs disks with identical IDs to the same machine. To avoid this issue, use the helper appliance option or a different Linux-based server that supports Btrfs. However, note that Btrfs disks can be restored for backups created by Veeam Agent for Linux. During the backup process, Veeam Agent for Linux changes disk IDs in the backup file.
- Restore from a ZFS pool using the original host as the helper host is not possible. The issue occurs due to restriction of mounting two ZFS pools with identical UUIDs to the same machine. To avoid this issue, use the helper appliance option or a different Linux-based server that supports ZFS pools.
- Hardened repositories cannot be selected as helper hosts.

# Source for Data Recovery

- You can restore VM guest OS files from a backup, VM replica, Cloud Director replica or CDP replica that has at least one successfully created restore point.
- You can restore files whose names are written in all locales in the UTF-8 encoding. If the encoding is other than UTF-8, you can restore only files whose names are written in the English locale.
- The multi-OS guest OS file restore wizard does not support restore of deduplicated volumes (for example, Microsoft Windows volumes with data deduplication enabled).
- You cannot restore files from a backup created in the reverse incremental mode if the backup job is being performed. If the backup is created in the incremental backup mode and the backup job is being performed, you can restore files from any available restore point.
- You cannot restore guest OS files from a running VM replica or if the replication job with the necessary VM is being performed. However, restore is possible for CDP replicas if the CDP policy is running.
- [For backups of BSD, macOS and Solaris VMs] You cannot restore files directly to the original location. Use the **Copy to** option instead.
- [For backups of Linux workloads] You can restore files from basic disks, Linux LVM (Logical Volume Manager) and ZFS pools. Encrypted, RAID1 and mirrored LVM volumes are not supported.
- [For backups of Linux workloads] RAID mounts and restores are not supported.
- [If you restore from backups with guest file system indexing disabled] To properly show the file system tree in the Veeam Backup browser, check that the `/etc/fstab` file lists disk UUIDs or labels. Disks listed using device names are shown outside the file system tree as separate disks.

# CDP Replicas

- You can launch restore for a CDP replica if its CDP policy is currently running. CDP will continue working.
- During restore, the CDP policy does not create new long-term restore points and does not delete the existing ones. Short-term restore points are still created.

- You cannot launch guest OS file restore, SureReplica, application item restore and failover in parallel for one VM or replica.

# Target for Data Recovery

- If you plan to restore VM guest OS files to their original location or to another VMware vSphere VM, VMware Tools or OpenVM Tools must be installed and running on the target VM.
- The following applies if guest processing components are not installed on the target VM:
  - Make sure that the `/tmp` directory on the target VM is mounted with the `exec` option. Otherwise, you will get an error with the permission denial.
  - If you restore files using the account for which only the Elevate account privileges automatically check box is selected, the `/home/<username>` directory also must be mounted with the `exec` option.
- [For restore to a new location] You can restore items only to Linux-based VMs. BSD, macOS and Solaris VMs are not supported.
- [For restore to a new location] If you want to restore files over network, make sure that the SSH daemon is configured and SCP utility is available on the target VM.
- [For restore to a new location] Veeam Backup & Replication can restore ACL for recovered VM guest OS files. To let Veeam Backup & Replication detect the target Linux system architecture and kernel version, the following utilities must be present in the minimal configuration of the system: *arch* and *uname*.
- For backups created by Veeam Agents, you can use the Restore to (restore to a new location) command only if the target computer is available over the network. This applies to backups created by Veeam Agent for Linux, Veeam Agent for IBM AIX or Veeam Agent for Oracle Solaris.

# Linux Firewalls in Helper Host and Target for Data Recovery

If you want to use the helper host option or if you want to restore files to a new location on another Linux VM, consider the following.

Veeam Backup & Replication automatically opens ports used for the restore process on the helper host and on the target Linux VM. Generally, Veeam Backup & Replication automatically opens ports for most popular firewalls (iptables, ufw, firewall-cmd). However, if for some reason the ports are not opened, you can open the ports manually. You can also specify these ports at the SSH Connection step of the **New Linux Server** wizard. Note that ports are opened dynamically: if 10 concurrent jobs are running, Veeam Backup & Replication opens ports 2500-2509.

If you use the `firewalld` tool, you can configure firewall rules to open ports only in necessary zones. By default, Veeam Backup & Replication opens ports in all active `firewalld` zones. If your firewall is configured for different zones, and you want to minimize security holes, you can configure Veeam Backup & Replication to open the ports only for certain zones. To do this, perform the following:

1. On the helper host or target Linux host, create the `/etc/VeeamNetConfig` file and define the following parameter:

   ```
   FirewalldZones=zone_name_1, zone_name_2
   ```

   where `zone_name_1, zone_name_2` is a list of zone names where the ports must be open. Veeam Backup & Replication will skip the zones that are not in this list.

2. [Only for helper host] If you select a Linux host that is already added to the Veeam Backup & Replication infrastructure, you should also add required zones to the `/opt/veeam/transport/VeeamTransportConfig` file.

```
FirewalldZones=zone_name_1, zone_name_2
```

**NOTE**

Veeam Backup & Replication opens the port 2500 in all zones even if you have specified the required zones in configuration files.

# Storage Snapshots

Requirements for guest OS file restore from storage snapshots are listed in the Data Recovery from Storage Snapshots section in the Storage System Snapshot Integration Guide.

## Restoring VM Guest OS Files (Multi-OS)

To restore VM guest OS files and folders from Linux or Unix-based machines, use the **Guest File Restore** wizard.

## Step 1. Launch Guest File Restore Wizard

To launch the **Guest File Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from backup** > **Guest files restore** > **Linux and other**.
- Open the Home view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup, click the VM whose files you want to restore and click **Guest files** > **Linux and other** on the ribbon. Alternatively, right-click the VM whose files you want to restore and select **Restore guest files** > **Linux and other**.

Alternatively for restore from regular backups, you can double-click the VBK or VBM file (for example, in Microsoft Windows Explorer). In the displayed window, select the VM and click **Restore** > **Guest files (Linux and other)**. You can use this option if you perform restore on the backup server. You cannot use this option if you perform restore remotely over the Veeam Backup & Replication console.

Alternatively for restore from storage snapshots, you can open the **Storage Infrastructure** view. In the inventory pane, expand the storage system tree and select the necessary volume snapshot. In the working area, select a VM whose files you want to restore and click **Guest Files > Linux and other** on the ribbon. You can also right-click a VM and select **Restore guest files > Linux and other**.

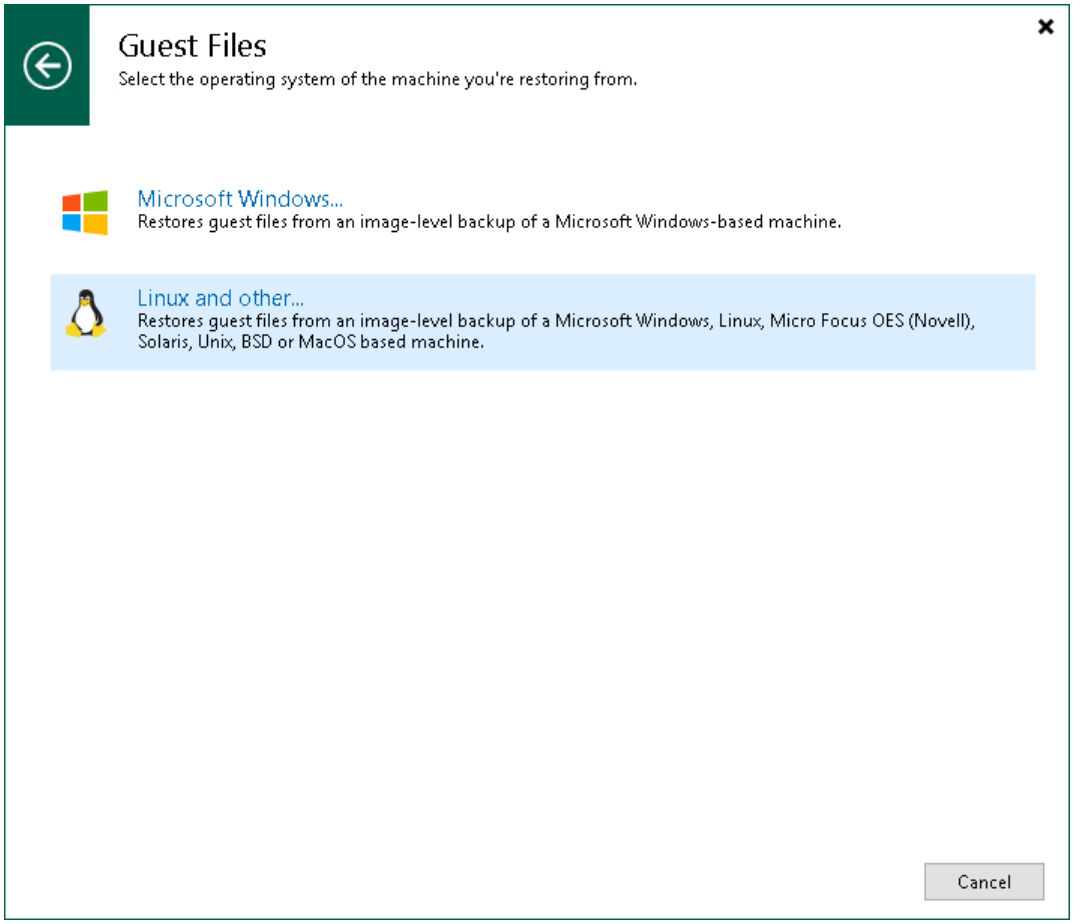

## Step 2. Select VM

At the **Virtual Machine** step of the wizard, select the VM whose guest OS files you want to restore:

1. In the **Virtual machine** list, expand the necessary backup.
2. Select the VM.

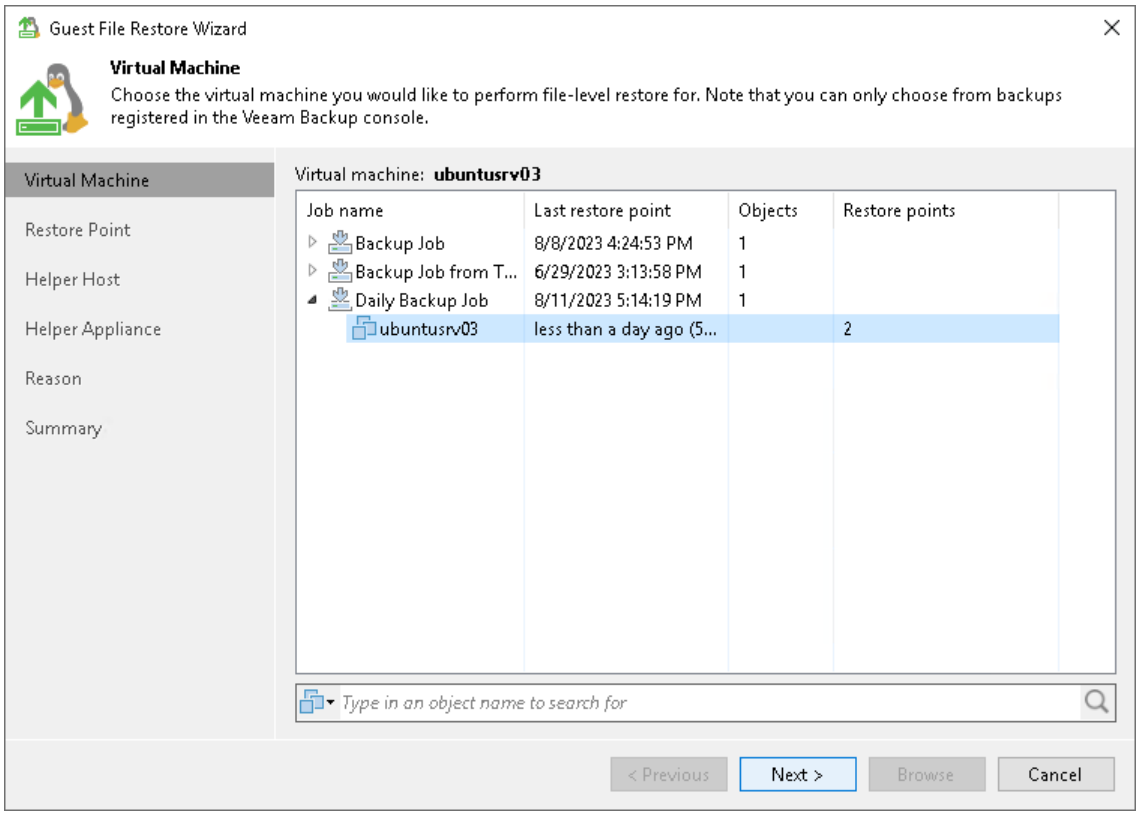

### Step 3. Select Restore Point

The UI of this step differs depending on the source for restore.

## Restore from Backups and VM Replicas

At the **Restore Point** step of the wizard, select the restore point from which you want to restore VM guest OS files.

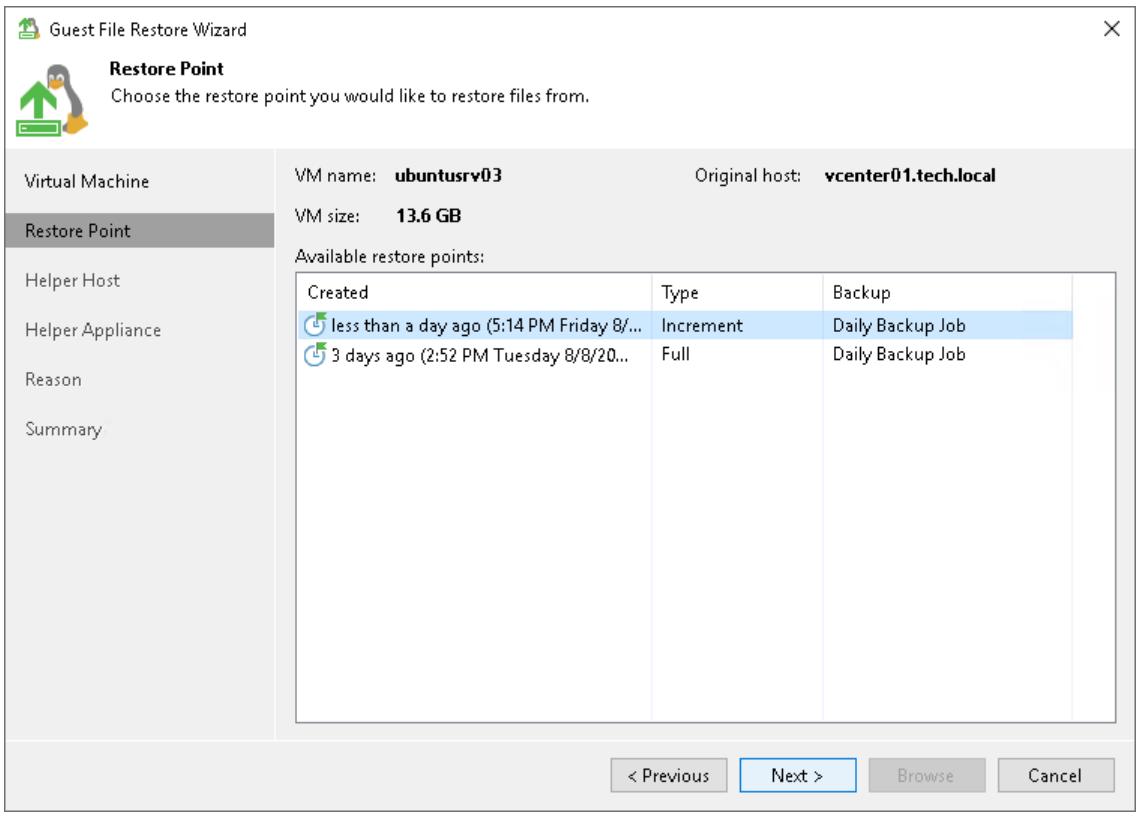


## Restore from CDP Replicas

At the **Restore Point** step of the wizard, select a short-term or long-term restore point from which you want to restore VM guest OS files. Use the right and left arrows on the keyboard to select the required restore point.

To restore to a short-term restore point, select a point in the green area. The darker the green, the more I/O load was produced on the source VM. To restore to a crash-consistent or application consistent long-term point, select a violet or turquoise vertical bar with a circle at the top.

To quickly find a long-term restore point, click a link that shows a date. In the opened window, you will see a calendar where you can select the necessary day. In the **Timestamp** section, you will see long-term restore points created during the selected day.

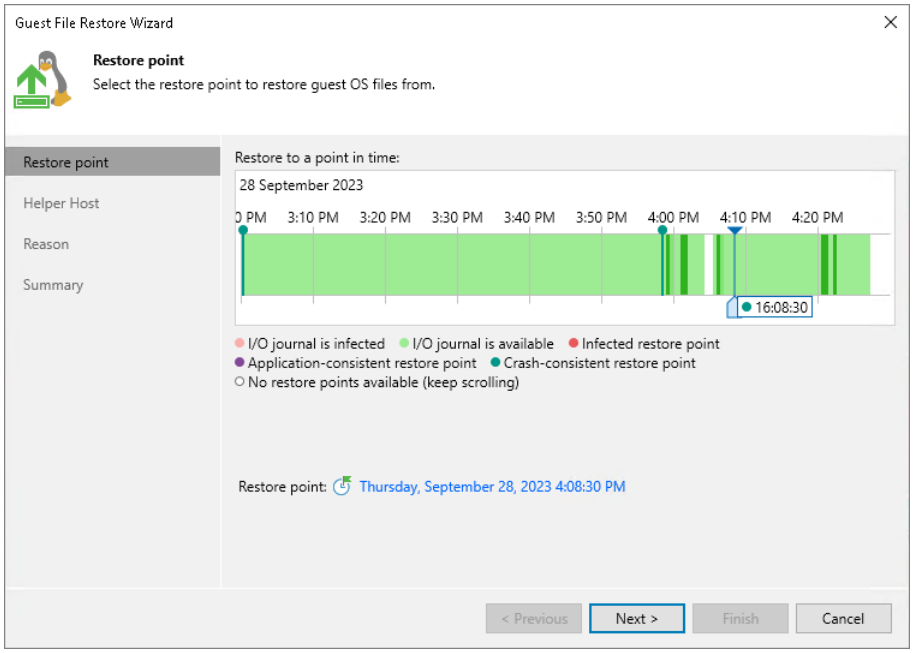

## Step 4. Configure Staging Host

This step of the wizard is available only if you restore VMs from storage snapshots.

At the **Staging Host** step of the wizard, select an ESXi host to which the clone or virtual copy of the storage snapshot will be mounted. On the selected ESXi host, Veeam Backup & Replication will create a temporary VM.

1. Click the **Customize** link at the bottom of the window.
2. In the **Host** field of the **Restore Configuration** window, select an ESXi host to which the snapshot clone must be mounted.
3. In the **Resource pool** field, specify a resource pool to which you want to place the temporary VM.
4. In the **VM folder** field, specify a VM folder to which you want to place the temporary VM.

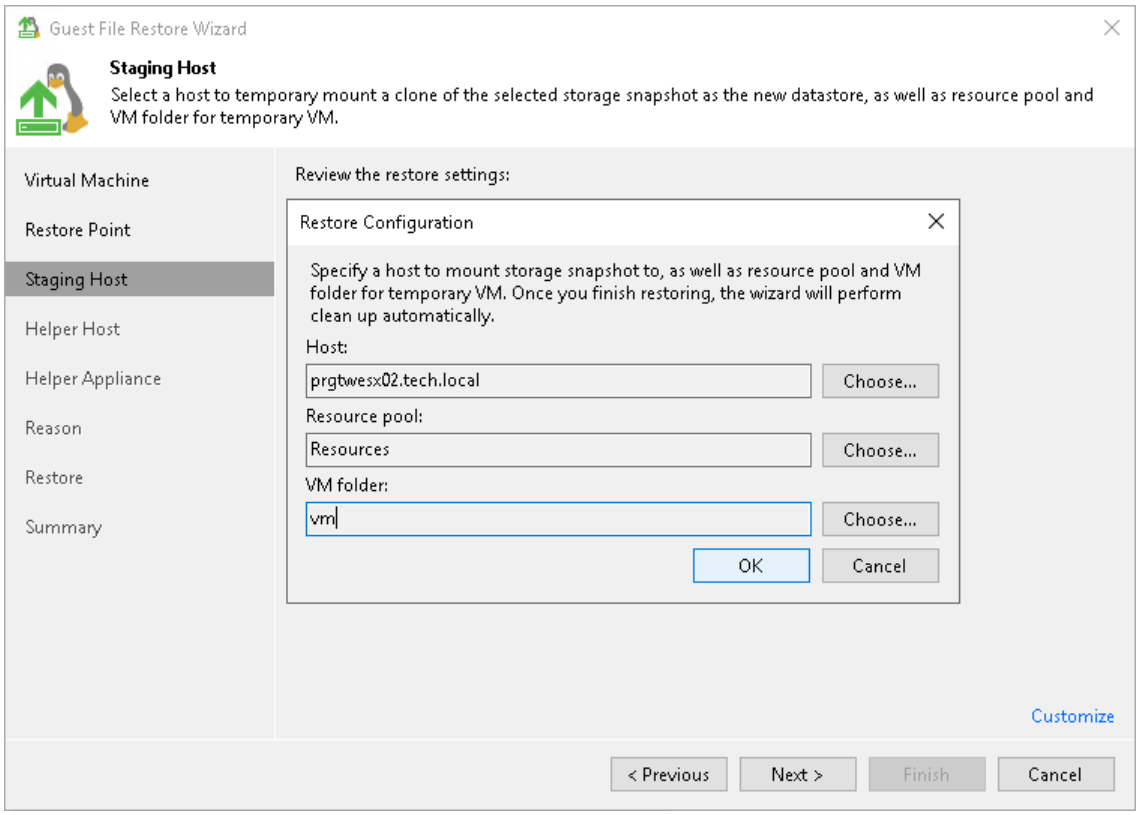

## Step 5. Specify Helper Host

At the **Helper Host** step of the wizard, you can select a machine that will be used as a helper host. Veeam Backup & Replication will mount disks of a machine from the backup or replica to the selected helper host. Alternatively, you can choose to configure a new helper appliance where Veeam Backup & Replication will mount VM disks. For details on helper hosts and appliances, see Restore from Linux, Unix and Other File Systems.

To specify a helper host, do the following:

1. From the **Linux server** list, select one of the following options:
   - **<Hostname> (original host)** — to specify the original Linux server as the helper host.
   - **<Hostname>** — to specify a Linux server as the helper host. The wizard displays all Linux servers added to the Veeam Backup & Replication infrastructure.
   - **Specify a different host** — to specify a Linux server that is not added to the Veeam Backup & Replication infrastructure.
   - **Use a temporary helper appliance** — to use a temporary helper appliance. If you select this option, the wizard will include a step for configuring the helper appliance.
2. [For **<Hostname> (original host)** option; optional] In the **Host credentials** window, specify credentials for the original Linux server.
3. [For **Specify different host** option] In the **Linux Server** window, specify the helper host name and connection setting:
   a. In the **Host name** field, specify the IPv4 or IPv6 address or the host name of the Linux server which will be used as a helper host. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
   b. From the **Credentials** list, select an account that will be used to connect to the helper host.

      If you have not added the account before, click the **Add** button on the right of the **Credentials** list and follow the instructions described in section Linux Accounts.
   c. If you want to change the default SSH settings, click **Advanced** and configure SSH settings as described in section Specify Credentials and SSH Settings.

d. Click **OK** to save the helper host settings.

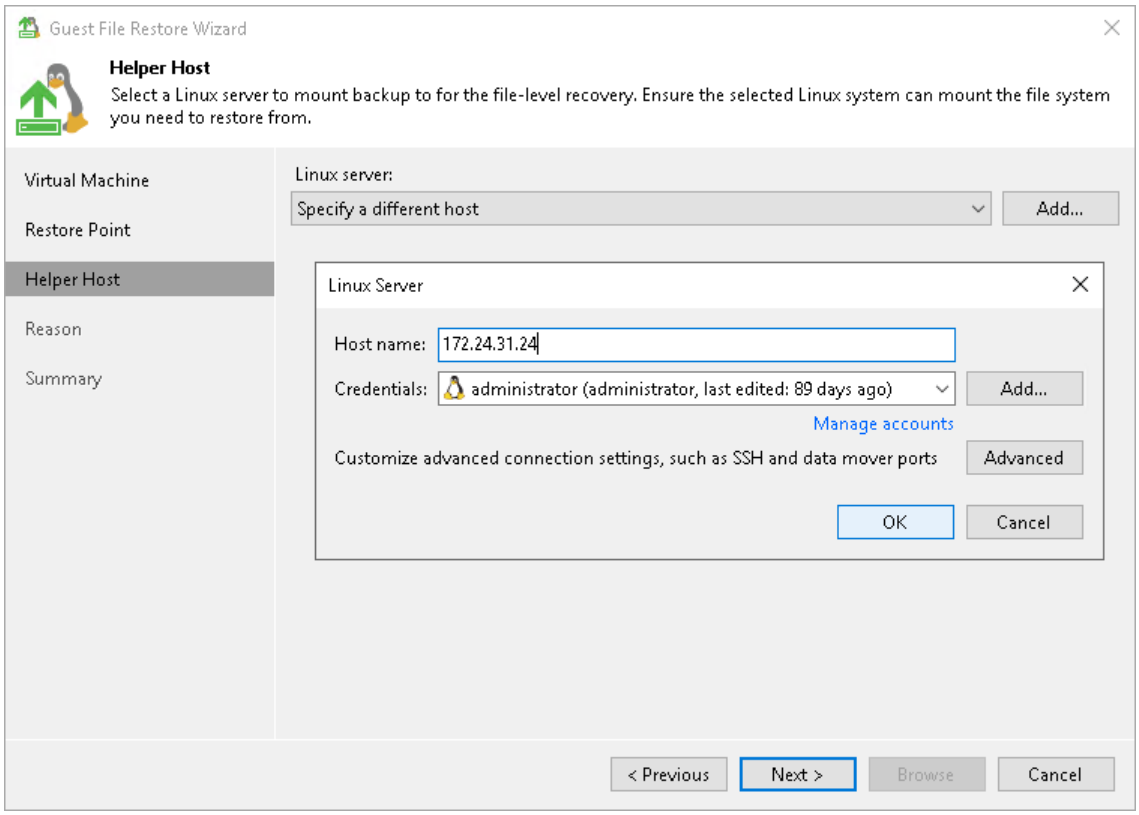

## Step 6. Specify Location for Helper Appliance

The **Helper Appliance** step of the wizard is available only if you have selected the **Use a temporary helper appliance** option at the Specify Helper Host step. This step is not available if you recover VMware vSphere VMs from backups or replicas.

At this step of the wizard, you specify an ESXi host to which you want to place the helper appliance.

To select a destination for the helper appliance:

1. At the bottom of the window, click **Customize**.
2. In the **Host** field, specify the ESXi host on which the helper appliance must be registered.
3. In the **Resource pool** field, specify a resource pool to which the helper appliance must be placed.

   > **NOTE**
   >
   > If you perform restore from storage snapshots, the Host and Resource pool fields cannot be changed. Their values are taken from the Configure Staging Host step of the wizard.

4. Select a network for the helper appliance:
   a. On the right of the **Network** field, click **Choose**.

      In the **Select Network** window, Veeam Backup & Replication will display a list of networks to which the specified host is connected.

   b. From the **Networks** list, select a network to which the helper appliance must be connected and click **OK**.

      Consider that the backup server and the mount server must have access to the helper appliance over the network.

5. Specify IP addressing settings for the helper appliance and DNS server:
   a. Click **Configure**.
   b. Switch to the **IPv4** or **IPv6** tab depending on which addresses you want to configure. Note that you can use IPv6 addresses only if IPv6 communication is enabled as described in section IPv6 Support.
   c. Configure IP settings for the helper appliance:
      - If you use a DHCP server in the network and want to obtain the IP address automatically, leave the **Obtain an IP address automatically** option selected.
      - To manually assign a specific IP address to the helper appliance, click **Use the following IP address** and specify the IP address settings.

   d. Configure IP settings for the DNS server:
      - If you use a DHCP server in the network and want to obtain the IP address automatically, leave the **Obtain DNS server address automatically** option selected.
      - To manually assign a specific IP address to the DNS server, click **Use the following DNS server address** and specify preferred and alternate addresses.

   e. Click **OK**.
6. To enable FTP access to the restored file system, select the **Enable FTP server on appliance (advanced)** check box. As a result, users will be able to access the helper appliance over FTP, browse the file system of the restored VM and download necessary files on their own.

7. If you are performing restore of a VM with the Novell Storage Services file system, select the **Restore from Novell Storage Services (NSS) file system** check box. Veeam Backup & Replication will deploy a specific appliance that supports the Novell Storage Services file system.

**IMPORTANT**

Consider the following:

- When choosing an ESXi host for the helper appliance used for file-level restore from the Novell Storage Services file system, make sure that it allows running VMs with 64-bit guest OSes.
- If you restore files from a non-NSS file system, check that the **Restore from Novell Storage Services (NSS) file system** check box is not selected. Otherwise, the restore process may work incorrectly.

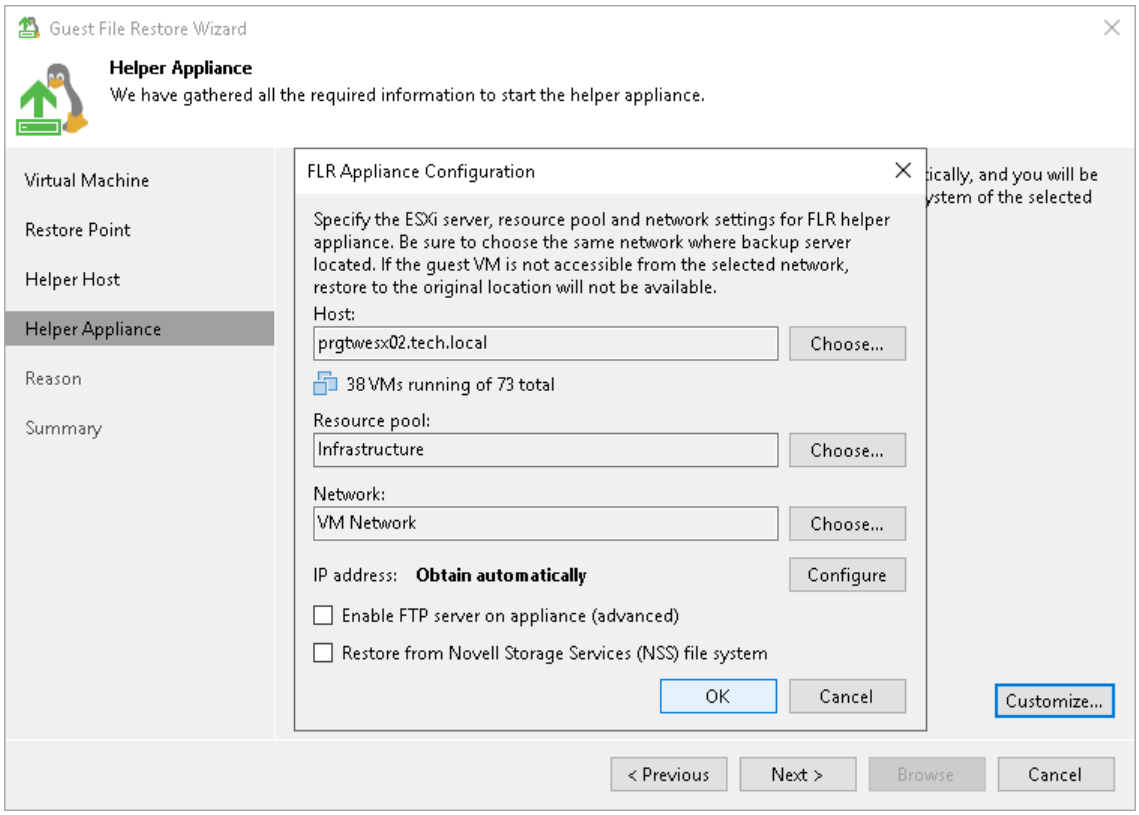

## Step 7. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring VM guest OS files. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

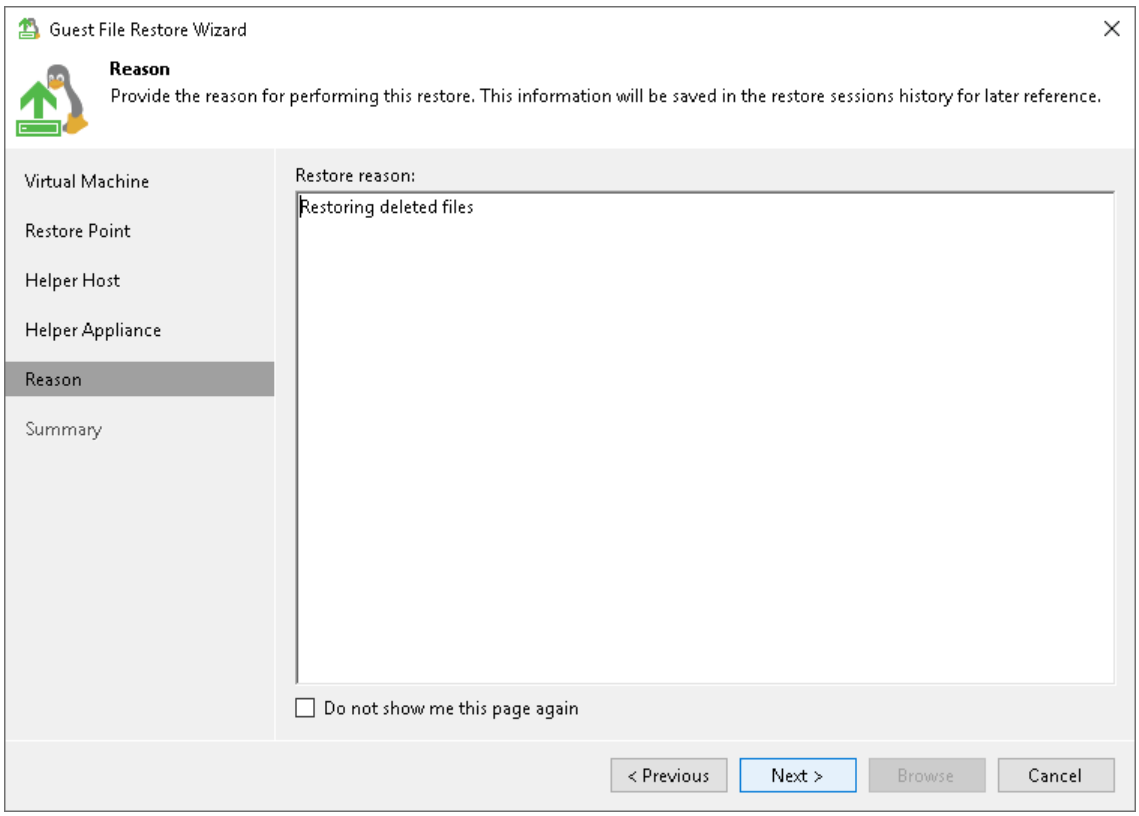

## Step 8. Start Mount Process

This step of the wizard is available only if you restore VMs from storage snapshots.

At the **Restore** step of the wizard, wait while Veeam Backup & Replication mounts VM disks. Then click **Next**.

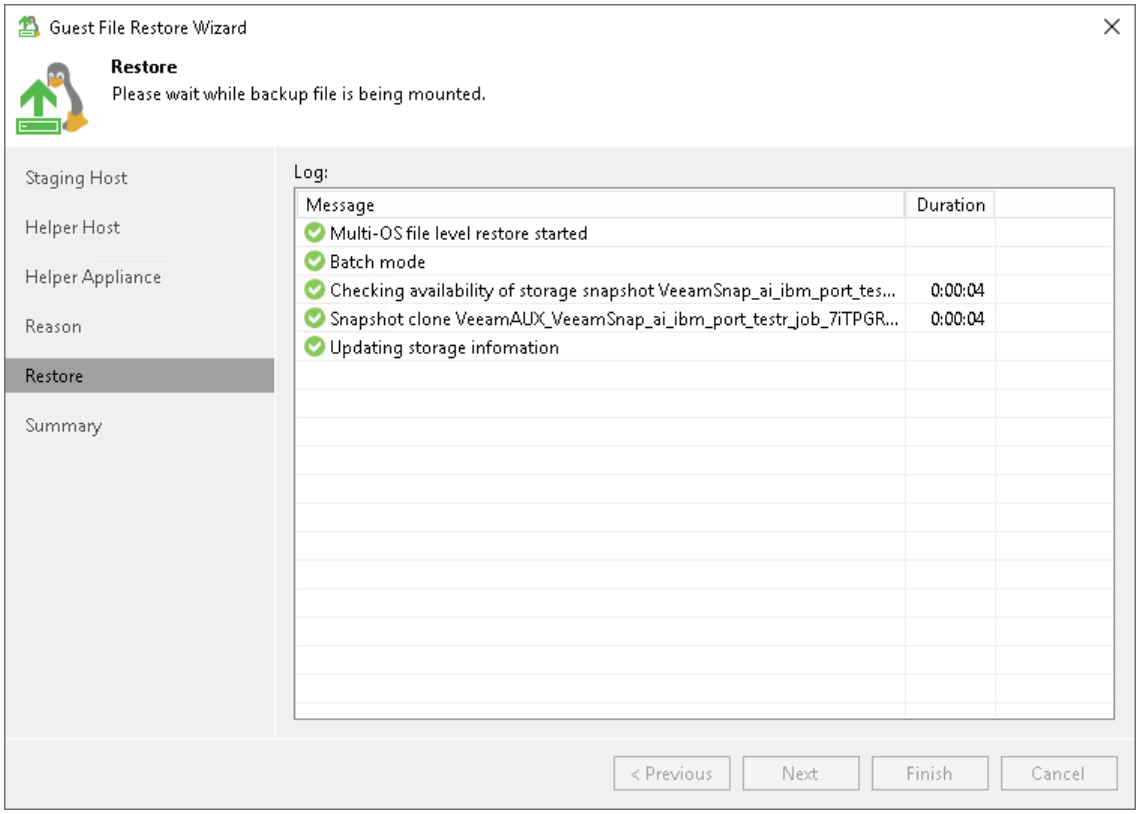

## Step 9. Verify Restore Settings

At the **Summary** step of the wizard, review the restore settings and click **Browse** to open the Veeam Backup browser.

If you have selected to mount disks to a helper appliance, it may take about 10-40 seconds to boot the helper appliance and open the browser.

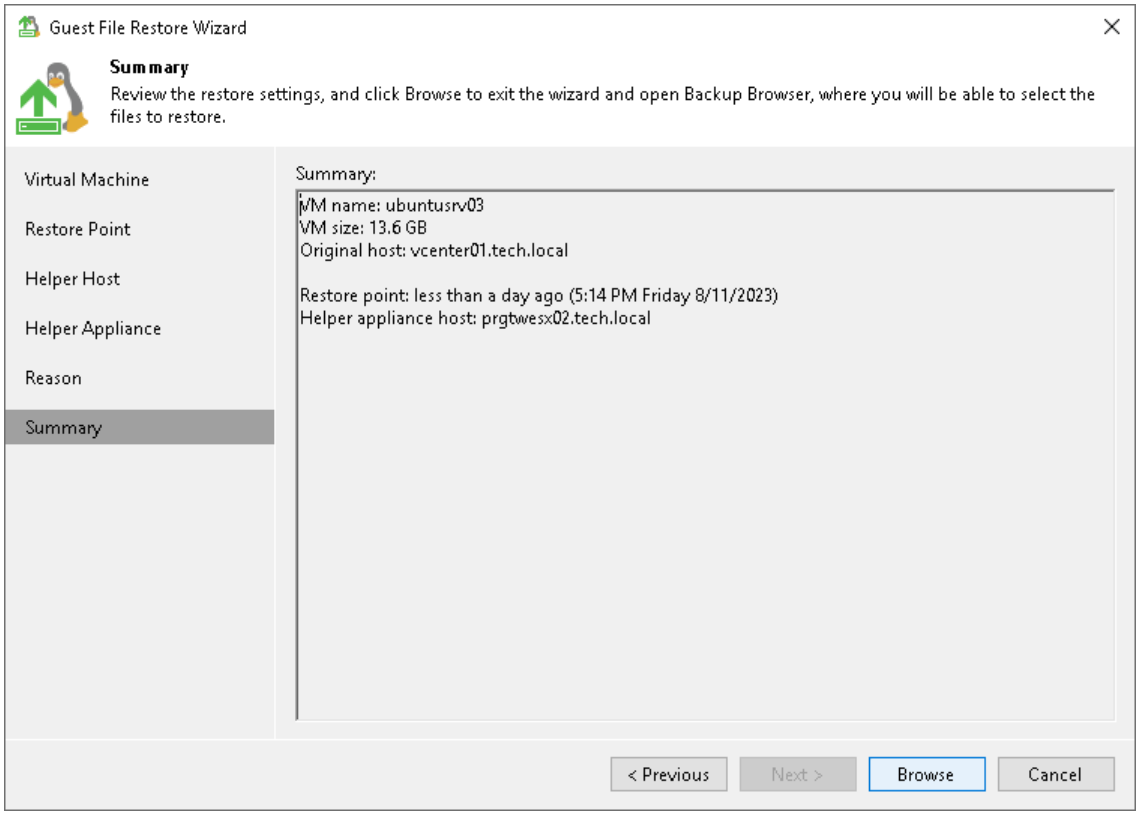

### Step 10. Finalize Restore

After the wizard is closed, Veeam Backup & Replication opens the Veeam Backup browser displaying the file system tree of the restored VM.

In the Veeam Backup browser, you can perform the following operations:

- Restore files to the original location (the **Restore** command)
- [For VMware vSphere, Cloud Director, Veeam Agent for Linux] Restore files to a new location (the **Restore to** command)
- Restore files to a new location over the network (the **Copy to** command)
- Access files over FTP
- Access helper appliance logs

> **NOTE**
>
> You can browse the VM guest OS files and access restored files on the FTP only while the Veeam Backup browser with the restored files is open. After the Veeam Backup browser is closed, Veeam Backup & Replication unmounts the VM disks from the helper appliance and removes the helper appliance from the ESXi host.

# Restoring Files to Original Location

To restore files and folders to the original location, do the following:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected items and select one of the following:
   - To overwrite the original files and folders with the ones restored from the backup, select **Restore** > **Overwrite**.
   - To save the restored files and folders next to the original ones, select **Restore** > **Keep**.

     Veeam Backup & Replication will add the *RESTORED_YYYYMMDD_HHMMSS* postfix to the original names and store the restored items in the same folder where the original items reside.

   Alternatively, you can select the same commands on the ribbon.

To restore files to the original location, Veeam Backup & Replication can use the following:

- The account for VM guest OS access specified in the backup job settings.
- The account specified for the original server used as the helper host.
- The credentials specified in the Protection Group.

If the account does not have sufficient rights to access the target VM, you will be prompted to enter credentials. In the **Credentials window**, specify a user account to access the destination location (server or shared folder). For more information on adding credentials, see Credentials Manager.

In some cases, you may remove the original VM and restore it from the backup by the time of guest OS file restore. If you then attempt to restore VM guest OS files to the original location, Veeam Backup & Replication will not be able to find the original VM by its reference ID, and will display a warning. Click **OK** and browse to the target VM in the virtual infrastructure to which you want to restore VM guest OS files.

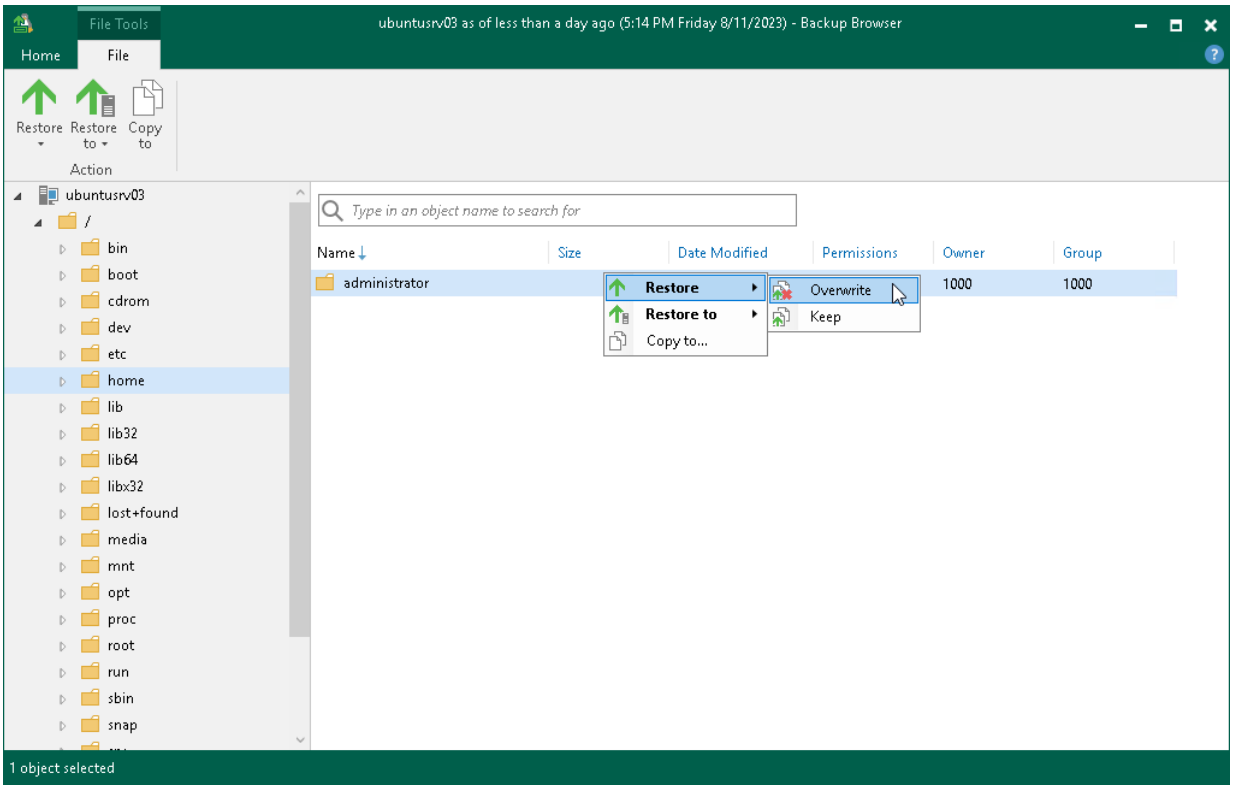


# Restoring Files and Folders to New Location

You can restore files and folders to a new machine of the same platform over the network or without the network. This functionality is available for the following types of backups:

- Backups of VMware vSphere VMs.
- Backups of VMware Cloud Director VMs created by Veeam Backup & Replication.
- Backups created by Veeam Agent for Linux, Veeam Agent for IBM AIX or Veeam Agent for Oracle Solaris.

To restore files and folders to a new location over the network or without the network, do the following:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected items and select one of the following:
   - If you want to overwrite files and folders with identical names in the target location, select **Restore to > Overwrite**.
   - If you want to keep files and folders with identical names in the target location, select **Restore to > Keep.**

     If there are items with identical names, Veeam Backup & Replication will add the *RESTORED_YYYYMMDD_HHMMSS* postfix to the original names and store the restored items in the target location.

   Alternatively, you can select the same commands on the ribbon.
3. In the **Select Virtual Machine** window, select the target workload.

4. In the **Credentials** window, provide credentials to connect to the target workload.
5. In the **Choose Target Folder** window, click **Browse** and select a folder where items will be restored.

   Note that the browse option is not available for backups created by Veeam Agents.

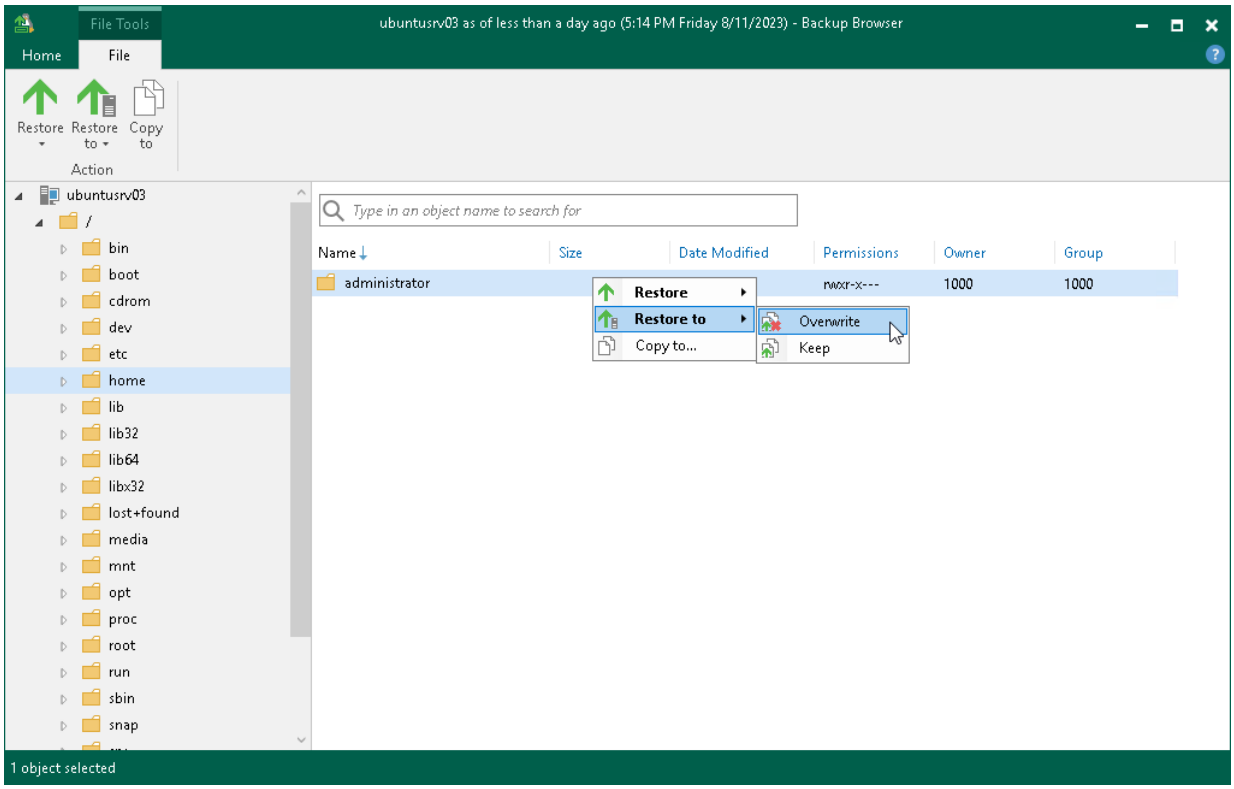


# Saving Files to New Location over Network

You can restore files and folders to components of the Veeam Backup & Replication infrastructure available over the network.

To save files and folders to a new location over the network:

1. Select the necessary files and folders in the file system tree or in the details pane on the right.
2. Right-click one of the selected files or folders and select **Copy to**.
3. In the **Select Destination** window, select the necessary destination:
   - To recover files to a server already added to the Veeam Backup & Replication infrastructure, select the server from the **Server** drop-down list and then specify the path to a folder where files will be recovered.
   - To recover files to a Linux server that is not added to the Veeam Backup & Replication infrastructure, select **Specify a different host** from the **Server** drop-down list and follow the steps of the wizard to add a Linux server that will be used as a target host. The server will be added ad-hoc. Then specify the path to a folder where files will be recovered.

     The server you add ad-hoc will not appear in the list of managed hosts in Veeam Backup & Replication: its purpose is to host the files that you recover. It will only remain visible in the Veeam Backup browser until all currently active file-level restore sessions are completed.

> **IMPORTANT**
>
> Only Linux servers can be added ad-hoc.

- To recover files to a shared folder, specify a path to the destination folder in the **Path to folder** field.

4. If you want to preserve original permissions and ownership for recovered files, select the **Preserve permissions and ownership** check box in the **Select Destination** window.

> **IMPORTANT**
>
> To restore original permissions and ownership settings, the user account you have specified must have privileges to change the owner on the selected server or shared folder.

5. If prompted, in the **Credentials** window, specify settings of the user account to access the destination location.

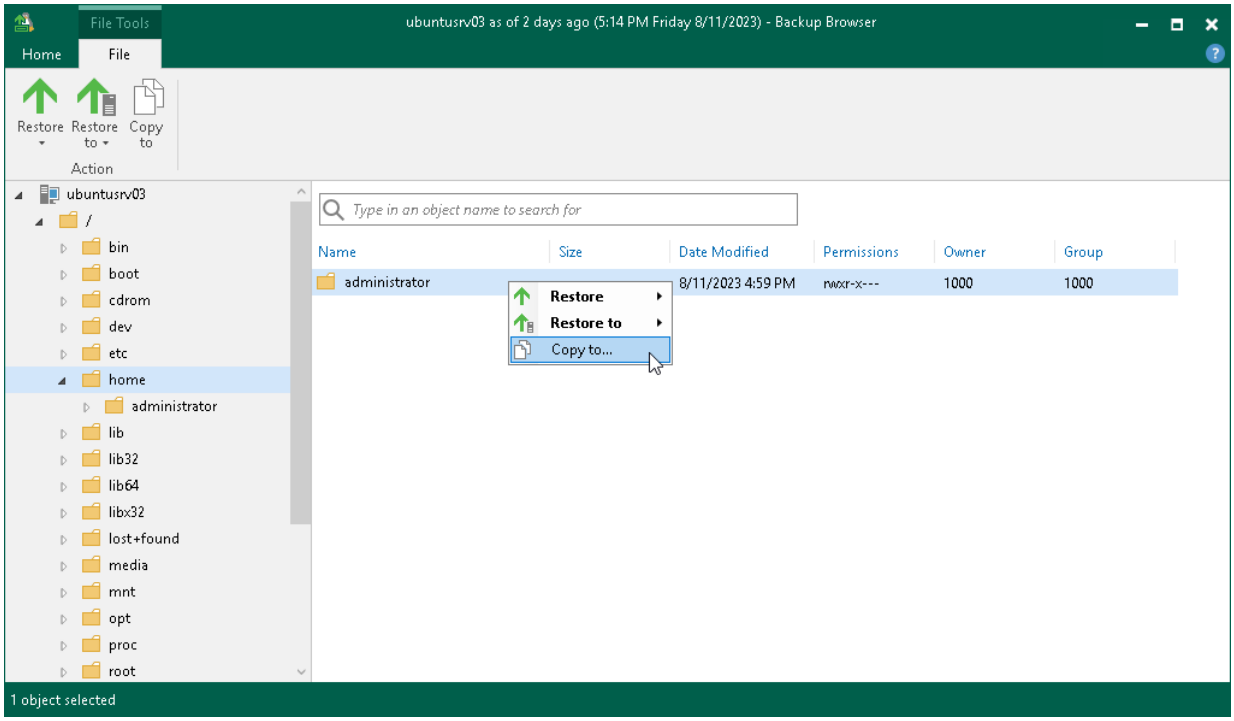


# Accessing Files over FTP

If you have chosen to enable FTP server on the helper appliance, the restored file system will also be available over FTP at *ftp://<FLR_appliance_IP_address>*. Other users in the same network can access the helper appliance to restore the files they need.

Accessing the appliance over FTP requires credentials. Use the Guest OS helper appliance credentials specified in managed credentials. If the password has not been updated, refer this Veeam KB article.

> **NOTE**
>
> [For restore from storage snapshots] You can browse to the VM guest OS files and access restored files on the FTP only while the Veeam Backup browser is open. After the Veeam Backup browser is closed, Veeam Backup & Replication will unmount VM disks from the helper appliance, remove the helper appliance and unmount the storage snapshot from the ESXi host (unless this storage snapshot is used by other restore operations).

# Accessing Helper Appliance Logs

If you need to access logs of the helper appliance, click **Support Logs** on the Veeam Backup browser ribbon. Veeam Backup & Replication will show the **Log** node under the file system tree. To hide this node, click **Support Logs** once again.

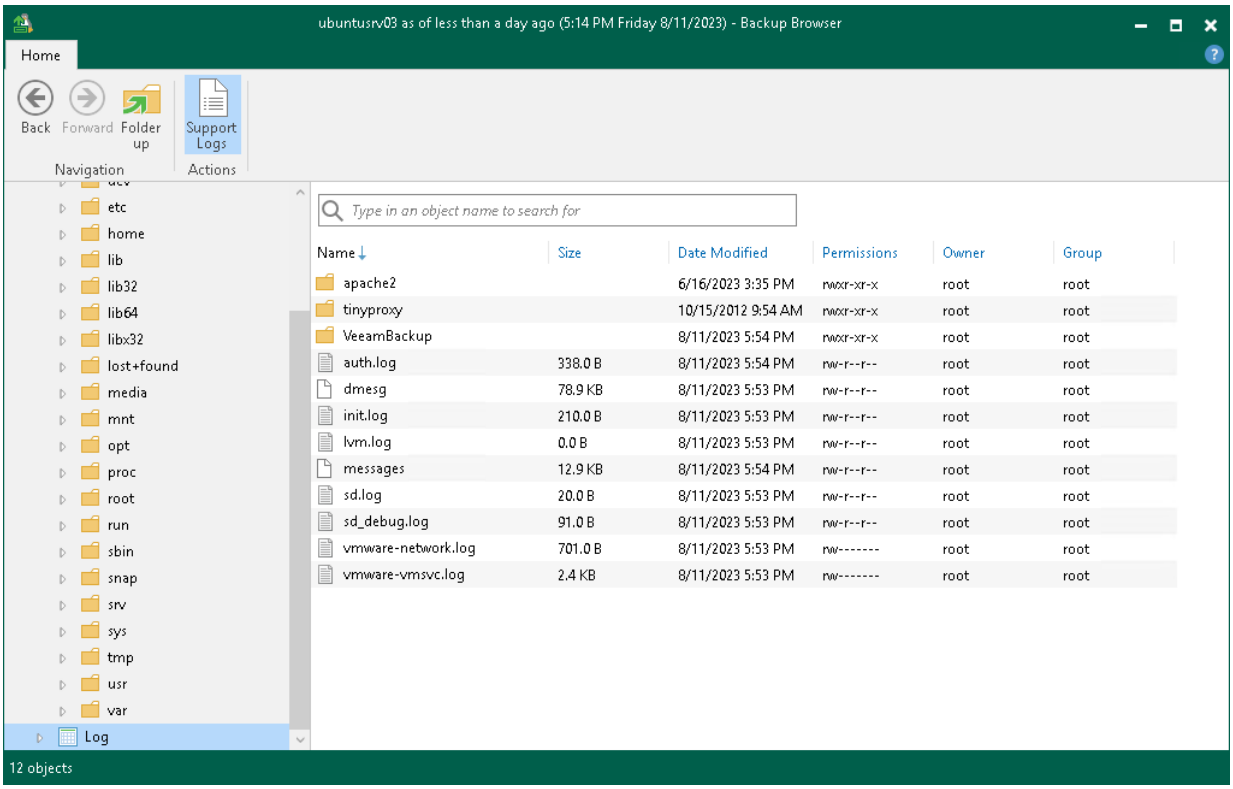


# Restore from Other File Systems

You can restore file systems other than those listed in the Guest OS File Restore section of Supported Platforms and Applications. For this purpose, use the Instant Disk Recovery technology.

To restore files and folders, do the following:

1. In the virtual infrastructure, find a VM that can read the file system of the original VM.
2. Use Instant Disk Recovery to mount a disk of the original VM to the VM that can read the file system.
3. Open the VM console, bring the disk to the online mode.
4. Copy the required files.
5. If you no longer need the disk, stop publishing it.

Alternatively, you can mount the VM disks to a Microsoft Windows VM and use file management tools.

# Viewing File Restore Session Statistics

You can view statistics about performed guest OS file restore sessions.

To view the restore session statistics, do one of the following:

- Open the **Home** view, in the inventory pane select **Last 24 hours**. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.
- Open the **History** view, in the inventory pane select **Restore**. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click **Statistics** on the ribbon or right-click the session and select **Statistics**.

The file restore statistics provides detailed data on file restore sessions:

- At the top of the **Restore Session** window, Veeam Backup & Replication shows general session statistics: a name of the machine whose guest OS files are restored during the session, a user name of the account under which the session was started, session status and duration details.
- The **Statistics** tab shows detailed information about the files restored during the session.
- The **Reason** tab shows the reason for the guest OS file restore that was specified at the **Reason** step of the **File Level Restore** wizard.
- The **Parameters** tab shows information about the restore point selected for the guest OS file restore at the **Restore Point** step of the **File Level Restore** wizard.
- The **Log** tab shows a list of operations performed during the session.

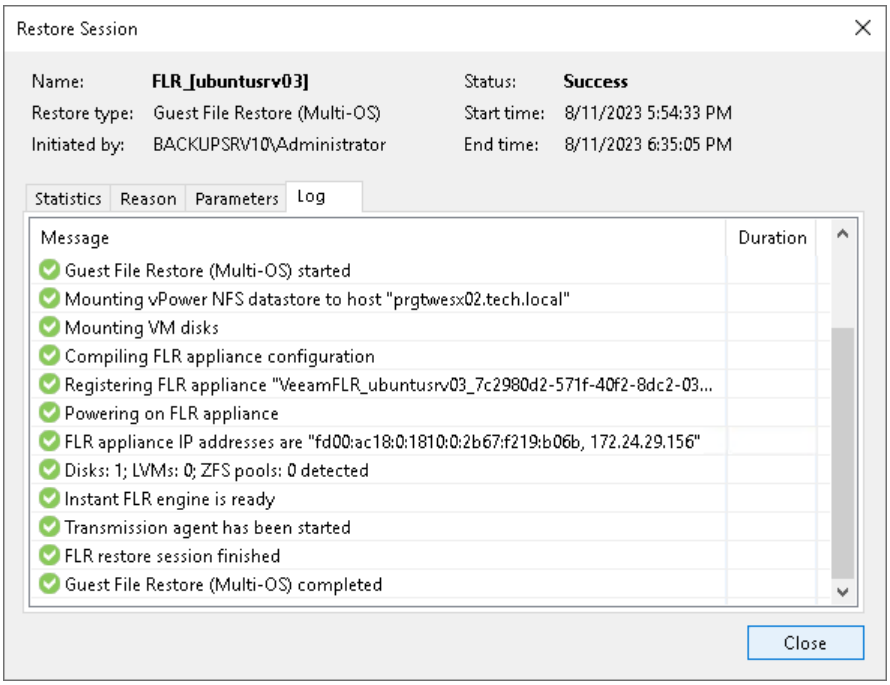

# Application Item Restore

Veeam Backup & Replication provides additional tools called Veeam Explorers. These tools help you restore application items directly from backups or replicas. These backups and replicas must be created with the enabled application-aware processing option.

You can restore items from the following applications:

- Microsoft Active Directory
- Microsoft SQL Server
- Oracle
- PostgreSQL
- Microsoft Exchange
- Microsoft SharePoint
- Microsoft OneDrive for Business*
- Microsoft Teams*

* Item restore for these applications is supported for backups created by Veeam Backup for Microsoft 365. For more information, see the Veeam Explorers User Guide for Veeam Backup for Microsoft 365.

For information on how to launch Veeam Explorers to restore application items, see Restoring Application Items.

Veeam Explorers are included into Veeam Backup & Replication, so you do not need to install them separately. You also do not need to purchase any additional license to use Veeam Explorers. For general information, see the Veeam Explorers Overview section of the Veeam Explorers User Guide.

# Restoring Application Items

To restore application items, use the wizard to restore the application items.

1. Check prerequisites.
2. Launch the wizard.
3. Select a VM or site.
4. Select a restore point or content database.
5. [For restore from storage snapshots] Select an ESXi host for snapshot mounting.
6. Specify a restore reason.
7. Open Veeam Explorer.

## Before You Begin

Before you recover application items, consider the following:

- The backup or replica you plan to restore from must be created with application-aware processing enabled.
- [For recovery from storage snapshots] Check requirements in the Data Recovery from Storage Snapshots section in the Storage System Snapshot Integration Guide.

- Check the Veeam Backup Explorers User Guide for requirements and limitations of applications.

# Step 1. Launch Application Item Restore Wizard

To launch the wizard, do one of the following:

- Open the **Home** view. In the inventory pane, select the **Backups** or **Replicas**. In the working area, expand the necessary backup or replica and select a workload whose application items you want to recover. Right-click the workload, select **Restore application items** and then select the type of application items:
  - Microsoft Active Directory objects
  - Microsoft SQL Server databases
  - Oracle databases
  - PostgreSQL instances
  - Microsoft Exchange Online mailbox items
  - Microsoft SharePoint Online content
  - Microsoft OneDrive for Business files
  - Microsoft Teams files
- Open the **Home** view. In the inventory pane, select **Backups** or **Replicas**. In the working area, expand the necessary backup or replica, select workload whose application items you want to recover, click **Application Items** on the ribbon and select the required application.
- For Microsoft Windows workloads, you can launch application item restore from the Veeam Backup browser. This browser opens after you launch guest OS file restore. For more information on performing guest OS file restore, see section the Guest OS File Restore.

Alternatively, to recover application items from a storage snapshot, you can open the **Storage Infrastructure** view. In the inventory pane, expand the storage system tree and select the necessary volume snapshot. In the working area, select the necessary VM and click **Application Items** on the ribbon and select an application.

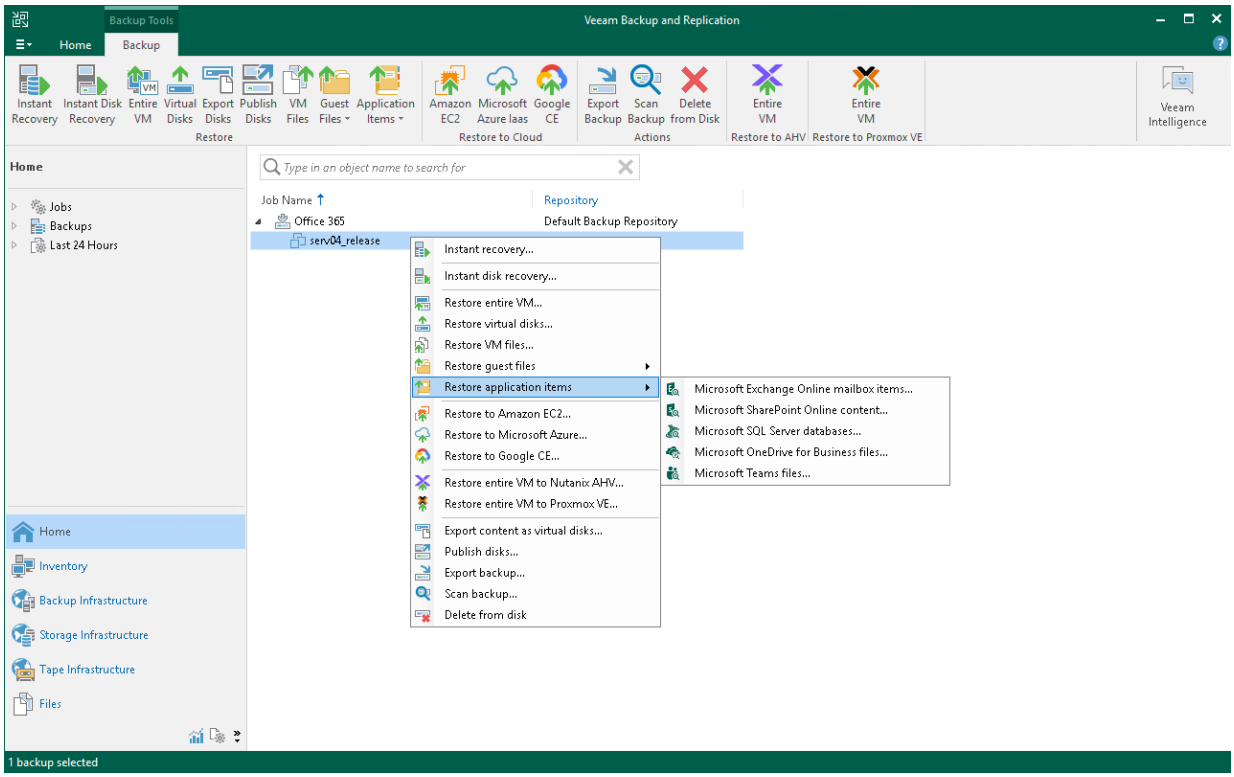

## Step 2. Select VM or Site

This step of the wizard differs for Microsoft SharePoint and other applications.

# Microsoft SharePoint

At the **Sites** step of the wizard, select a Microsoft SharePoint site from which you want to restore application items.

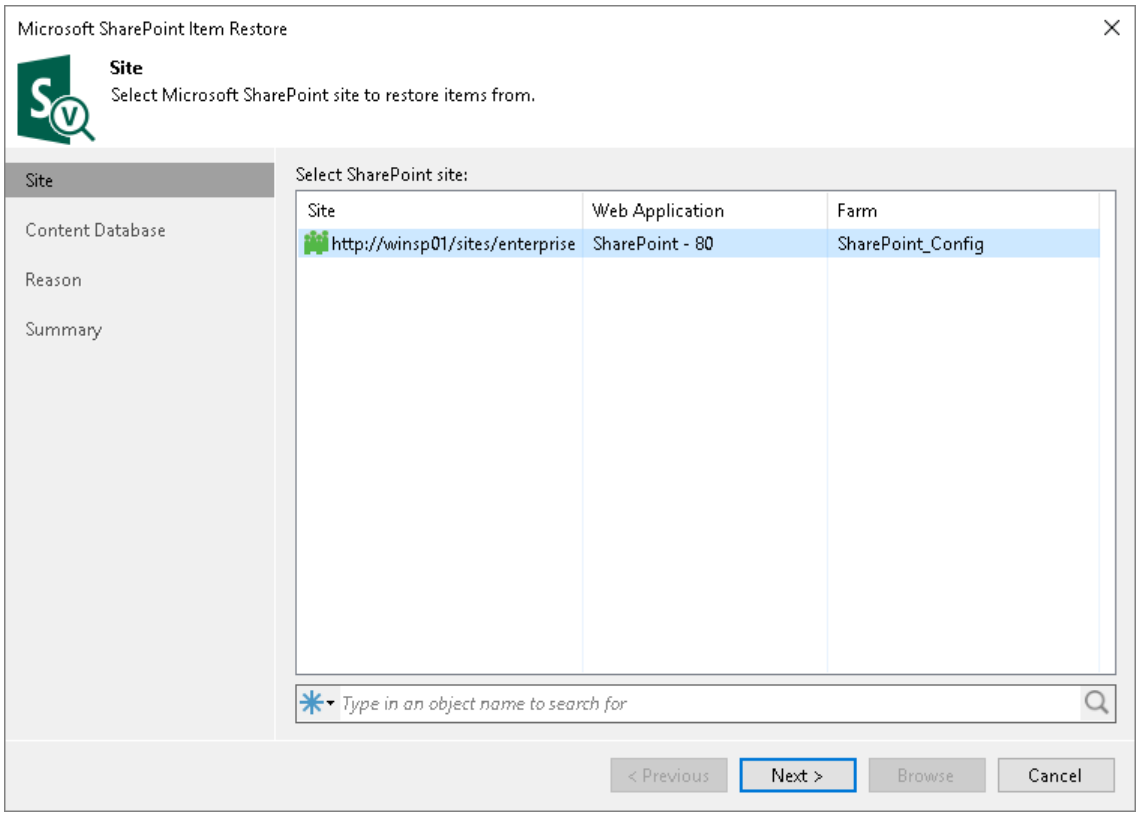

# Other Applications

At the **Machines** step of the wizard, select a workload with the application installed. If the necessary workload is not displayed in the list, select the **Show all VMs** check box.

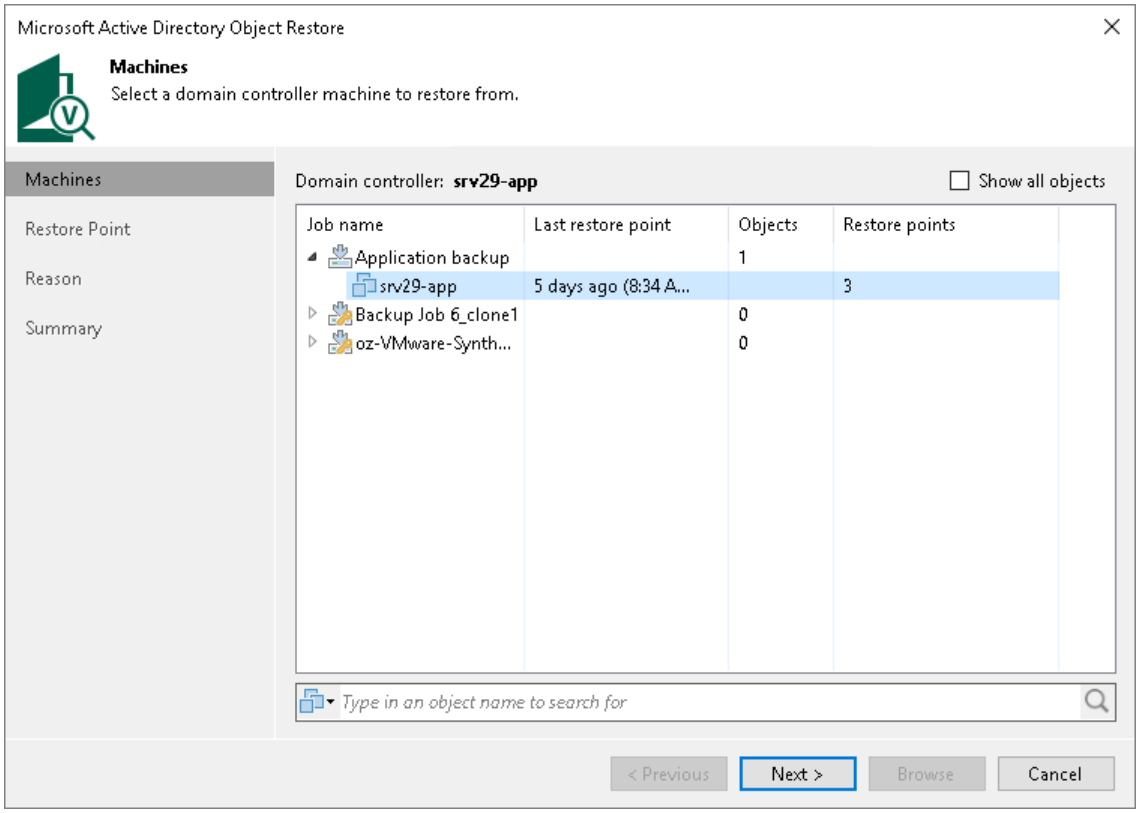

## Step 3. Select Restore Point or Content Database

This step of the wizard differs for Microsoft SharePoint and other applications.

# Microsoft SharePoint

At the **Content Database** step of the wizard, select a content database restore point from which you want to extract application items.

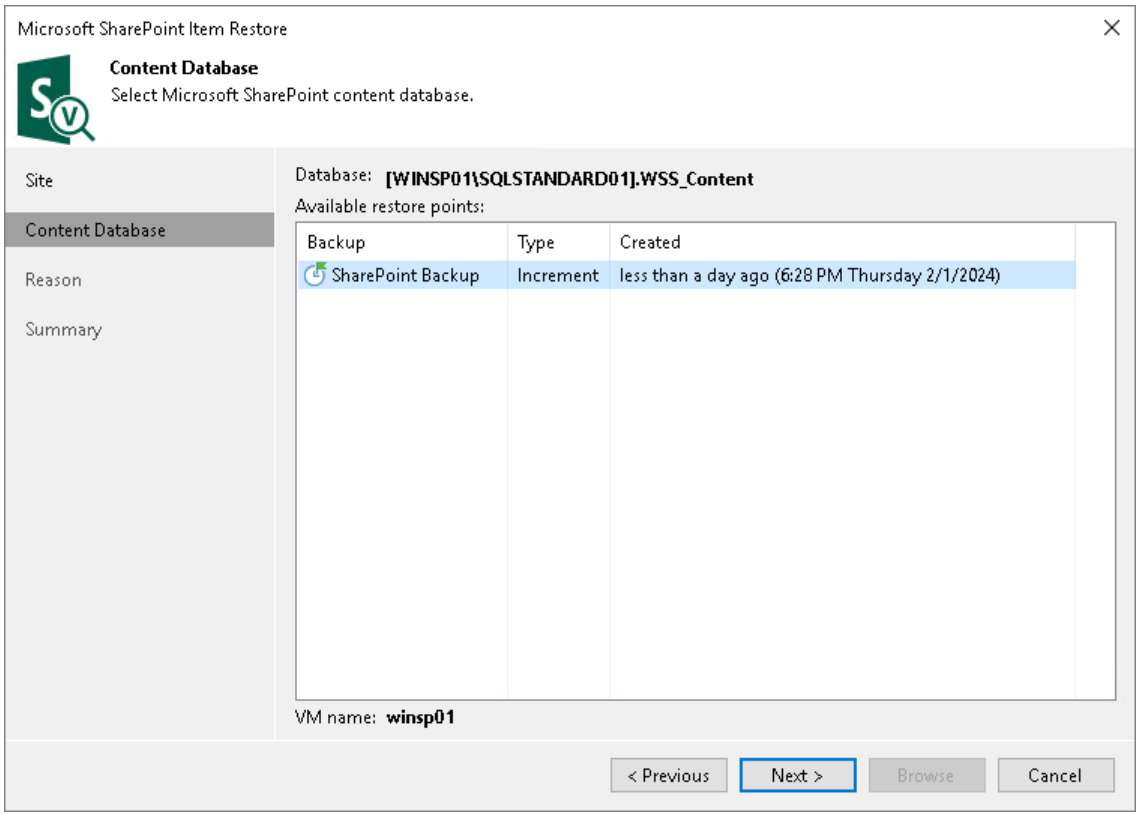

# Other Applications

At the **Restore Point** step of the wizard, select a restore point from which you want to restore application items.

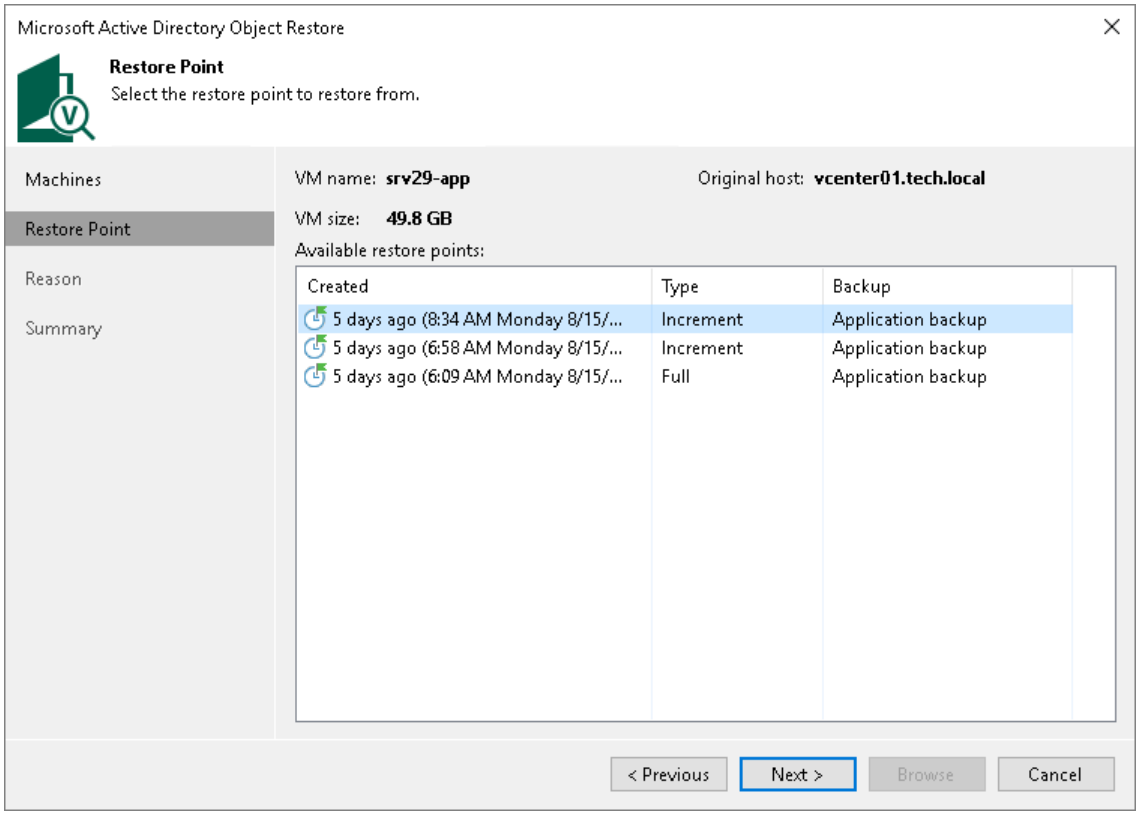

# Step 4. Select ESXi Host for Snapshot Mounting

This step is available only if you restore application items from a storage snapshot.

At the **Location** step of the wizard, select an ESXi host where the clone or virtual copy of the storage snapshot must be mounted. On the selected ESXi host, Veeam Backup & Replication will create a temporary VM and mount disks of the virtualized application to this temporary VM.

To specify a destination for the snapshot clone or virtual copy and temporary VM:

1. At the **Location** step of the wizard, click **Customize**.
2. Next to the **Host** field, click **Choose** and select an ESXi host where the snapshot clone or virtual copy must be mounted and where the temporary VM must be created.
3. Next to the **Resource pool** field, click **Choose** and select a resource pool where you want to place the temporary VM.
4. Next to the **Folder** field, click **Choose** and select a folder where you want to place the temporary VM.
5. Click **OK**.

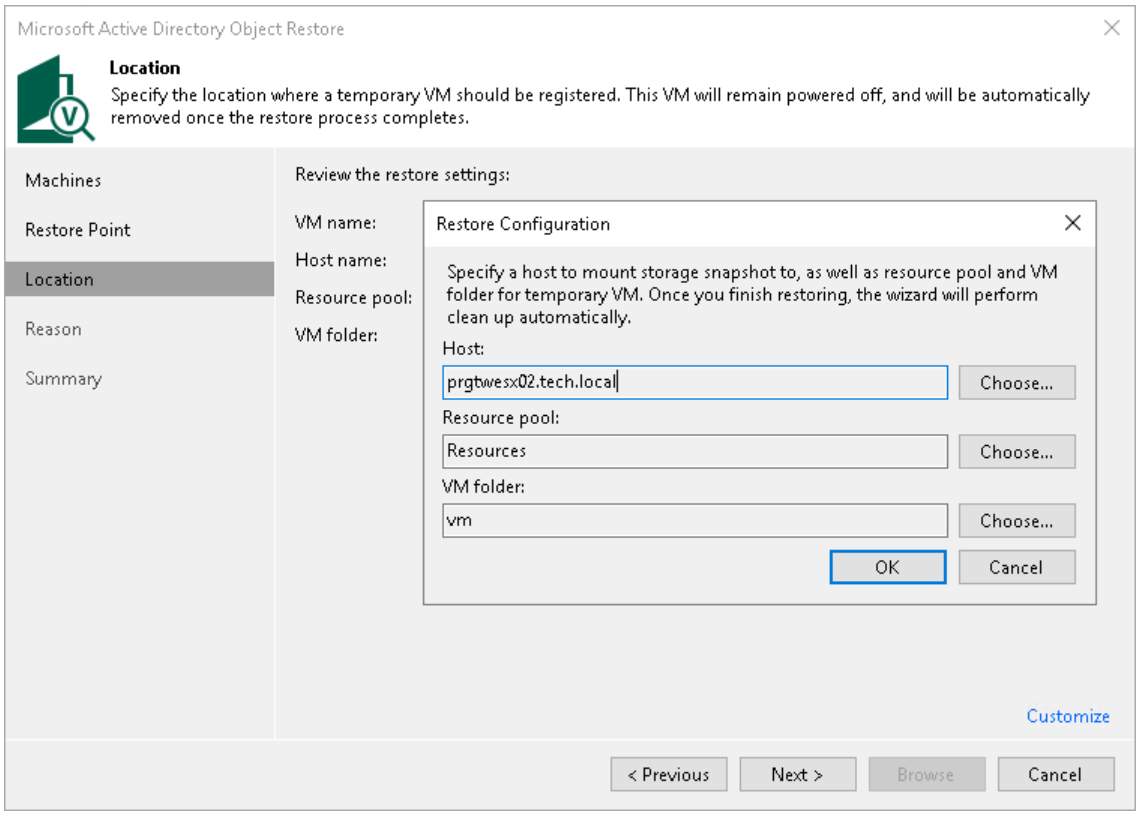

## Step 5. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the application item. The information you provide will be saved in the session history, and you will be able to view it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

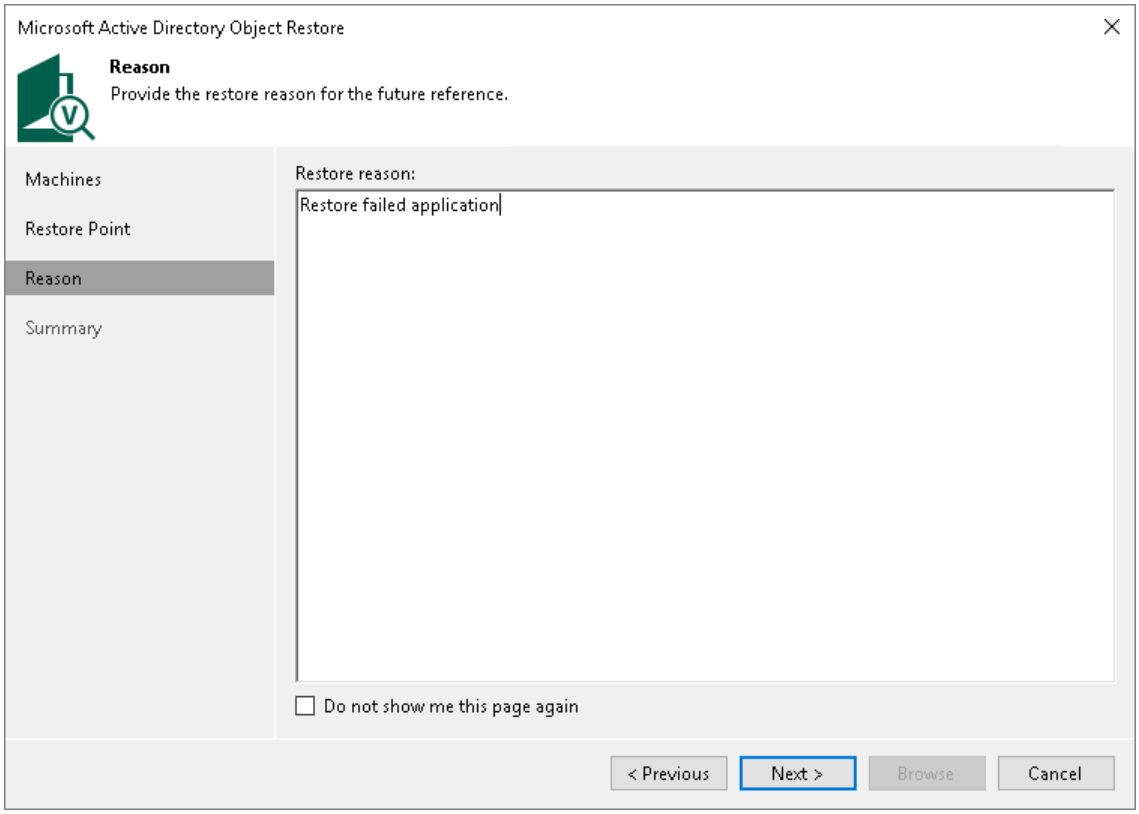

# Step 6. Open Veeam Explorer

At the **Summary** step of the wizard, click **Browse** to start the restore process. Veeam Backup & Replication will automatically locate the application item and open it in Veeam Explorer. In this explorer, you can browse, search, restore application items and so on.

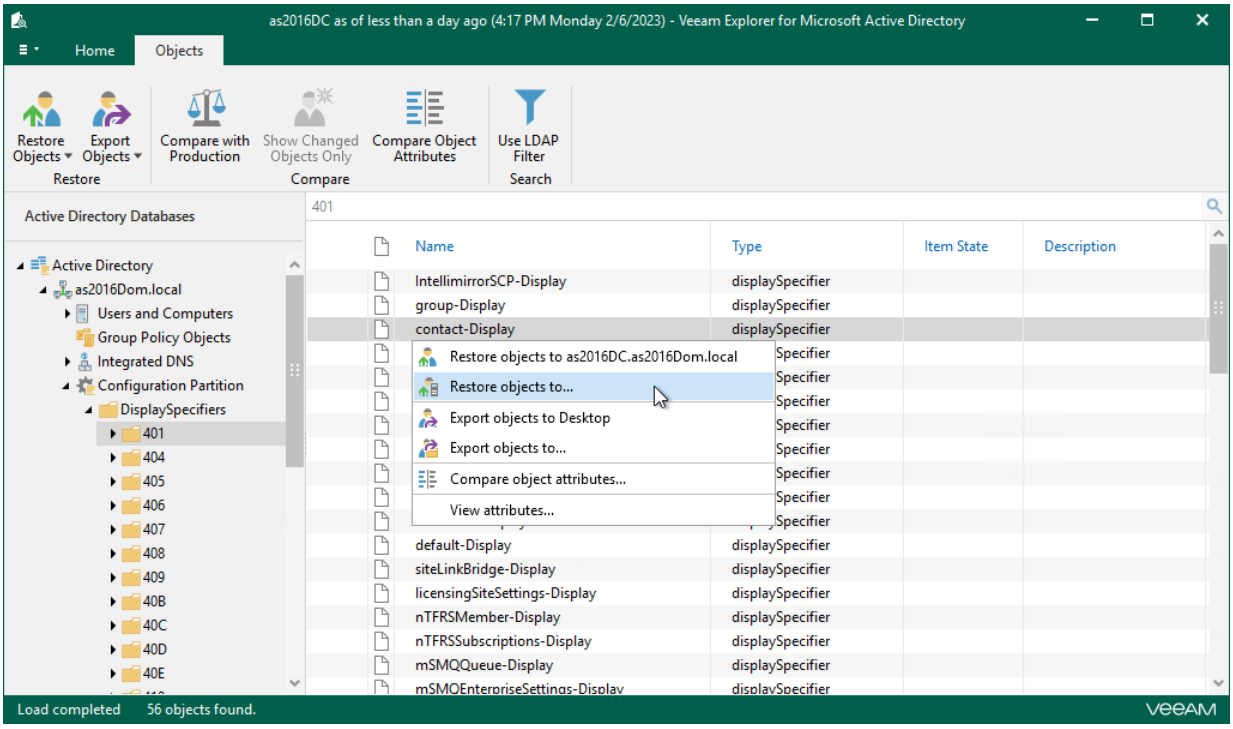


The detailed information on what you can do is provided in the Veeam Backup Explorers User Guide. For example:

- Veeam Explorer for Microsoft Active Directory: General Information, Browsing, Searching and Viewing Items, Data Restore.
- Veeam Explorer for Microsoft SQL Server: General Information, Viewing Database Information, Data Restore.
- Veeam Explorer for Oracle: General Information, Viewing Database Information, Data Restore.
- Veeam Explorer for PostgreSQL: General Information, Data Restore, Data Publishing.
- Veeam Explorer for Microsoft Exchange: General Information, Browsing, Searching and Viewing Items, Restore from Veeam Backup & Replication Backups.
- Veeam Explorer for Microsoft SharePoint: General Information, Browsing, Searching and Viewing Items, Restore from Veeam Backup & Replication Backups.
- Veeam Explorer for Microsoft OneDrive for Business: General Information, Browsing, Searching and Viewing Items, Restoring Microsoft OneDrive Data.
- Veeam Explorer for Microsoft Teams: General Information, Browsing, Searching and Viewing Items, Data Restore.

# Restoring Data from Encrypted Backups

To restore encrypted data, you need to decrypt them first. For more information, see Decrypting Backups.

Data restore from the decrypted backup does not differ from that from an unencrypted one. For more information, see Data Recovery.

# VMware Cloud Director Support

Veeam Backup & Replication provides support for VMware Cloud Director. Veeam Backup & Replication uses VMware Cloud Director API to back up vApps and VMs and restore them directly to the VMware Cloud Director hierarchy.

The main entity with which Veeam Backup & Replication works during backup is a vApp. A vApp is a virtual system that contains one or more individual VMs along with parameters that define operational details — vApp metadata. When Veeam Backup & Replication performs backup of VMs, it captures not only data of VMs being a part of vApps, but also vApp metadata. As a result, you can restore VMware Cloud Director objects back to the VMware Cloud Director hierarchy and do not need to perform any additional actions on import and VM configuration.

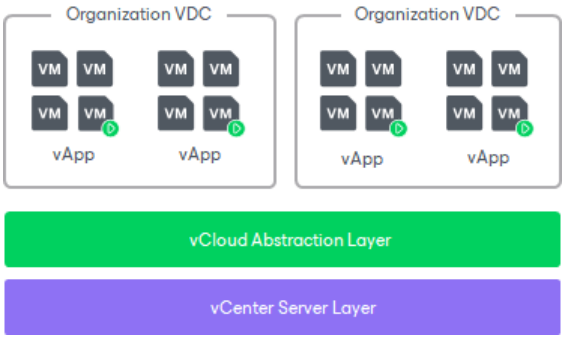


**NOTE**

For VMware Cloud Director objects, we recommend that you use features dedicated to Cloud Director objects, not to regular VMware vSphere VMs. For example, Cloud Director Backup Jobs.

# Viewing VMware Cloud Director VMs

After you add the VMware Cloud Director server to the backup infrastructure, you can view the VMware Cloud Director hierarchy in Veeam Backup & Replication and work with VMware Cloud Director VMs.

To open the VMware Cloud Director hierarchy:

1. Open the **Inventory** view.
2. Click the **View** tab on the ribbon.
3. On the **View** tab, click **Cloud Director View**.

The hierarchy of the VMware Cloud Director server will become available in the inventory pane. VMs managed by VMware Cloud Director will be displayed in the working area. You can work with these VMs just as if you worked with VMs managed by vCenter Servers or registered on ESXi hosts in your backup infrastructure.

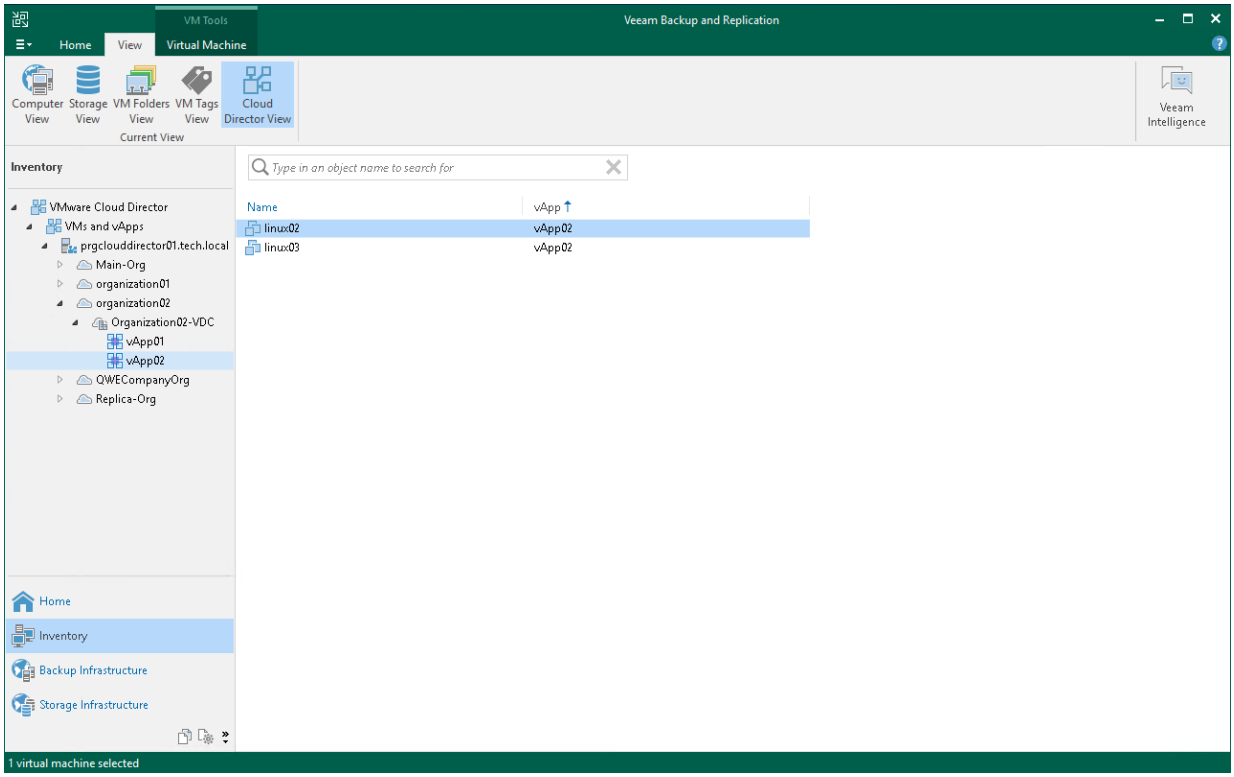

# Backup and Restore of vApps

Veeam Backup & Replication lets you back up VMware Cloud Director vApps and restore them back to the VMware Cloud Director hierarchy.

In terms of VMware Cloud Director, a vApp is a coherent system that includes one or more VMs. Every vApp is described with a set of operational details — vApp metadata. vApp metadata contains the following information:

- vApp owner settings
- Access rights settings
- vApp network settings: information about organization networks to which the vApp is connected
- Lease settings and so on

When Veeam Backup & Replication performs backup of a vApp, it backs up all VMs being a part of this vApp along with the vApp metadata. Backup of the vApp is performed with the VMware Cloud Director backup job. The VMware Cloud Director backup job may contain one or several vApps. If necessary, you can exclude specific VMs and VM disks from the backup when configuring a VMware Cloud Director backup job.

After Veeam Backup & Replication creates backups, you can scan backup data using different malware detection methods. For more information, see Malware Detection.

Veeam Backup & Replication also offers the following restore options for backed-up vApps:

- VM Recovery
- vApp Recovery
- Item Recovery

> **NOTE**
>
> Veeam Backup & Replication shows restore points of a VM that is no longer processed, for example excluded from the job or deleted, as long as incremental backup files with this VM exist in the active backup chain. After Veeam Backup & Replication deletes the files according to the retention policy, it also hides the VM restore points from the list of restore points created by this backup job. Although the full backup file with the VM may still exist, restore from this file into the VMware Cloud Director is not supported.

# Backup for VMware Cloud Director

Veeam Backup & Replication lets you perform backup for vApps and VMs, as well as VM containers in VMware Cloud Director such as organization VDC, organization and even the VMware Cloud Director instance.

When Veeam Backup & Replication performs backup of vApps and VMs, it additionally captures vApp metadata.

vApp metadata includes:

- General information about the vApp where VMs reside, such as: vApp name, description, VMs descriptions
- Information about vApp networks and organization networks to which the vApp is connected
- VMs startup options
- User information
- Lease
- Quota
- Storage template and so on

vApp metadata is stored together with the VM content. Capturing vApp metadata is extremely important for restore: without it, you will not be able to restore vApps and VMs back to VMware Cloud Director.

# Data to Back Up

Veeam Backup & Replication lets you back up the following types of VMware Cloud Director VMs:

- Regular VMs that are part of vApps
- Standalone VMs that were created in the VMware Cloud Director tenant portal
- Linked clone VMs that are associated with vApps

## Backup of Regular and Standalone VMs

When you back up regular or standalone VMs, Veeam Backup & Replication captures and stores to the backup file the following data:

- VM disk content
- vApp metadata
- VM metadata

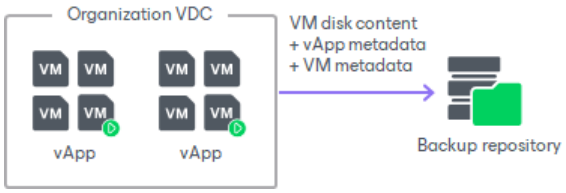


## Backup of Linked Clone VMs

When you back up linked clone VMs, Veeam Backup & Replication captures and stores to the backup file the following data:

- Content of the template to which the VM is linked
- Content of the VM user disk — delta disk
- vApp metadata
- VM metadata

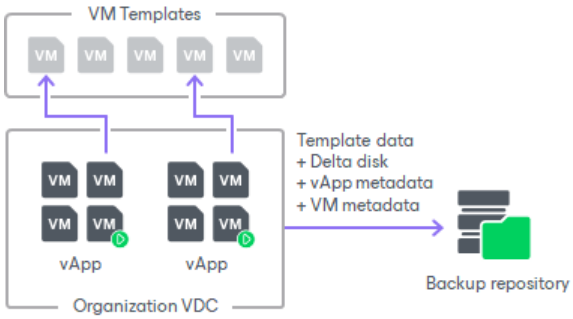


During full backup of linked clone VMs, Veeam Backup & Replication consolidates data of the VM template and delta disk and saves it as a regular VM disk in the backup file. Data merging guarantees proper VM restore: even if a VM template is lost by the time of recovery, you will still be able to restore the linked clone VM from the backup.

During incremental backup, Veeam Backup & Replication saves only changed data of the delta file.

## Limitations for Backup of Linked Clone VMs

Before you back up linked clone VMs, consider the following:

- [For vCenter 6.5 or later] If you back up a linked clone VM that has snapshots, Veeam Backup & Replication may fail to produce a valid restore point. To overcome this issue, do one of the following:
  - Disable CBT (Change Block Tracking) in the backup job settings.
  - Ensure that CBT is enabled on the VM template to which the VM is linked.

  For details on how to enable CBT on the VM template, contact Veeam Customer Support.
- Backup of linked clone VMs that were created by services other than VMware Cloud Director may cause snapshot-related problems. To overcome this issue, disable Veeam Snapshot Hunter. For details, see this Veeam KB article.

# Cloud Director Backup Jobs

For VMs managed by VMware Cloud Director, Veeam Backup & Replication offers a special type of the backup job — VMware Cloud Director backup job. VMware Cloud Director backup jobs process VMware Cloud Director objects, ensure their proper restore and support of Cloud Director-specific features.

It is recommended that you use VMware Cloud Director backup jobs to back up VMs managed by VMware Cloud Director. If you back up VMs managed by VMware Cloud Director using a regular backup job, Veeam Backup & Replication will perform backup at the level of the underlying vCenter Server and will not capture vApp metadata. As a result, you will not be able to restore a fully-functioning VM to VMware Cloud Director.

# Performing Backup of VMware Cloud Director VMs

The VMware Cloud Director backup is practically the same as a regular VM backup. The VMware Cloud Director backup job aggregates main settings for the backup task and defines when, what, how and where to back up VMware Cloud Director VMs.

You can perform the VMware Cloud Director backup job for single VMs and for VM containers:

- vApp
- Organization VDC
- Organization
- VMware Cloud Director instance

Just like a regular backup job, the VMware Cloud Director backup job can be scheduled or run manually. To create a VMware Cloud Director backup job, do one of the following:

- On the **Home** tab, click **Backup Job** > **Virtual Machine** and select **VMware Cloud Director**.
- Open the **Home** view, in the inventory pane right-click **Jobs** and select **Backup** > **Virtual Machine** > **VMware Cloud Director**.
- Open the **Inventory** view, click the **View** tab and click **Cloud Director View** on the ribbon. In the inventory pane expand the **VMware Cloud Director** hierarchy, in the working area select one or more VMs, click **Add to Backup** on the ribbon and select **New job**. Alternatively, you can right-click one or several VMs and select **Add to backup job** > **New job**. In this case, the selected VMs will be automatically added to the new VMware Cloud Director backup job. You can add other VMs to the job when passing through the wizard steps.

You can quickly include VMs to already existing VMware Cloud Director backup jobs. To do this, in the **Inventory** view, in the working area right-click necessary VMs and select **Add to backup job** > <*name of the job*>.

The backup wizard for VMware Cloud Director offers the same options as the backup wizard for VMware virtual machines. For more information, see Creating Backup Jobs.

**IMPORTANT**

Consider the following:

- Veeam Backup & Replication supports the VMware Cloud Director multisite deployment. To back up vApps and VMs, you must add all Cloud Director servers to Veeam Backup & Replication separately.
- We recommend adding VMs of one vApp into one Cloud Director backup job. Splitting VMs into different jobs can cause additional load on the system during backup and restore. This can happen because Veeam Backup & Replication must handle multiple vApp metadata copies.
- Limitation of concurrent tasks applies in the same way as for backup jobs.
  Veeam Backup & Replication ignores vApps. For more information, see Limitation of Concurrent Tasks.
- If you run a VMware Cloud Director backup job for the vApp, the job is considered to finish with the *Success* status and the complete restore point for the vApp is created only if all VMs in the vApp are successfully backed up. If any VM in the job fails, the restore point for the vApp will be in the incomplete status, and you will not be able to restore the whole vApp from such restore point. However, you will be able to perform restore to the original vApp for VMs that were partially or successfully backed up and whose data is available in the incomplete restore point.
  Veeam Backup & Replication will restore all data that is available for such VMs in the backup.

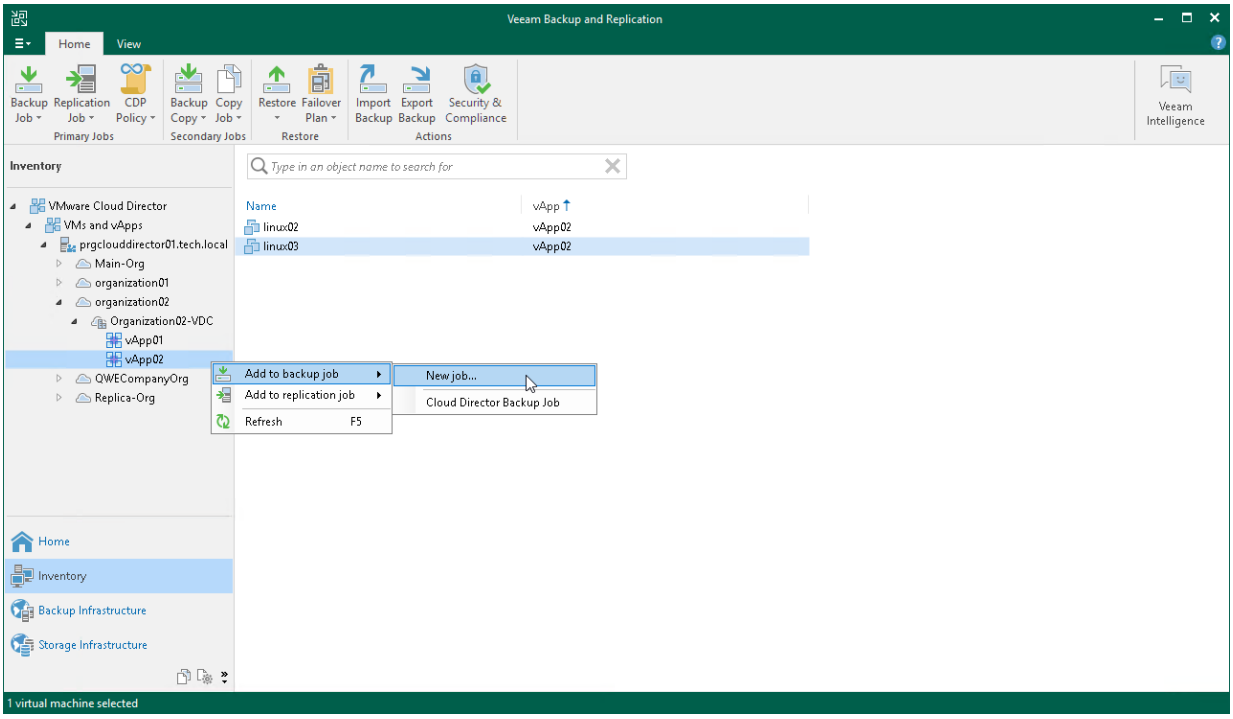

# Managing Cloud Director Backups and Jobs

Managing for Cloud Director backups and jobs is practically the same as for regular backups. You can view backup properties, copy, move, export and delete backups from disk. For more information, see the following sections:

- Viewing Backup Properties
- Moving Backups
- Copying Backups
- Exporting Backups
- Deleting Backups from Disk
- Detaching Backups from Jobs
- Editing Job Settings
- Cloning Jobs
- Retrying Jobs
- Disabling and Deleting Jobs
- Starting and Stopping Jobs
- Performing Active Full Backup

  You can perform active full backup for the whole job as described in Performing Active Full Backup or create a partial active full backup as described further in this section.

> **NOTE**
>
> Consider the following:
>
> - When you launch a retry for a vApp, Veeam Backup & Replication performs retry for the vApp and its failed VMs. When you launch the retry for a VM, Veeam Backup & Replication performs retry only for this VM.
> - You cannot retry a backup job targeted at an immutable repository. That is because the last session of the preceding automatic retry session makes immutable not only VM backup files but also the vApp backup file. Then this manual retry tries to change this immutable vApp backup file and fails. Instead of the manual retry, you can create an active full backup. This operation can help because it creates new vApp backup files. Note that automatic job retry can still be used because it makes the vApp backup file immutable only on the last retry session.
> - The move and copy operations are available at the backup job level and at the vApp level. You cannot use these operations for individual VMs.

## Performing Partial Active Full Backup

You can create a partial active full backup, that is, create a backup for an individual vApp or an individual VM. When you launch the active full backup for an individual vApp, Veeam Backup & Replication creates the active full backup for the vApp and its VMs that the job processes. When you launch the active full backup for an individual VM, Veeam Backup & Replication creates the active full backup for the VM and its parent vApp.

To create a partial active full backup:

1. Open real-time statistics or sessions results of the job.
2. Select a vApp or VM for which you want to perform active full backup.
3. Right-click one of the selected workload and click **Active full**. Note that you will be able to launch active full backup for other workloads only after the running session finishes.

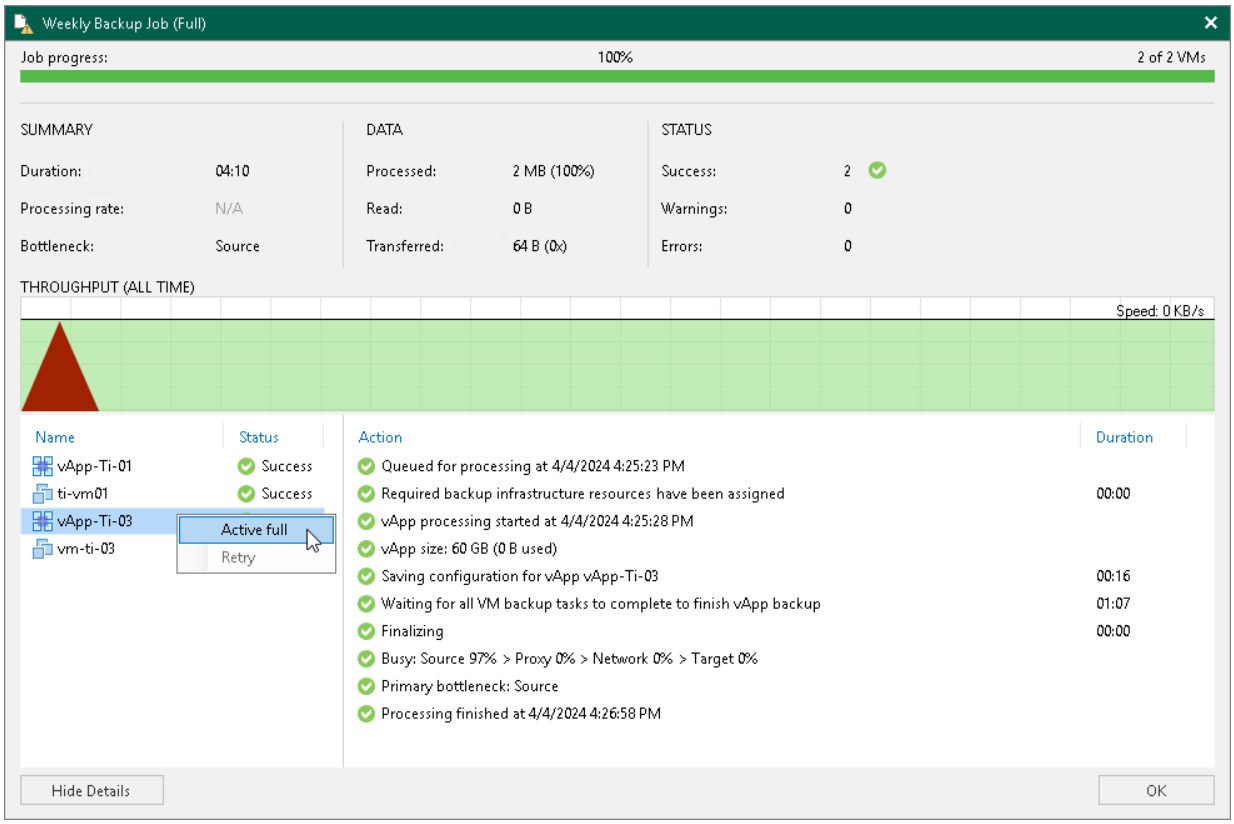

# Creating VeeamZIP Files for VMware Cloud Director VMs

You can create a VeeamZIP file for one or more VMware Cloud Director VMs. When Veeam Backup & Replication creates a VeeamZIP file for a VMware Cloud Director VM, it backs up VMs and the vApp metadata.

The process of VeeamZIP files creation for VMware Cloud Director VMs does not differ from that for regular VMware VMs. For more information, see Creating VeeamZIP Backups.

> **IMPORTANT**
>
> Creating Cloud Director VeeamZIP backups to the object storage repository is not supported.

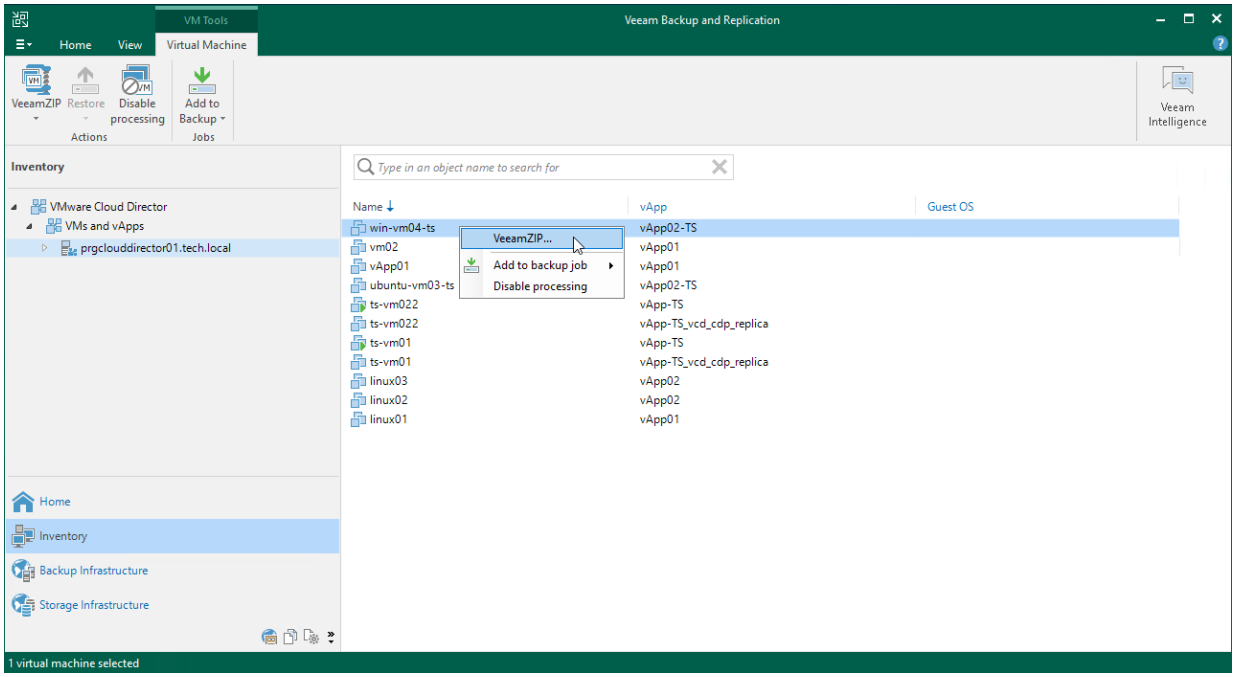

# CDP for VMware Cloud Director

Continuous data protection (CDP) for VMware Cloud Director is a technology that helps you protect mission-critical vApps and VMs when data loss for seconds or minutes is unacceptable. CDP also provides minimum recovery time objective (RTO) in case a disaster strikes because CDP replicas are in a ready-to-start state.

CDP for VMware Cloud Director utilizes a similar mechanisms as CDP for VMware vSphere VMs. The main difference is that the unit for replication and failover is a vApp, not a VM. For more information on CDP for VMware vSphere VMs, see Continuous Data Protection (CDP).

To protect vApps using CDP for Cloud Director, you first need to configure the infrastructure and then to create a Cloud Director CDP policy. For more information, see Backup Infrastructure for Cloud Director CDP and Creating Cloud Director CDP Policies. To recover vApps to a short-term or long-term restore point, you need to fail over to its replica. For more information, see Failover and Failback for Cloud Director CDP.

# Backup Infrastructure for Cloud Director CDP

The following backup infrastructure components are required for VMware Cloud Director CDP:

- Backup server
- Source and Target organization VDCs
- VMware CDP proxies

## Backup Server

The backup server is the configuration, administration and management core of the backup infrastructure. The backup server runs the Veeam CDP Coordinator Service. This service coordinates replication and data transfer tasks, and controls resource allocation. We recommend that you place the backup server in the target site or as a separate unit.

For more information on the backup server, see Backup Server.

## Source and Target Organization VDCs

The source and the target organization VDCs are two terminal points between which replicated vApp and VM data is moved. The source and target organization VDCs must be a part of one VMware Cloud Director server or two different servers. In the underlying VMware vSphere infrastructure, hosts must be added to clusters managed by vCenter Servers. For more information on the requirements to the infrastructure and how to add Cloud Director servers, see Requirements and Limitations and Adding VMware Cloud Director Servers.

The source and target VDC organizations perform the following tasks:

- The source organization VDC reads vApp and VM disk data, reads and processes I/O operations and sends data to source VMware CDP proxies. The data is sent uncompressed.
- The target organization VDC receives data from target VMware CDP proxies and saves this data to replicas on the datastore. Also, the target VDC manages replicas: creates replicas, retains restore points and so on.

## I/O Filter on Organization VDCs

To be able to use organization VDCs for CDP, you must install the I/O filter on each VDC where vApps reside. After you install the I/O filter on the organization VDCs, Veeam Backup & Replication automatically installs the filter on all underlying clusters and their hosts. For more information on how to install the filter, see Installing I/O Filter on VDCs.

It is the I/O filter that reads and processes I/O operations between the protected vApps and their underlying datastore and that sends and receives data from VMware CDP proxies. Also, the filter communicates with the Veeam CDP Coordinator Service on the backup server and notifies the service that the backup infrastructure must be reconfigured if any proxy becomes unavailable. This I/O filter is built on the basis of vSphere API for I/O filtering (VAIO).

## VMware CDP Proxies

A VMware CDP proxy is a component that operates as a data mover and transfers data between the source and target organization VDCs. We recommend that you configure at least two VMware CDP proxies: one (source proxy) in the production site and one (target proxy) in the disaster recovery site.

The source and target VMware CDP proxies perform the following tasks:

- The source proxy prepares data for short-term restore points from data received from the source host, compresses and encrypts the data (if encryption is enabled in the network traffic rules). Then sends it to the target proxy.
- The target proxy receives the data, decompresses and decrypts it, and then sends to the target host.

For more information on VMware CDP proxies, their requirements, limitations and deployment, see VMware CDP Proxies.

# Requirements and Limitations

Consider the following for CDP for VMware Cloud Director:

- As CDP for VMware Cloud Director is based on the CDP mechanism for VMware vSphere VMs, requirements and limitations are also actual. For more information, see Requirements and Limitations for CDP.
- Check the supported Cloud Director versions in Supported Platforms and Applications.
- The version of target hosts must be 7.0 or later. Target hosts of version 6.5 and 6.7 are not supported.
- The *Any* storage policy is not supported.
- Most configurations require that you disable the VM discovery option in VMware Cloud Director global settings. For more information on where you can change the option, see VMware Docs.
- Fast provisioned VMs cannot be protected with CDP.
- The maximum number of VMs in a vApp that can be protected with CDP is 128.
- If you add a new organization VDC to the Cloud Director server after the I/O filter is installed on the existing VDCs, you need to install the I/O filter manually on the newly added VDC. To do that, open the I/O Filter Management wizard, make sure that check boxes are selected near the organization VDCs where the I/O filter must be present and finish the wizard.

# How Cloud Director CDP Works

Cloud Director CDP works in a similar way as CDP for VMware vSphere VMs. For more information, see the following sections:

- How CDP Works
- CDP Replication Chain
- Retention Policies
- Guaranteed Delivery
- Replica Seeding and Mapping

# Installing I/O Filter on VDCs

To be able to protect VM containers with CDP, you must install the I/O filter on each organization VDC that you plan to protect and where Cloud Director replicas will reside. For more information on the filter, see Backup Infrastructure for Cloud Director CDP.

To install the I/O filter on a VMware Cloud Director server and all the connected organization VDCs, use the **I/O Filter Management** wizard.

# Step 1. Launch I/O Filter Management Wizard

To launch the **I/O Filter Management** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane, navigate to the **Managed Servers** > **VMware Cloud Director**. In the working area select a VMware Cloud Director server on which you want to install the filter. Right-click the servers and select **Manage I/O filters**. Alternatively, click **Manage I/O Filters** on the ribbon.
- Open the **Inventory** view. Click the **View** tab on the ribbon and click **Cloud Director View**. In the inventory pane, navigate to the **VMware Cloud Director** > **VMs and vApps** > *<Cloud Director Server >* node. Right-click the node and select **Manage I/O filters**. Alternatively, click **Manage I/O Filters** on the ribbon.

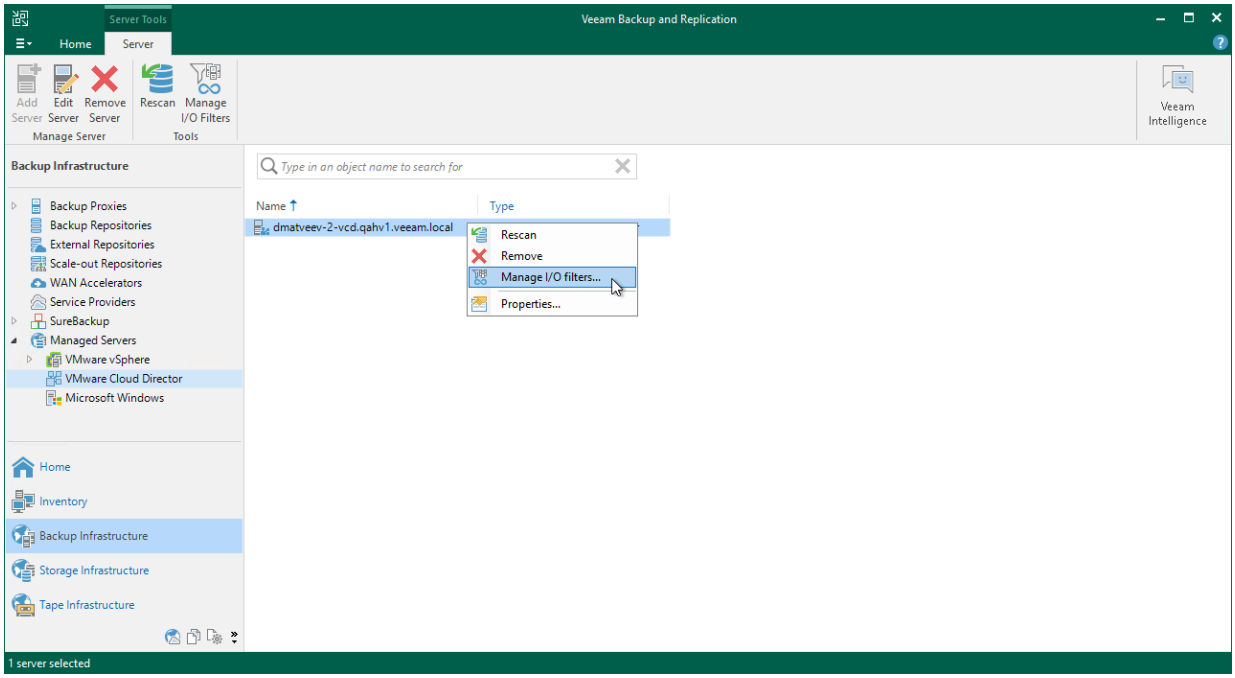

# Step 2. Select VDC Organizations

At the **Organization VDC** step of the wizard, select check boxes near organization VDCs on which you want to install the I/O filter.

If you select check boxes near VDCs where the filters are installed, Veeam Backup & Replication will update the filters. If you clear check boxes, Veeam Backup & Replication will delete the I/O filter from these VDCs. For more information, see Updating and Uninstalling I/O Filter.

> **NOTE**
>
> If you add a new organization VDC to the Cloud Director server after the I/O filter is installed on the existing VDCs, you need to install the I/O filter manually on the newly added VDC. To do that, open the **I/O Filter Management** wizard, make sure that check boxes are selected near the organization VDCs where the I/O filter must be present and finish the wizard.

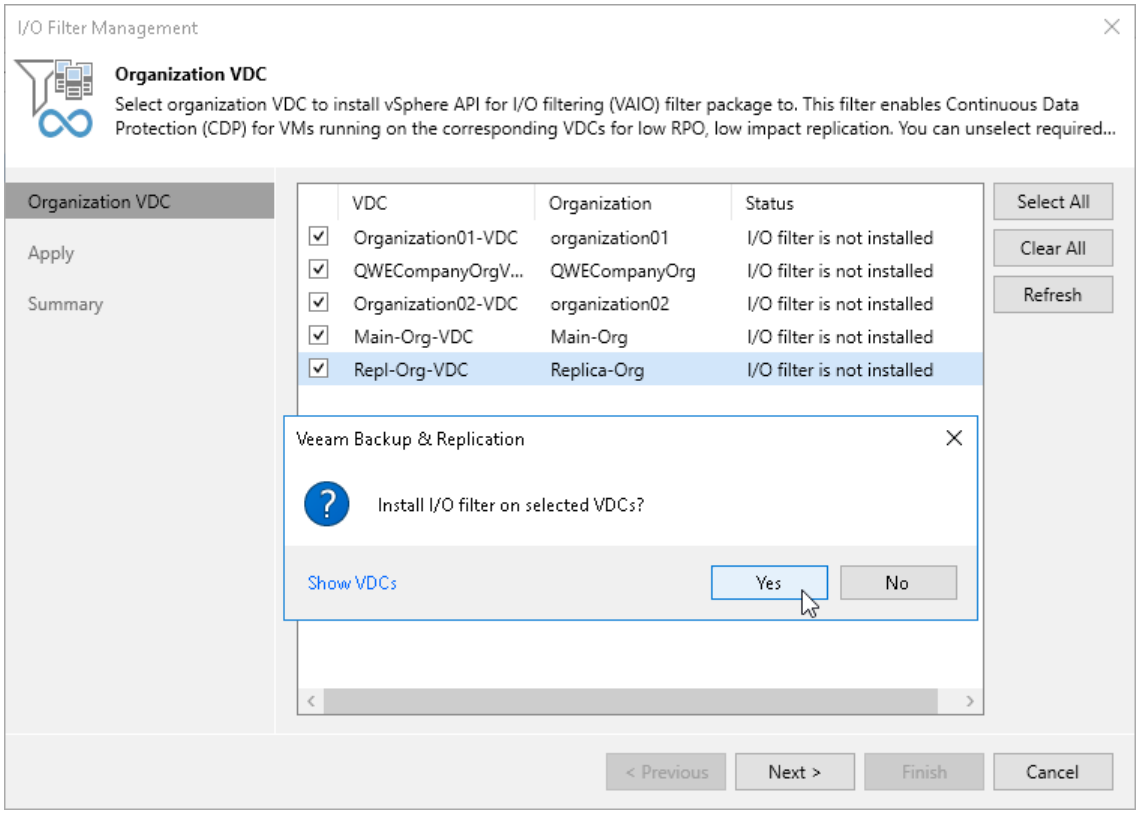

# Step 3. Apply Filter Settings

At the **Apply** step of the wizard, wait untill Veeam Backup & Replication installs the I/O filter. Click **Next**.

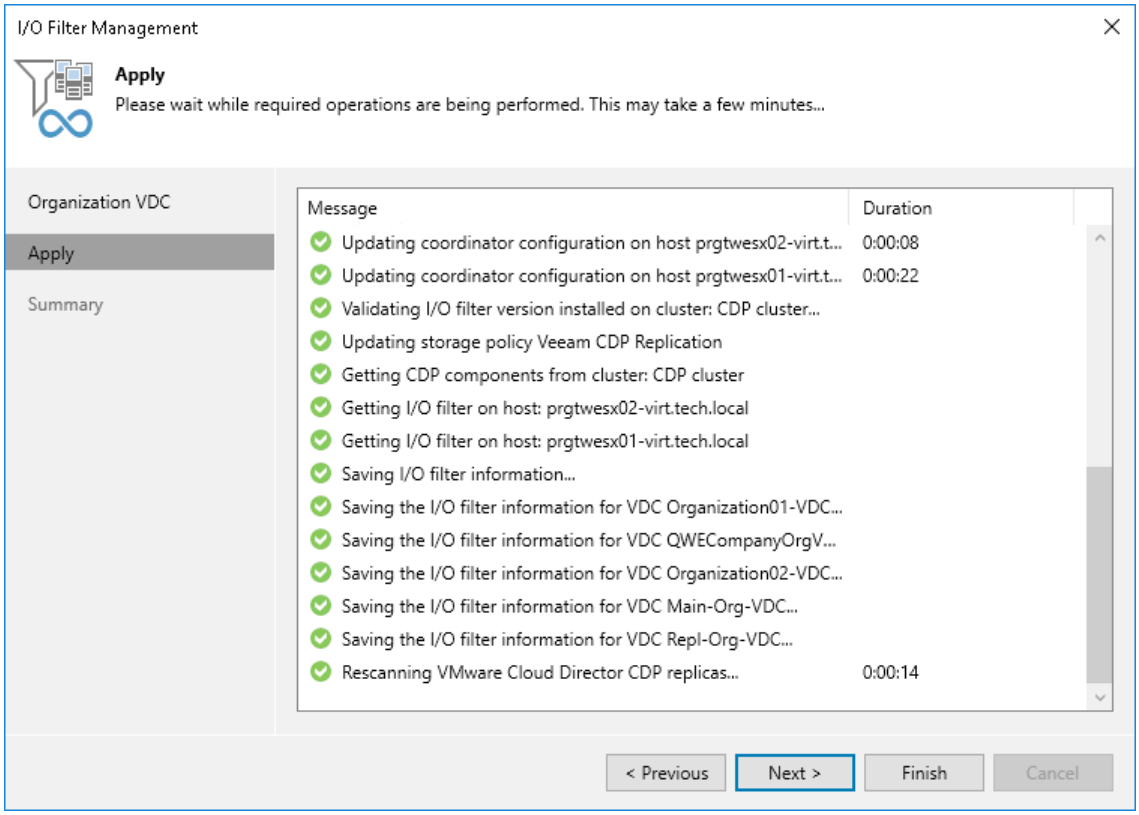

# Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review on which organization VDCs the I/O filter is installed and click **Finish** to exit the wizard.

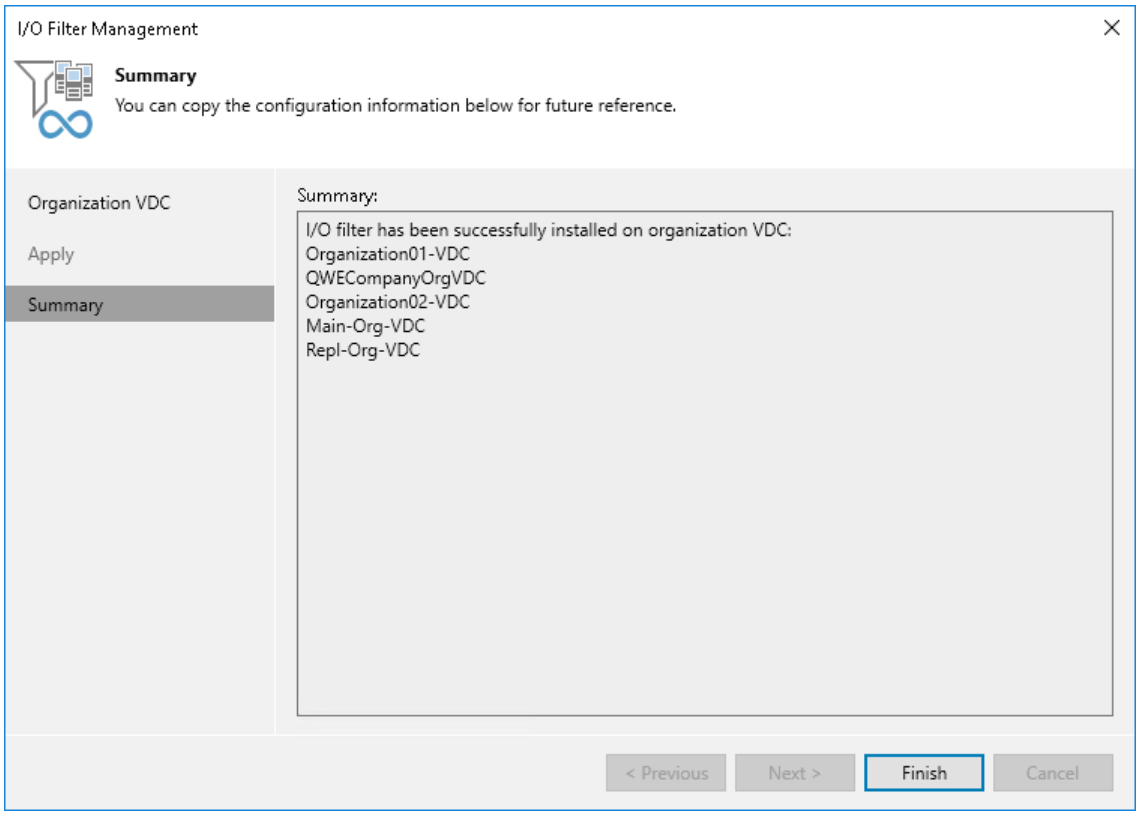

# Updating and Uninstalling I/O Filter

Veeam Backup & Replication allows you to update or uninstall the I/O filter from organization VDCs using the Veeam Backup & Replication console.

## Requirements

Consider the following:

- Make sure that you have disabled or deleted all VMware Cloud Director CDP policies as described in section Disabling and Deleting Policies.
- [When uninstalling the filter] If you have manually assigned the Veeam CDP Replication storage policy to VMs that are parts of organization VDCs to which the filter was installed, or replicas are still present in the Veeam Backup & Replication configuration database, you must change the storage policy for these VMs.

  For more information on how to change storage policies, see VMware Docs. To see the list of replicas, open the **Home** view. In the inventory pane, click the **Replicas** node.

## Updating or Uninstalling I/O Filter

To update or uninstall the I/O filter, do the following:

1. Launch the **I/O Filter Management** wizard as described in section Launch VeeamCDP Filter Management Wizard.
2. At the **Organization VDC** step of the wizard, do the following:
   - To update the filter, make sure that check boxes are selected near the necessary organization VDCs.
   - To uninstall the filter, clear the check boxes near the necessary organization VDCs.

3. Proceed to the last step of the wizard and close the wizard.

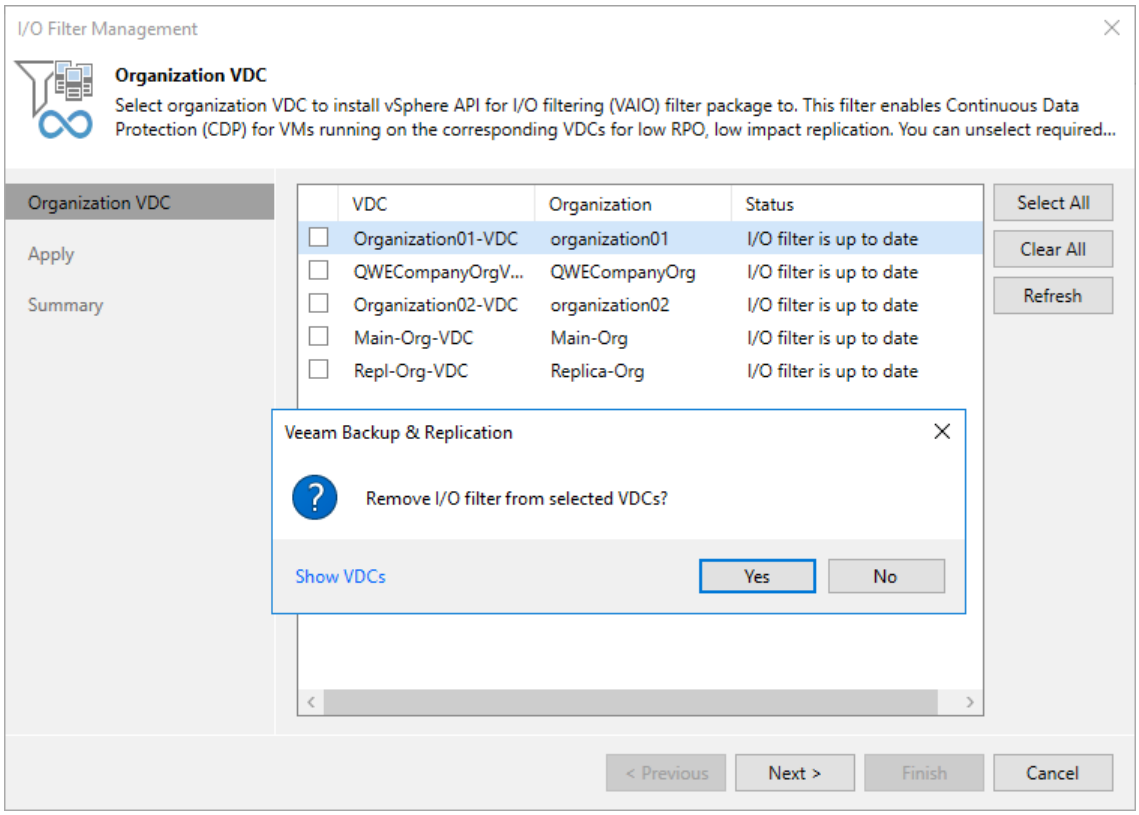

# Creating Cloud Director CDP Policies

To protect vApps with CDP, you must configure a VMware Cloud Director CDP policy. The CDP policy defines which vApps to protect, where to store replicas, how often create short-term and long-term restore points, and so on. One CDP policy can process one or multiple vApps.

To create a CDP policy, use the **New VMware Cloud Director CDP Policy** wizard.

## Before You Begin

Before you create a CDP policy, check the following prerequisites:

- Check that all the required components are added to the backup infrastructure. For more information on the required components, see Backup Infrastructure for Cloud Director CDP.
- The I/O filter must be installed on each VDC where VMs that you plan to protect reside. For more information on how to install the filter, see Installing I/O Filter on VDCs.
- If you plan to use replica seeding, you must create a seed as described in section Creating vApp Replica Seeds.

# Step 1. Launch New VMware Cloud Director CDP Policy Wizard

To launch the **New VMware Cloud Director CDP Policy** wizard, do one of the following:

- Open the **Home** view. On the ribbon, click **CDP Policy** > **VMware Cloud Director**.
- Open the **Home** view. In the inventory pane, right-click **Jobs** and select **CDP Policy** > **VMware Cloud Director**.
- Open the **Inventory** view. In the inventory pane, right-click workloads that you want to replicate. Select **Add to CDP policy** > **New job** if you want to create a new VMware Cloud Director CDP policy, or **Add to CDP policy** > *<Policy Name>* if you want to add workloads to an existing VMware Cloud Director CDP policy.

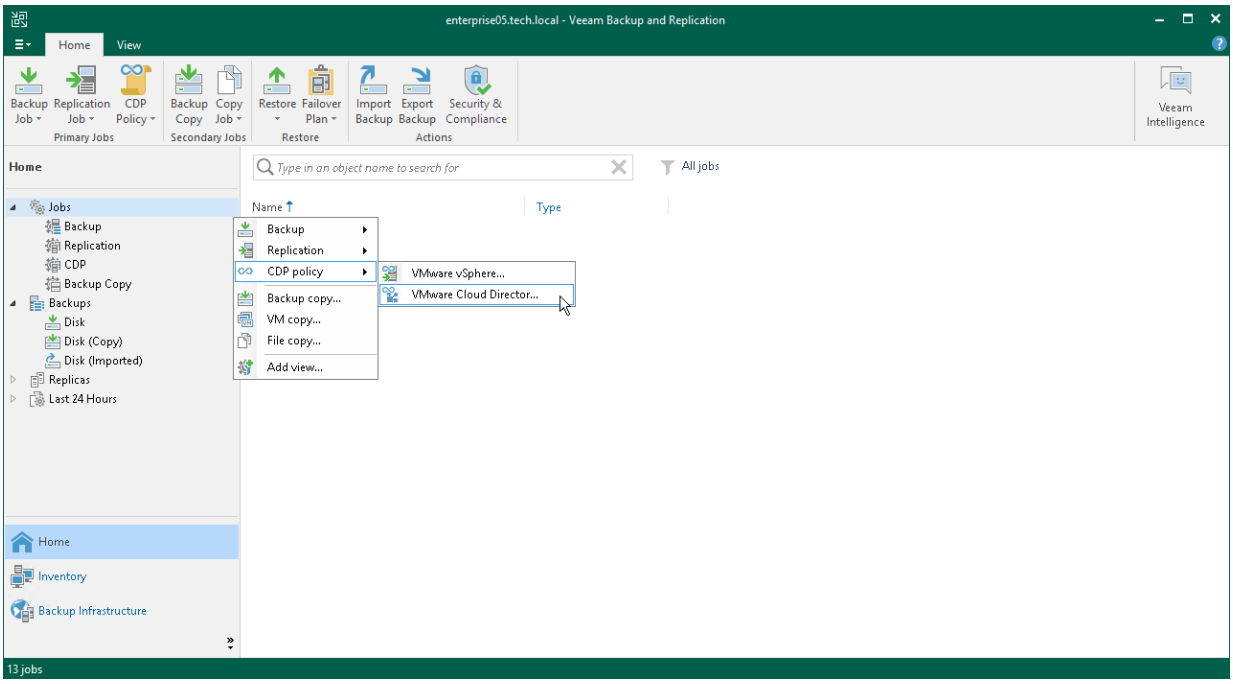

# Step 2. Specify Policy Name and Advanced Settings

At the **Name** step of the wizard, specify a name and description for the CDP policy, and choose whether you want to use replica seeding or network mapping:

1. In the **Name** field, enter a name for the CDP policy.
2. In the **Description** field, provide a description for future reference.
3. If a network between your production and disaster recovery (DR) sites has low bandwidth, and you want to reduce the amount of traffic sent during the initial synchronization of the CDP policy, select the **Replica seeding (for low bandwidth DR sites)** check box.

   When selected, this check box enables the **Seeding** step where you will have to configure replica seeding and mapping.
4. If your DR site networks do not match your production site networks, select the **Network remapping (for DR sites with different virtual networks)** check box.

   When selected, this check box enables the **Network** step where you will have to configure a network mapping table.

> **NOTE**
>
> Cloud Director CDP policies do not support network mapping of the vApp networks. You can configure a mapping table for organization VDC network only.

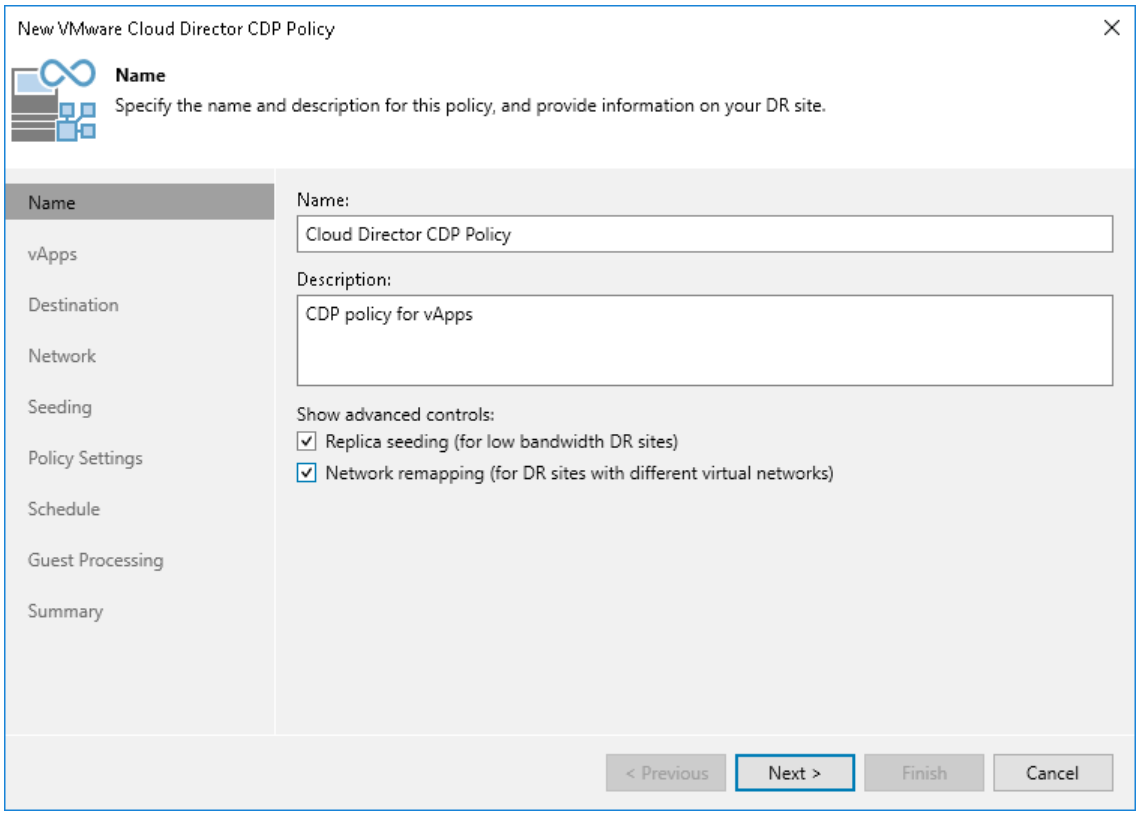

# Step 3. Select vApps to Replicate

At the **vApps** step of the wizard, select VM containers (vApps, organization or organization VDCs) that you want to replicate.

1. Click **Add**.
2. In the **Selects Objects** window, select the necessary VM containers. Click **Add**.

   When you add new items to VM containers, Veeam Backup & Replication updates settings automatically to process these new items.

**IMPORTANT**

Consider the following:

- You cannot replicate standalone VMs. You can replicate only VM containers (vApps, organization VDCs and so on).
- If you replicate VM containers and add new VMs to these container in future, Veeam Backup & Replication will update policy settings automatically to include these VMs. However, note that only vApps and VMs are replicated. VM templates, logs, folders and so on are not replicated.
- You can replicate only turned on VMs added to vApps. The turned off VMs will be skipped from processing.
- You cannot add to a CDP policy vApps that were already added to other CDP policies created on the same backup server.

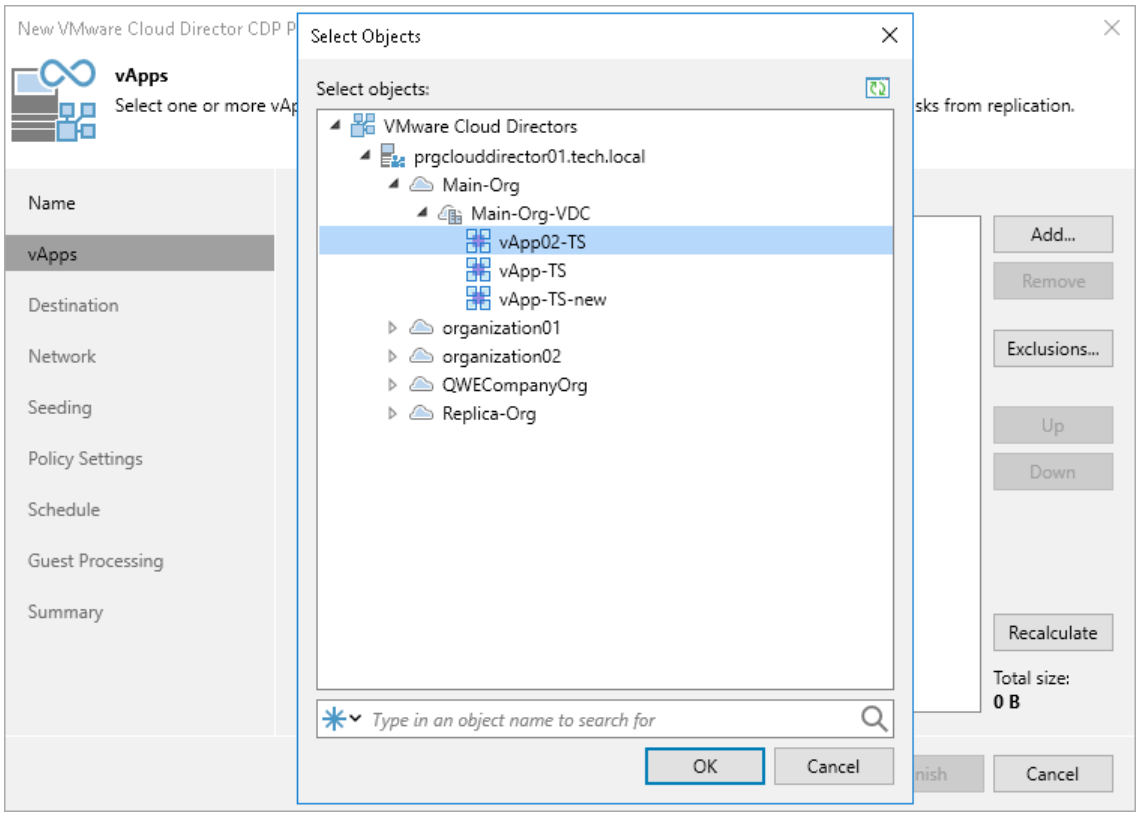

# Step 4. Exclude Objects

After you have added vApps and VM containers to the replication job, at the **vApps** step of the wizard you can specify which objects you want to exclude from the replication job.

> **NOTE**
>
> You can exclude only child objects that are added to a VM container. For example, if you add an organization VDC to a replication job, you can exclude only vApps that are available in this organization VDC.

# Excluding VMs and VM Containers

To exclude VMs and VM containers (vApps, organizations, organization VDCs and so on):

1. At the **vApps** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, check that the **VMs** tab is selected. Click **Add**.
3. In the **Select Objects** window, select VMs and VM containers that you want to exclude. Click **OK**.

   Select the **Show full hierarchy** check box to display the hierarchy of all VMware Cloud Director servers added to the backup infrastructure.

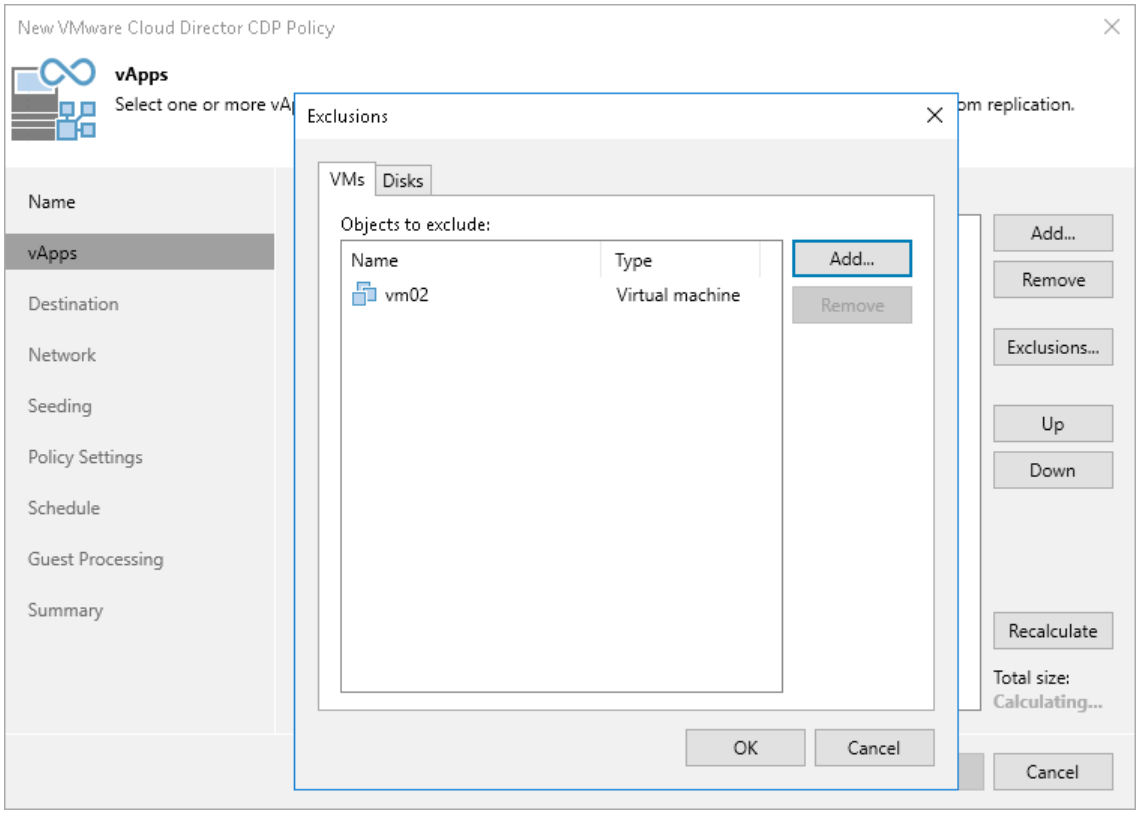


# Excluding Disks

To exclude disks:

1. At the **vApps** step of the wizard, click **Exclusions**.

2. In the **Exclusions** window, do the following:
   a. Switch to the **Disks** tab.
   b. To exclude disks of VMs, click **Add**. In the **Add Objects** window, select the necessary VMs and click **Add**. Veeam Backup & Replication will include these VMs in the list as standalone objects.
   c. In the **Disks to process** list, select the necessary VMs.
   d. Click **Edit**.
3. In the **Select Disks** window, select disks that you want to replicate: all disks, 0:0 disks (as a rule, system disks) or specific IDE, SCSI, SATA or NVMe disks. Disks that you do not select will be excluded from processing. Click **OK**.
4. In the **Exclusions** window, click **OK**.

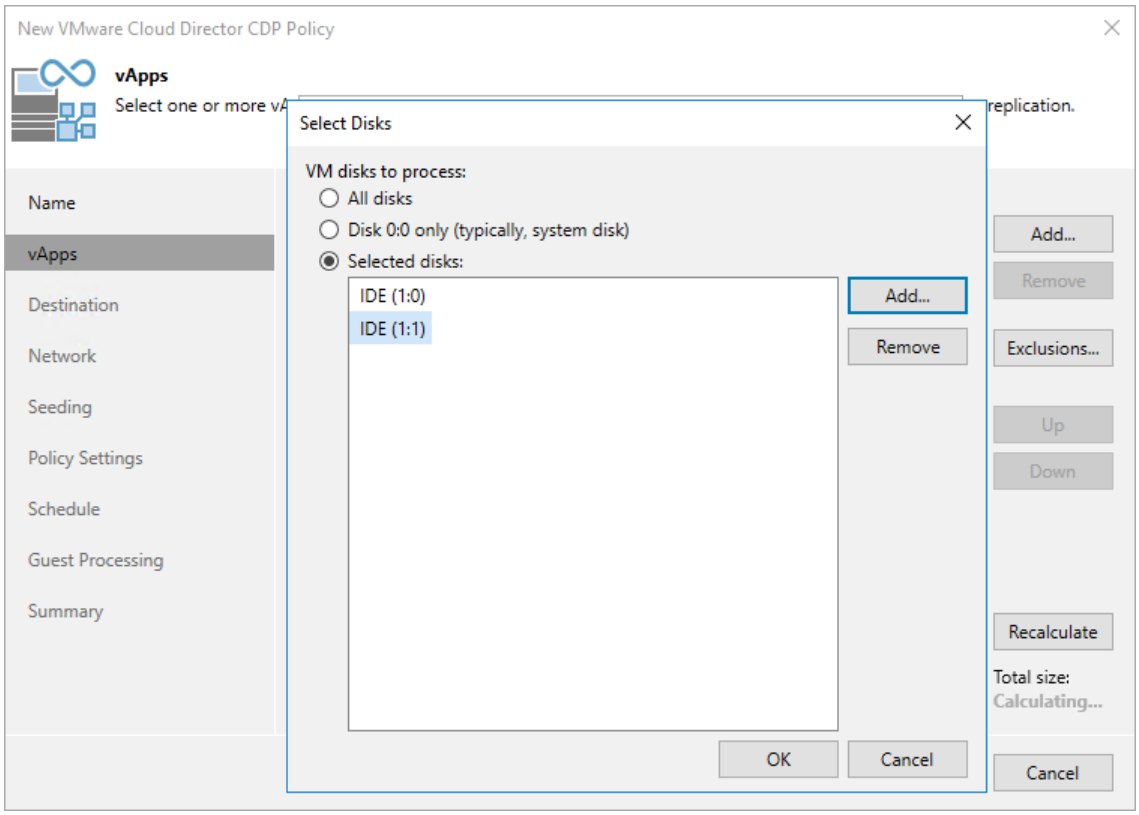

# Step 5. Specify vApp Processing Order

At the **vApps** step of the wizard, click **Up** and **Down** to change the processing order. VM containers (vApps, organization, organization VDCs, and so on) at the top of the list have a higher priority and will be processed first.

> **NOTE**
>
> Consider the following:
>
> - VMs inside a VM container are processed at random.
> - To ensure that vApps are processed in the specified order, you must add them as standalone vApps, not as a part of containers.
> - The processing order may differ from the order that you have specified. For example, if resources of a vApp that is higher in the priority are not available, and resources of a vApp that is lower in the priority are available, Veeam Backup & Replication will process the vApp with the lower priority first.

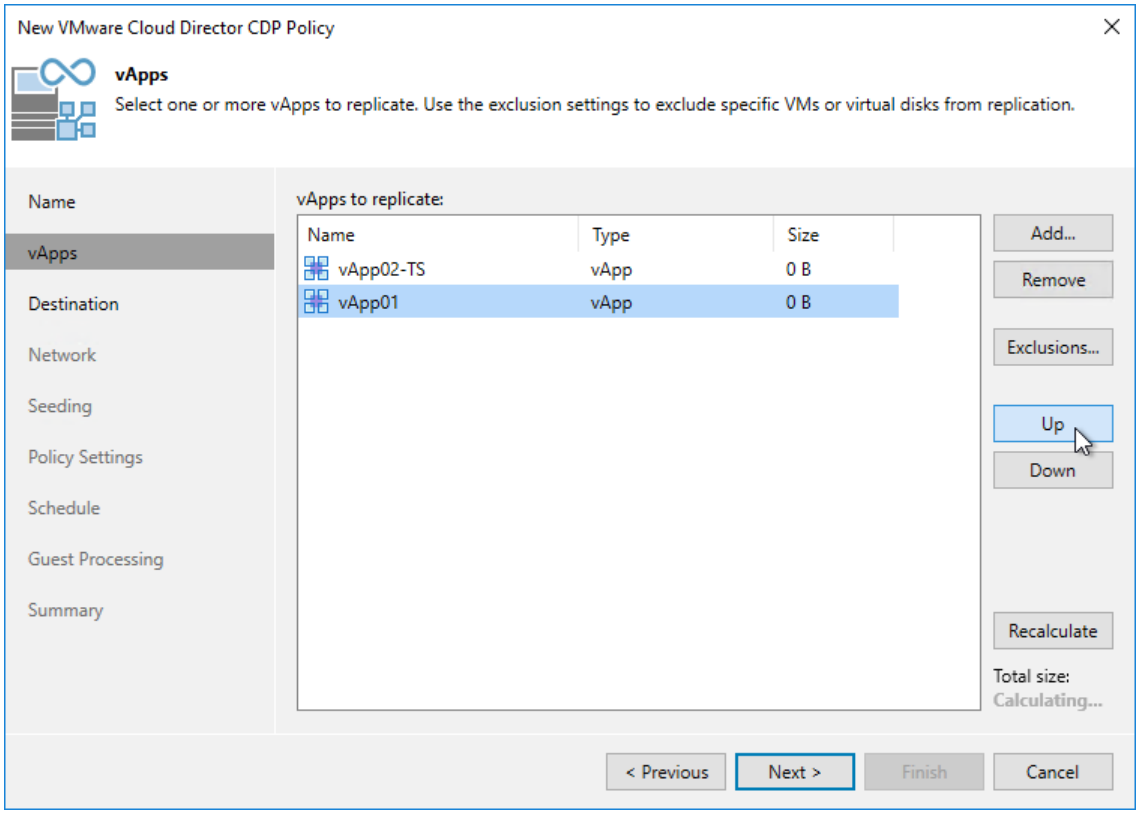

# Step 6. Select Replica Destination

At the **Destination** step of the wizard, select an organization VDC where you want to store replicas and the storage policy that you want to apply to the replicas.

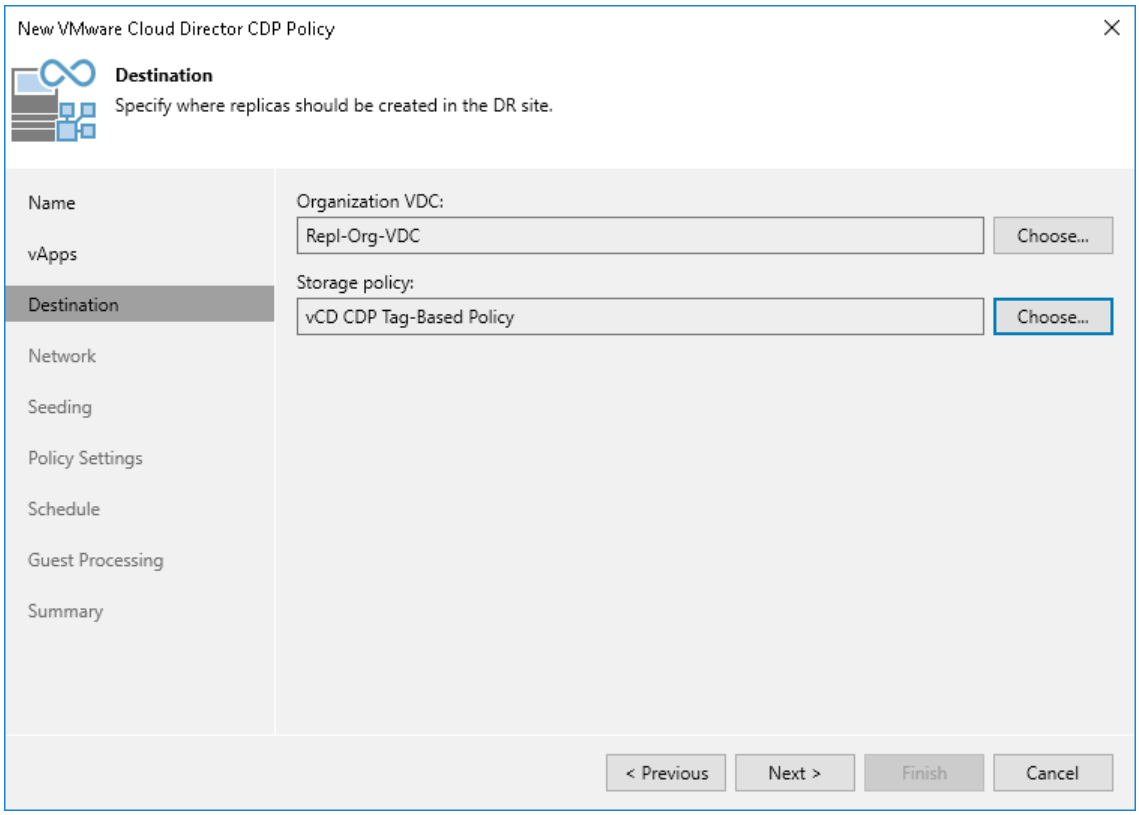

# Step 7. Configure Network Mapping

The **Network** step of the wizard is available if you have selected the **Network remapping** option at the **Name** step of the wizard.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the production site to networks in the disaster recovery (DR) site. When the replication session starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update replica configuration to replace the production networks with the specified networks in the DR site. As a result, you will not have to re-configure network settings manually.

To add a row to the network mapping table:

1. Click **Add**.
2. In the **Network Mapping** window, click **Browse** next to the **Source network** field.
3. In the **Select Network** window, select the production network to which the source workloads are connected and click **OK**.
4. In the **Network Mapping** window, click **Browse** next to the **Target network** field.
5. In the **Select Network** window, select a network in the DR site to which replicas will be connected and click **OK**.
6. In the **Network Mapping** window, click **OK**.

NOTE

Consider the following:

- Cloud Director CDP policies do not support network mapping of vApp networks. You can configure a mapping table for organization VDC networks only.
- If your replica has a configured network remapping and IP addresses are set as a static pool on the target, Veeam Backup & Replication will assign these IP addresses to VM NICs on a "first come, first served" basis: the first available IP address will be assigned to the first NIC, the second available IP address will be assigned to the second NIC and so on.

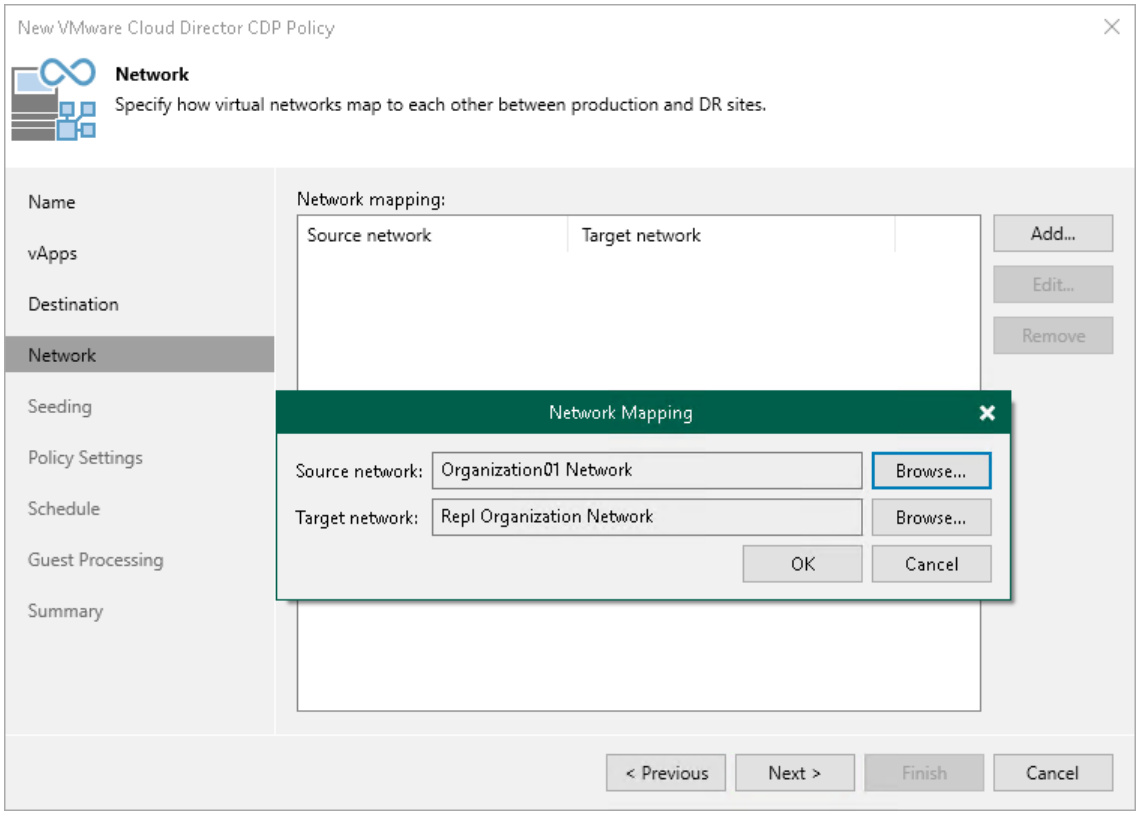
New VMware Cloud Director CDP Policy
Network
Specify how virtual networks map to each other between production and DR sites.
Name
vApps
Destination
Network
Seeding
Policy Settings
Schedule
Guest Processing
Summary
Network mapping:
Source network
Target network
Add...
Edit...
Remove
Network Mapping
Source network:
Organization01 Network
Browse...
Target network:
Repl Organization Network
Browse...
OK
Cancel
< Previous
Next >
Finish
Cancel

# Step 8. Configure vApp Seeding and Mapping

The **Seeding** step is available if you have selected the **Replica seeding** check box at the **Name** step of the wizard.

At the **Seeding** step of the wizard, configure replica seeding and mapping. Seeding and mapping help reduce the amount of traffic sent during the initial replica synchronization. For more information on when to use seeding and mapping, see vApp Seeding and Mapping.

> **IMPORTANT**
>
> If the **Replica seeding** check box is enabled in a policy, all vApps in the policy must be covered with seeding or mapping. If a vApp neither has a seed, nor has mapping to an existing vApp, it will be skipped from processing.

# Configuring Replica Seeding

To configure replica seeding:

1. Make sure that you have backups of replicated vApps in a backup repository in the DR site. If you do not have the backups, create them as described in section Creating vApp Replica Seeds.

   > **IMPORTANT**
   >
   > Consider the following:
   >
   > - Backups must be created by Veeam Backup & Replication.
   > - When you start replication, Veeam Backup & Replication will attempt to restore all VMs added to replication from the vApp seed that you have specified. If a VM is not found in the vApp seed, the VM will be skipped from replication.
   > - Backups must not reside in a scale-out backup repository.

2. In the **Initial seeding** section, select the **Get seed from the following backup repository** check box.
3. From the list of available backup repositories, select the repository where your replica seeds are stored.

> **NOTE**
>
> If a vApp has a seed and is mapped to an existing replica, replication will be performed using replica mapping because mapping has a higher priority.

# Configuring Replica Mapping

To configure replica mapping:

1. Select the **Map replicas to existing vApps** check box.
2. If you want Veeam Backup & Replication to scan the DR site to detect existing copies of vApps that you plan to replicate, click **Detect**.

   If any matches are found, Veeam Backup & Replication will populate the mapping table. If Veeam Backup & Replication does not find a match, you can map a vApp to its copy manually. Note that the mapping list does not display vApps added to the list of exclusions for replication.

3. If you want to map a VM manually, select a source VM from the list, click **Edit** and select the copy of this VM on the target host in the DR site.

> **TIP**
>
> If there is no existing replica in the DR site, you can restore a vApp from the backup and map it to the source vApp.

To remove a mapping association, select a VM in the list and click **Remove**.

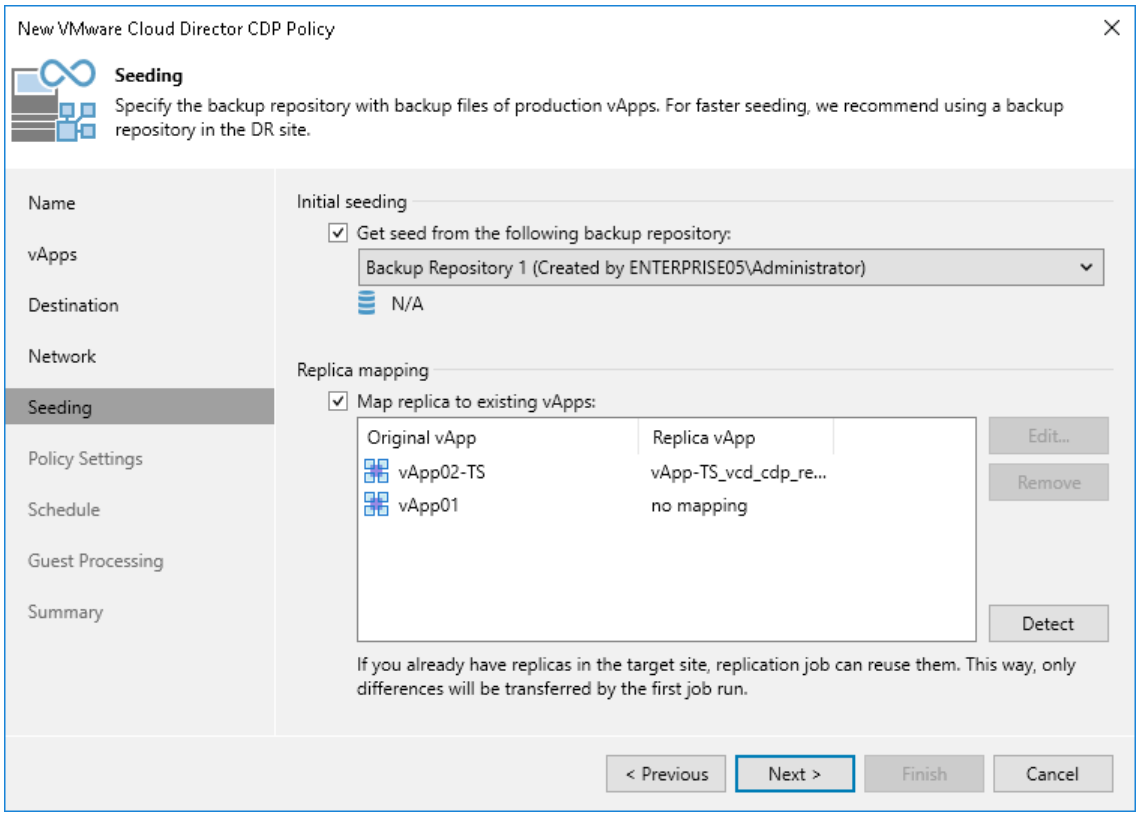

# Step 9. Specify Data Transfer and Replica Settings

At the **Policy Settings** step of the wizard, select VMware CDP proxies that must be used for the CDP policy, specify which suffix to add to replica names and VMware CDP proxies availability:

1. Specify which VMware CDP proxies you want to use:
   - If you want Veeam Backup & Replication to select proxies automatically, leave **Automatic selection** in the **Source proxy** and **Target proxy** fields.

     Veeam Backup & Replication will assign VMware CDP proxies for VM processing one by one. Before processing a new VM from the list, Veeam Backup & Replication will check available VMware CDP proxies.
   - If you want to select VMware CDP proxies manually, do the following:
     i. Click **Choose** next to the **Source proxy** field if you want to select VMware CDP proxies in the production site, or next to the **Target proxy** field if you want to select VMware CDP proxies in the disaster recovery site.
     ii. In the **Backup Proxy** window, click **Use the selected backup proxy servers only**. Select proxies that you want to use and click **OK**.

   > **NOTE**
   >
   > We recommend that you deploy at least two VMware CDP proxies: one CDP proxy in the production site and one CDP proxy in the disaster recovery site.

2. To test whether VMware CDP proxies available in the backup infrastructure can handle replication, click **Test**.

   Veeam Backup & Replication will analyze available CPU on all source and all target VMware CDP proxies, the maximum VM disk write speed during the last hour, and will calculate approximate requirements for VMware CDP proxies. In the **CDP Infrastructure Assessment** window, you will see the calculated values:
   - The **CPU** rows show CPU cores available on all proxies (source or target).
   - The **Proxy RAM** rows show RAM required for CDP and, in parenthesis, RAM available on all proxies (source or target). If values in the parentheses and near the paranthesis are the same, you need to upgrade proxies for which values coincide to provide more resources. For example, you can double up the amount of RAM.
   - The **Proxy Bandwidth** rows show the maximum disk write speed during the last hour and, in parenthesis, available bandwidth based on available cores of source or target proxies.

3. In the **Replica name suffix** field, specify a suffix that will be added to names of replicas.

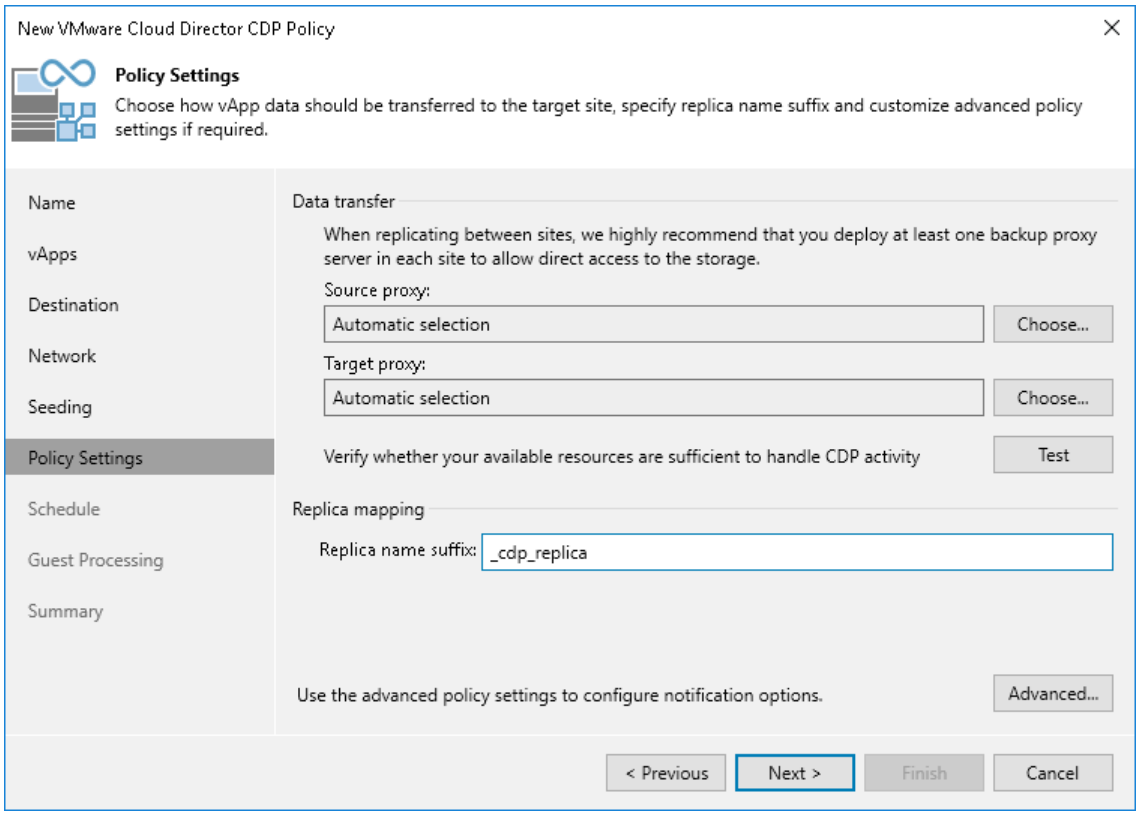

# Step 10. Specify Notification Settings

Veeam Backup & Replication can send by email two types of notifications for CDP policies: session reports and RPO reports.

Veeam Backup & Replication sends session reports after a CDP policy session stops. This report includes information on the policy during the session, for example read and transferred data. The session report is configured using the Global Email Notification Settings. Veeam Backup & Replication sends RPO reports after the configured RPO period ends. This report contains information on the maximum delay, SLA and other information. The RPO report is configured at the **Policy Settings** step of the wizard.

At the **Policy Settings** step of the wizard, specify RPO notification settings:

1. At the lower right corner, click **Advanced**.
2. To receive SNMP traps on the CDP policy, select the **Send SNMP notifications for this job** check box.

   SNMP traps will be sent if you configure global SNMP settings in Veeam Backup & Replication and configure software on recipient machine to receive SNMP traps. For more information, see Specifying SNMP Settings.
3. To receive notifications by email in case of policy failure, success or warning, select the **Send email notifications to the following recipients** check box. Then configure notification settings:
   a. Check that you have configured global email notification settings as described in section Configuring Global Email Notification Settings.
   b. In the text field, specify a recipient email address. If you want to specify multiple addresses, separate them by a semicolon.
   c. To use global notification settings, select **Use global notification settings**.
   d. To specify a custom notification subject and redefine at which time notifications must be sent, select **Use custom notification settings specified below.** Then specify the following settings:
      i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%JobResult%, %JobName%*, *%ObjectCount%* (number of VMs in the policy) and *%Issues%* (number of VMs in the policy that have been processed with the *Warning* or *Failed* status).
      ii. Select the **Notify on success**, **Notify on error** or **Notify on warning** check boxes to receive email notification if the policy gets the *Warning*, *Success* of *Error* status.

> **NOTE**
>
> A CDP policy will get the *Warning* or *Error* status according to the reporting settings configured at the **Schedule** step of the wizard. The policy will get the *Success* status after the initial configuration succeeds and every day at 8 A.M. if no error or warning occurs.

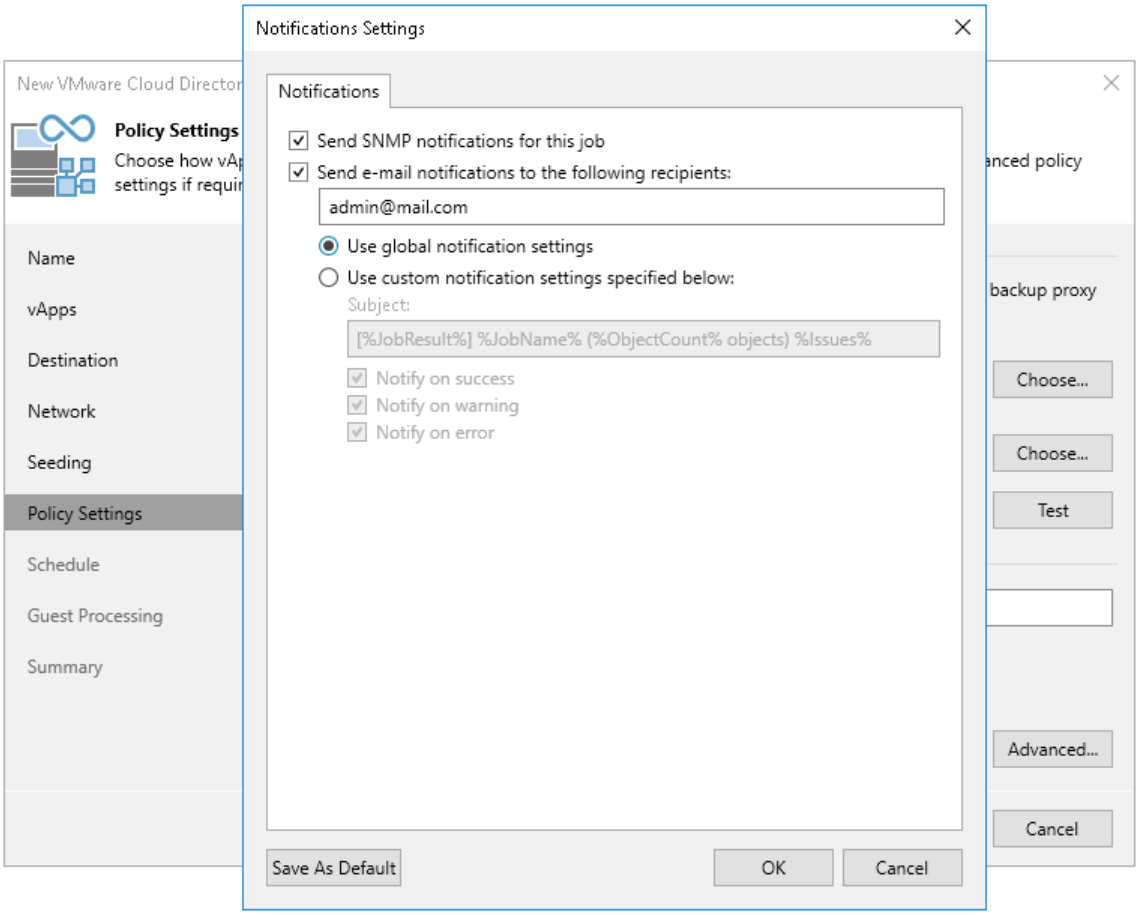
New VMware Cloud Director
Policy Settings
Choose how vA
settings if requi
Name
vApps
Destination
Network
Seeding
Policy Settings
Schedule
Guest Processing
Summary
anced policy
backup proxy
Choose...
Choose...
Test
Advanced...
Cancel
Notifications Settings
Notifications
Send SNMP notifications for this job
Send e-mail notifications to the following recipients:
admin@mail.com
Use global notification settings
Use custom notification settings specified below:
Subject:
[%JobResult%] %JobName% (%ObjectCount% objects) %Issues%
Notify on success
Notify on warning
Notify on error
Save As Default
OK
Cancel

# Step 11. Configure Schedule

At the **Schedule** step of the wizard, configure the schedule and retention policies:

1. Specify scheduling options:

   a. In the **Recovery point objective** field, specify the necessary RPO in seconds or minutes, that is, how often to create short-term restore points.

      The minimum RPO is 2 seconds, however it can be not optimal if your CDP policy contains many VMs with high workload. The optimal RPO is not less than 15 seconds. The maximum RPO is 60 minutes.

      During every specified period, Veeam Backup & Replication will prepare data for short-term restore points for VM replicas and send this data to the target destination. Note that short-term restore points are crash-consistent.

   b. If you want to prohibit the policy to run at specific time intervals, click **Schedule**. In the schedule box, click **Denied** and select the necessary time area.

2. To instruct the CDP policy to display a warning or error if a newly created restore points are not transferred to the target within the set RPO, click **Reporting**. Then specify when the policy must display error and warning.

   If you have configured email notification settings, Veeam Backup & Replication will mark the policy with the *Warning* or *Error* status and will also send email notifications.

3. In the **Short-term retention** section, configure the short-term retention policy, that is, specify for how long to store short-term restore points.

4. In the **Long-term retention** section, specify when to create long-term restore points and for how long to store them:

   a. In the **Create additional restore points every** field, specify how often you want to create long-term restore points.

   b. In the **Keep restore points for** field, specify for how long to retain these long-term restore points.

c. To specify time periods when Veeam Backup & Replication must create application-consistent and crash-consistent long-term restore points, click **Schedule**. In the schedule box, click **Crash-consistent** or **Application-consistent** and select the necessary time area. By default, Veeam Backup & Replication creates application-consistent backups if you enable application-aware processing at the **Guest Processing** step of the wizard. If you do not enable application-aware processing, Veeam Backup & Replication will create crash-consistent long-term restore points.

   If you want to shift the schedule, specify the offset in the **Start time within an hour** field. For example, you schedule creation of crash-consistent restore points from 00:00 to 01:00, and set the offset value to 25. The schedule will be shifted forward, and the crash-consistent restore points will be created from 0:25 and to 01:25.

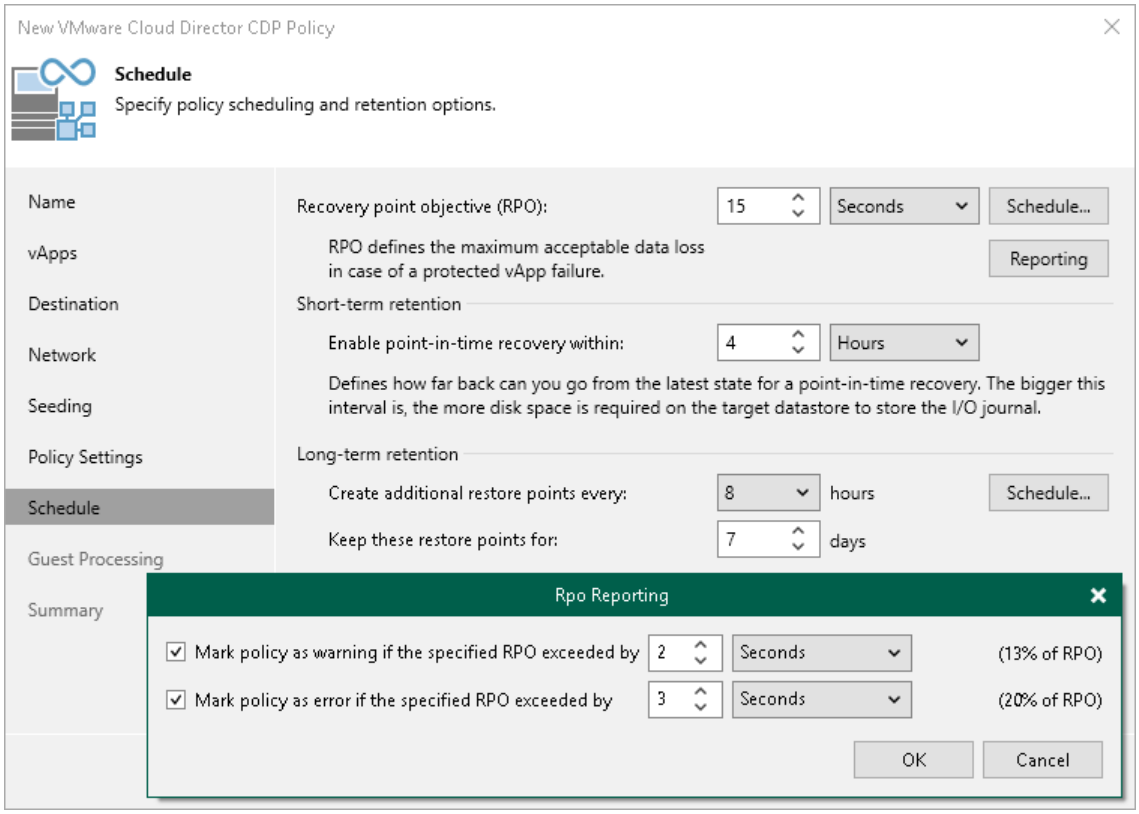

# Step 12. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, enable and configure guest OS processing.

Guest OS processing involves application-aware processing that allows creation of transactionally consistent replicas and guest file system indexing (however, indexing is not available for replicas). In its turn, application-aware processing includes log truncation, execution of custom scripts and guest OS file exclusions. For more information on guest processing, see the Guest Processing section.

To be able to use guest processing, you must also configure user accounts to access guest OSes and guest interaction proxies.

To enable guest OS processing and start configuring it (accounts and guest interaction proxies):

1. Select **Enable application-aware processing**.

   When you select this option, Veeam Backup & Replication enables application-aware processing with the default settings for all VMs. You can further disable application-aware processing for individual VMs and reconfigure the default settings.

2. If you have added Microsoft Windows VMs to be processed, specify which guest interaction proxy Veeam Backup & Replication can use to perform different guest processing tasks:

   - If you want Veeam Backup & Replication to select the guest interaction proxy automatically, leave **Automatic selection** on the **Guest interaction proxy** field.
   - If you want to explicitly specify which servers will perform the guest interaction proxy role, click **Choose**. In the **Guest Interaction Proxy** window, click **Prefer the following guest interaction proxy server**, and select the necessary proxies.

   For more information on the guest interaction proxy, requirements and limitations for it, see Guest Interaction Proxies.

3. From the **Guest OS credentials** list, select a user account that will be used to connect to guest OSes and that has enough permissions. For more information on the permissions and requirements for the user account, see Permissions for Guest Processing.

   [For Microsoft Windows VMs] Veeam Backup & Replication will also use this account to deploy the non-persistent runtime components or use (if necessary, deploy) persistent agent. For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

   [For Linux VMs] If you installed persistent agent components for VMs running Linux or Unix operating systems, select *Use management agent credentials* from the list. For more information, see Persistent Agent Components.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. For more information on adding credentials, see the Credentials Manager section.

   > **NOTE**
   >
   > If you plan to use Kerberos authentication, check limitations and requirements listed in section Guest Processing.

4. To specify credentials for individual workloads, click **Credentials**. Then select the necessary workload and set user credentials for it.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose a VM whose settings you want to customize. Then select the VM in the list and define the necessary settings.

5. To check whether Veeam Backup & Replication can connect to VMs using the specified guest OS credentials and can deploy the non-persistent runtime components or connect to persistent agent components on the guest OSes, click **Test Now**.

After you have enabled application-aware processing for all VMs and configured other settings required for guest processing, you can disable application-aware processing for individual VMs and change the default settings. For more information, see the following sections:

- Application-aware Processing
- Microsoft SQL Server transaction log settings
- Oracle archived log settings
- Script settings

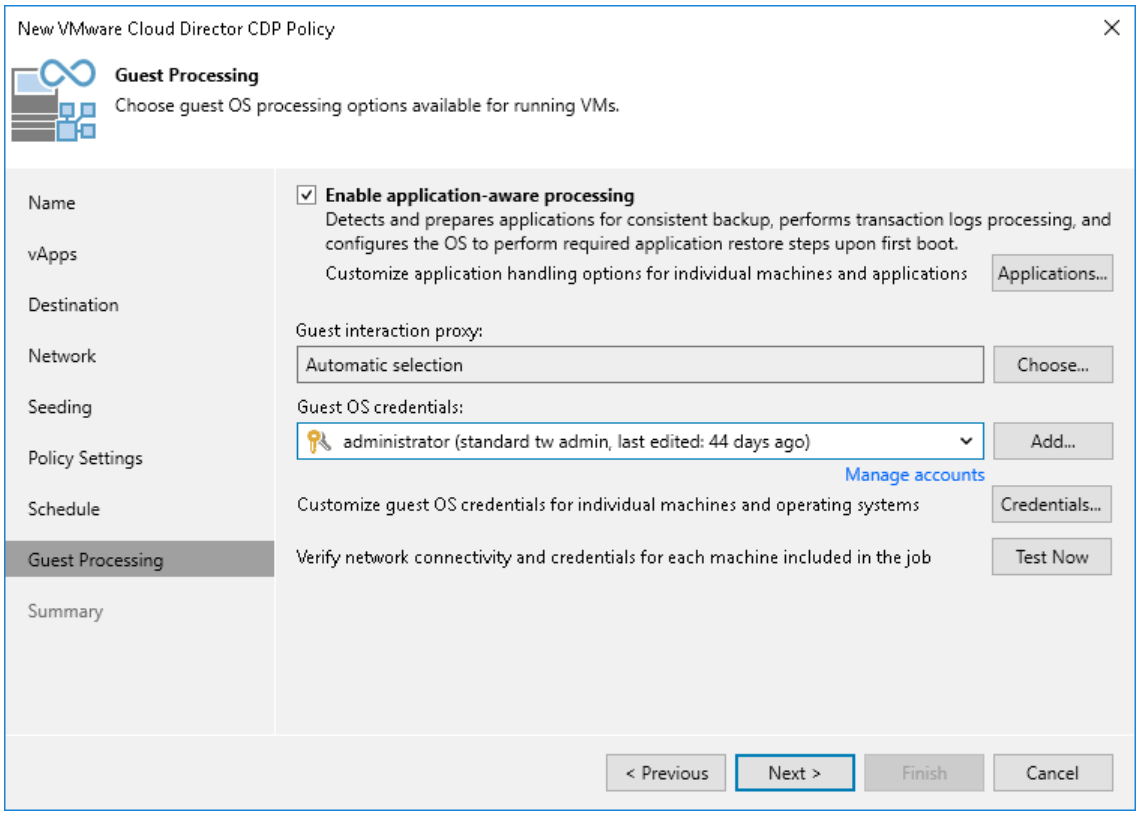


## Application-Aware Processing and Transaction Logs

Application-aware processing helps create transactionally consistent replicas. The transactionally consistent replicas guarantee proper recovery of applications without data loss. For more information on application-aware processing, see Application-Aware Processing.

To configure general application-aware processing settings and specify whether Veeam Backup & Replication processes transaction logs or creates copy-only replicas:

1. At the **Guest Processing** step of the wizard, check that you have selected the **Enable application-aware processing** check box and configured guest OS credentials.
2. At the **Guest Processing** step of the wizard, click **Applications**.

3. In the **Application-Aware Processing Options** window, select workloads for which you want to configure application-aware processing, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.

4. On the **General** tab, in the **Applications** section, specify the behavior scenario for application-aware processing:

   - Select **Require successful processing** if you want Veeam Backup & Replication to stop the replication process if any error occurs during application-aware processing.
   - Select **Try application processing, but ignore failures** if you want to continue the replication process even if an error occurs during application-aware processing. This option guarantees that replication will continue working. However, the resulting replica will be crash consistent, not transactionally consistent.
   - Select **Disable application processing** if you want to disable application-aware processing for the workload.

5. [For Microsoft Exchange and Microsoft SQL Server] In the **VSS Settings** section, specify if Veeam Backup & Replication must process transaction logs or create copy-only replicas:

   a. Select **Process transaction logs with this job** if you want Veeam Backup & Replication to process transaction logs.

      [For Microsoft Exchange VMs] With this option selected, the non-persistent runtime components or persistent components running on the VM guest OS will wait for replication to complete successfully and then trigger truncation of transaction logs. If the replication job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

      [For Microsoft SQL Server] You will need to configure how to process transaction logs.

      > **TIP**
      >
      > To configure log processing for Oracle and PostgreSQL databases, switch to the Oracle and PostrgeSQL tabs.

   b. Select **Perform copy only** if you use another tool to perform guest level processing, and this tool maintains consistency of the database state. Veeam Backup & Replication will create a copy-only replica for the selected VMs. The copy only replica preserves the chain of full and differential files and transaction logs on the VM. For more information, see Microsoft Docs.

6. [For Microsoft Windows VMs] In the **Persistent guest agent** section, select the **Use persistent guest agent** check box to use for application-aware processing persistent guest agents on each protected VM.

   By default, Veeam Backup & Replication uses non-persistent runtime components.
   Veeam Backup & Replication deploys runtime components on each protected VM when the replication job starts, and removes the runtime components as soon as the replication job finishes.

   For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

[For Linux VMs] To use persistent guest agents, you must install Management Agent on protected VMs. For more information, see Persistent Agent Components.

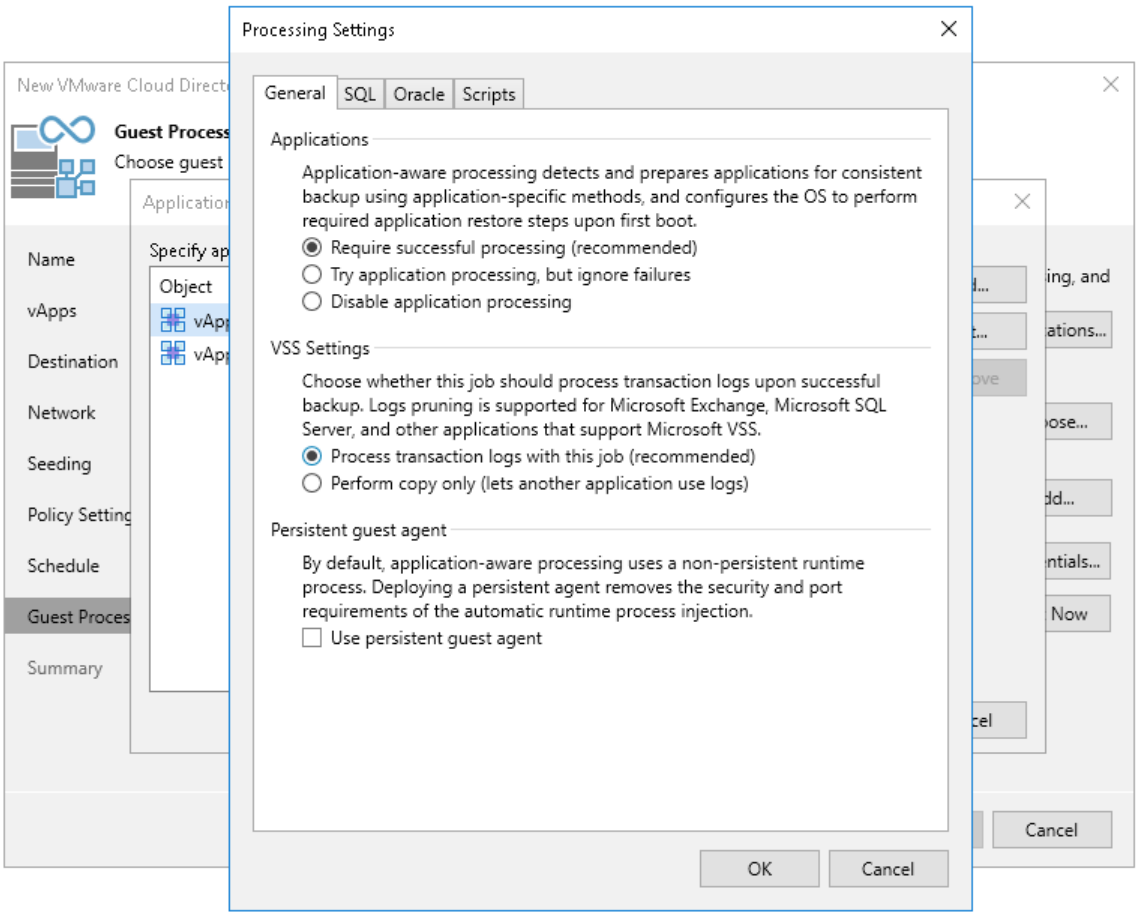


## Microsoft SQL Server Transaction Log Settings

The **SQL** tab is available for VMs that run Microsoft SQL Server and if you have selected **Process transaction logs with this job** when configuring application-aware processing.

To create transactionally consistent backups of an Microsoft SQL Servers, you must check that application-aware processing is enabled and then specify settings of transaction log processing.

# Enabling Application-Aware Processing

Before configuring transaction log processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Microsoft SQL Server and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

# Specifying Transaction Log Settings

In the **Processing Settings** window, switch to the **SQL** tab and specify how transaction logs must be processed:

- If you want Veeam Backup & Replication to trigger truncation of transaction logs after the CDP policy creates a long-term restore point, select **Truncate logs**.

  In this case, transaction logs will be truncated after the CDP policy creates a long-term restore point. If the creation fails, the logs will remain untouched until the next start of the long-term restore point creation.

- If you do not want Veeam Backup & Replication to truncate logs at all, select **Do not truncate logs**.

  This option is recommended if you use another tool to perform VM guest-level replication, and this tool maintains consistency of the database state.

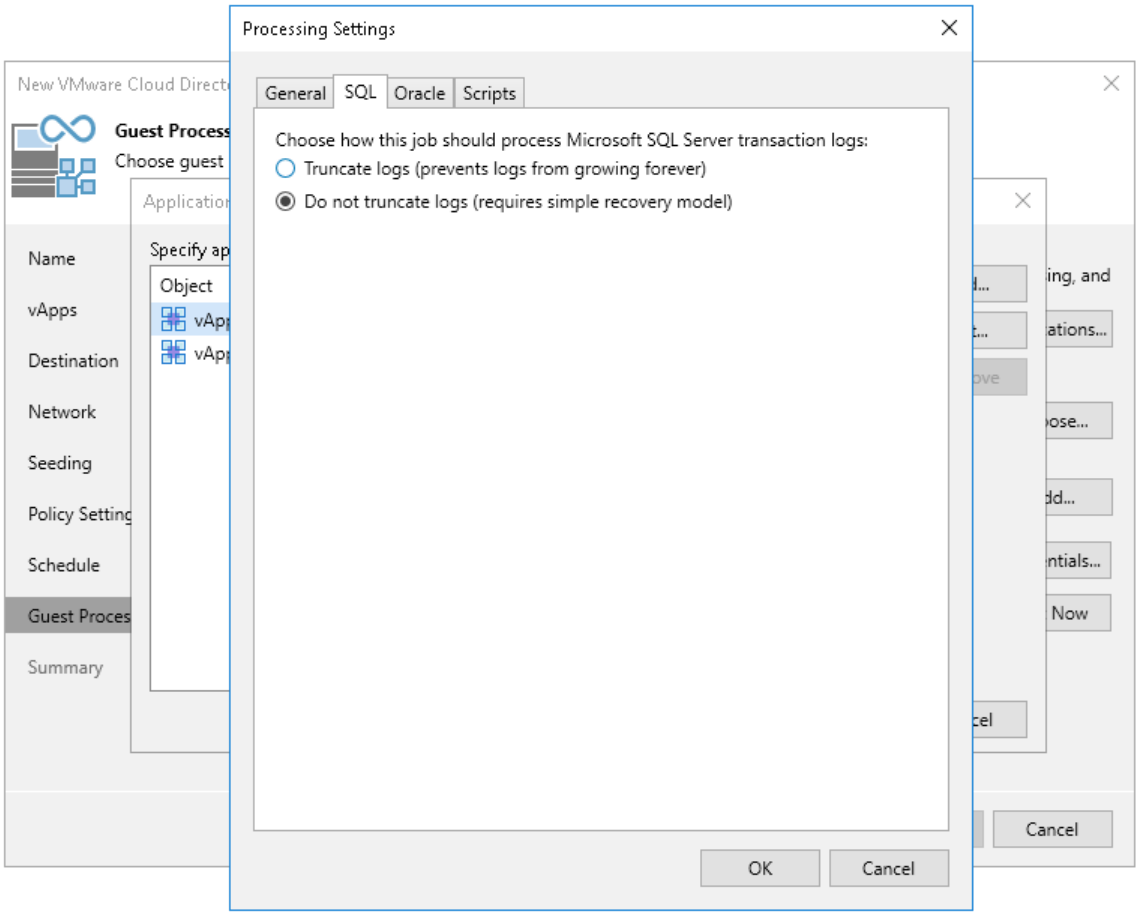


## Oracle Archived Log Settings

The **Oracle** tab applies to Linux VMs that run Oracle.

To create transactionally consistent backups of an Oracle server, you must check that application-aware processing is enabled and then specify settings of archive log processing.

# Enabling Application-Aware Processing

Before configuring archive log processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.

2. Click **Applications**.
3. In the displayed list, select the Oracle server and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

# Specifying Archive Log Settings

To configure how Veeam Backup & Replication must process archive logs of an Oracle server:

1. In the **Processing Settings** window, switch to the **Oracle** tab.
2. From the **Specify Oracle account with SYSDBA privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the Oracle databases. The account that you plan to use must have privileges described in section Permissions.

   You can select **Use guest credentials** from the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle databases.
3. In the **Archived logs** section, specify how to process archived logs:
   - If you want to preserve archived logs on the VM guest OS, select **Do not delete archived logs**. When the replication job completes, the non-persistent runtime components or persistent components will not truncate transaction logs.

     It is recommended that you select this option for databases where the ARCHIVELOG mode is turned off. If the ARCHIVELOG mode is turned on, archived logs may grow large and consume all disk space.
   - If you want to delete archived logs older than <N> hours, select **Delete logs older than <N> hours** and specify the number of hours.

- If you want to delete archived logs larger than <N> GB, select **Delete logs over <N> GB** and specify the size. The specified size refers to the log size of each database, not all databases on the selected Oracle server.

Transaction logs will be deleted using Oracle Call Interface after the CDP policy creates a long-term restore point. If the creation fails, the logs will remain untouched until the next start of the long-term restore point creation.

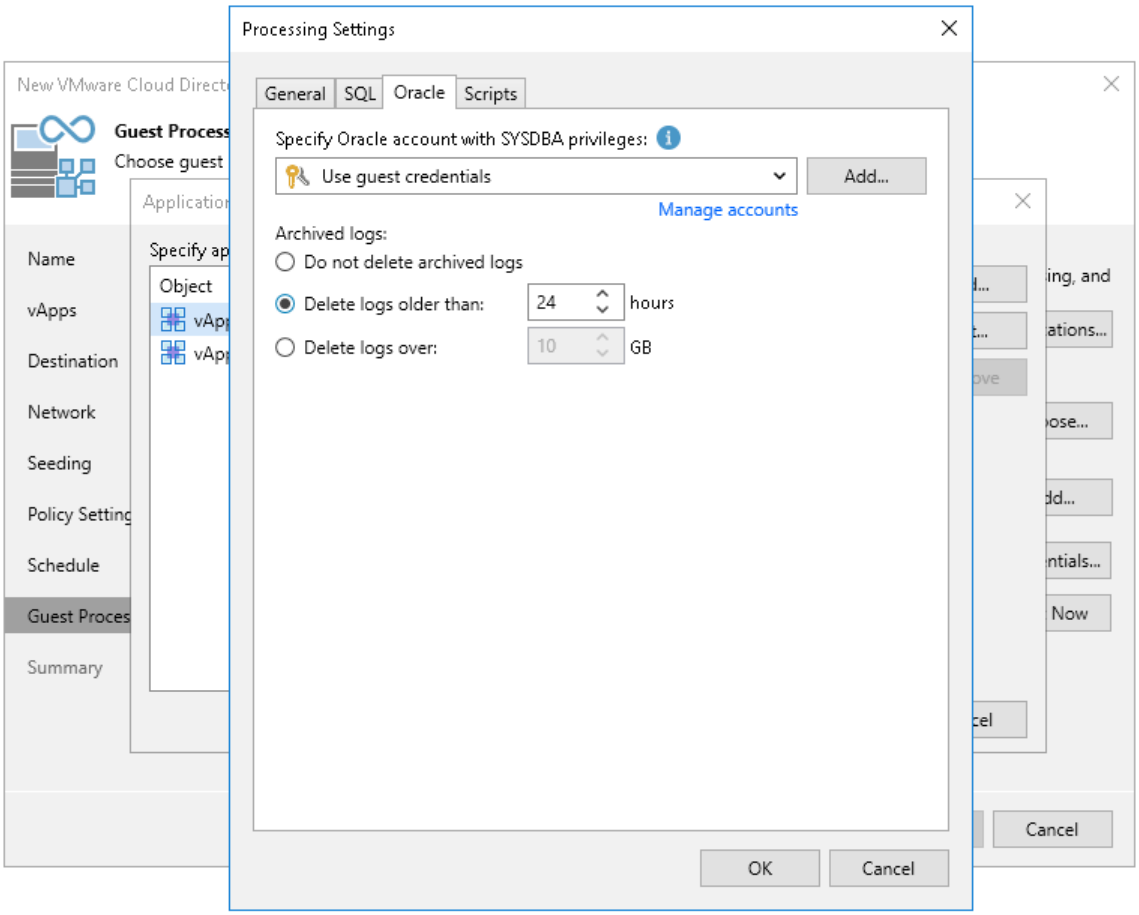


## Script Settings

You can instruct Veeam Backup & Replication to run custom scripts before the CDP policy starts the creation of a long-term restore point and after the policy finishes the creation. For example, these can be pre-freeze and post-thaw scripts for a VM that does not support VSS. The scripts will quiesce the VM file system and application data to bring the VM to a consistent state before the creation of a restore point, and bring the VM and applications to their initial state after the creation finishes.

To specify pre-freeze and post-thaw scripts:

1. At the **Guest Processing** step of the wizard, check that you have selected the **Enable application-aware processing** check box and configured guest OS credentials.
2. At the **Guest Processing** step of the wizard, click **Applications**.
3. In the **Application-Aware Processing Options** list, select workloads for which you want to configure scripts, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.

4. Click the **Scripts** tab.
5. In the **Script processing mode** section, select a scenario for script execution:
   - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the replication process if scripts fail.
   - Select **Ignore script execution failures** if you want to continue the replication process even if script errors occur.
   - Select **Disable script execution** if you do not want to run scripts for the VM.
6. In the **Windows scripts** section, specify paths to scripts for Microsoft Windows VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.
7. In the **Linux scripts** section, specify paths to scripts for Linux VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

   If you plan to replicate a VM container with Microsoft Windows and Linux VMs, you can select to execute both Microsoft Windows and Linux scripts. When replication starts, Veeam Backup & Replication will automatically determine which OS type is installed on the VM and use the correct scripts for this VM.

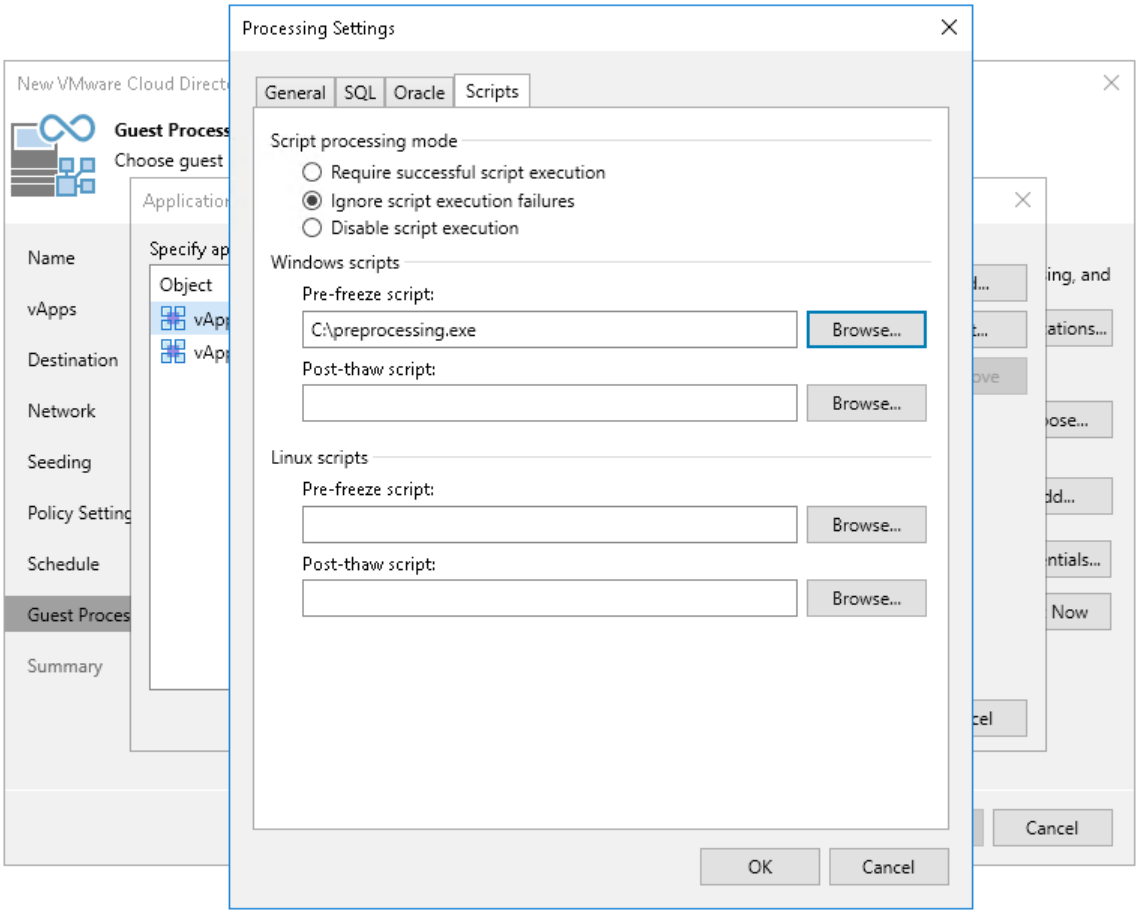

# Step 13. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured settings. If you want to start the policy right after you close the wizard, leave the **Enable the policy when I click Finish** check box selected, otherwise clear the check box. Then click **Finish** to close the wizard.

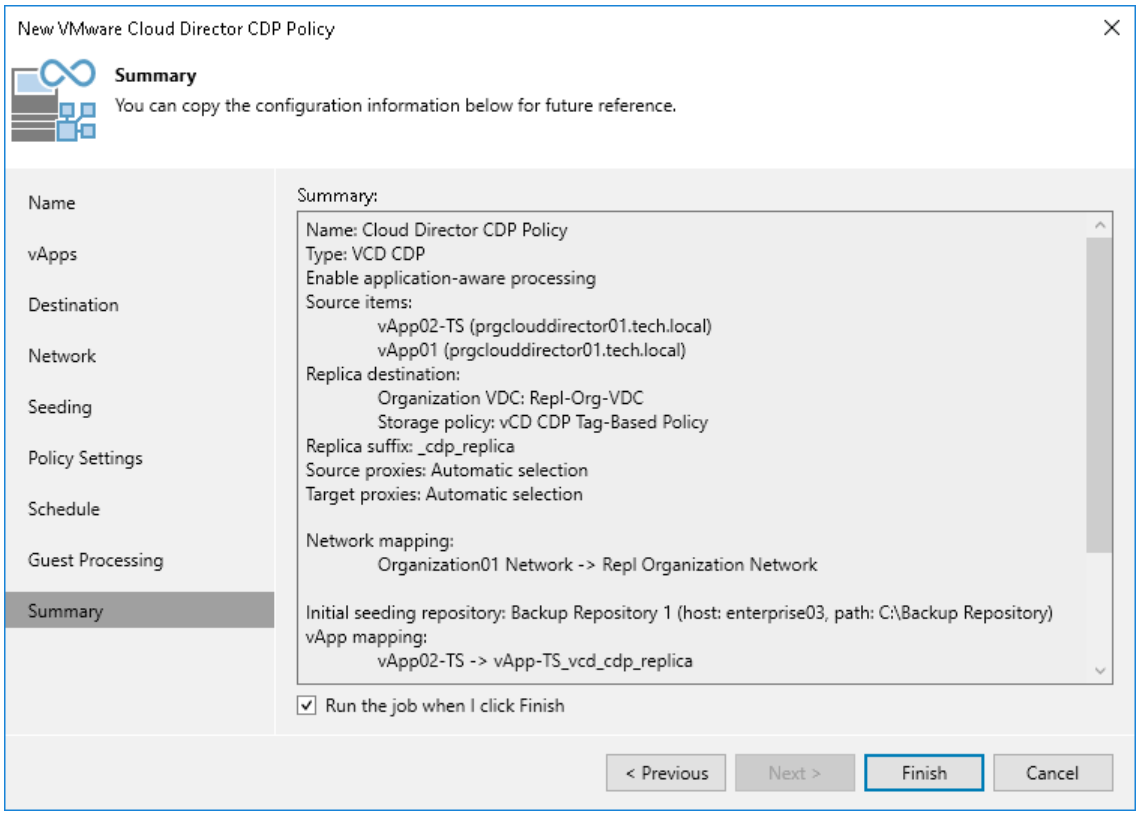

# Managing Cloud Director CDP Policies

After you create VMware Cloud Director CDP policies, you can view statistics, edit, disable and delete policies.

To view all created CDP policies, open the **Home** view and navigate to the **Jobs** > **CDP** node. The working area displays the full list of the created policies. Here, you can manage the policies.

## Viewing Session Statistics and Results

Veeam Backup & Replication allows you to view statistics and session results for a vApp and VMs added to a Cloud Director CDP policy:

- To view real-time statistics for the policy, open the **Home** view. In the inventory pane select **Jobs** > **CDP**. In the working area, double-click the necessary policy. Alternatively, right-click the policy and select **Statistics**. In the opened window, you will be able to switch between statistics of individual vApps and VMs.
- To view real-time statistics for an individual vApp, open the **Home** view. In the inventory pane select **Jobs, Last 24 hours or Running**. In the working area, double-click the necessary vApp. Alternatively, right-click the vApp and select **Statistics**.
- To view statistics on the finished policy sessions, open the **History** view. In the inventory pane select **CDP**. In the working area, double-click the necessary policy session.

The statistics provides detailed data on policy sessions: duration, performance bottlenecks, amount of processed data, read and transferred data.

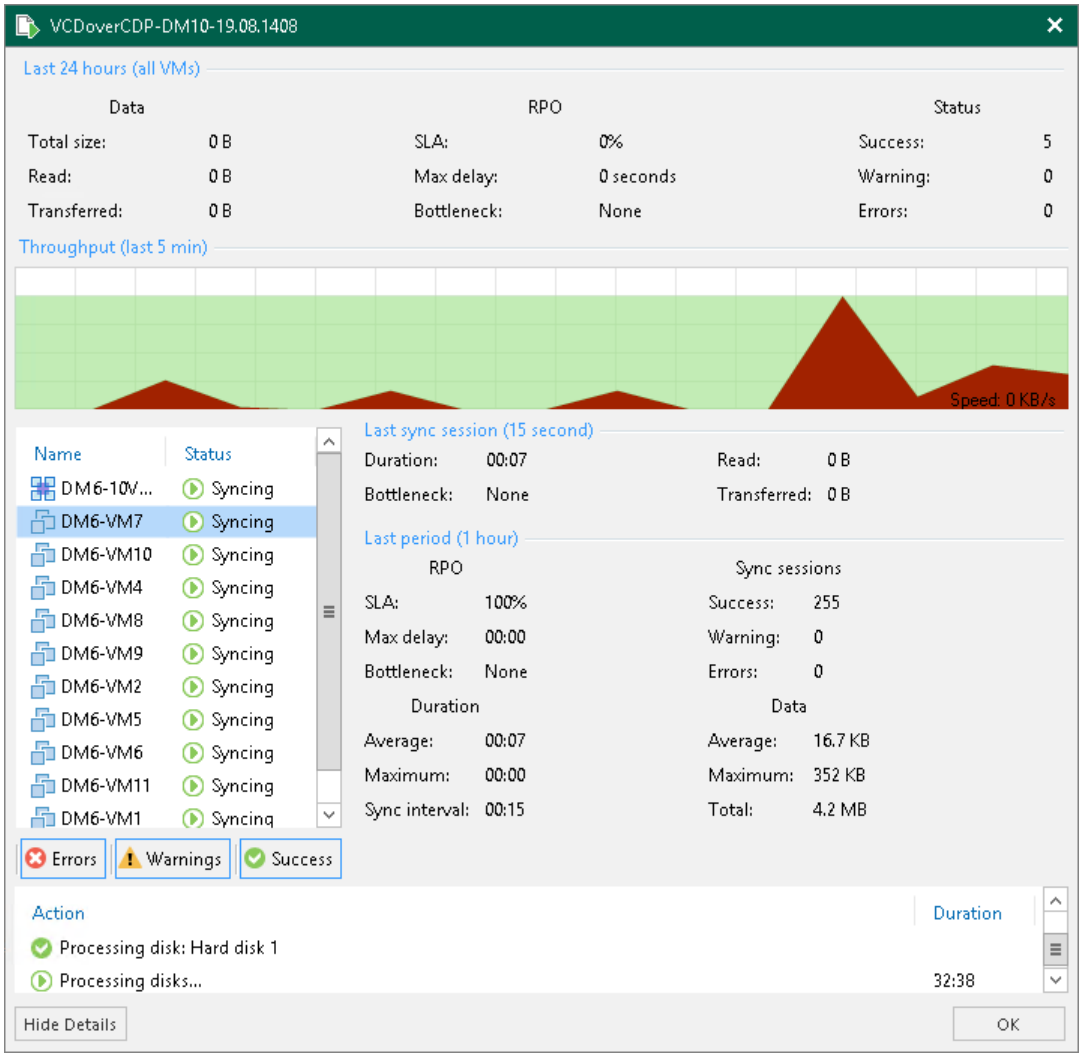


# Statistics Counters

Veeam Backup & Replication displays policy statistics for the following counters:

- The **Last 24 hours (all VMs)** section shows the general information for all vApp and VMs added to vApps for which you have opened the statistics. This information is collected during 24 hours.
  - The **Data** box shows information about processed VM data:
    - **Total size** – total size of the processed data.
    - **Read** – amount of data read from the datastore prior to applying compression and deduplication. The value of this counter is typically lower than the value of the **Total size** counter. Veeam Backup & Replication reads only data blocks that have changed since the last policy session, processes and copies these data blocks to the target.
    - **Transferred** – amount of data transferred from the source VMware CDP proxy to the target VMware CDP proxy. The data is transferred after compression and deduplication.
  - The **RPO** box shows information related to RPO:
    - **SLA** – percentage of sessions completed within the desired RPO.

        - **Max delay** — difference between the configured RPO and time required to transfer and save data.
        - **Bottleneck** — bottleneck in the data transmission process.
    - The **Status box** shows information about task session results. This box informs how many task sessions have completed with the Success, Warning and Error statuses during the 24-hour session. One task session lasts for the period between the creation of two long-term restore points.
- The **Last sync session** section shows general information collected during the last synchronization session of the 24-hour session. One synchronization session lasts for the period between the creation of the two short-term restore points.
    - **Duration** — time period during which data was collected and sent to the target host.
    - **Bottleneck** — bottleneck in the data transmission process.
    - **Read** — amount of data read from the datastore prior to applying compression and deduplication.
    - **Transferred** — amount of data transferred from the source VMware CDP proxy to the target VMware CDP proxy. The data is sent after compression and deduplication.
- The **Last period** section shows general information collected during the last task session of the 24-hour session. One task session lasts for the period between the creation of two long-term restore points.
    - The **RPO** box shows information related to RPO:
        - **SLA** — percentage of sessions completed within the desired RPO.
        - **Max delay** — difference between the configured RPO and time required to transfer and save data.
        - **Bottleneck** — bottleneck in the data transmission process.
    - The **Sync session** box shows information about synchronization session results. This box informs how many synchronization sessions have completed with the *Success*, *Warning* and *Error* statuses during the last task session. One synchronization session lasts for the period between the creation of the two short-term restore points.
    - The **Duration** box shows information about duration of synchronization sessions:
        - **Average** — average duration of a synchronization session.
        - **Maximum** — maximum duration of a synchronization session.
        - **Sync interval** — duration of a synchronization session configured in the policy, that is, the specified RPO.
    - The **Data** box shows information about processed vApp data:
        - **Average** — average amount of data processed within one synchronization session.
        - **Maximum** — maximum amount of data processed within one synchronization session.
        - **Total** — total size of data sent during the task session.
- The pane at the left shows the list of vApps and VMs included into the session statistics. VMs that belong to a vApp are showed right under the vApp.
- The pane at the bottom shows the list of operations performed during the session. If you open statistics for a policy, you can see the list of operations for the whole policy or an individual vApp or VM. To see the list of operations for an individual vApp or VM, click the vApp or VM in the pane on the left. To see the list of operations for the whole policy, click anywhere on the blank area in the left pane.

# Colored Graph

To visualize the data transfer process, Veeam Backup & Replication displays a colored graph in the statistics window:

- The green color defines the amount of data read from the source.
- The brown color defines the amount of data transferred to the target.
- The horizontal line defines the current data processing speed.

If the session is still being performed, you can click the graph to view data rate for the last 5 minutes or the last 24 hours. If the session has already ended, the graph will display information for the last 24 hours only.

The colored graph is displayed only for the currently running session or the latest finished session. If you open statistics for past sessions other than the latest one, the colored graph will not be displayed.

# Editing Policies

To edit a CDP policy:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **CDP** node.
3. In the working area, select the necessary policy and select **Edit** on the ribbon. As an alternative, right-click the necessary policy and select **Edit**.
4. Follow the instructions provided in the Creating CDP Policies section.

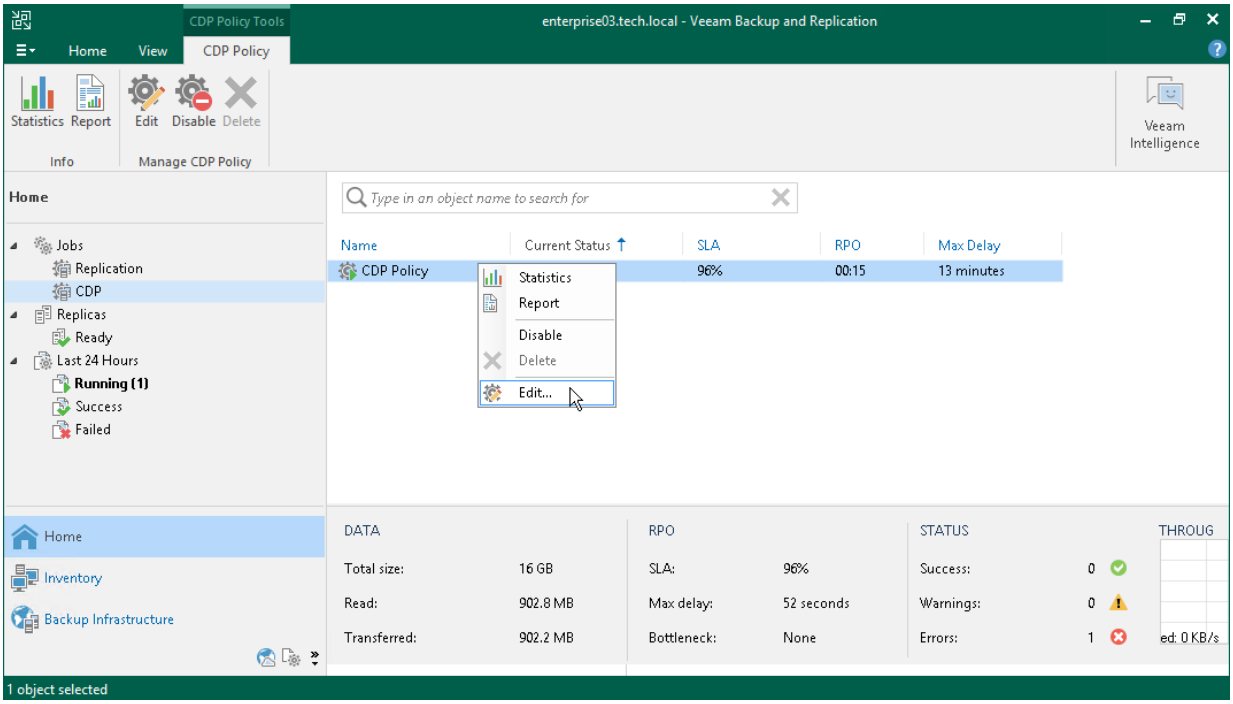


# Cloning Policies

You can create a new CDP policy by cloning an existing one. Policy cloning allows you to create an exact copy of any policy with the same policy settings.

The name of the cloned policy is formed by the following rule: *<policy_name_clone1>*, where *policy_name* is the name of the original policy and *clone1* is a suffix added to the original policy name. If you clone the same policy again, the number in the name will be incremented, for example, *policy_name_clone2*, *policy_name_clone3*, and so on.

# Considerations and Limitations

When cloning a policy, Veeam Backup & Replication can change some policy settings so that cloned policies do not hinder original policies.

- Veeam Backup & Replication disables the cloned policy. To enable the cloned policy, select it in the policy list and click **Disable** on the ribbon. Alternatively, right-click the policy and select **Disable**.
- Before enabling the policy, you must edit the list of protected workloads. One workload can be protected only by one CDP policy.
- In the cloned policy, you cannot change the target destination.

# Cloning Policy

To clone a CDP policy:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the policy and click **Clone** on the ribbon. Alternatively, right-click the policy and select **Clone**.

   After the policy is cloned, you can edit its settings, including the policy name.

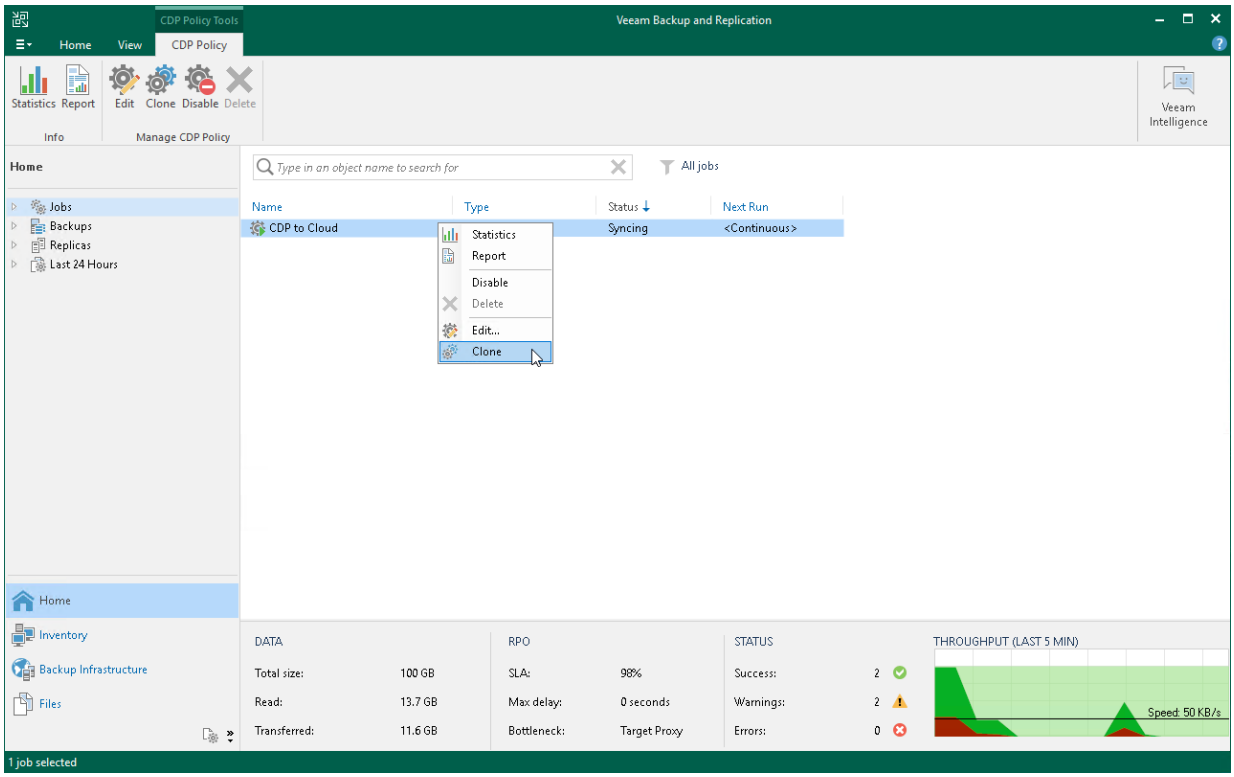

# Disabling and Deleting Policies

Veeam Backup & Replication allows you to temporarily disable or permanently delete created CDP policies.

# Disabling CDP Policies

To disable a CDP policy:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **CDP** node.
3. In the working area, select the necessary policy and select **Disable** on the ribbon. Alternatively, right-click the necessary policy and select **Disable**.

> **TIP**
>
> To enable a disabled policy, select it and click **Disable** once again.

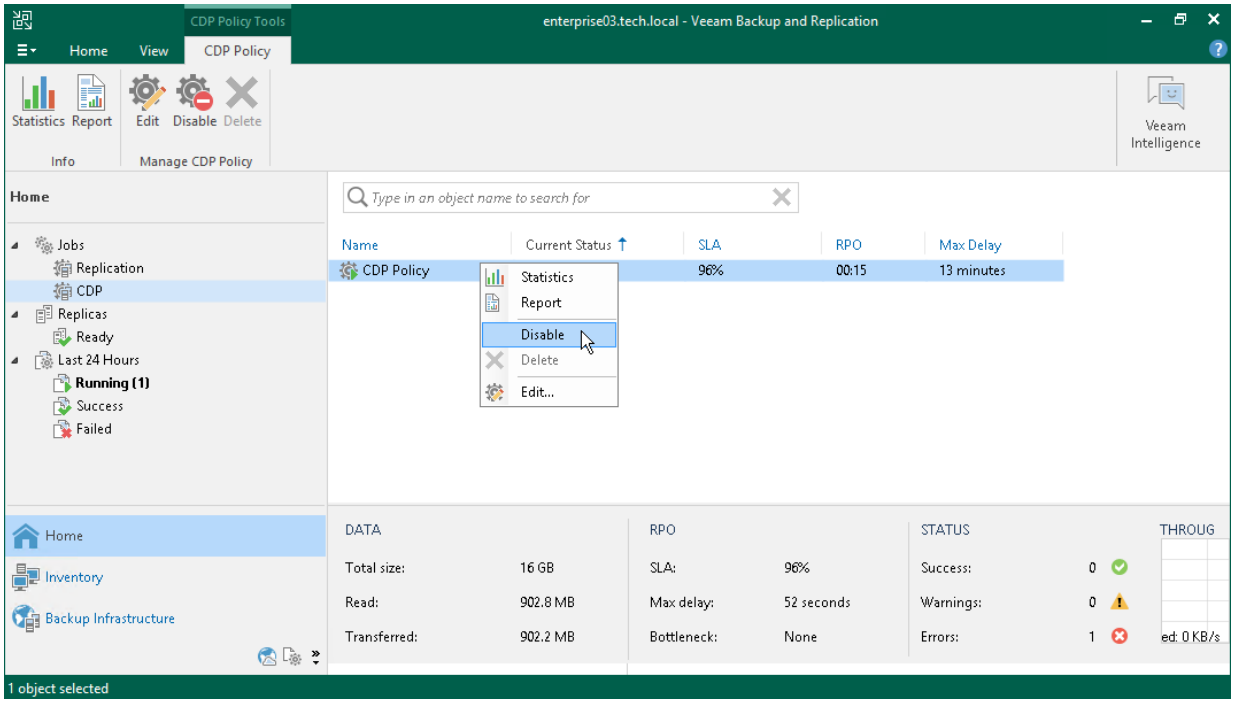


# Deleting CDP Policies

Veeam Backup & Replication allows you to delete only disabled policies. To delete a CDP policy:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **CDP** node.
3. Check that you have disabled the policy that you want to delete.

4. In the working area, select the necessary policy and select **Delete** on the ribbon. Alternatively, right-click the necessary policy and select **Delete**.

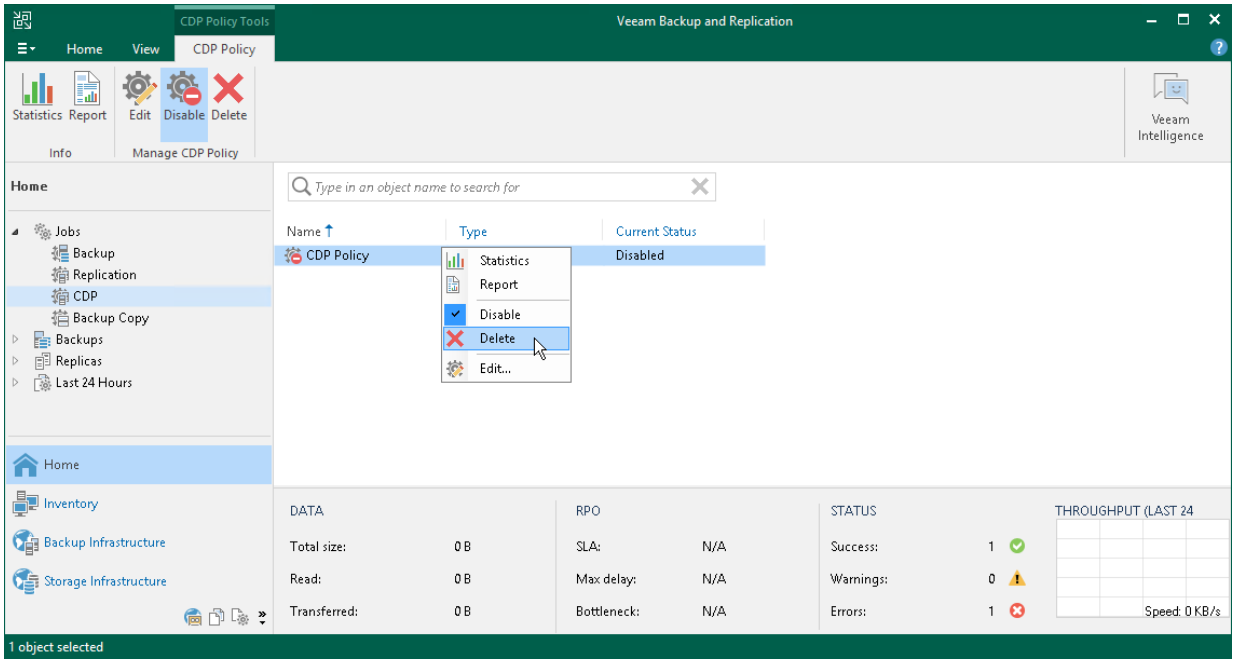

# Managing Cloud Director CDP Replicas

To view all created replicas, open the **Home** view and navigate to the **Replicas** node. The working area displays the full list of the created replicas. Here, you can view replica properties and delete replicas from the configuration database or disk.

## Viewing Replica Properties

You can view replica properties that provide the following information:

- Available restore points
- Date of restore points creation
- Data size and replica status

To view replica properties:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, right-click the necessary replica and select **Properties**. Alternatively, select **Properties** on the ribbon.

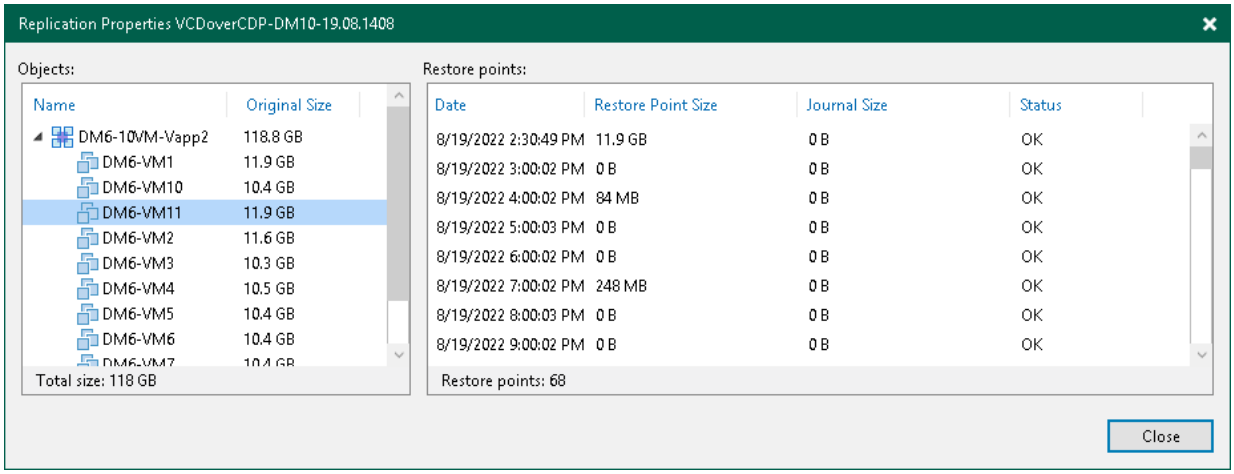


## Rescanning Replicas

Rescan is a process that allows Veeam Backup & Replication to verify that the following data written to configuration database is up to date: information about the disaster recovery (DR) site, its components and VM containers. The replica rescan process is performed the following way:

1. Veeam Backup & Replication gathers information on replicas that are currently available in the DR site.
2. Veeam Backup & Replication compares this information with information stored in the configuration database about replicas from this DR site.
3. If information about replicas from the DR site differs from information stored in the configuration database about these replicas, Veeam Backup & Replication updates the configuration database.

To rescan replicas, do the following:

1. Open the **Home** view.

2. In the inventory pane, right-click the **Replicas** node and select **Rescan replicas**.

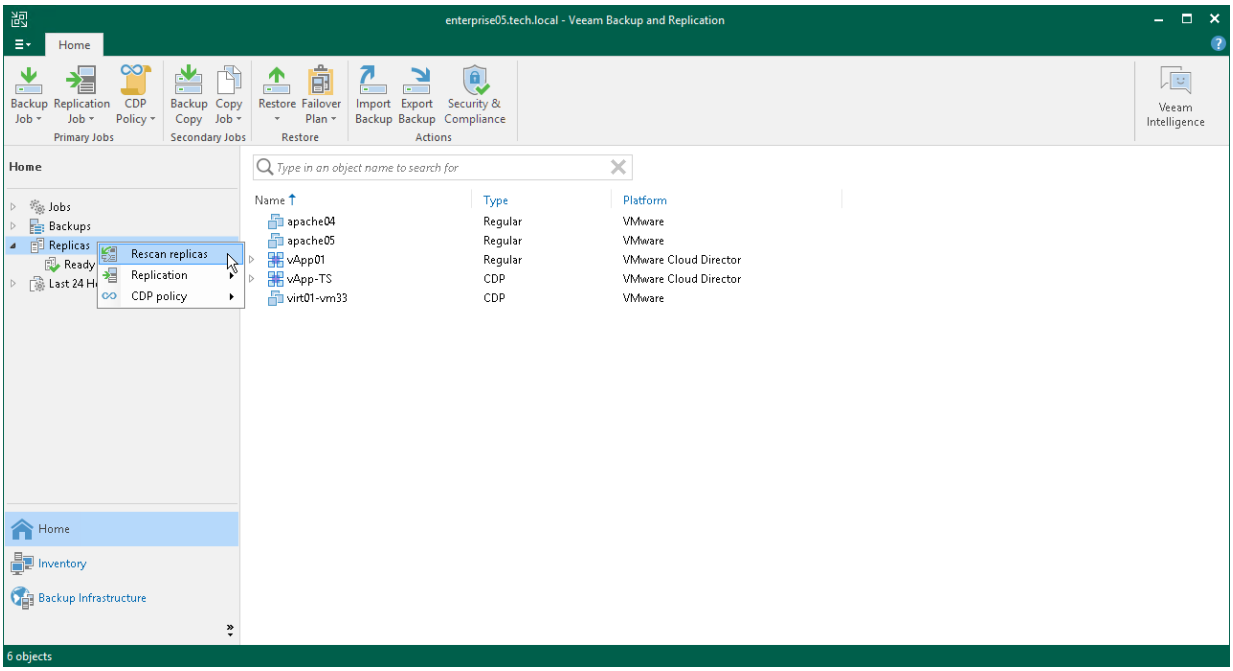


# Removing from Configuration

When you remove replicas from the configuration, Veeam Backup & Replication removes records about the replicas from the configuration database, stops showing the replicas in Veeam Backup & Replication console and also stops synchronizing their state with the state of the source VMs. However, the actual replicas remain on hosts.

To remove records about replicas from the Veeam Backup & Replication console and configuration database:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.
3. In the working area, select replicas in the *Ready* state and click **Remove from** > **Configuration** on the ribbon. Alternatively, right-click one of the selected replicas and select **Remove from configuration**.

**NOTE**

Consider the following:

- The **Remove from configuration** operation can be performed only for replicas in the *Ready* state. If the replica is in the *Failover* or *Failback* state, this option is disabled.
- When you perform the **Remove from configuration** operation for a vApp that is replicated as a standalone object, Veeam Backup & Replication removes this vApp from the initial Cloud Director CDP policy. When you perform the **Remove from configuration** operation for a vApp that is replicated as part of a container (organization VDC or Cloud Director server), Veeam Backup & Replication adds this vApp to the list of exclusions in the initial policy. For more information, see Exclude Objects.
- When you remove from the configuration a replica, Veeam Backup & Replication consolidates the replica into a regular VM. Veeam Backup & Replication removes technical files, such as transaction log and delta disk files, and commits data from them into the VM.

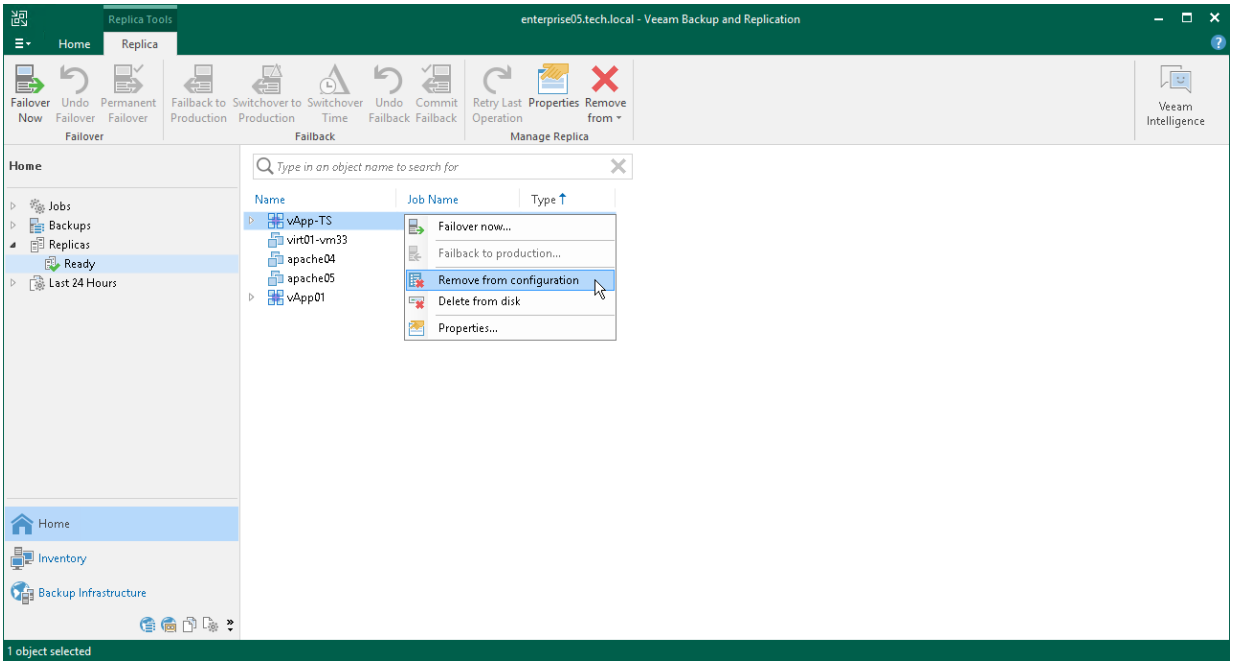


# Deleting from Disk

When you delete replicas from disks, Veeam Backup & Replication removes the replicas not only from the Veeam Backup & Replication console and configuration database, but also from host storage.

> **NOTE**
>
> Consider the following:
>
> - You can delete records only about replicas that are in the *Ready* state.
> - Do not delete replica files from the destination storage manually, use the **Delete from disk** option instead. If you delete replica files manually, subsequent replication sessions will fail.
> - Unlike the **Remove from configuration** operation, the **Delete from disk** operation does not remove the processed workload from the initial replication job. This means that the replication process will restart for this workload. To avoid this, you can exclude the workload from the replication job or disable the job.

To delete replica files from disks:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.

3. In the working area, select the necessary replica and click **Remove from** > **Disk** on the ribbon. As an alternative, right-click the replica and select **Delete from disk**.

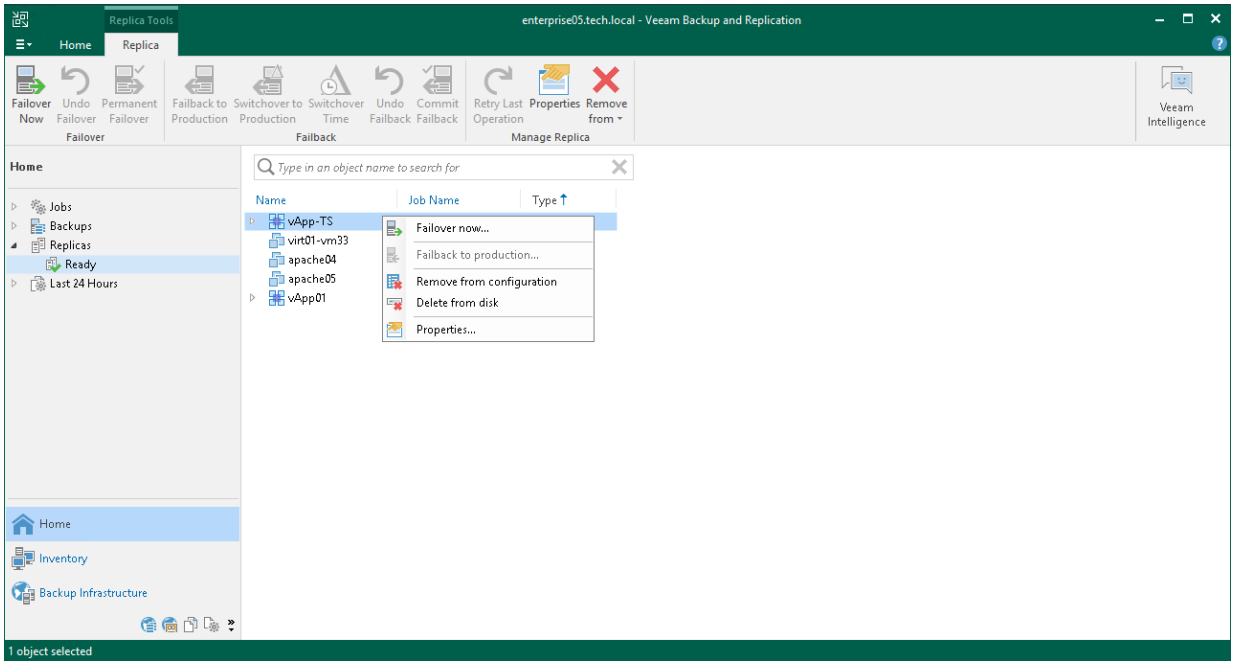

# Failover and Failback for Cloud Director CDP

Failover and failback are operations that allow you to manage your production and disaster recovery (DR) sites if a disaster strikes. Failover is a process of switching from the vApp on the source organization VDC to its replica on a target organization VDC that is set up as the disaster recovery (DR) site. Failback is a process of returning from the replica to the source vApp or a new vApp.

Veeam Backup & Replication provides the following failover and failback operations:

- **Perform failover**

  When you perform failover, you shift all processes from the source vApp in the production organization VDC to the replica in the DR organization VDC. Failover is an intermediate step that needs to be finalized: you can perform permanent failover, perform failback or undo failover.

- **Perform permanent failover**

  When you perform permanent failover, you permanently switch from the source vApp to a replica and use this replica as the production vApp. The source vApp is excluded from VMware Cloud Director replica processing.

- **Undo failover**

  When you undo failover, you switch back to the source vApp and discard all changes made to the replica while it was running. For example, you can use the undo failover scenario if you have failed over to the replica for testing and troubleshooting purposes, and you do not need to synchronize the source vApp state with the current state of the replica.

- **Perform failback**

  When you perform failback, you switch back to the source vApp and send to the source vApp all changes that took place while the replica was running. If the source organization VDC is not available, you can recover a vApp with the same configuration as the source vApp and switch to it.

  When you perform failback, changes are only sent to the source or recovered vApp but not published. You must test whether the source or recovered vApp works with these changes. Depending on the test results, you can do the following:

  - **Commit failback**. When you commit failback, you confirm that the source or recovered vApp works as expected and you want to get back to it.
  - **Undo failback**. When you undo failback, you confirm that the source or recovered vApp is not working as expected and you want to get back to the replica.

The following scheme can help you decide which steps are preferable when you fail over to a replica.

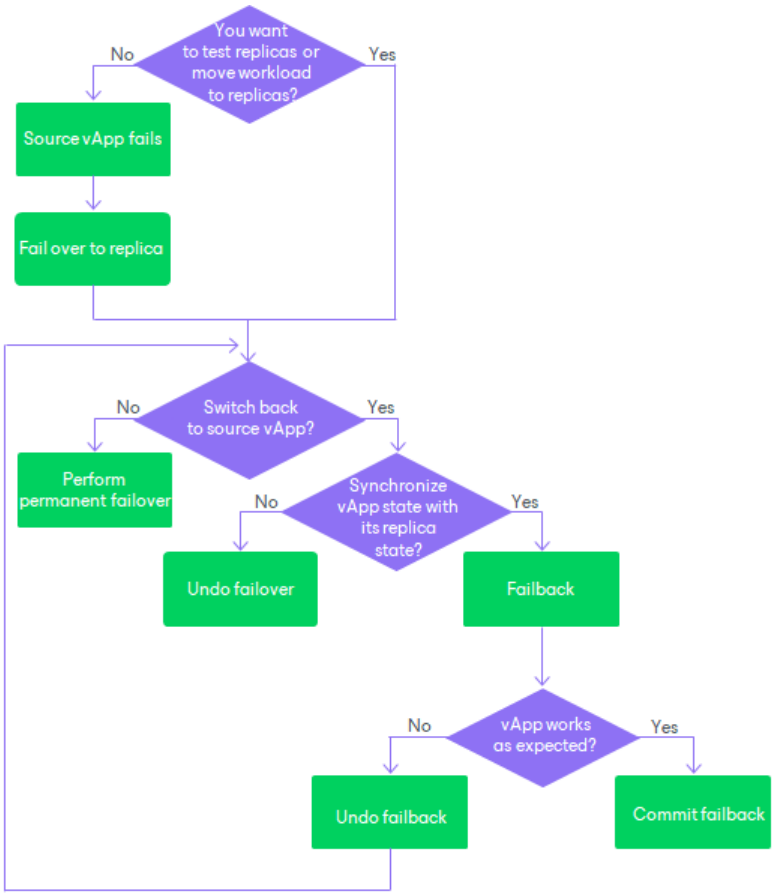


# Failover

Failover is a process when Veeam Backup & Replication switches processes from the source vApp in the production organization VDC to its replica in the disaster recovery organization VDC. During failover, Veeam Backup & Replication recovers the replica to the required restore point and shifts all I/O processes from the source vApp to its replica. As a result, you have a fully functional vApp within several minutes, and your users can access services and applications with minimum disruption.

You can fail over to replicas not only when a disaster strikes the production organization VDC, but also to test replicas for recoverability. If the source vApps and replicas are located in the same network, consider temporary disconnecting the source vApps from the network to avoid IP address or machine name conflicts.

# How Failover Works

The failover operation is performed in the following way:

1. Veeam Backup & Replication puts all replication activities on hold.
2. The state of the replica is changed from *Ready* to *Processing*.
3. Veeam Backup & Replication recovers a replica to the required restore point.
4. Veeam Backup & Replication powers on the replica.

   The source vApp still exists and failover does not change the vApp state: if the vApp is powered on when you perform failover, it remains powered on when failover completes; if the vApp is powered off, it remains in this state.

5. All changes made to the replica while it is running in the *Failover* state are written to the delta file and stored in the target host.
6. After failover completes successfully, the vApp state is changed from *Processing* to *Failover*.

# Finalizing Failover

Failover is an intermediate step that needs to be finalized. You can perform the following operations:

- Undo failover
- Perform permanent failover
- Perform failback

## Performing Failover

For more information on failover, see Failover and Failback for Cloud Director CDP and Failover.

To perform failover, do the following.

## Step 1. Launch Failover Wizard

To launch the **Failover** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vCloud Director** > **Restore from replica** > **Entire vApp** > **Failover vApp to a replica**.
- Open the **Home** view, in the inventory pane select **Replicas**. In the working area, select the necessary replica, right-click one of the selected replica and click **Failover Now**. Alternatively, click **Failover Now** on the ribbon.

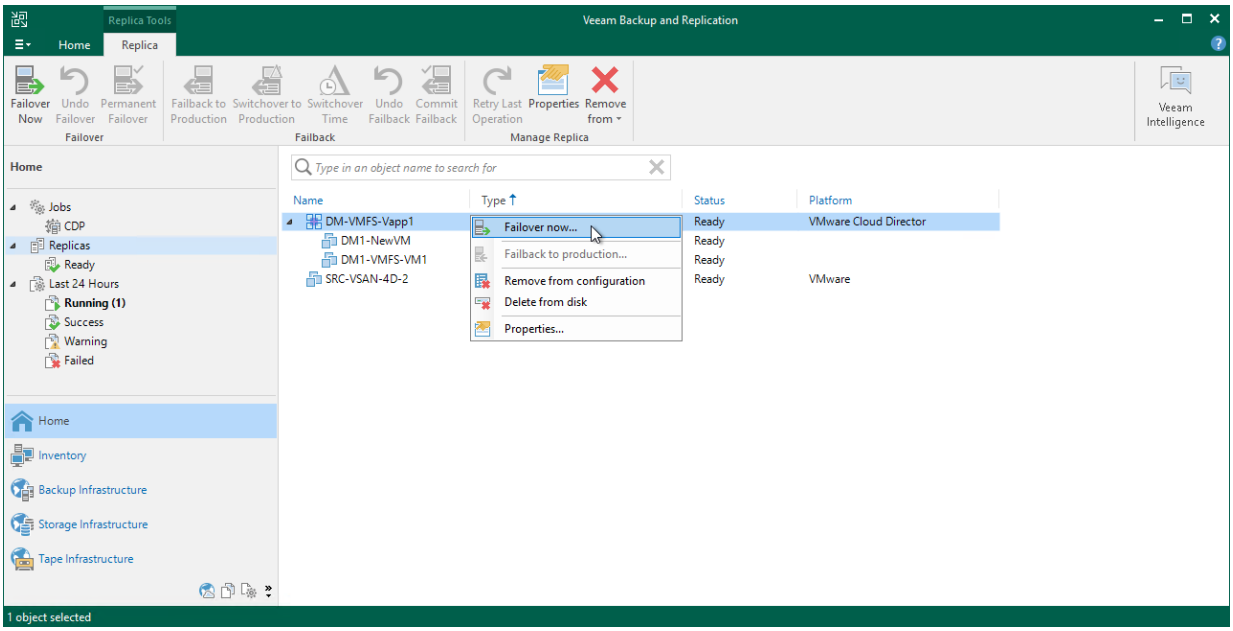

## Step 2. Select vApps

At the **vApp** step of the wizard, you can modify a list of vApps from which you fail over. To add vApps, click **Add** > **From infrastructure** if you want to add vApps from the virtual infrastructure, or **Add** > **From replicas** if you want to add vApps from existing replicas. Then select the necessary vApps. If you select organizations or organization VDCs, Veeam Backup & Replication will expand them to a vApp list.

> **NOTE**
>
> Make sure that vApps you select from the virtual environment have been successfully replicated at least once.

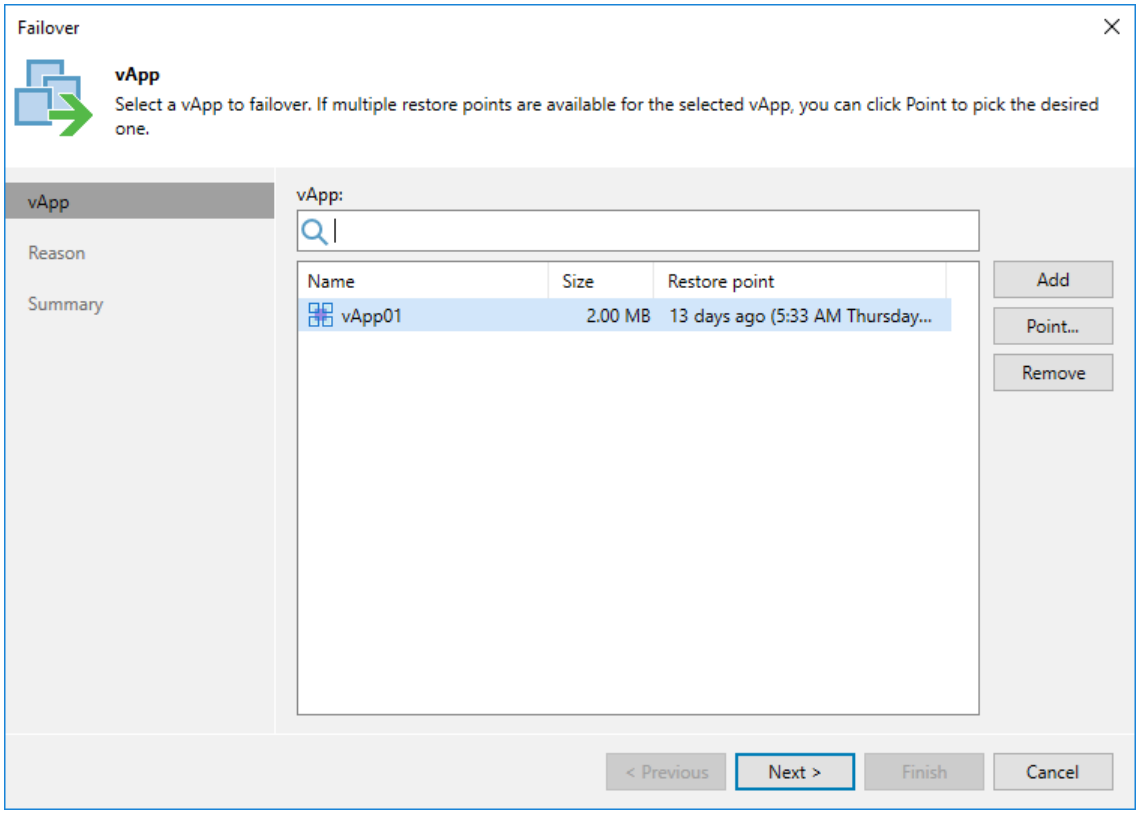

## Step 3. Select Restore Points

At the **vApp** step of the wizard select to which state of replicas you want to fail over:

1. In the **vApp** list, select the necessary vApp and click **Point**.
2. In the **Restore Points** window, select to which restore point you want to fail over:
   - Latest crash-consistent restore point. In this case, you will lose the least amount of data.
   - Latest long-term application-consistent restore point. In this case, you will be able to restore specific application data.
   - Specific point in time. In this case, you will restore to any selected restore point.

     To restore to a short-term restore point, select a point in the green area. To restore to a crash-consistent or application consistent long-term point, select a violet or turquoise vertical bar with a circle at the top. A yellow restore point (a short-term restore point in the yellow area or a long-term restore point as a vertical bar) indicates that the point is in the mixed state. This means that some VMs were powered off, not processed successfully, have different states — crash-consistent and application-consistent, and so on. In this case, Veeam Backup & Replication will suggest to restore the VMs on the nearest short-term restore point.

     Use the right and left arrows on the keyboard to select the required restore point. To quickly find a long-term restore point, click a link that shows a date. In the opened window, you will see a calendar where you can select the necessary day. In the **Timestamp** section, you will see long-term restore points created during the selected day.

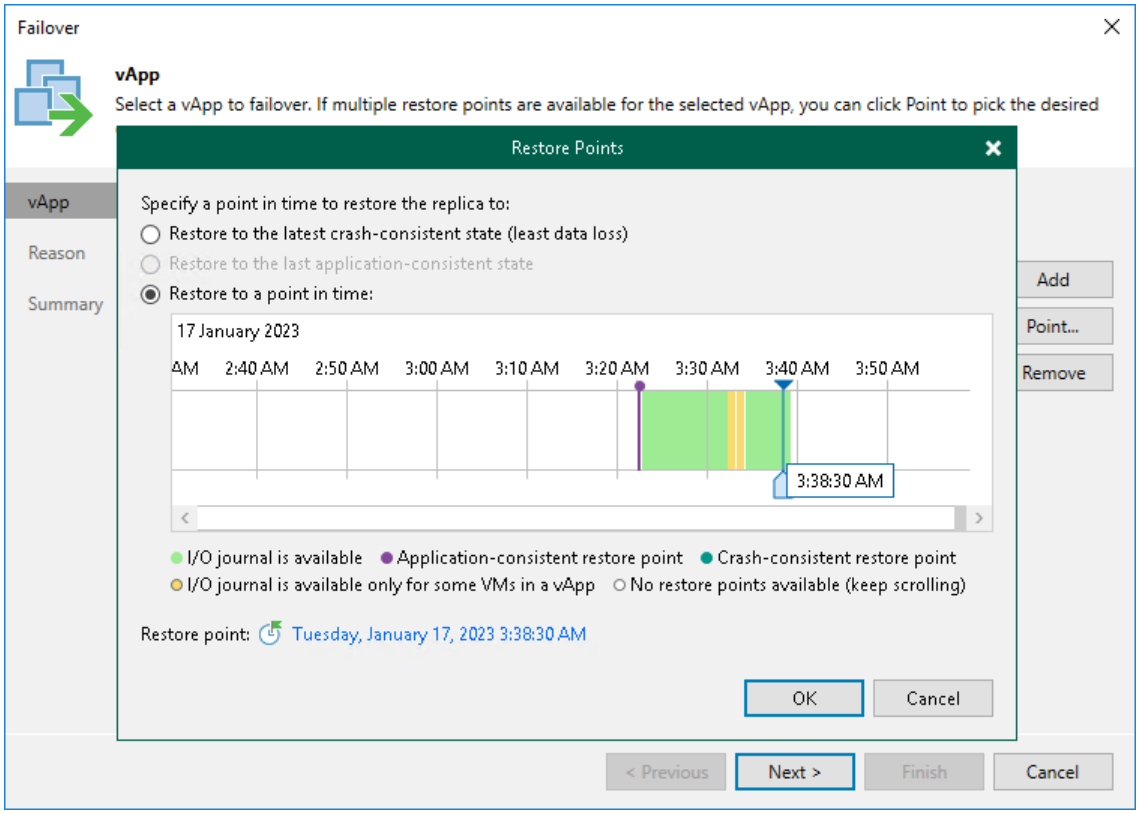

## Step 4. Specify Failover Reason

At the **Reason** step of the wizard, enter a reason for failing over to the replicas. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

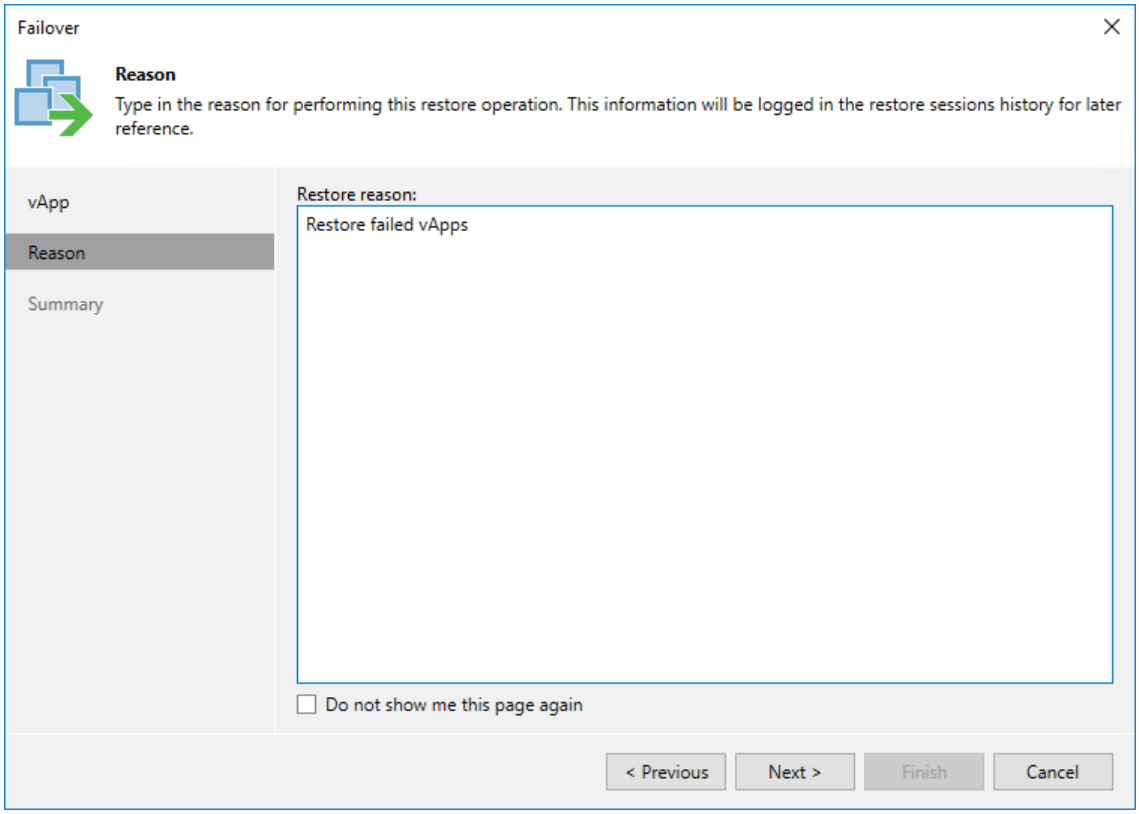

### Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the failover task. Then click **Finish** to start the failover process.

When the failover process is complete, the replicas will be started on the target hosts.

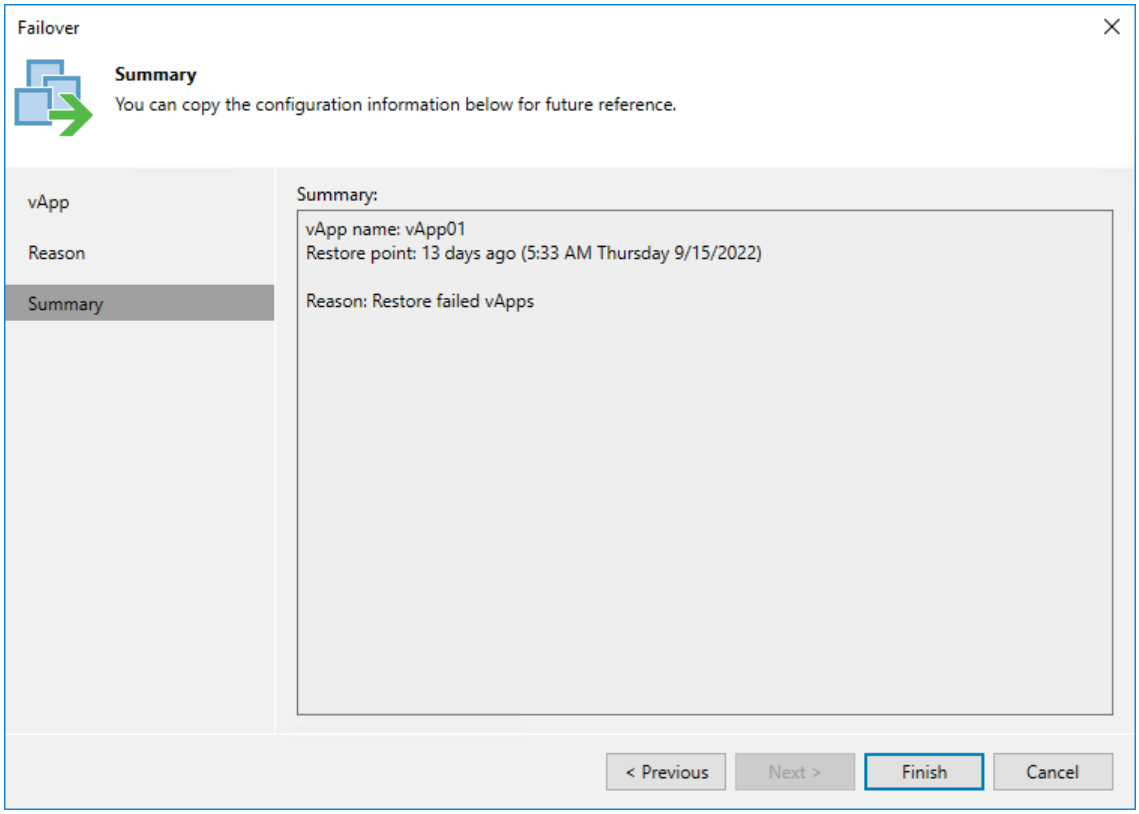


# What You Do Next

Failover is an intermediate step that needs to be finalized. You can finalize failover in the following ways:

- Perform permanent failover
- Undo failover
- Perform failback

## Performing Failover Retry

The failover retry option is necessary when failover of vApps fails with the *Incomplete* state. When you perform a retry, Veeam Backup & Replication restarts failover only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, failover takes less time and does not consume as many resources as when processing a whole vApp.

To retry failover:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary vApp and select **Retry Failover** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry failover**.

# Permanent Failover

Permanent failover is one of the ways to finalize failover. When you perform permanent failover, you permanently switch processes from the source vApp to its replica. As a result, the replica stops acting as a replica and starts acting as the production vApp.

The permanent failover operation is performed in the following way:

1. Veeam Backup & Replication removes restore points of the replica from the replication chain and deletes associated files from the datastore. Changes that were written to the delta disks are committed to the replica to bring the replica to the most recent state.
2. Veeam Backup & Replication removes the replica from the list of replicas in the Veeam Backup & Replication console and the replica becomes the production vApp.
3. To protect the replica from corruption after permanent failover is complete, Veeam Backup & Replication reconfigures the job or policy and adds the source vApp to the list of exclusions. When replication job starts, the source vApp is skipped from processing. As a result, no data is written to the working replica.

## Performing Permanent Failover

For more information on permanent failover, see Failover and Failback for Cloud Director CDP and Permanent Failover.

To perform permanent failover, do one of the following:

- Open the **Home** view, in the inventory pane select **Replicas** > **Active**. In the working area, select the necessary vApp and click **Permanent Failover** on the ribbon.
- Open the **Home** view, in the inventory pane select **Replicas** > **Active**. In the working area, right-click the necessary vApp and select **Permanent Failover**.

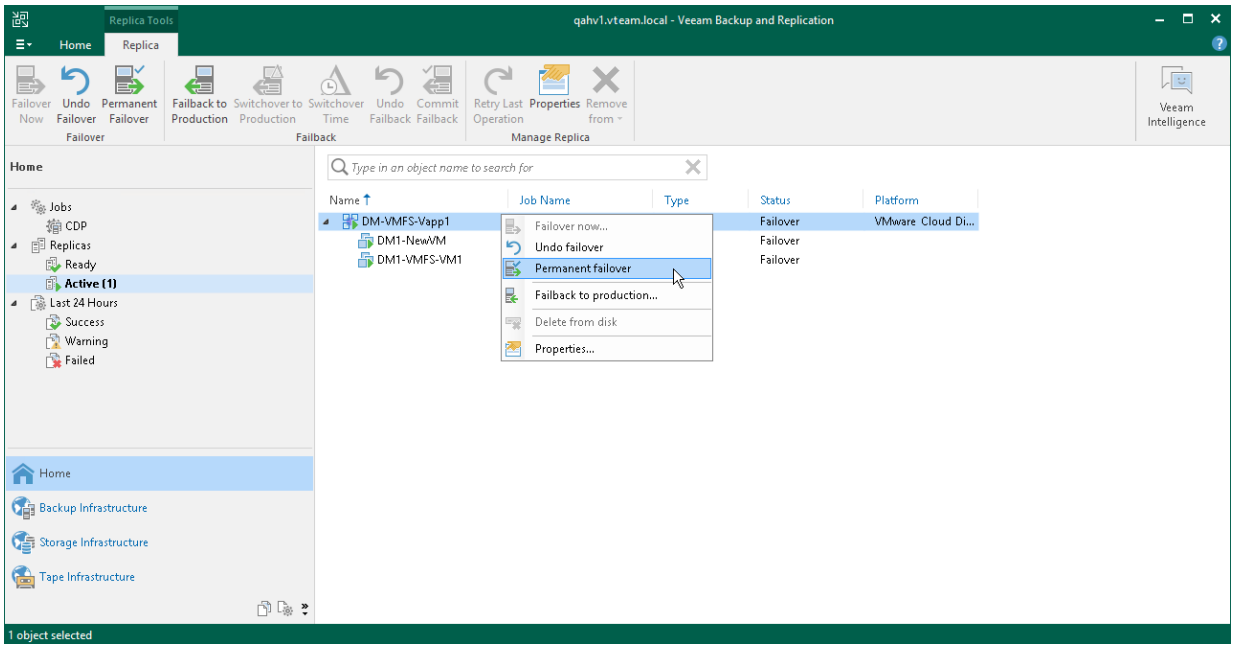

## Performing Permanent Failover Retry

If permanent failover failed, you can retry this operation. When you perform a retry, Veeam Backup & Replication restarts permanent failover only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, permanent failover takes less time and does not consume as many resources as when processing a whole vApp.

To perform a retry:

1. Open the **Home** view, in the inventory pane, navigate to the **Replicas** > **Active** node.
2. In the working area, select the necessary vApp and select **Retry Permanent Failover** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry permanent failover**.

# Failover Undo

Failover undo is one of the ways to finalize failover. When you undo failover, you switch back from a vApp replica to the original vApp. Veeam Backup & Replication discards all changes made to the vApp replica while it was in the *Failover* state.

The failover undo operation is performed in the following way:

1. Veeam Backup & Replication reverts the replica to its pre-failover state. To do this, Veeam Backup & Replication powers off the vApp replica and gets it back to the latest restore point in the replication chain.
2. The state of the replica gets back to *Ready*, and Veeam Backup & Replication resumes replication activities for the original vApp on the source host.

## Undoing Failover

For more information onfailover undo, see Failover and Failback for Cloud Director CDP and Failover Undo.

To undo failover:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, select the necessary replica and click **Undo Failover** on the ribbon. Alternatively, right-click the necessary replica and select **Undo Failover**.

4. In the displayed window, click **Yes** to confirm the operation.

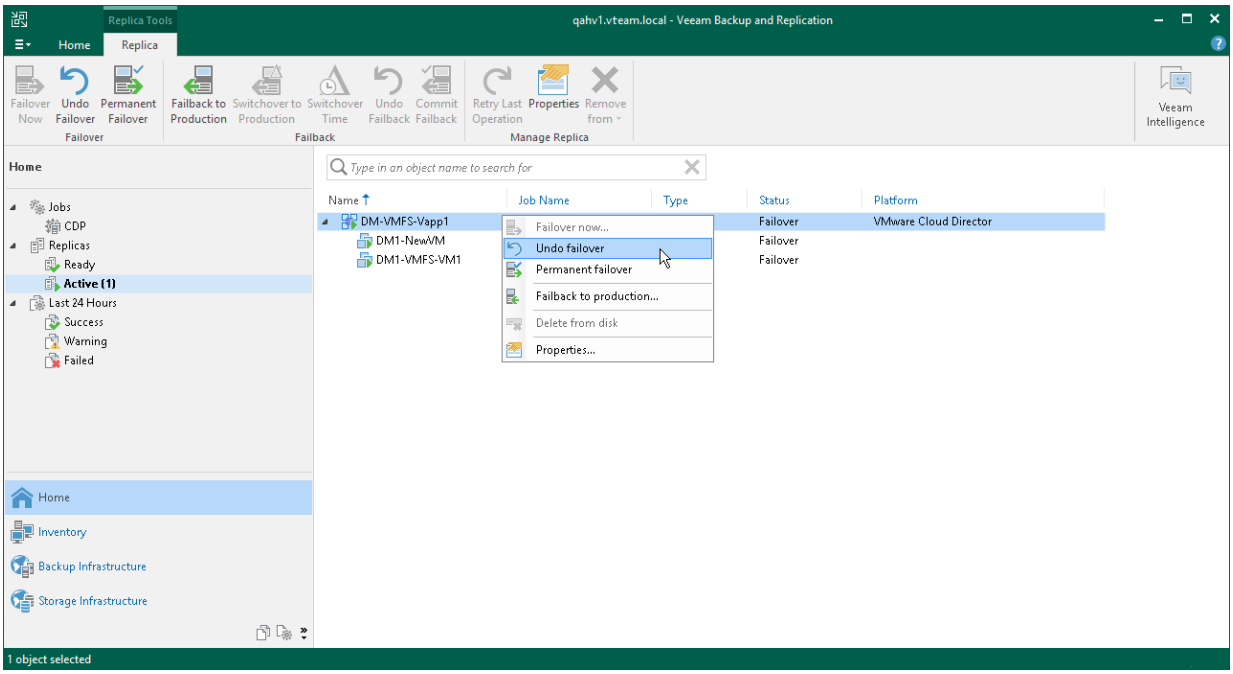


## Performing Failover Undo Retry

If the failover undo operation failed, you can retry this operation. When you perform a retry, Veeam Backup & Replication restarts the failover undo operation only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, the failover undo operation takes less time and does not consume as many resources as when processing the whole vApp.

To perform a retry:

1. Open the **Home** view, in the inventory pane, navigate to the **Replicas** > **Active** node.
2. In the working area, select the necessary vApp and select **Retry Undo Failover** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry undo failover**.

# Failback

Failback is an option that allows you to switch back from the replicated vApp on a disaster recovery organization VDC to the production vApp. When you perform failback, you switch back to the production VM from a VM replica, shift I/O processes from the disaster recovery site to the production site.

You can perform failback in the following ways:

- Fail back to the source vApp in the original location.
- Fail back to a vApp already recovered to a new location. This vApp must be recovered before you perform failback. For example, you can recover the VM from a backup.
- Fail back to a vApp recovered from a replica to a new location, or to any location but with different settings. The vApp will be recovered from the replica during the failback process.

The first two options help you decrease recovery time and the use of the network traffic because Veeam Backup & Replication needs to transfer only differences between the source or recovered vApp and replica. For the third option, Veeam Backup & Replication needs to transfer the whole vApp data, including its configuration and virtual disk content. Use the third option if there is no way to use the source vApp or restore it from a replica.

Veeam Backup & Replication performs failback in two phases:

- **First phase**: Veeam Backup & Replication synchronizes the state of the production vApp (the source vApp, an already recovered vApp or a vApp recovered from the replica) with the current state of the replica. This phase may take a lot of time especially if vApp is large. While Veeam Backup & Replication performs the first phase of failback, VMs from replicas are still up and running, users can access these VMs and perform daily routine tasks as normal. All changes made to vApps during the first phase of a failback a written to a delta file.
- **Second phase**: Veeam Backup & Replication transfers all changes made to the replica during the first phase of failback to the production vApp, switches all processes from the replica to the production vApp and turns off the replica.

  The time when the second phase starts depends on how you want to switch from the replica to the production vApp. You can switch to the production vApp automatically, at the scheduled time or manually. If you select to switch automatically, the second phase will start right after the first phase finishes. If you select to switch at the scheduled time or manually, the second phase will start at the time you want.

The process of failing back to the source vApp or a source vApp restored in a different location differs from the process of failing back to a specific location:

- How failback to the source vApp or a source vApp restored in a different location works
- How failback to a specific location works

# How Failback to Source vApp or Source vApp Restored in Different Location Works

When you fail back to the source vApp or an already recovered vApp, Veeam Backup & Replication performs the following operations during the first phase:

1. Veeam Backup & Replication calculates the difference between disks of the production vApp and disks of the replica in the *Failover* state. The calculation of the difference helps Veeam Backup & Replication understand what data needs to be transferred to the production vApp and to synchronize its state with the state of the replica.

   [For ESXi hosts prior to version 7.0] If you fail back to the source vApp in the source location and you have enabled the **Quick rollback** option, this calculation can be performed much faster. For more information on the **Quick rollback** option, see Quick Rollback.

2. Veeam Backup & Replication transfers the data that was detected as different to the production vApp. The transferred data is written to the production vApp.

3. Veeam Backup & Replication changes the state of the replica from *Failover* to *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the following operations:

1. The guest OS of the replica is shut down or the replica is powered off.

   If VMware Tools are installed on the VM added to the replica, Veeam Backup & Replication tries to shut down the replica guest OS. If nothing happens in 15 minutes, Veeam Backup & Replication powers off the vApp replica. If VMware Tools are not installed on the VM added to the replica or the vApp is suspended, Veeam Backup & Replication powers off the vApp. The replica remains powered off until you commit failback or undo failback.

2. Veeam Backup & Replication calculates the difference between disks of the production vApp and disks of the replica. The calculation of the difference helps Veeam Backup & Replication understand what data was changed while the replica was in the *Ready to switch* state.

3. Sends data changed on the replica while it was in the *Ready to switch* state to the production vApp.

4. The state of the replica is changed from *Ready to switch* to *Failback*.

5. [If you fail back to a recovered vApp] Veeam Backup & Replication updates the ID of the source vApp in the Veeam Backup & Replication configuration database. The ID of the source vApp is replaced with the ID of the recovered vApp.

6. If you have selected to power on the production vApp after failback, Veeam Backup & Replication powers on the production vApp on the host.

# How Failback to Specific Location Works

When you fail back to a vApp recovered from a replica, Veeam Backup & Replication performs the following operations during the first phase:

1. Veeam Backup & Replication requests VMware Cloud Director to create on the target organization VDC an empty vApp with the same configuration as the replica. VMware Cloud Director server registers the created production vApp.

2. Veeam Backup & Replication transfers data of the replica to the production vApp to update the production vApp state to the replica state.

3. Veeam Backup & Replication changes the state of the replica from *Failover* to the *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the same operations as described in section How Failback to Source vApp or Already Recovered vApp Works except for the step 2.

Failback is an intermediate step that needs to be finalized. If the production vApp works as expected and you want to get back to it, commit failback. If the vApp does not work as expected, undo failback.

## Quick Rollback

Quick rollback helps you significantly reduce the failback time. You can use quick rollback if you fail back from a replica to the source vApp in the original location.

During failback, Veeam Backup & Replication calculates differences between VM disks of the source vApp and disks of the replica. With the quick rollback option enabled, Veeam Backup & Replication compares only those disk sectors that have changed during the replica was in the *Failover* state instead of comparing entire disks. To get information about the changed disk sectors, Veeam Backup & Replication uses VMware vSphere Changed Block Tracking (CBT).

As a result of enabling quick rollback, difference calculation becomes much faster. After the differences are calculated, Veeam Backup & Replication performs failback in a regular way: transport changed blocks to the source vApp, powers off the replica and synchronizes the source vApp with the replica once again.

# Requirements for Quick Rollback

To perform quick rollback, make sure that the following requirements are met:

- You fail back to the source vApp in the original location.
- Do not use quick rollback if the problem occurred at the vApp hardware level, storage level or due to a power loss.

  Use quick rollback if you fail back to the source vApp that had a problem at the guest OS level — for example, there was an application error or a user accidentally deleted a file on the source VM guest OS.
- CBT must be enabled for the source vApp.

# Limitations for Quick Rollback

During the first replication job session after failback with quick rollback, CBT on the source vApp is reset. Due to that Veeam Backup & Replication will read data of the entire vApp.

## Performing Failback

For more information on failback, see Failover and Failback for Cloud Director CDP and Failback.

To perform failback, do the following:

## Step 1. Launch Failback Wizard

To launch the **Failback** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware Cloud Director** > **Restore from replica** > **Entire vApp** > **Failback to production**.
- Open the **Home** view, in the inventory pane select **Replicas** > **Active**. In the working area, right-click the necessary replica and select **Failback to production**. Alternatively, click **Failback to Production** on the ribbon.

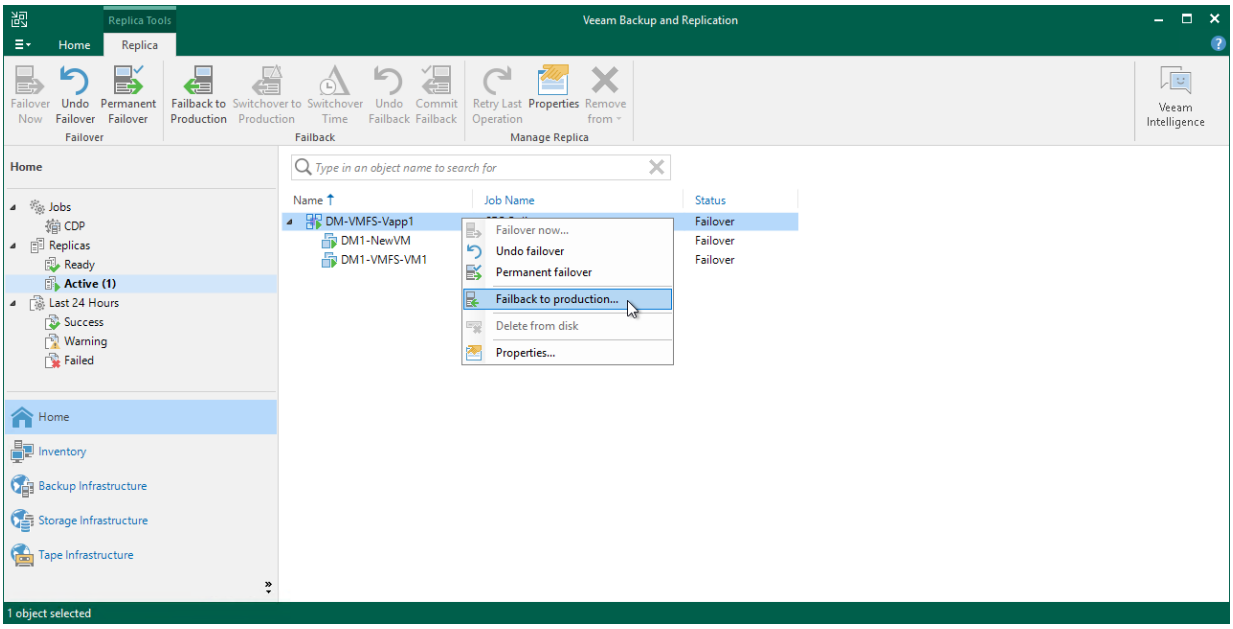

## Step 2. Select Replicas

At the **Replica** step of the wizard, select replicas from which you want to fail back.

To update the list of replicas that are ready for failback (replicas in the *Failover* state), click **Populate**.

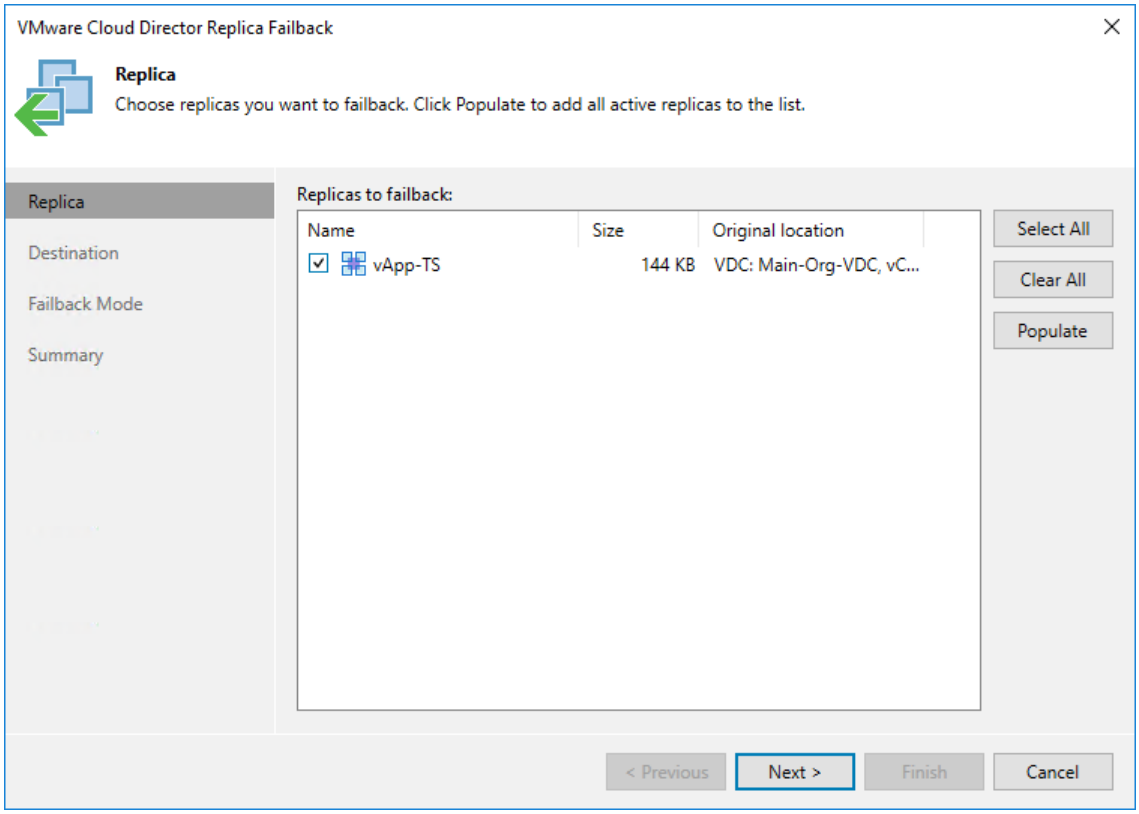

## Step 3. Select Failback Destination

At the **Destination** step of the wizard, select the failback destination and backup proxies for vApp data transport during failback:

1. Select a destination for failback. Veeam Backup & Replication supports the following options:
   - **Failback to the original vApp** – select this option if you want to fail back to the source vApps that reside on the source hosts. Veeam Backup & Replication will synchronize the state of the source vApps with the current state of their replicas to apply any changes that occurred to the replicas while running in the disaster recovery (DR) site.

     If this option is selected, you will proceed to the **Failback Mode** step of the wizard.
   - **Failback to the original vApp restored in a different location** – select this option if the source vApps have already been recovered to a new location, and you want to switch to the recovered vApps from their replicas. Veeam Backup & Replication will synchronize the state of the recovered vApps with the current state of the vApp replicas to apply any changes that occurred to the replicas while running in the DR site.

     If this option is selected, you will proceed to the **Target vApp** step of the wizard.

     > **TIP**
     >
     > You can restore to another VMware Cloud Director.
   - **Failback to the specified location** – select this option if you want to recover the source vApps from replicas. You can recover vApps to a new location, or to any location but with different settings (such as network settings, virtual disk type, configuration file path and so on). Select this option if there is no way to fail back to the source vApp or an already recovered vApp.

     If you select this option, the wizard will include additional steps.

   If you select one of the first two options, Veeam Backup & Replication will send to the source/recovered vApps only differences between the existing virtual disks of VMs included in the vApps.
   Veeam Backup & Replication will not send replica configuration changes such as different IP address or network settings (if replica Re-IP and network mapping were applied), new hardware or virtual disks added while the replicas were in the *Failover* state.

   If you select **Failback to the specified location**, Veeam Backup & Replication will send to the specified location whole replica data, including configurations and virtual disk content.

2. To select which backup proxies will be used for data transfer, click **Pick backup proxies for data transfer**.

   By default, Veeam Backup & Replication selects proxies automatically. Before processing a new vApp in the vApp list, Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication selects the most appropriate proxy basing on the following information: transport modes that the backup proxies can use and the current workload on the backup proxies.

   If want to select proxies manually and if vApps and their replicas reside in different sites, select at least one backup proxy in the production site and one proxy in the disaster recovery site. If vApps and replicas reside in the same site, you can use the same backup proxy as the source and target one.

   We recommend that you select at least two backup proxies in each site to ensure that failback will be performed in case one proxy fails or looses the network connection.

4. [For failback to the source vApps] If you want to fasten failback, and the source vApps had problems at the guest OS level, select the **Quick rollback** check box.

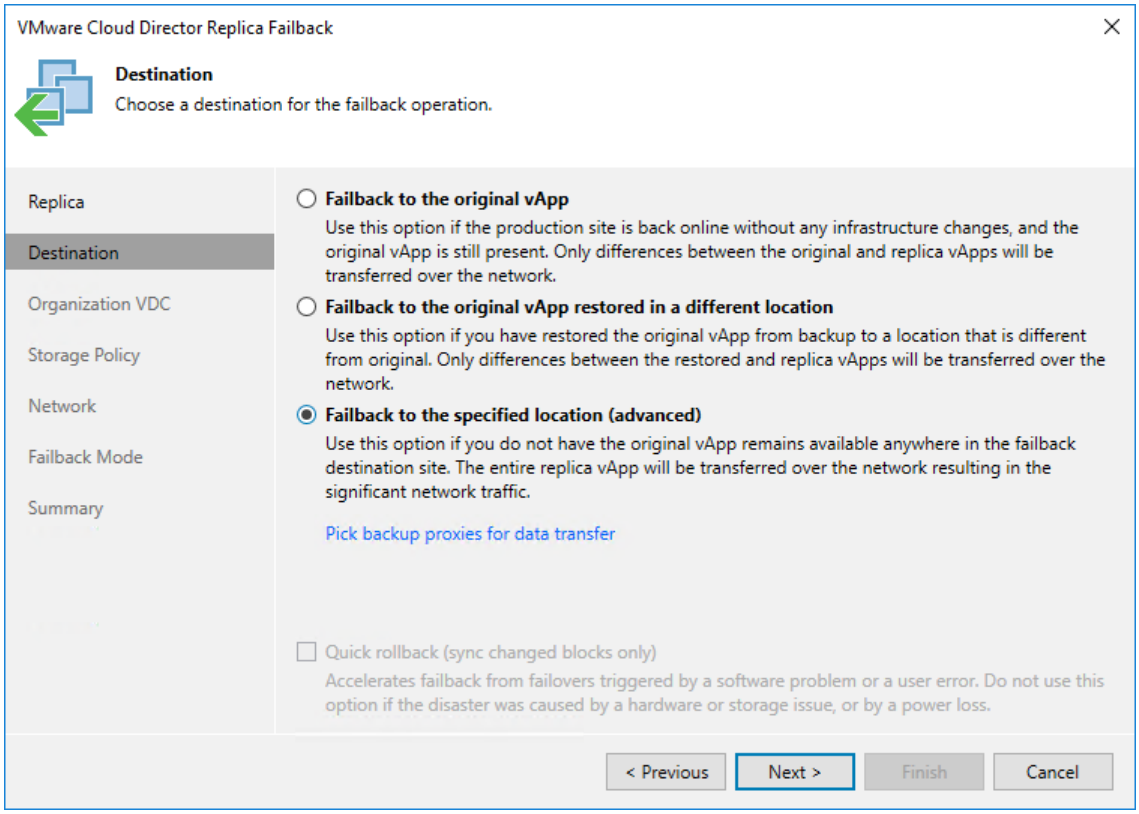


# Restoring Storage Policies

If the replicated vApp was associated with the storage policy, in the failback to original location scenario, Veeam Backup & Replication will associate the restored vApp with this storage policy.

When you click **Next**, Veeam Backup & Replication will check storage policies in the virtual environment and compare this information with the information about the replica storage policy. If the original storage policy has been changed or deleted, Veeam Backup & Replication will display a warning. You can select one of the following options:

- **Current** – the restored VM will be associated with the profile with which the source VM in the production environment is currently associated.
- **Default** – the restored VM will be associated with the profile that is set as default for the target datastore.
- **Stored** – the restored VM will be associated with the profile that was assigned to the source VM at the moment of replication.

For more information, see Storage Profiles.

## Step 4. Specify Organization VDCs

The **Organization VDC** step is available if you have selected the **Failback to the specified location** option at the **Destination** step.

At the **Organization VDC** step of the wizard, specify names for the restored vApps and the organization VDCs to which Veeam Backup & Replication will add restored vApps. To do this, select the necessary vApp and use the Name and VDC buttons.

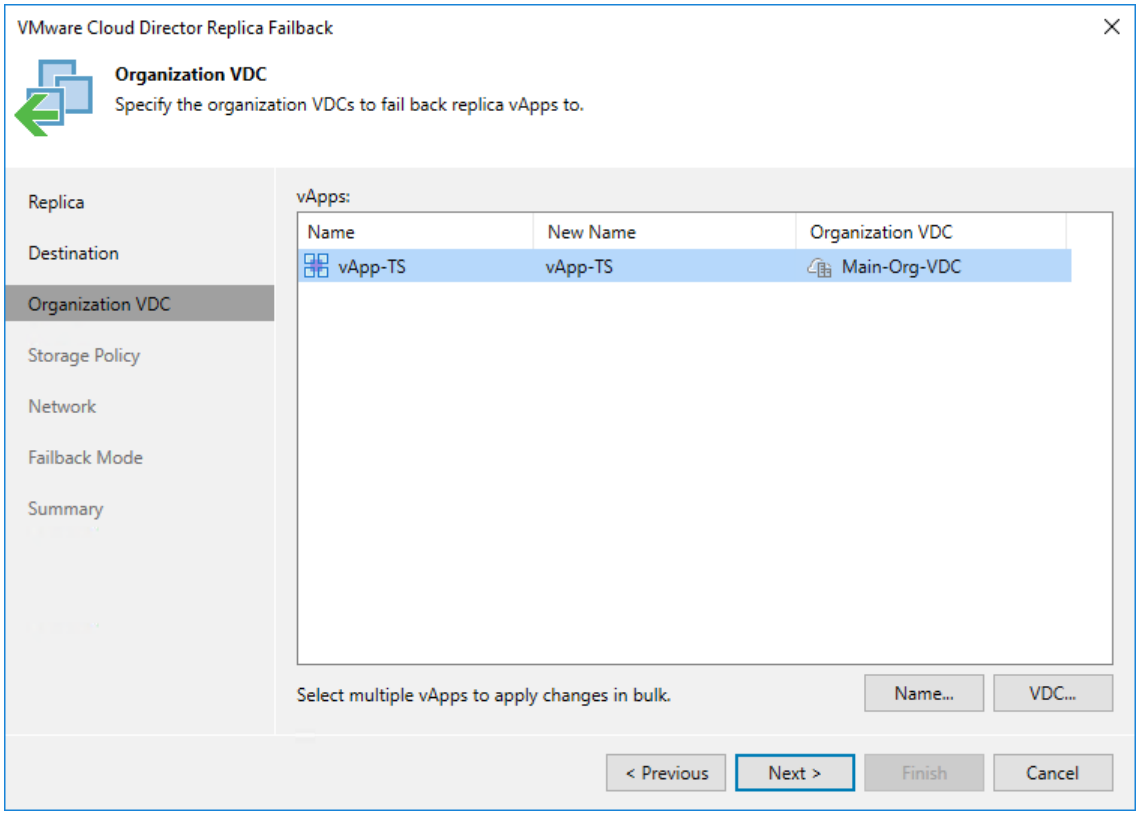

## Step 5. Specify Storage Policies

The **Storage Policy** step is available if you have selected the **Failback to the specified location** option at the **Destination** step.

At the **Storage Policy** step of the wizard, specify storage policies that Veeam Backup & Replication will apply to vApps that you want to restore:

1. In the **vApps** list, select vApps for which you want to change the policy and click **Policy**.
2. In the **Select storage policy** window select the policy that you want to apply.

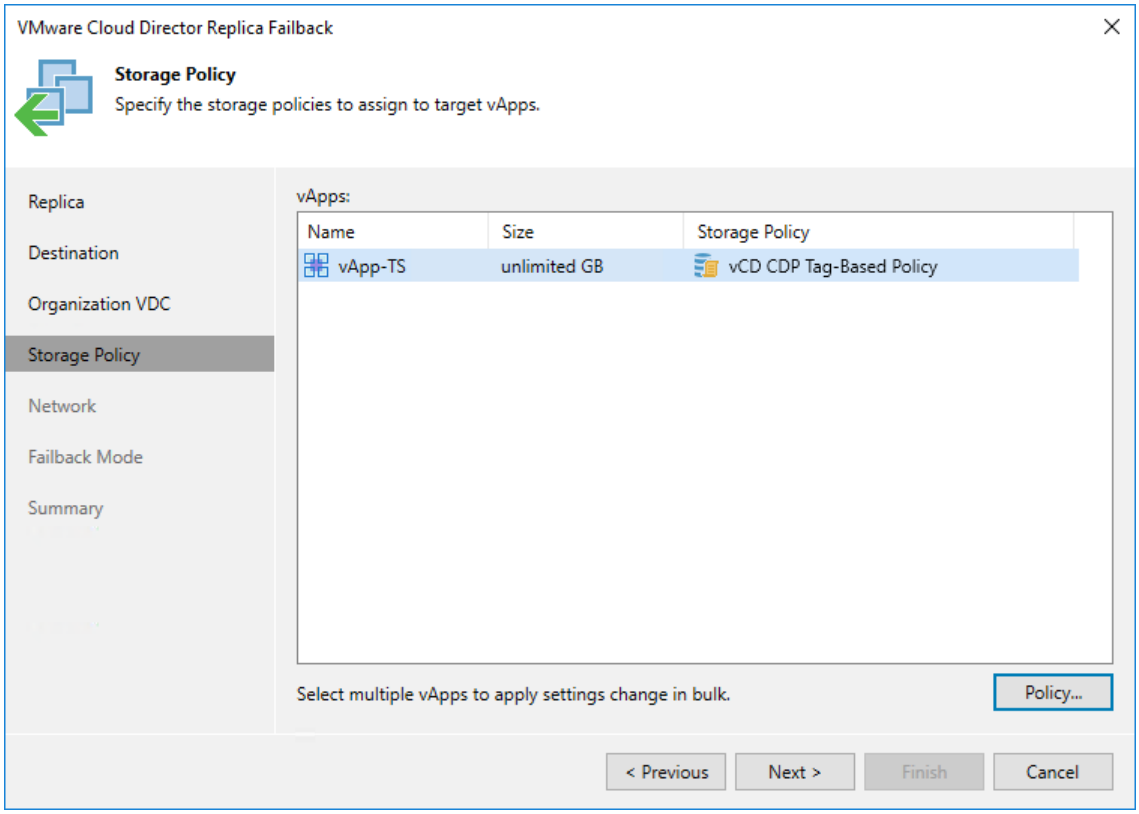

## Step 6. Configure Network Mapping

The **Network** step is available if you have selected the **Failback to the specified location** option at the **Destination** step.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the DR site to networks in the site where recovered vApps reside. Veeam Backup & Replication will use the network mapping table to update configuration files of VMs added to vApps on the fly, during the failback process.

To change networks to which the restored vApps will be connected:

1. In the **Network connections** list, select the necessary vApps and click **Network**.

   If vApps are connected to multiple networks, select the necessary network and click **Network**.

2. In the **Select network** window, select networks to which vApps must be connected after restore.

If you do not want to connect restored vApps to any virtual network, select the necessary vApps and click **Disconnect**.

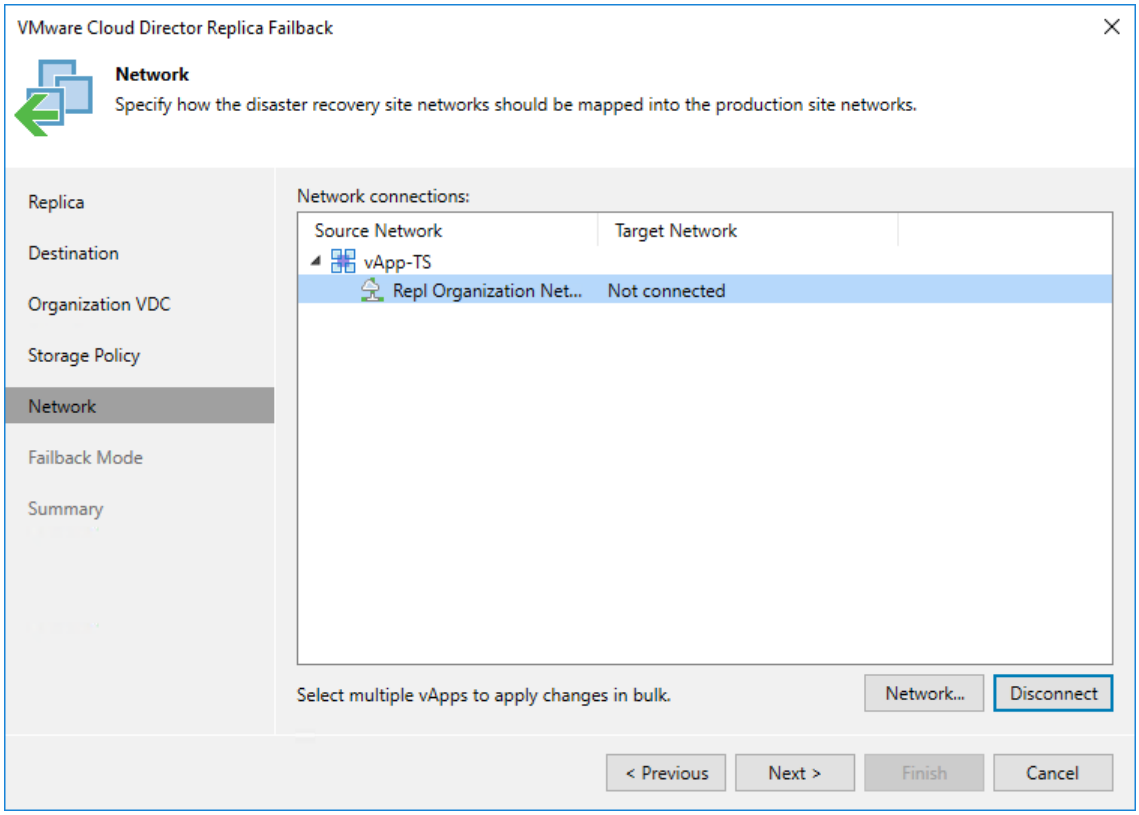

## Step 7. Select Target vApp

The **Target vApp** step is available if you have selected the **Failback to the original VM restored in a different location** option at the **Destination** step.

At the **Target vApp** step of the wizard, specify to which vApps you want to fail back from replicas. These vApps must be already restored from backups in the required location.

By default, Veeam Backup & Replication fails back the replica to the source vApps. If you want to specify the target vApp manually, perform the following steps:

1. Select a replica and click **Edit**.
2. In the **Selects Objects** window, select a vApp or vApp container to which you want to fail back.
3. Click **Add**.

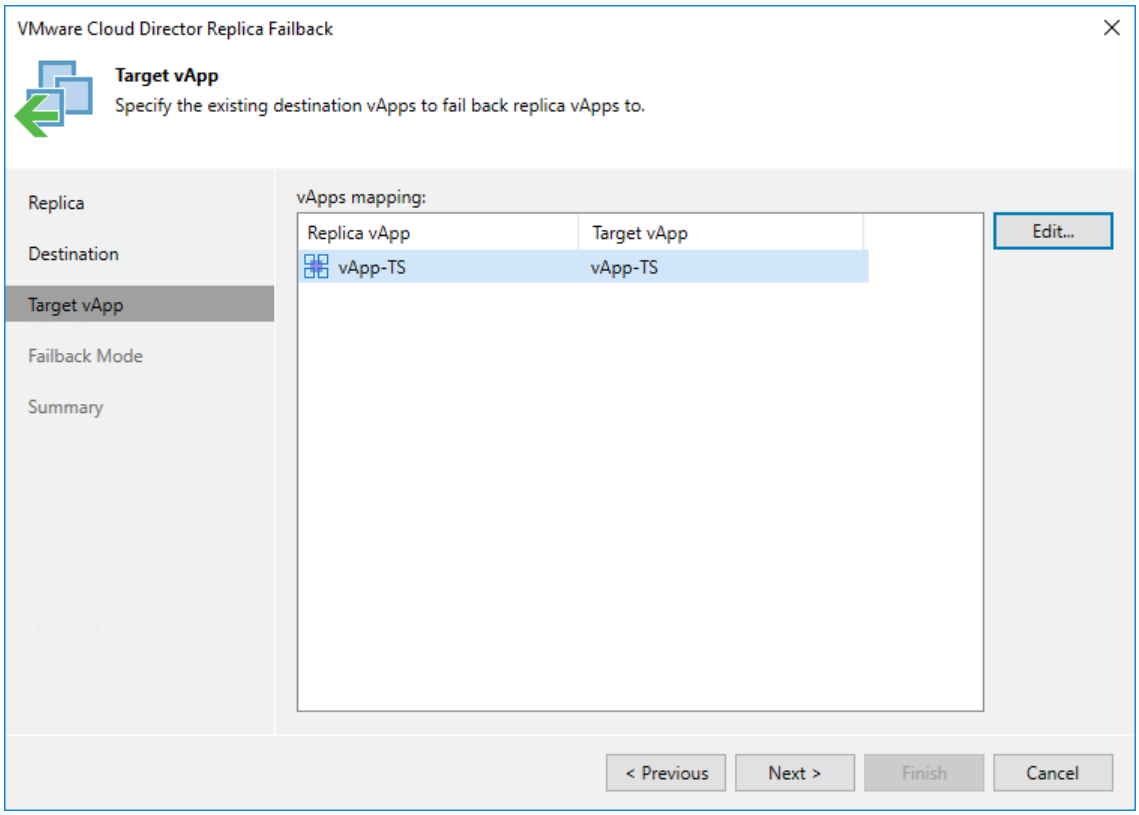

## Step 8. Schedule Switch to Production vApps

At the **Failback Mode** step of the wizard, specify when switch from replicas to production vApps must be performed:

- Select **Auto** if you want Veeam Backup & Replication to perform the switch automatically right after the state of the production vApps is synchronized with the state of their replicas.
- Select **Scheduled** if you want Veeam Backup & Replication to perform the switch at a specific time.
- Select **Manual** if you want to perform the switch manually.

If you select the **Scheduled** or **Manual** option, you can further reset or set the scheduled time or switch to the production VM manually. For more information, see Changing Switching Time and Switching to Production vApps Manually.

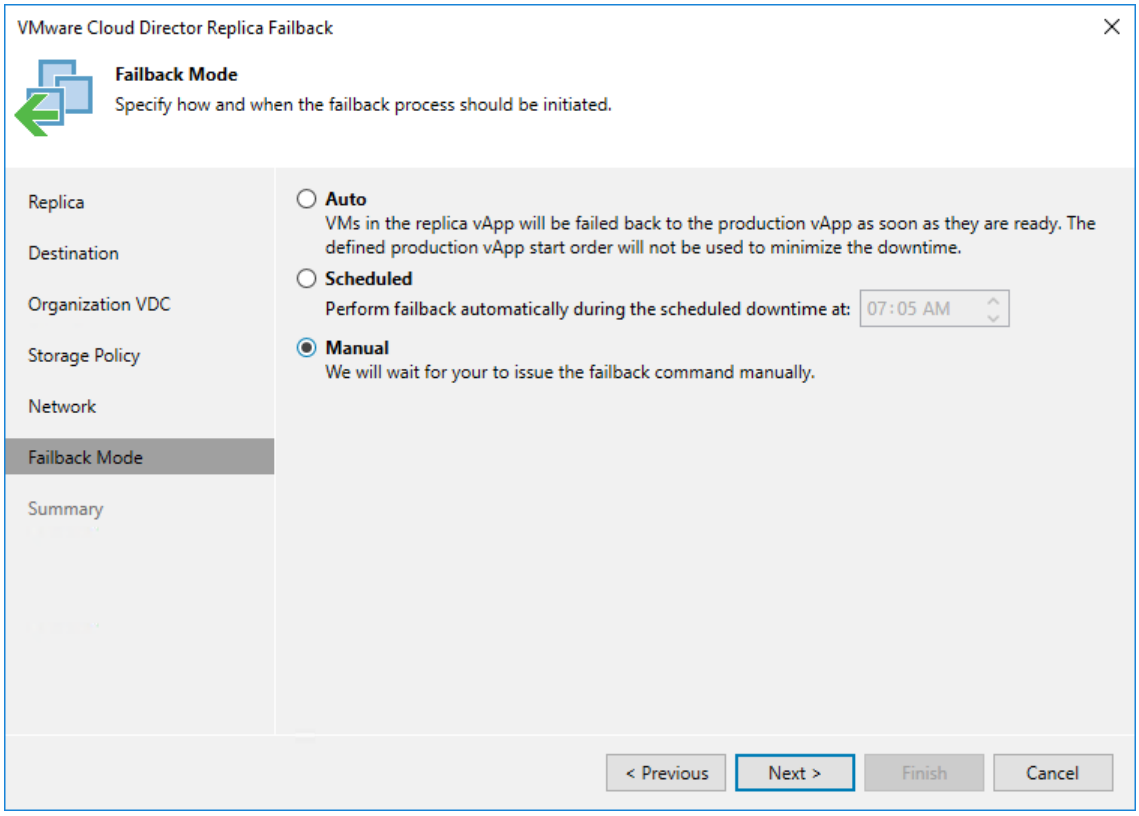

### Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured failback settings. If you want to power on the production VMs right after the switch to production operation is performed, select the **Power on target vApp after failback** check box. Then click **Finish**.

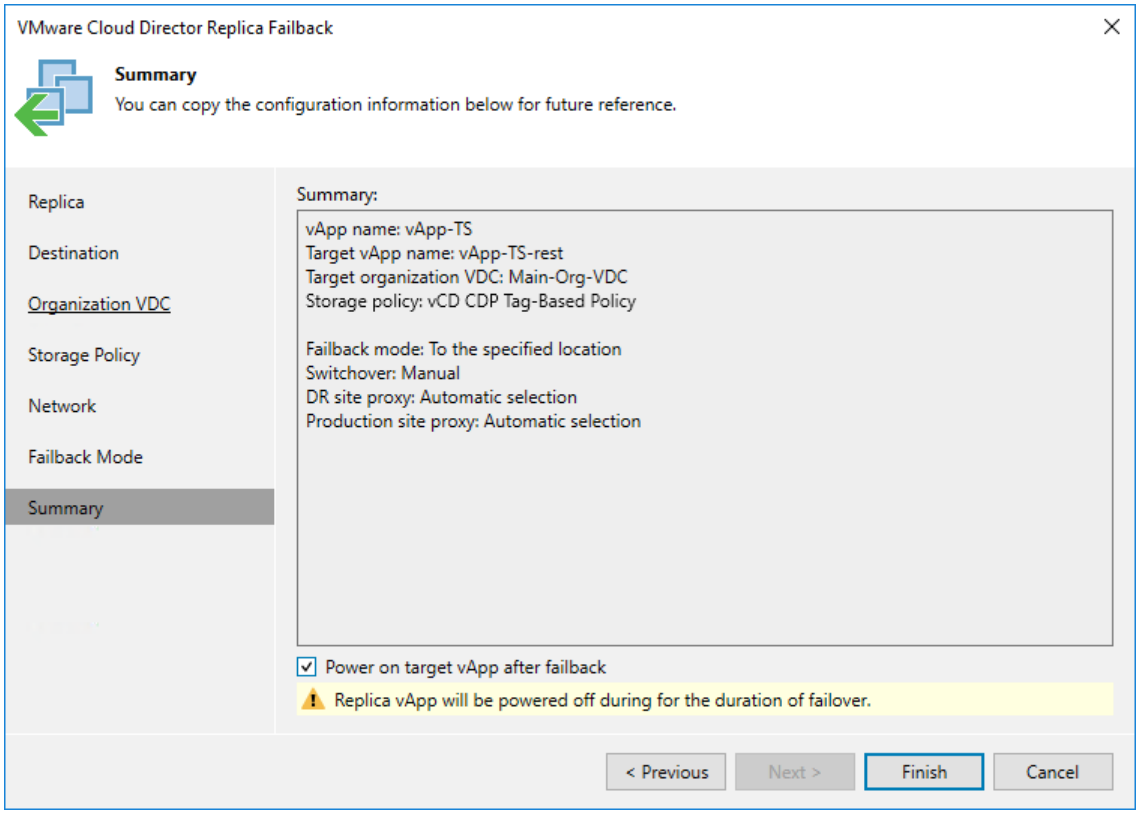


## What You Do Next

Failback is an intermediate step that needs to be finalized. You can finalize failback in the following ways:

- Commit failback
- Undo failback

## Changing Switching Time

The following instructions apply if you have selected to switch from replicas to production vApps manually or at the scheduled time at at the **Failback Mode** step of the **Failback** wizard.

To change the time when the switch from replicas to production vApps must be performed:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the vApp in the *Ready to switch* state and select **Switchover Time** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Change switchover time.**

If the switching time operation failed, you can retry this operation again. To perform a retry, in the working area, select the necessary vApp and select **Retry Switchover Time** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry switchover time.**

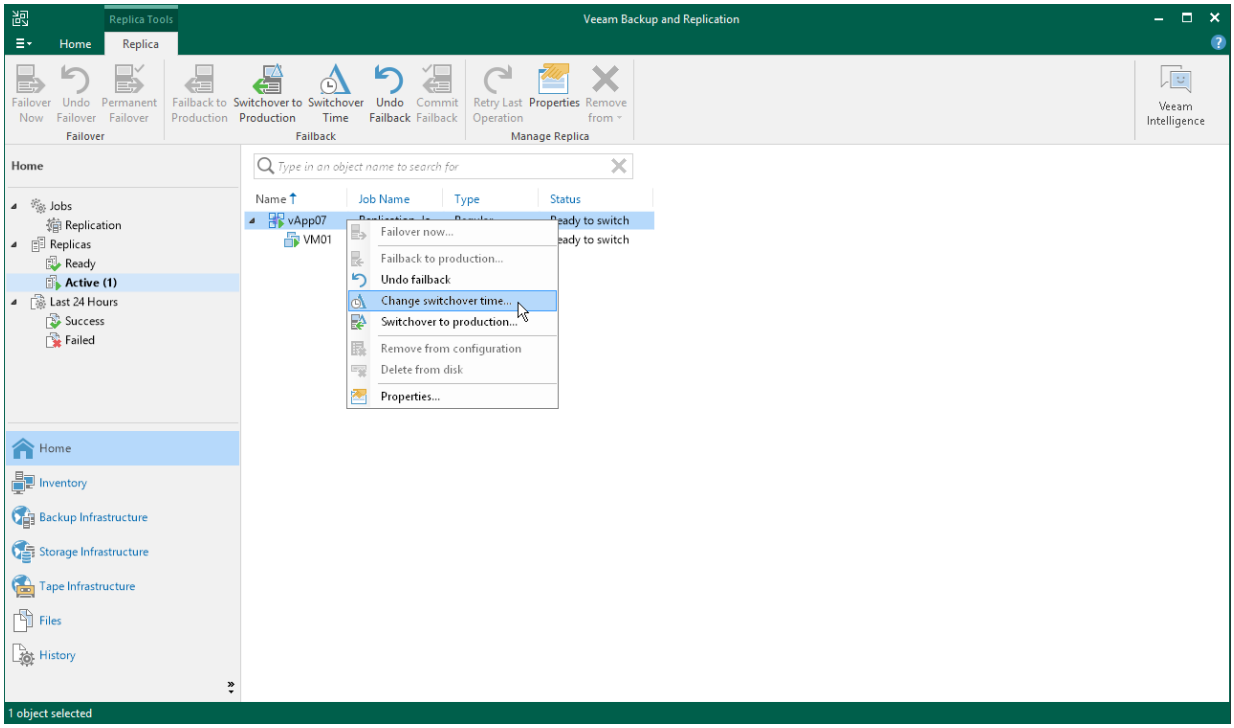


## Switching to Production vApps Manually

The following instructions apply if you have selected to switch from replicas to production vApps manually at the **Failback Mode** step of the **Failback** wizard.

To switch to a production vApp from its replica, do the following:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary vApp and select **Switchover to production** on the ribbon. As an alternative, you can right-click the necessary vApp and select **Switchover to production**.

If the switch to production operation failed, you can retry this operation again. To perform a retry, in the working area, select the necessary vApp and select **Retry Switchover to Production** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry switchover to production**.

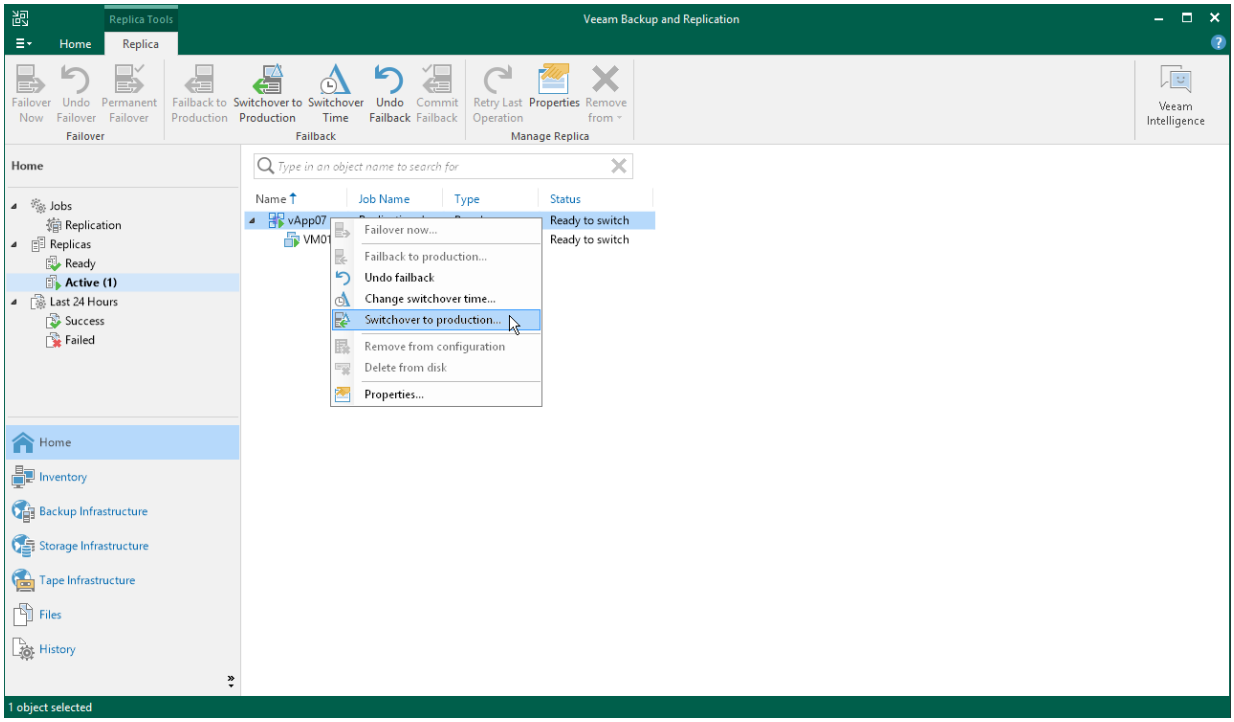


# What You Do Next

After you switch to the production VM, you must finalize failback. You can finalize failback in the following ways:

- Commit failback
- Undo failback

## Performing Failback Retry

The failback retry option is necessary when failback of vApps fails with the *Incomplete* state. When you perform a retry, Veeam Backup & Replication restarts failback only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, failback takes less time and does not consume as many resources as when processing the whole vApp.

To retry failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary vApp and select **Retry Failback** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry failback**.

# Failback Undo

Failback undo is one of the ways to finalize failback. You can use this option if the vApp to which you failed back (the production vApp) works in a wrong way and you want to get back to the replica.

The failback undo operation is performed in the following way:

1. Veeam Backup & Replication powers off the production vApp.
2. Veeam Backup & Replication reverts the replica to its pre-failback state.
3. Veeam Backup & Replication powers on the replica and changes the replica state from *Failback* or *Ready to Switch* to *Failover*.

## Undoing Failback

For more information on failback undo, see Failover and Failback for Cloud Director CDP and Failback Undo.

To undo failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary replica and click **Undo Failback** on the ribbon. Alternatively, you can right-click the necessary replica and select **Undo Failback**.

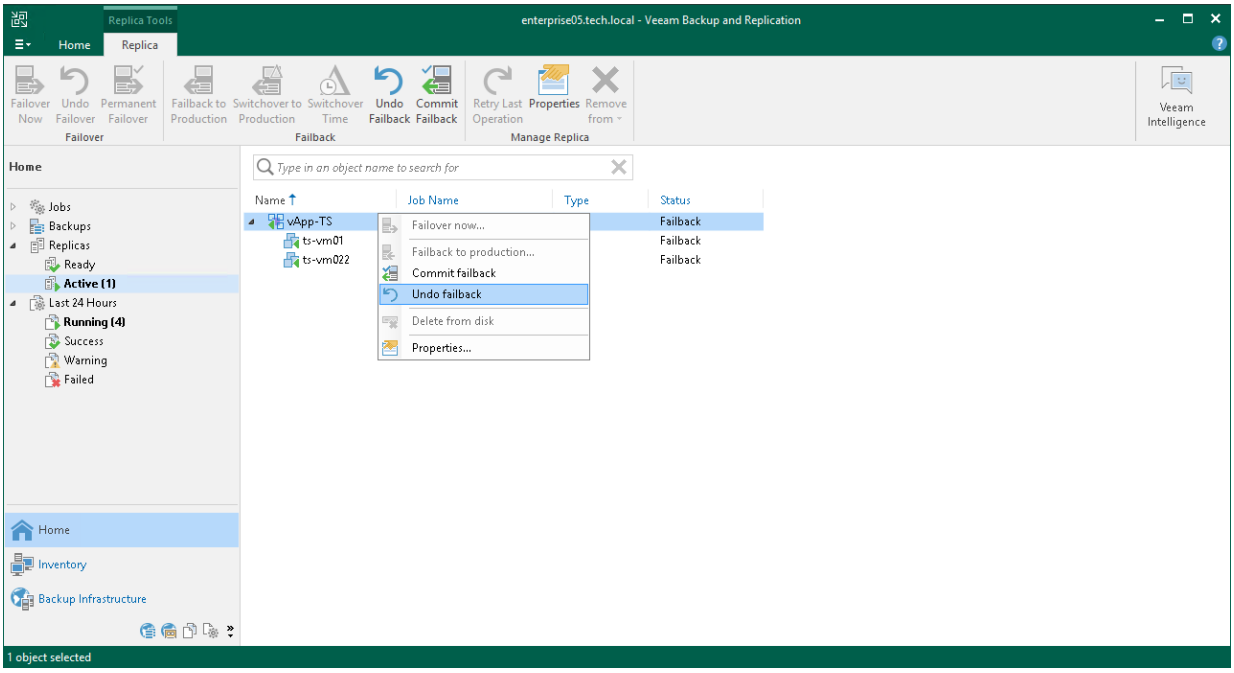


## Performing Failback Undo Retry

If the failback undo operation failed, you can retry this operation. When you perform a retry, Veeam Backup & Replication restarts the failback undo operation only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, the failback undo operation takes less time and does not consume as many resources as when processing the whole vApp.

To perform a retry:

1. Open the **Home** view, in the inventory pane, navigate to the **Replicas** > **Active** node.
2. In the working area, select the necessary vApp and select **Retry Undo Failover** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry undo failback**.

# Failback Commit

Failback commit is one of the ways to finalize failback. When you commit failback, you confirm that the vApp to which you failed back (the production vApp) works as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production vApp.

> **NOTE**
>
> If during failback, you have selected to switch to the production VM manually, you must first perform the switchover.

The failback commit operation is performed in the following way:

1. Depending on whether you have failed back to the source vApp or recovered vApp:
   - If you have failed back to a vApp recovered from a backup or replica, Veeam Backup & Replication reconfigures all existing jobs where the source vApp is present and adds the source vApp to the list of exclusions. The recovered vApp takes the role of the source vApp and is included into all jobs instead of the excluded vApp. When the VMware Cloud Director replication process starts, Veeam Backup & Replication processes the recovered vApp instead of the former source vApp.
   - If you have failed back to the source vApp, the replication job or policy is not reconfigured. When the replication process starts, Veeam Backup & Replication still processes the source vApp.
2. Veeam Backup & Replication changes the state of the replica from *Failback* to *Ready*.

During failback commit, the failback delta disk that saves the pre-failback state of a replica is not deleted. Veeam Backup & Replication uses this delta disk as an additional restore point for replica. With the pre-failback delta disk, Veeam Backup & Replication needs to transfer fewer changes and therefore puts less load on the network when replication activities are resumed.

## Committing Failback

For more information on failback commit, see Failover and Failback for Cloud Director CDP and Failback Commit.

To commit failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. In the working area, select the necessary replica and click **Commit Failback** on the ribbon. As an alternative, you can right-click the replica and select **Commit failback**.

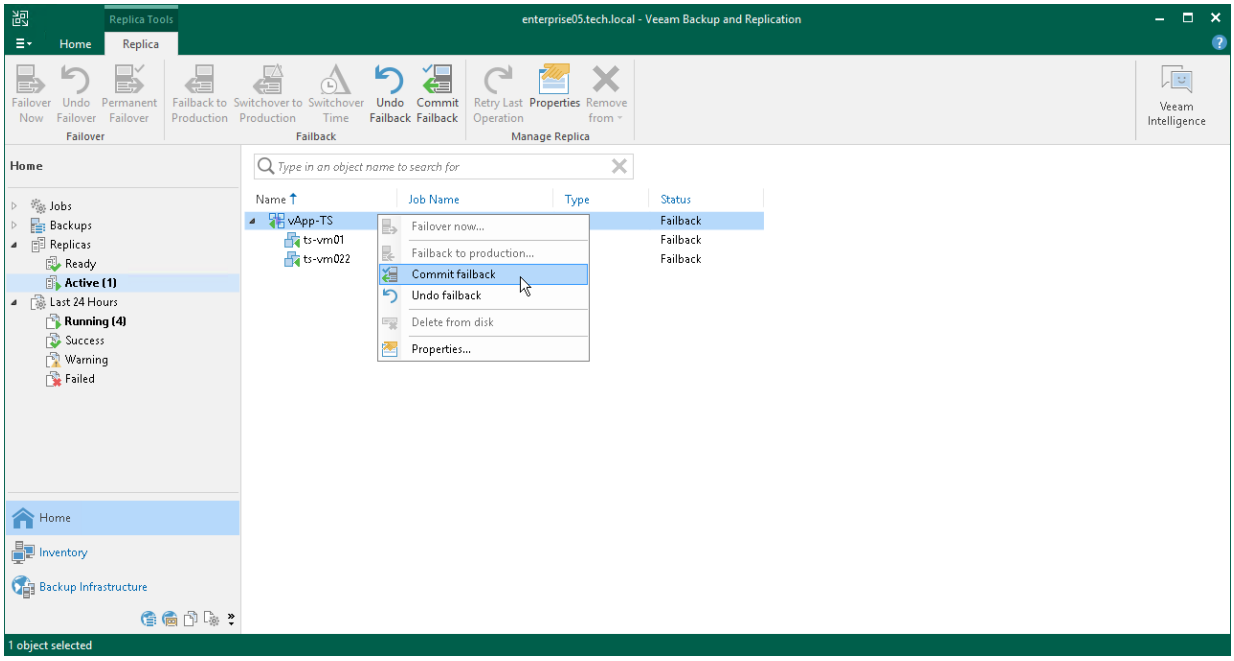


# Performing Failback Commit Retry

If the failback commit operation failed, you can retry this operation. When you perform a retry, Veeam Backup & Replication restarts the failback commit operation only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, the failback commit operation takes less time and does not consume as many resources as when processing the whole vApp.

To perform a retry:

1. Open the **Home** view, in the inventory pane, navigate to the **Replicas** > **Active** node.
2. In the working area, select the necessary vApp and select **Retry Commit Failback** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry commit failback**.

# Replication for VMware Cloud Director

VMware Cloud Director replication is a technology that allows you to replicate VM containers between production and disaster recovery VMware Cloud Director environments, to the source organization VDC or to another organization VDC added to a different VMware Cloud Director. Using the VMware Cloud Director replication you can replicate the following VM containers:

- vApps
- Cloud Director organizations
- Cloud Director organization VDC
- VMware Cloud Director instances

The VMware Cloud Director replication technology utilizes the same mechanisms as the VM replica and follows the same recovery scenarios. The main differences from VM replica are the following:

- The VMware Cloud Director replica target is an organization VDC and it must be set up beforehand in your VMware Cloud Director infrastructure.
- A minimal unit of a VMware Cloud Director replication is a vApp, you cannot replicate a single VM that is added to a vApp.
- The VMware Cloud Director replica snapshots are created on the VMs that are added to the target vApp.
- A single restore point of the VMware Cloud Director replica contains snapshots of all VMs added to a vApp.
- Veeam Backup & Replication assigns a new CLOUD.UUID to the replicated VMs when a VMware Cloud Director replication job runs.

The VMware Cloud Director replication technology supports functionality similar to normal VM replication: you can create VMware Cloud Director replication jobs, perform failover and failback operations with replicated VM containers and manage these VM containers.

## Related Topics

- Creating Cloud Director Replication Job
- Managing Cloud Director Replication Jobs
- Managing Cloud Director Replicas
- Failover and Failback for Cloud Director

# vApp Replica Seeding and Mapping

Replica seeding and mapping are technologies that help reduce the amount of traffic sent over the network. With these technologies, Veeam Backup & Replication does not have to transfer all of vApp data from the source host to the target host across the sites during the initial synchronization.

You can use seeding and mapping in the following scenarios:

- **Seeding**

  Configure replica seeding if, in a backup repository located in the disaster recovery (DR) site, you have backups of vApps that you plan to protect. During replication, Veeam Backup & Replication will restore image of VMs that are added to vApps from these backups and will synchronize the state of the restored vApps with the latest state of the source vApps. Then Veeam Backup & Replication will use these restored vApps as replicas.

  For more information on how to create backups that can be used as "seeds" for replica, see Creating vApp Replica Seeds.

- **Mapping**

  Configure replica mapping if, on the host in the DR site, you have ready-to-use copies of the source vApps. These can be restored vApps or replicas. Veeam Backup & Replication will synchronize the state of these ready-to-use vApps with the latest state of the source vApps and will use these vApps as replicas.

You can also configure both replica seeding and replica mapping in the same job. For example, if a job includes 2 vApps, you can use seeding for one vApp and map the other vApp to an existing vApp.

> **IMPORTANT**
>
> If seeding or mapping is enabled in a job, all vApps in the job must be covered with seeding or mapping. If a vApp neither has a seed, nor is mapped to an existing vApp, it will be skipped from processing.

## Algorithm for Seeding

Replica seeding includes the following steps:

1. As a preparatory step for replica seeding, you need to create a backup of a vApp that you plan to protect. For more information on how to create a backup that will be used as a "seed" for replica, see Creating vApp Replica Seeds.
2. When you create a replication job, you should point it to a backup repository in the DR site. During the initial synchronization, Veeam Backup & Replication accesses the backup repository where the replica seed is located, and restores the image of VMs that are added to vApp from the backup. The restored vApp is registered on the target host in the DR site. Files of the restored vApp are placed to the location you specify as the replica destination datastore.

   Virtual disks of a replica restored from the backup preserve their format (that is, if the source VM used thin provisioned disks, virtual disks of the replica are restored as thin provisioned).
3. Veeam Backup & Replication synchronizes the restored vApp with the latest state of the source vApp.

   After successful synchronization, in the **Home** view in the Veeam Backup & Replication console, under **Replicas** node you will see a vApp replica with two restore points. One point will contain the state of the vApp from the backup file; the other point will contain the latest state of the source vApp you want to replicate.
4. During incremental synchronization, Veeam Backup & Replication transfers only incremental changes in a regular manner.

Replica seeding dramatically reduces traffic sent over WAN or slow connections because Veeam Backup & Replication does not send the full contents of the vApp. Instead, it transmits only differential data blocks.

# Algorithm for Mapping

Replication to a mapped vApp is performed in the following way:

1. The first step differs depending on which vApp you have selected for mapping:
   - If you have selected a regular vApp, Veeam Backup & Replication calculates the differences between the source and mapped vApp.
   - If you have selected a snapshot replica, Veeam Backup & Replication deletes all restore points and delta disks and then calculates the differences between the source and mapped vApp.
   - If you have selected a CDP replica, Veeam Backup & Replication imports all restore points of this replica and then calculates the differences between the source and mapped vApp. Note that if disk sizes of the source and mapped vApp differ, Veeam Backup & Replication will delete all restore points of the mapped vApp.
2. To synchronize the state of the mapped vApp with the state of the source vApp, Veeam Backup & Replication sends the calculated changes to the mapped vApp.

   The first and second steps take place during the initial synchronization.
3. During the incremental synchronization, Veeam Backup & Replication transfers only incremental changes in a regular manner.

After the successful initial synchronization, in the **Home** view of Veeam Backup & Replication, under **Replicas** node you will see a replica with restore points. If you have selected for mapping a regular vApp or snapshot replica, you will see two restore points: one restore point will contain the latest state of the mapped vApp, the other will contain the state of the source vApp. If you have selected a CDP replica, you will see all restore points of the mapped vApp plus one restore point that will contain the state of the source vApp.

# Network Mapping

To establish a connection between a production VM and its replica, Veeam Backup & Replication uses network mapping.

Network mapping is implemented with the help of a network mapping table. This table contains networks in the production site and suitable networks in the disaster recovery (DR) site. When the replication session starts, Veeam Backup & Replication checks the network mapping table. Then Veeam Backup & Replication updates replica configuration to replace the production networks with the specified networks in the DR site.

Veeam Backup & Replication supports the following types of network mapping:

- **Manual network mapping**.

  During manual network mapping, Veeam Backup & Replication uses the network mapping table that a user creates when configuring a replication job.

- **Automatic network mapping**.

  During automatic network mapping, Veeam Backup & Replication analysis networks in the DR site and searches for the networks with the same names as in the production site. If such networks are found, Veeam Backup & Replication tries to configure network connections similar to the production site. This mapping applies to networks for which manual mapping is not configured.

> **NOTE**
>
> Cloud Director replication jobs do not support network mapping of vApp networks. You can configure a mapping table for organization VDC networks only.

# Replica Network Settings Configuration

How Veeam Backup & Replication configures replica network settings depends on how IP addresses are allocated for the source VM NICs.

## DHCP Allocation Mode

If DHCP allocation mode is used on the source VM, replica NIC settings are the following:

- If mapping (manual or automatic) was performed — NIC is connected, the IP mode is DHCP and new MAC address is assigned.
- If mapping (manual or automatic) was not performed — NIC is disconnected, the IP mode is not set and new MAC address is assigned.

## Static-Manual IP Allocation Mode

If static-manual IP allocation mode is used on the source VM, replica NIC settings are the following:

- If mapping (manual or automatic) was performed — NIC is disconnected, the IP mode is not set and new MAC address is assigned.
- If mapping (manual or automatic) was not performed — NIC is disconnected, the IP mode is not set and new MAC address is assigned.

If manual mapping is used, you can connect a replica NIC after the first replication job run. However, note that to stay connected the following NIC settings must be the same as on the source VM: connection state, primary NIC value and network adapter type. Also, the replica NIC network must be the same as the mapped network (the target network from the mapping table). If any of these parameters differs, the NIC will be disconnected after the next job run.

## Static IP Pool Allocation Mode

If static IP pool allocation mode is used on the source VM, Veeam Backup & Replication assigns replica NIC settings in the following way:

- If mapping (manual or automatic) was performed — NIC is connected, the IP mode is DHCP and new MAC address is assigned.
- If mapping (manual or automatic) was not performed — NIC is disconnected, the IP mode is not set and new MAC address is assigned.

# Creating Cloud Director Replication Job

To replicate a vApp or another VM container, you must configure a VMware Cloud Director replication job. The VMware Cloud Director replication job defines a scope of VM containers to replicate, where to store replicated VM containers and how often to replicate VM containers. Every time the VMware Cloud Director replication job runs, replica snapshots are created on the VMs that are added to the target vApp. If a disaster strikes, you can fail over to the necessary restore point of your VM container. During failover, Veeam Backup & Replication shifts from the source VM container to their replicas. As a result, you have fully functional vApps within a couple of seconds, and your users can access services and applications with minimum disruption.

> **IMPORTANT**
>
> VMware Cloud Director replication job does not support replication of a single VM that is added to a vApp. You can replicate only vApps or VM containers.

To create a VMware Cloud Director replication job, do the following:

1. Check prerequisites.
2. Launch the New Replication Job wizard.
3. Specify a job name and advanced settings.
4. Select vApps to replicate.
5. Exclude objects from a replication job.
6. Specify the vApp processing order.
7. Select a destination for replicas.
8. Configure network mapping.
9. Specify replication job settings.
10. Specify advanced replica settings.
11. Specify data transfer and replica settings.
12. Configure seeding and mapping settings.
13. Specify guest processing settings.
14. Define a job schedule.
15. Finish working with the wizard.

## Before You Begin

Before you create a VMware Cloud Director replication job, check the following prerequisites:

- Check the supported Cloud Director versions in Supported Platforms and Applications.
- You must add VMware Cloud Director server to the backup infrastructure. The backup server must be able to resolve short names and connect to source and target virtualization hosts.
- You must set up the organization VDC that will keep the replicas.
- You must disable the VM discovery option in VMware Cloud Director settings. For more information on where you can change the option, see VMware Docs.

- VMware Cloud Director replication does not support replication of a vApp to a Cloud Provider.

  > **NOTE**
  >
  > It is possible to replicate vSphere VMs to your Cloud Provider Cloud. For more information, see Creating Replication Jobs.

- If you plan to replicate VM containers using WAN accelerators, source and target WAN accelerators must be added to the backup infrastructure and properly configured. For more information, see Adding WAN Accelerators.
- If you plan to replicate VM containers using WAN accelerators, it is recommended that you pre-populate global cache on the target WAN accelerator before you start the VMware Cloud Director replication job. Global cache population helps reduce the amount of traffic transferred over WAN. For more information, see Manually Populating Global Cache.
- If you plan to use pre-job and post-job scripts or pre-freeze and post-thaw scripts, you must create scripts before you configure the VMware Cloud Director replication job. Veeam Backup & Replication supports script files in the following formats: EXE, BAT, CMD, JS, VBS, WSF, PS1, SH. For more information, see Pre-Freeze and Post-Thaw Scripts.
- You must check limitations for replication. For more information, see Considerations and Limitations.
- Due to VMware vSphere limitations, if you change the size of VM disks on the source VM, all VM snapshots that were created previously, will not be available for failover, failback and other subsequent operations. For more information, see this VMware KB article.

# Step 1. Launch New Replication Job Wizard

To launch the **New Replication Job** wizard, do one of the following:

- On the **Home** tab, click **Replication Job** > **VMware vCloud Director**.
- Open the **Home** view, in the inventory pane right-click **Jobs** and select **Replication** > **VMware vCloud Director**.

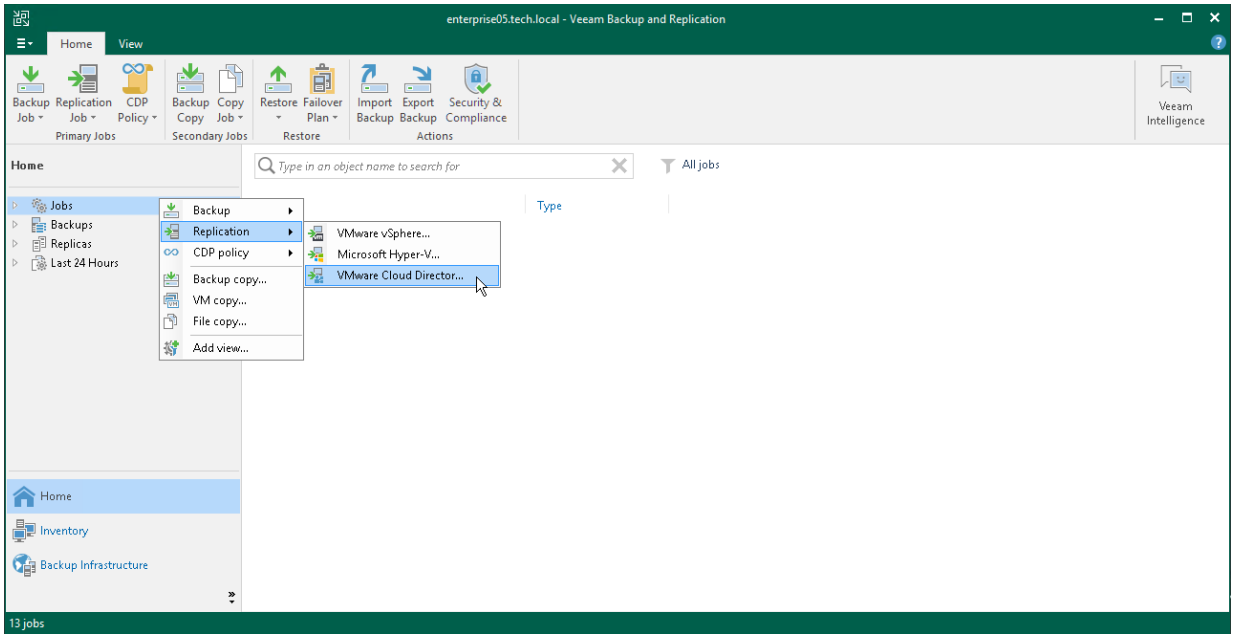

# Step 2. Specify Job Name and Advanced Settings

At the **Name** step of the wizard, specify a job name and description, and configure advanced settings for the VMware Cloud Director replication job.

1. In the **Name** field, enter a name of the replication job.
2. In the **Description** field, provide a description for future reference.
3. Depending on your DR site configuration, you can select the following advanced settings for the job:
   - Select the **Replica seeding (for low bandwidth DR sites)** check box to enable the Seeding step in the wizard. Replica seeding can be used if you plan to replicate vApps and want to reduce the amount of traffic sent over the network during the first run of the replication job.
   - Select the **Network remapping (for DR sites with different virtual networks)** check box to enable the Network step in the wizard. If the network in the DR site does not match the production network, you can resolve this mismatch by creating a network mapping table.

     > **NOTE**
     >
     > VMware Cloud Director replication jobs do not support network mapping of the vApp networks. You can configure a mapping table for organization VDC network only.

4. Select the **High priority** check box if you want Veeam Backup & Replication to prioritize this job higher than other similar jobs. Veeam Backup & Replication will allocate resources to this job in the first place. For more information on job priorities, see Job Priorities.

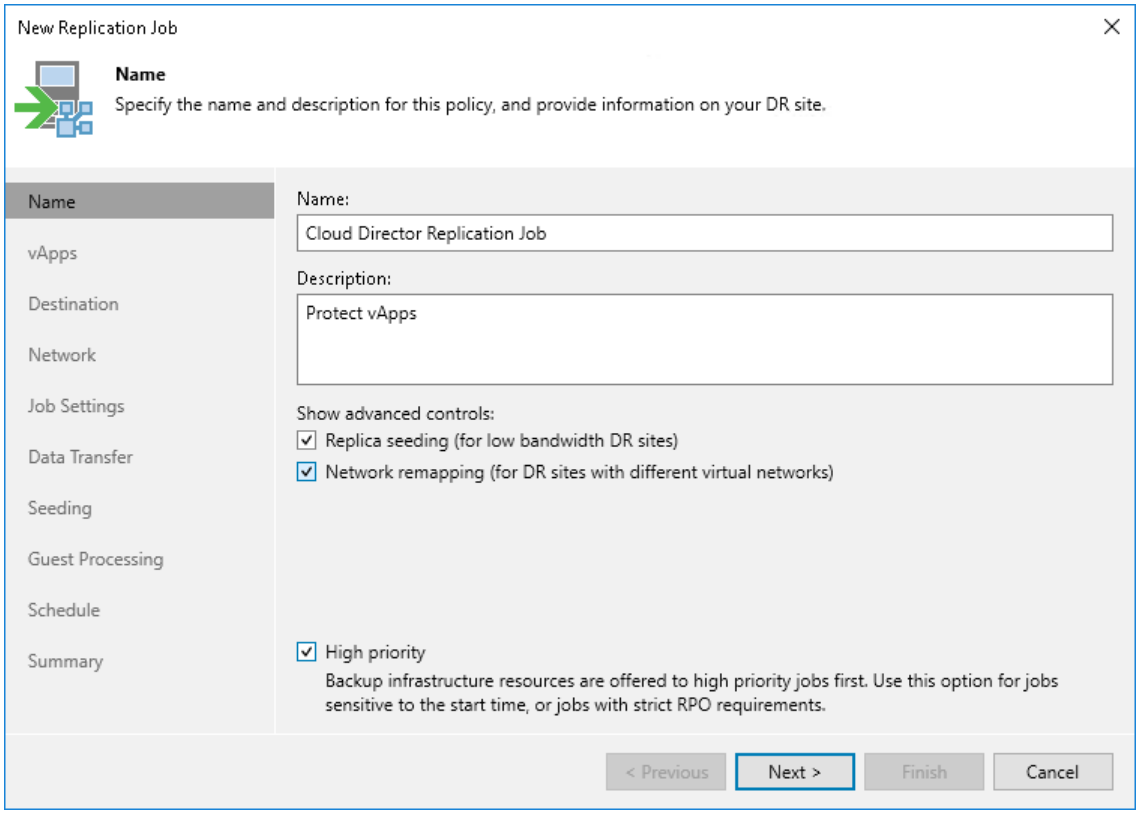

# Step 3. Select vApps to Replicate

At the **vApps** step of the wizard, select VM containers (vApps, organization or organization VDCs) that you want to replicate.

1. Click **Add**.
2. In the **Selects Objects** window, select the necessary VM containers. Click **Add**.

   When you add new items to VM containers, Veeam Backup & Replication updates settings automatically to process these new items.

**IMPORTANT**

VMware Cloud Director replication job does not support replication of a single VM that is added to a vApp. You can replicate only vApps or VM containers.

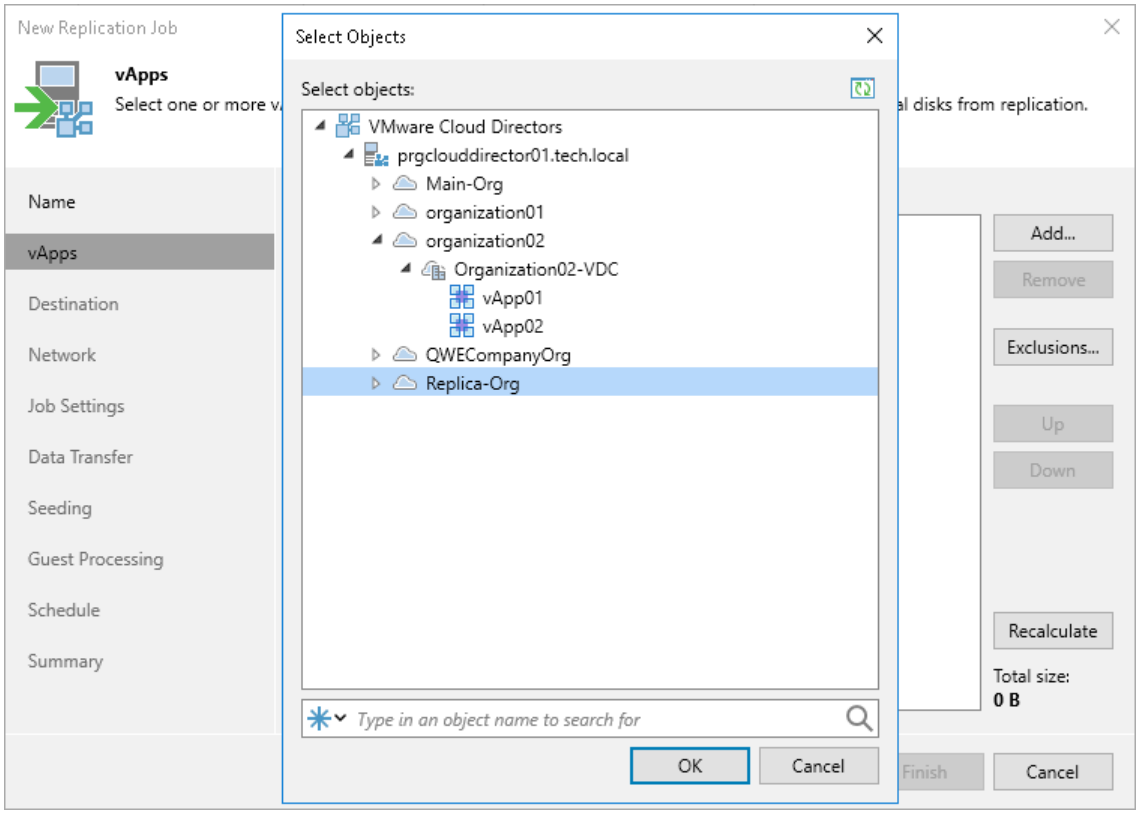

# Step 4. Exclude Objects

After you have added vApps or VM containers to the replication job, at the **vApps** step of the wizard you can specify which objects you want to exclude from the replication job.

> **NOTE**
>
> You can exclude only child objects that are added to a VM container. For example, if you add an organization VDC to a replication job, you can exclude only vApps that are available in this organization VDC.

# Excluding VMs and VM Containers

To exclude VMs and VM containers (vApps, organizations, organization VDCs and so on):

1. At the **vApps** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, check that the **VMs** tab is selected. Click **Add**.
3. In the **Select Objects** window, select VMs and VM containers that you want to exclude. Click **OK**.

   Select the **Show full hierarchy** check box to display the hierarchy of all VMware Cloud Director servers added to the backup infrastructure.

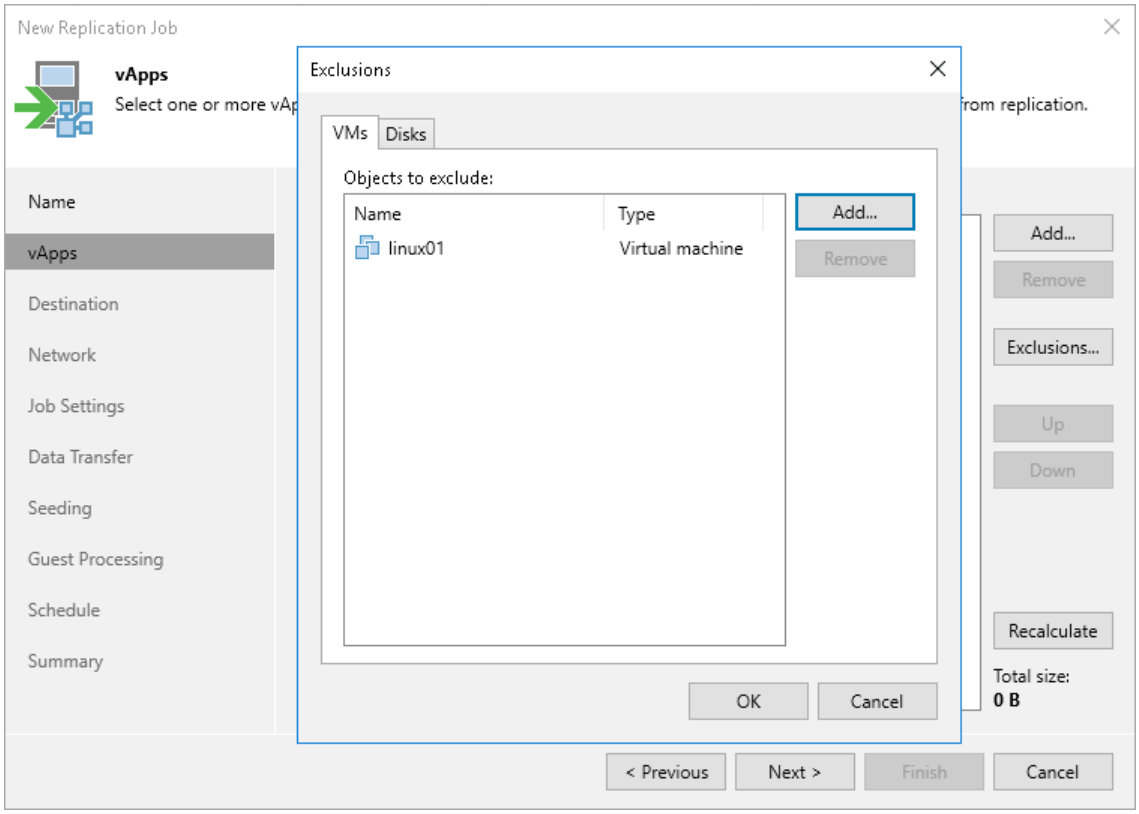


# Specifying Disks

To exclude disks:

1. At the **vApps** step of the wizard, click **Exclusions**.

2. In the **Exclusions** window, do the following:
   a. Switch to the **Disks** tab.
   b. To exclude disks of VMs, click **Add**. In the **Add Objects** window, select the necessary VMs and click **Add**. Veeam Backup & Replication will include these VMs in the list as standalone objects.
   c. In the **Disks to process** list, select the necessary VMs.
   d. Click **Edit**.
3. In the **Select Disks** window, select disks that you want to replicate: all disks, 0:0 disks (as a rule, system disks) or specific IDE, SCSI, SATA or NVMe disks. Disks that you do not select will be excluded from processing. Click **OK**.
4. In the **Exclusions** window, click **OK**.

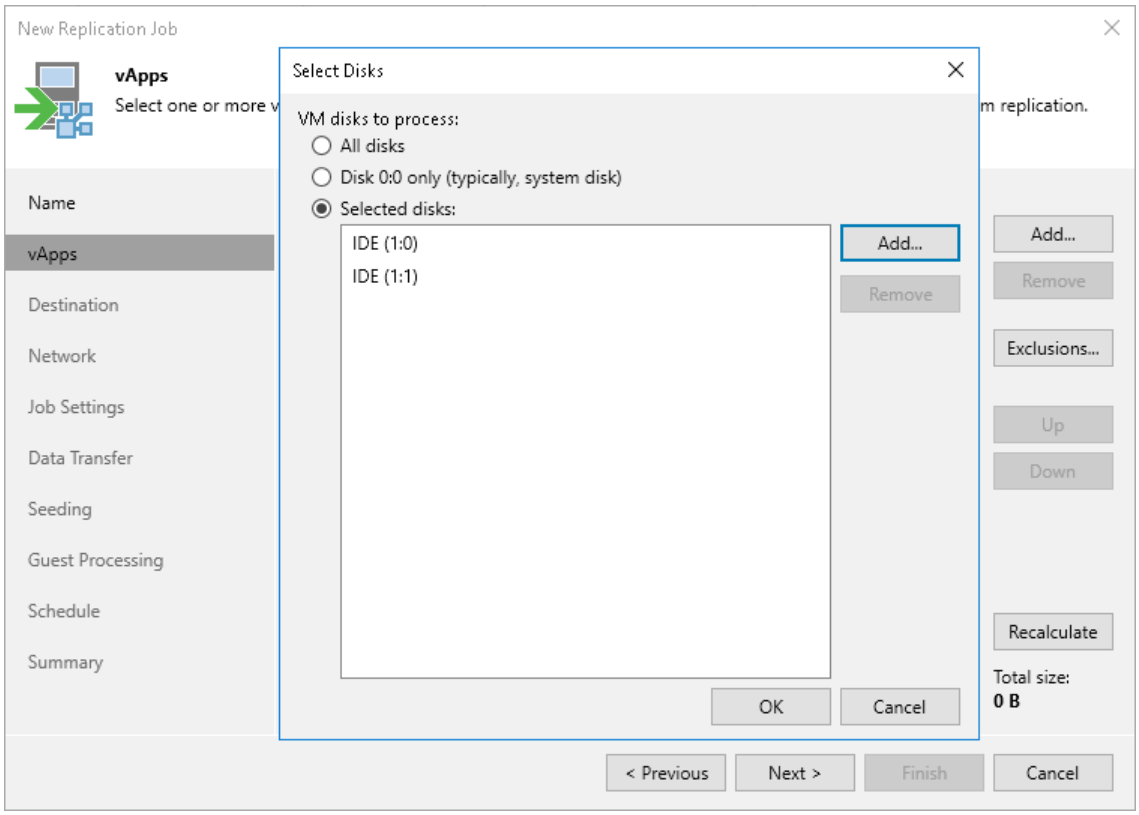

# Step 5. Specify vApp Processing Order

At the **vApps** step of the wizard, click **Up** and **Down** to change the processing order. VM containers (vApps, organization, organization VDCs, and so on) at the top of the list have a higher priority and will be processed first.

> **NOTE**
>
> Consider the following:
>
> - VMs inside a VM container are processed at random.
> - To ensure that vApps are processed in the specified order, you must add them as standalone vApps, not as a part of containers.
> - The processing order may differ from the order that you have specified. For example, if resources of a vApp that is higher in the priority are not available, and resources of a vApp that is lower in the priority are available, Veeam Backup & Replication will process the vApp with the lower priority first.

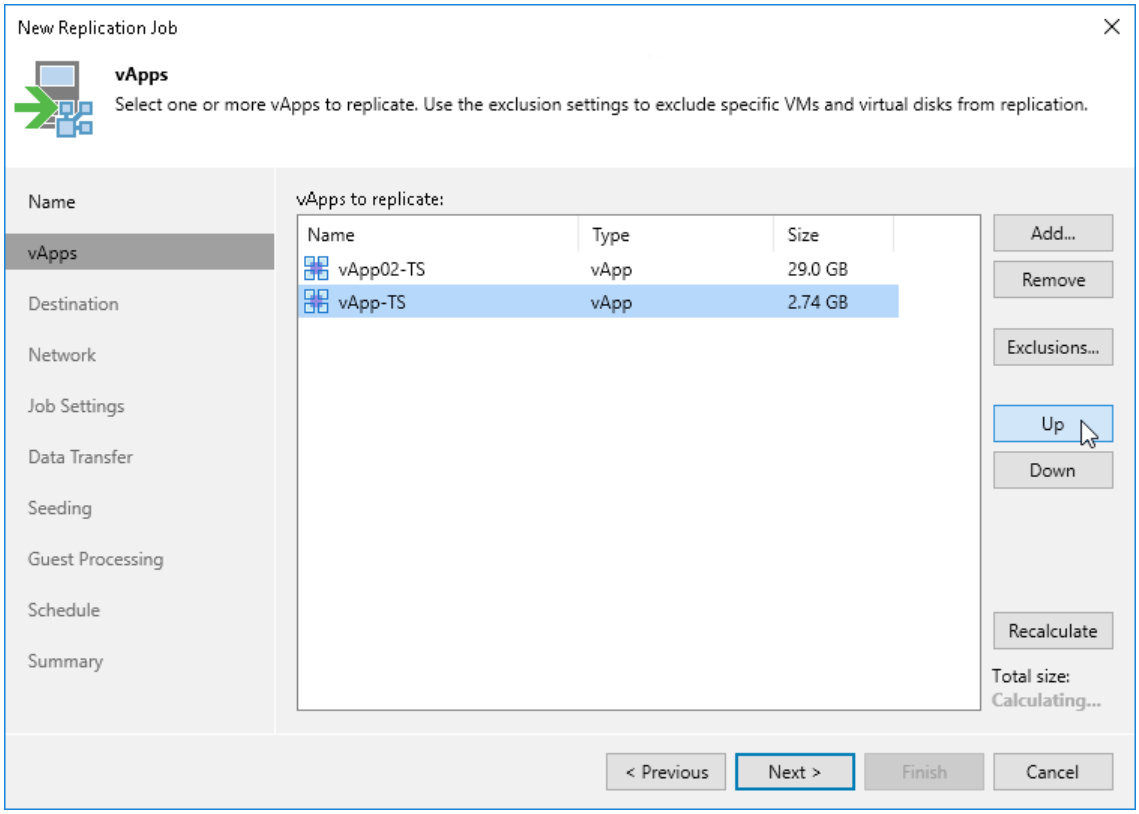

# Step 6. Select Replica Destination

At the **Destination** step of the wizard, select an organization VDC where you want to store replicas and the storage policy that you want to apply to the replicas. You can assign a storage policy to all objects of your VMware Cloud Director infrastructure and also specify a storage policy for a particular vApp or VM.

1. Next to the **Organization VDC** field, click **Choose** and select an organization VDC where replicas must be registered.
2. Next to the **Storage policy VDC** field, click **Choose** and select the storage policy.
   Veeam Backup & Replication will apply this storage policy to the replicas.

   If you want to assign a specific storage policy to individual replicas:

   a. Click the **Select storage policy for individual replicas** link.

   b. In the **Storage policy** window, click **Add** on the right.

   c. In the **Select vApp(s)** window, select necessary vApps.

   d. Select the added vApp in the **Storage policy** list and click **Policy** at the bottom of the window.

   e. From the list of available storage policies, select the necessary storage policy.

      You can select a storage policy for a single VM added to a VM container.

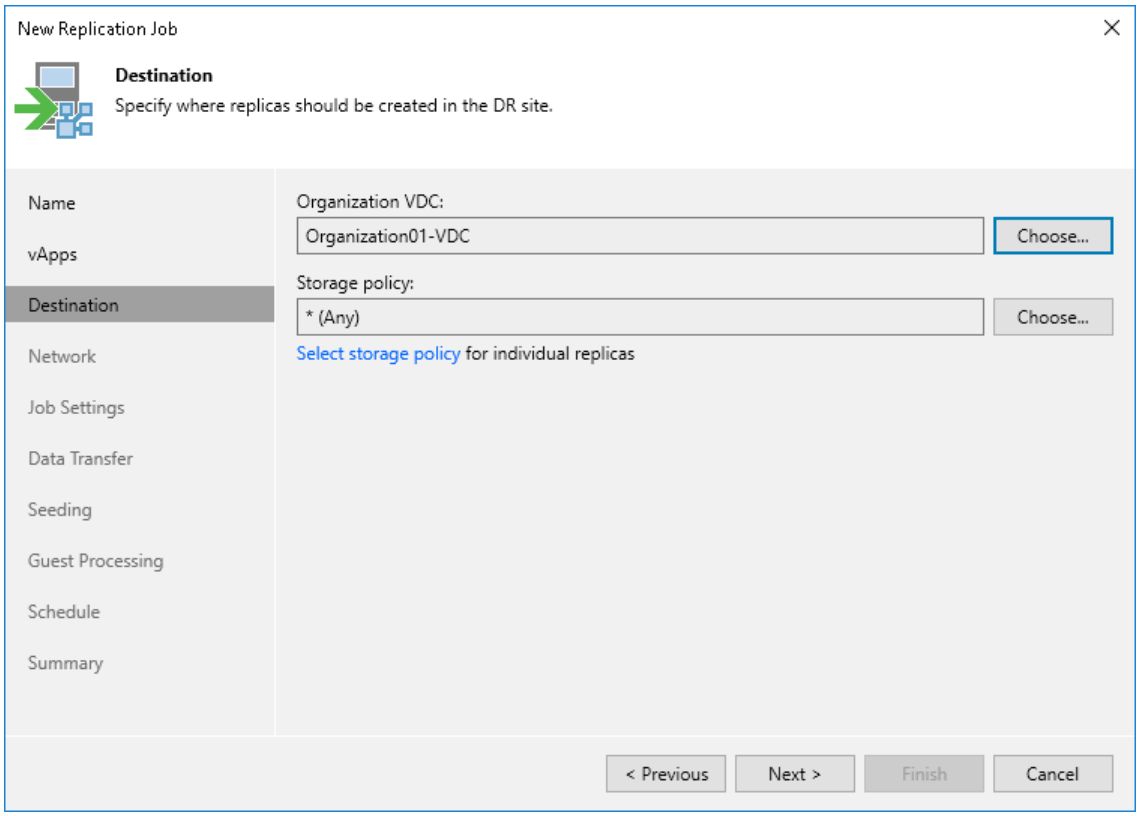

# Step 7. Configure Network Mapping

The **Network** step of the wizard is available if you have selected the **Network remapping** option at the **Name** step of the wizard.

At the **Network** step of the wizard, configure a network mapping table. For more information on network mapping and how replica network settings are configured, see Network Mapping.

To add a row to the network mapping table:

1. Click **Add**.
2. In the **Network Mapping** window, click **Browse** next to the **Source network** field.
3. In the **Select Network** window, select the production network to which the source workloads are connected and click **OK**.
4. In the **Network Mapping** window, click **Browse** next to the **Target network** field.
5. In the **Select Network** window, select a network in the DR site to which replicas will be connected and click **OK**.
6. In the **Network Mapping** window, click **OK**.

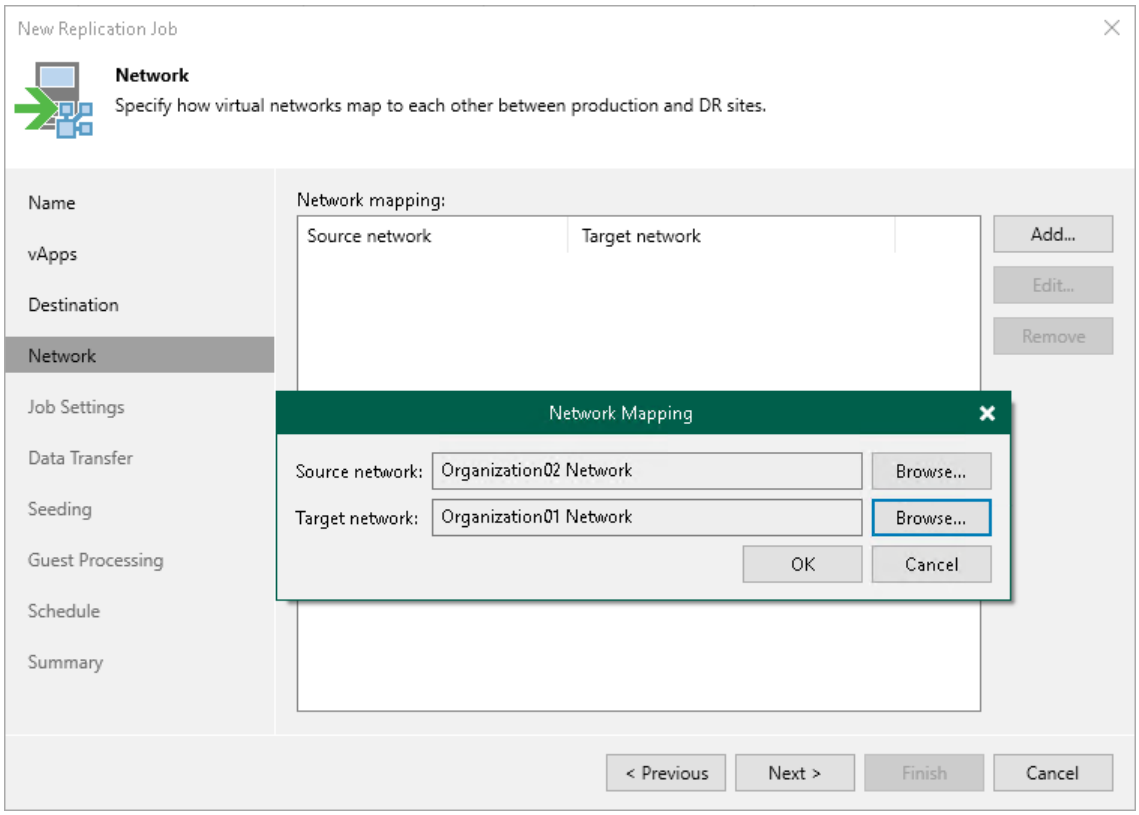

# Step 8. Specify Replication Job Settings

At the **Job Settings** step of the wizard, define VMware Cloud Director replication job settings.

1. From the **Repository for replica metadata** list, select a backup repository. Veeam Backup & Replication uses this backup repository to keep metadata for replicas. Metadata is required when you perform failback with the Quick Rollback option. For more information on metadata, see Changed Block Tracking.
2. In the **Replica name suffix** field, enter a suffix for the name of replicas. Veeam Backup & Replication will add this suffix to the name of the target vApps and will register replicas with this suffix.
3. In the **Restore points to keep** field, specify the number of restore points the replication job must maintain. If this number is exceeded, Veeam Backup & Replication removes the earliest restore point.

   When you specify the retention policy settings for the replication job, consider available space on the target datastore. Due to VMware restrictions on the number of VM snapshots, the maximum number of restore points for replicas is limited to 28. An excessive amount of restore points may overfill the target datastore.

**IMPORTANT**

Consider the following:

- You cannot store replica metadata on deduplicating storage appliances. During replication jobs, Veeam Backup & Replication frequently reads and writes small portions of metadata from/to the backup repository. Frequent access to metadata causes low performance of deduplicating storage appliances, which may result in low performance of replication jobs.
- You cannot store replica metadata in a scale-out backup repository, object storage repository or NFS backup repository.
- At least one successful run of a VMware Cloud Director replication job is required to apply a retention policy to a VMware Cloud Director replica and to delete previous restore points.

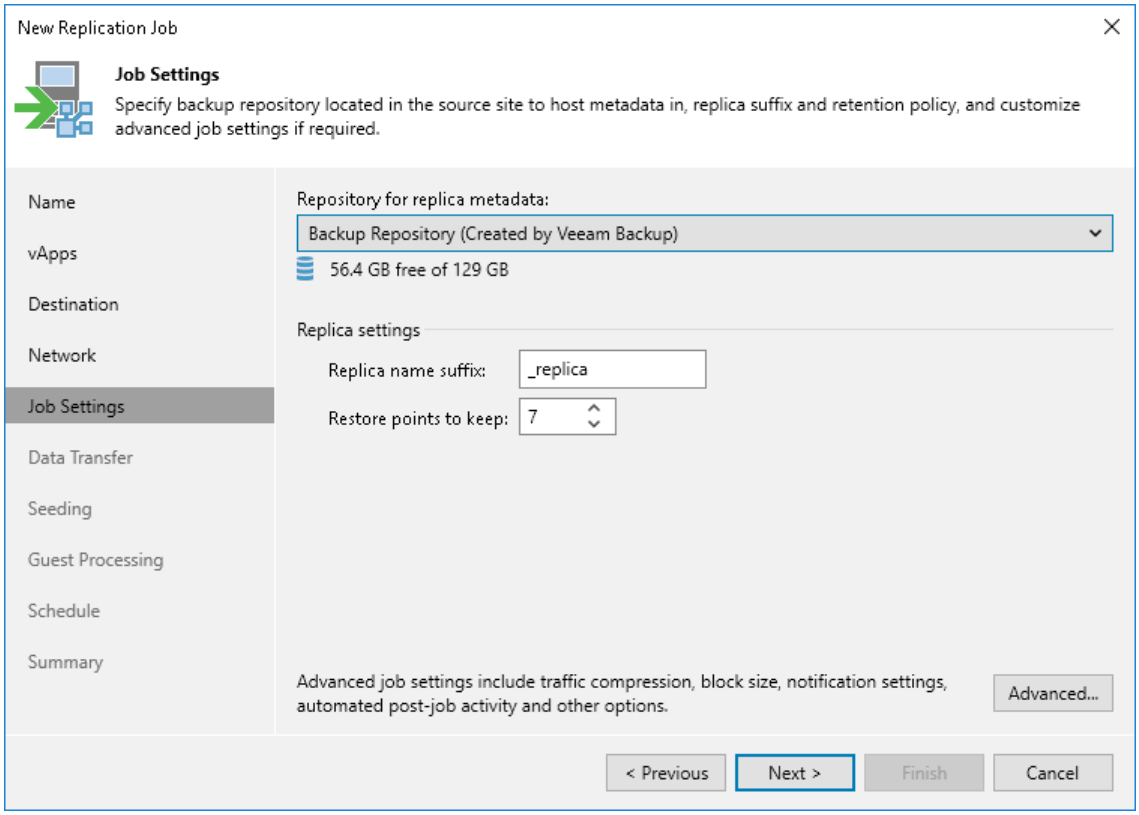

# Step 9. Specify Advanced Replica Settings

At the **Job settings** step of the wizard, you can specify the following settings for the VMware Cloud Director replication job:

- Traffic Settings
- Notification Settings
- vSphere Settings
- Integration Settings
- Script Settings

> **TIP**
>
> After you specify necessary settings for the replication job, you can save them as default settings. To do this, click Save as Default at the bottom left corner of the Advanced Settings window. When you create a new replication job, Veeam Backup & Replication will automatically apply the default settings to the new job.

## Traffic Settings

To specify traffic settings for a VMware Cloud Director replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. Verify that the **Traffic** tab is selected.
3. By default, Veeam Backup & Replication checks the NTFS MFT file on VMs with Microsoft Windows OS to identify data blocks of the `hiberfil.sys` file (file used for the hibernate mode) and `pagefile.sys` file (swap file), and excludes these data blocks from processing. The swap file is dynamic in nature and changes intensively between replication job sessions, even if the VM itself does not change much. Processing of service files reduces the job performance and increases the size of incremental data.

   If you want to include data blocks of the `hiberfil.sys` file and `pagefile.sys` file to the replica, clear the **Exclude swap file blocks** check box. For more information, see Swap Files.
4. By default, Veeam Backup & Replication does not copy deleted file blocks ("dirty" blocks on the VM guest OS) to the target location. This option lets you reduce the size of the replica and increase the job performance. If you want to copy dirty blocks, clear the **Exclude deleted file blocks** check box. For more information, see Deleted File Blocks.
5. From the **Compression level** list, select a compression level for the created replica: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*. Compression is applicable only if data is transferred between two backup proxies. If one backup proxy acts as the source and target backup proxy, data is not compressed at all.

6. In the **Storage optimization** section, select block size that will be used to process VMs. For more information on the data blocks sizes and how they affect performance, see Storage Optimization.

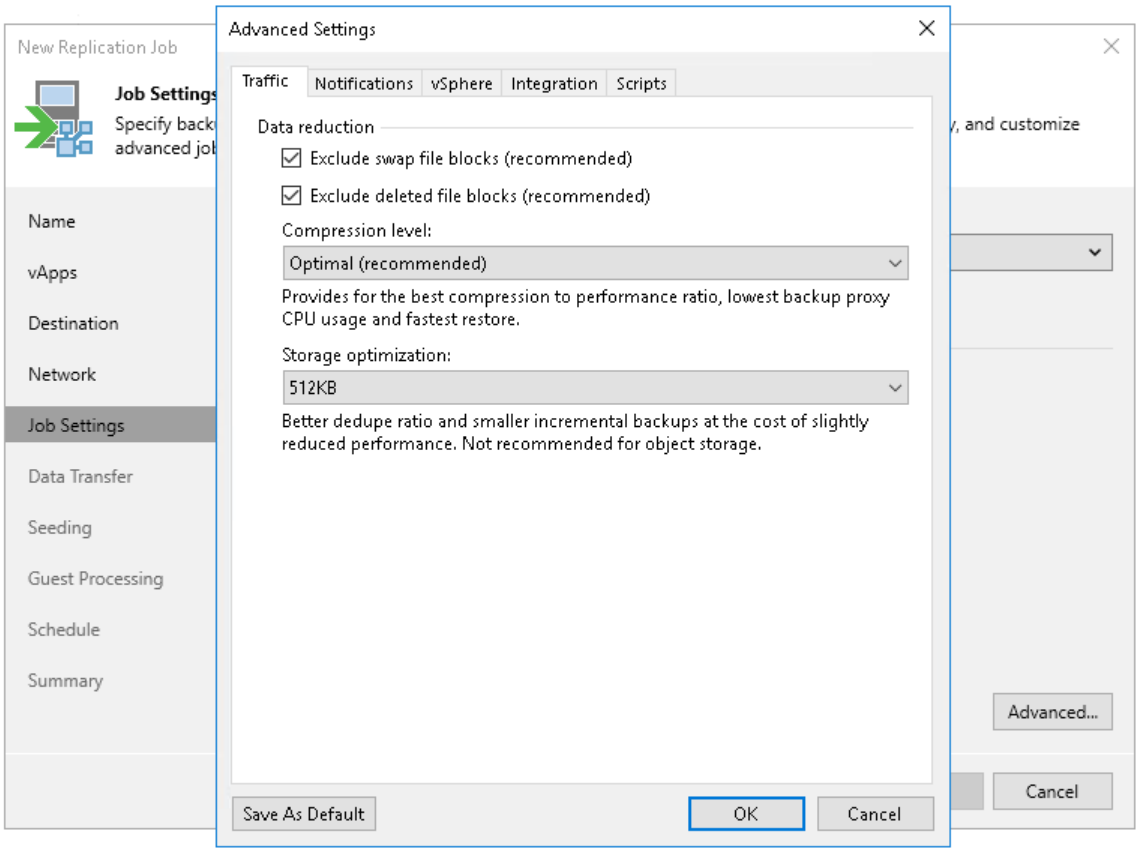


## Notification Settings

To specify notification settings for a VMware Cloud Director replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully.

   SNMP traps will be sent if you configure global SNMP settings in Veeam Backup & Replication and configure software on recipient's machine to receive SNMP traps. For more information, see Specifying SNMP Settings.
4. Select the **Send email notifications to the following recipients** check box if you want to receive notifications by email if a job completes with a failure, warning or success. In the field under the check box, specify the recipient email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication. For more information, see Configuring Global Email Notification Settings.
5. You can choose to use global notification settings or specify custom notification settings.
   - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see Configuring Global Email Notification Settings.
   - To configure a custom notification for a job, select **Use custom notification settings specified below**. You can specify the following notification settings:

i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%Time%* (completion time), *%JobName%*, *%JobResult%*, *%ObjectCount%* (number of VMs in the job) and *%Issues%* (number of VMs in the job that have been processed with the *Warning* or *Failed* status).

ii. Select the **Notify on success**, **Notify on warning** or **Notify on error** check boxes to receive email notification if the job completes successfully, fails or completes with a warning.

iii. Select the **Suppress notifications until the last retry** check box to receive a notification about the final job status. If you do not enable this option, Veeam Backup & Replication will send one notification per every job retry.

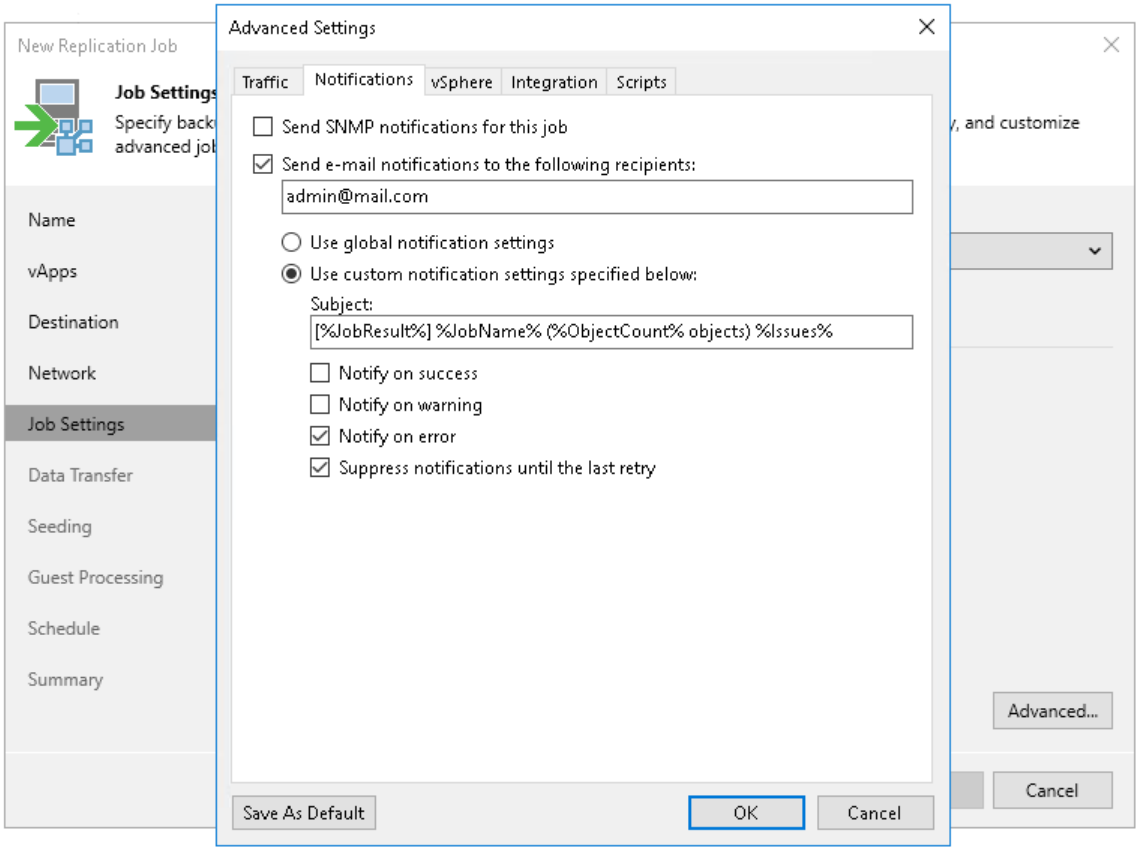


## vSphere Settings

To specify vSphere settings for a VMware Cloud Director replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. Click the **vSphere** tab.
3. Select the **Enable VMware tools quiescence** check box to freeze the file system of processed VMs during replication. Depending on the VM version, Veeam Backup & Replication will use VMware FileSystem Sync Driver (vmsync) or VMware VSS component in VMware Tools for VM snapshot creation. These tools are responsible for quiescing the VM file system and bringing the VM to a consistent state suitable for replication.
4. In the **Changed block tracking** section, configure VMware vSphere Changed Block Tracking CBT:
   a. Make sure that the **Use changed block tracking data** check box is selected if you want to enable CBT.
   b. Make sure that the **Enable CBT for all processed VMs automatically** check box is selected if you want to force using CBT even if CBT is disabled in VM configuration.

c. Make sure that the **Reset CBT on each Active Full backup automatically** check box is selected if you want to reset CBT after the replication job starts for the first time.

CBT reset helps avoid issues, for example, when CBT returns incorrect changed data.

For more information on CBT, see Changed Block Tracking.

**IMPORTANT**

You can use CBT for VMs with virtual hardware version 7 or later. These VMs must not have existing snapshots.

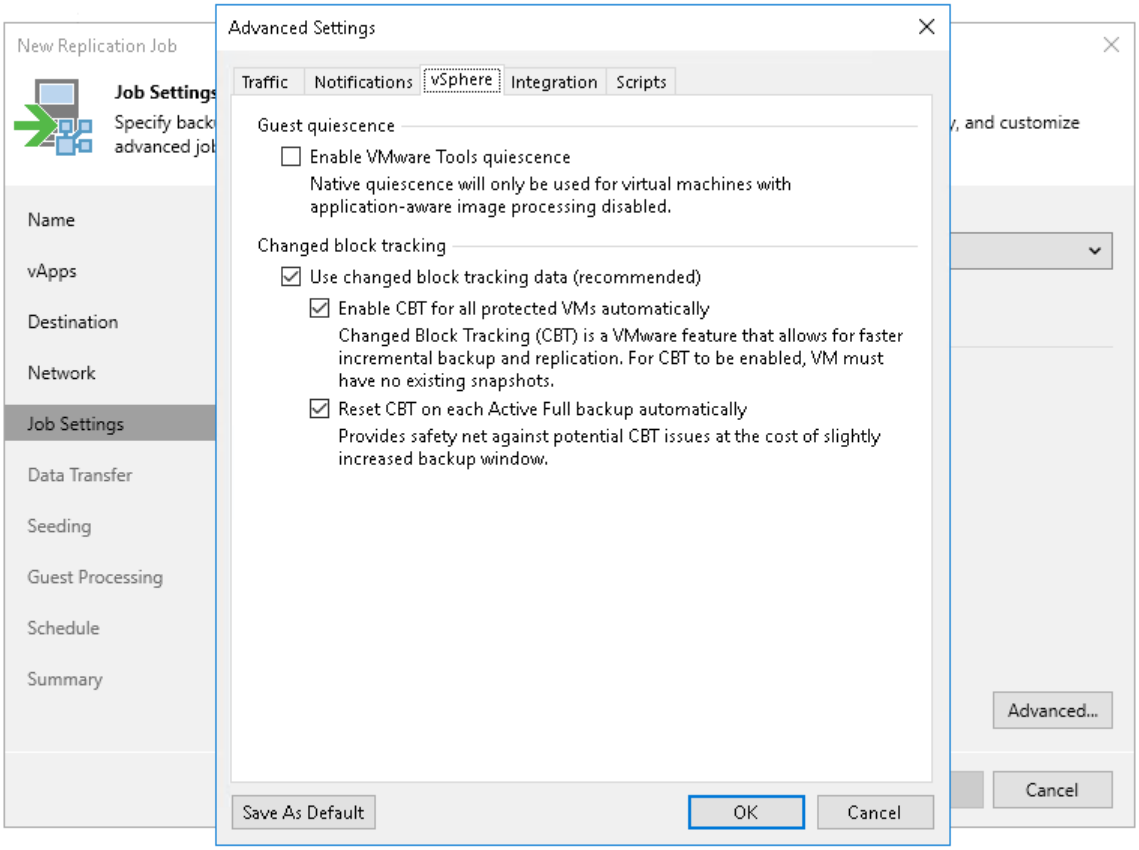


## Integration Settings

On the **Integration** tab, define whether you want to use the Backup from Storage Snapshots technology to create a VMware Cloud Director replica. Backup from Storage Snapshots lets you leverage storage snapshots for VM data processing. The technology improves RPOs and reduces the impact of replication activities on the production environment. For more information, see Backup from Storage Snapshots section.

To specify storage integration settings for the replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. Click the **Integration** tab.
3. By default, the **Enable backup from storage snapshots** option is enabled. If you do not want to use Backup from Storage Snapshots, clear this check box.

4. If you want to replicate vApps with multiple VMs whose disks are located on the same volume or LUN, select the **Limit processed VM count per storage snapshot to** check box and specify the number of VMs for which one storage snapshot must be created. In a regular job processing course, Veeam Backup & Replication creates a VMware snapshot for every VM added to the job and then triggers one storage snapshot for all VMs. In some situations, creating VMware snapshots for all VMs may require a lot of time. If you limit the number of VMs per storage snapshot, Veeam Backup & Replication will divide VMs into several groups, trigger a separate storage snapshot for every VM group and read VM data from these snapshots. As a result, the job performance will increase.

   For example, you add to the job vApps with 30 VMs whose disks are located on the same volume and set the **Limit processed VM count per storage snapshot to** option to 10. Veeam Backup & Replication will divide all VMs into 3 groups and create 3 storage snapshots from which it will read VM data.

5. If the backup infrastructure is configured incorrectly, for example, the backup proxy does not meet the necessary requirements, Backup from Storage Snapshots will fail and VMs residing on the storage systems will not be processed by the job at all. To fail over to the regular VM processing mode and process such VMs in any case, select the **Failover to standard backup** check box.

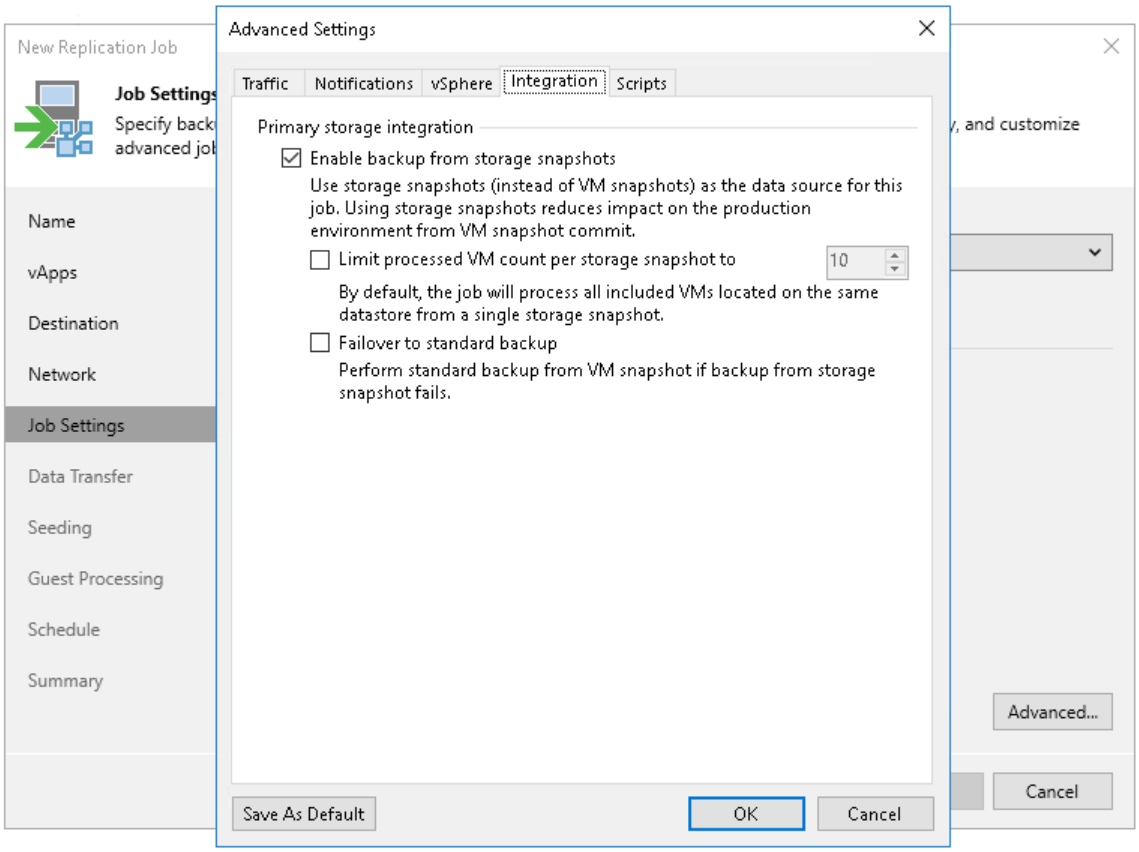


## Script Settings

To specify script settings for the replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. Click the **Scripts** tab.

3. If you want to execute custom scripts before and after the replication job, select the **Run the following script before the job** and **Run the following script after the job** check boxes and click **Browse** to choose executable files from a local folder on the backup server. The scripts are executed on the backup server.

   You can select to execute pre- and post-replication actions after a number of job sessions or on specific week days.

   - If you select the **Run scripts every... backup session** option, specify the number of the replication job sessions after which scripts must be executed.
   - If you select the **Run scripts on selected days only** option, click **Days** and specify week days on which scripts must be executed.

**NOTE**

Consider the following:

- Custom scripts that you define in the advanced job settings relate to the replication job itself, not the VM quiescence process. To add pre-freeze and post-thaw scripts for VM image quiescence, use the **Guest Processing** step of the wizard.
- If you select the **Run scripts on the selected days only** option, Veeam Backup & Replication executes scripts only once on each selected day – when the job runs for the first time. During subsequent job runs, scripts are not executed.
- To run the script, Veeam Backup & Replication uses the Service Account under which the Veeam Backup Service is running.

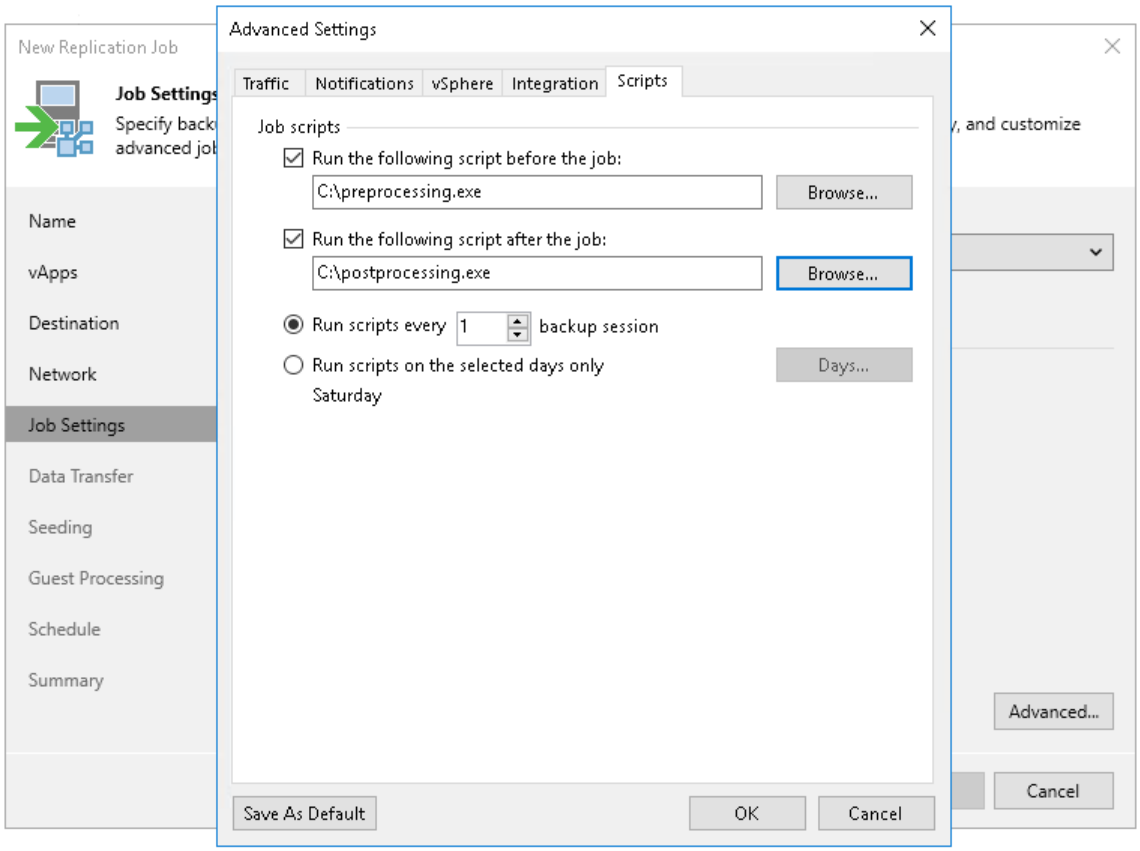

# Step 10. Specify Data Transfer Settings

At the **Data Transfer** step of the wizard, select backup infrastructure components that Veeam Backup & Replication will use for the replication process and choose a path for vApp data transfer.

1. If you plan to replicate vApp data within one site, the same proxy can act as the source and target proxy. For off-site replication, you must deploy at least one proxy in each site to establish a stable connection for vApp data transfer across sites.

   Click **Choose** next to the **Source proxy** and **Target proxy** fields to select backup proxies for the job. In the **Backup Proxy** window, you can choose automatic proxy selection or assign backup proxies explicitly.

   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that have access to the source and target datastores and automatically assign optimal proxy resources for processing vApp data.

     Veeam Backup & Replication assigns resources to vApp included in the replication job one by one. Before processing a new vApp from the list, Veeam Backup & Replication checks available backup proxies. If more than one proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use and the current workload on the backup proxies to select the most appropriate proxy for vApp processing.

   - If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that the job can use. It is recommended that you select at least two backup proxies to ensure that the job will be performed if one of backup proxies fails or loses its connectivity to the source datastore.

2. Select a path for vApp data transfer:

   - To transport vApp data directly using backup proxies to the target datastore, select **Direct**.

   - To transport vApp data using WAN accelerators, select **Through built-in WAN accelerators**. From the **Source WAN accelerator** list, select the WAN accelerator configured in the source site. From the **Target WAN accelerator** list, select the WAN accelerator configured in the target site.

     You should not assign one source WAN accelerator to several replication jobs that you plan to run simultaneously. The source WAN accelerator requires a lot of CPU and RAM resources, and does not process multiple replication tasks in parallel. Alternatively, you can create one replication job for all vApps you plan to process over one source WAN accelerator.

The target WAN accelerator, however, can be assigned to several replication jobs. For more information, see Adding WAN Accelerators.

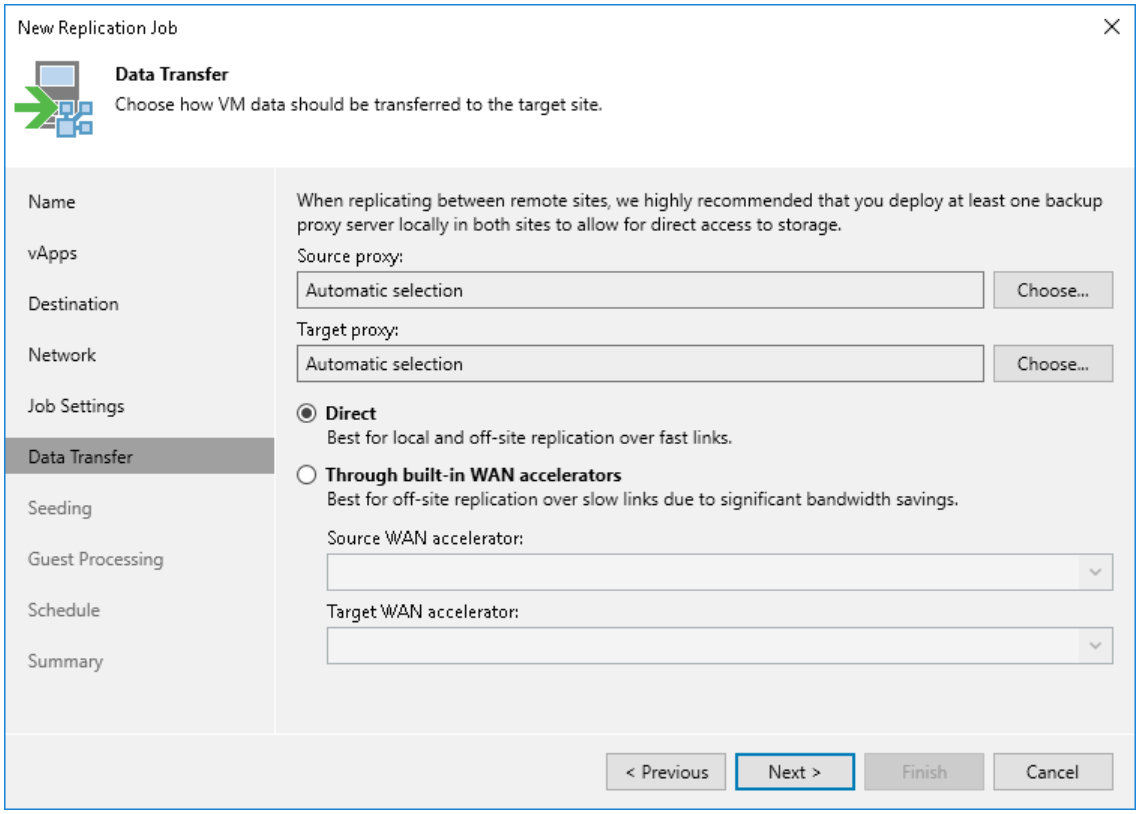

# Step 11. Configure Seeding and Mapping Settings

The **Seeding** step is available if you have selected the **Replica seeding** option at the **Name** step wizard. You can use this step to configure replica seeding and mapping for a VMware Cloud Director replication job. For more information on when to use seeding and mapping, see vApp Seeding and Mapping.

If you use replica seeding or mapping, make sure that you select the following backup infrastructure components for the job properly: source-side backup repository to keep metadata and proxies. It is recommended that you explicitly assign proxies in the production site and disaster recovery (DR) site. For more information, see Specify Data Transfer Settings.

> **IMPORTANT**
>
> If the **Replica seeding** check box is enabled in a policy, all vApps in the policy must be covered with seeding or mapping. If a vApp neither has a seed, nor has mapping to an existing vApp, it will be skipped from processing.

# Configuring Replica Seeding

To configure replica seeding:

1. Make sure that you have backups of replicated vApps in a backup repository in the DR site. If you do not have the backups, create them as described in section Creating vApp Replica Seeds.

   > **IMPORTANT**
   >
   > Consider the following:
   >
   > - Backups must be created by Veeam Backup & Replication.
   > - When you start replication, Veeam Backup & Replication will attempt to restore all VMs added to replication from the vApp seed that you have specified. If a VM is not found in the vApp seed, the VM will be skipped from replication.
   > - Backups must not reside in a scale-out backup repository.

2. In the **Initial seeding** section, select the **Get seed from the following backup repository** check box.
3. From the list of available backup repositories, select the repository where your replica seeds are stored.

> **NOTE**
>
> If a vApp has a seed and is mapped to an existing replica, replication will be performed using replica mapping because mapping has a higher priority.

# Configuring Replica Mapping

To configure replica mapping:

1. Select the **Map replicas to existing vApps** check box.
2. If you want Veeam Backup & Replication to scan the DR site to detect existing copies of vApps that you plan to replicate, click **Detect**.

   If any matches are found, Veeam Backup & Replication will populate the mapping table. If Veeam Backup & Replication does not find a match, you can map a vApp to its copy manually. Note that the mapping list does not display vApps added to the list of exclusions for replication.

3. If you want to map a VM manually, select a source VM from the list, click **Edit** and select the copy of this VM on the target host in the DR site.

> **TIP**
>
> If there is no existing replica in the DR site, you can restore a vApp from the backup and map it to the source vApp.

To remove a mapping association, select a VM in the list and click **Remove**.

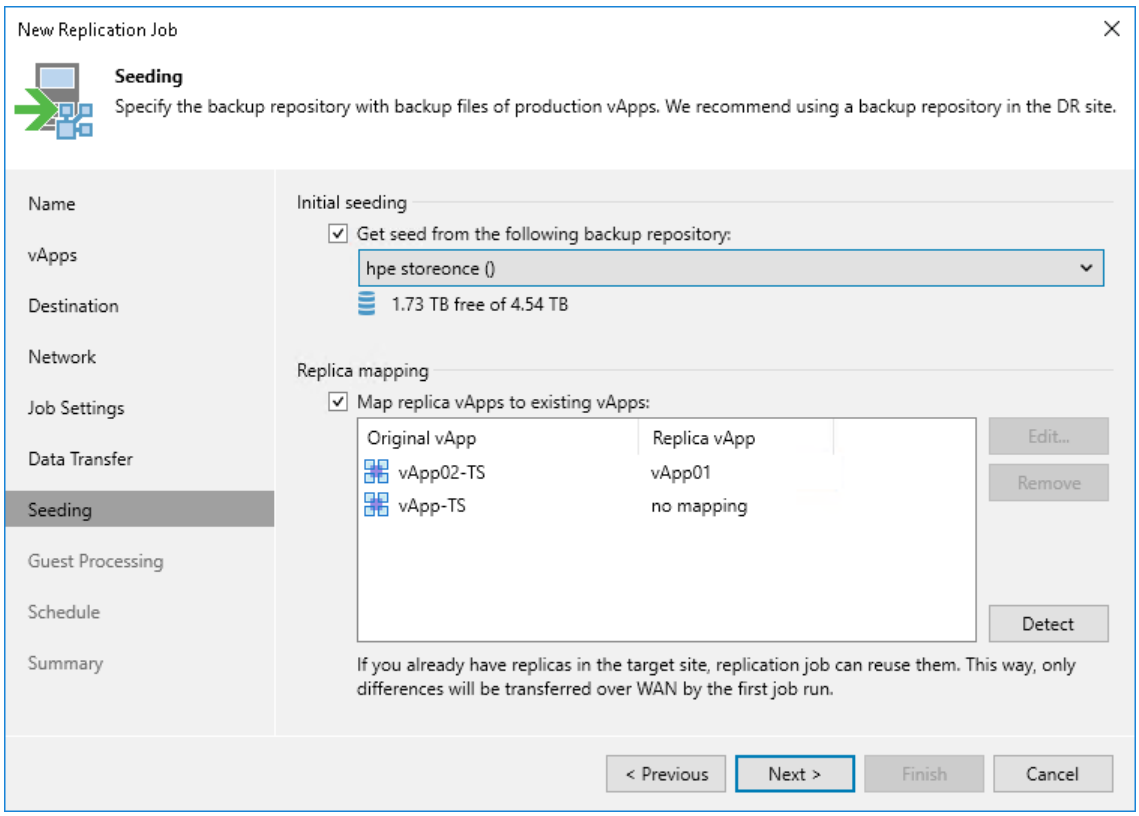

# Step 12. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, enable and configure guest OS processing.

Guest OS processing involves application-aware processing that allows creation of transactionally consistent replicas and guest file system indexing (however, indexing is not available for replicas). In its turn, application-aware processing includes log truncation, execution of custom scripts and guest OS file exclusions. For more information on guest processing, see the Guest Processing section.

To be able to use guest processing, you must also configure user accounts to access guest OSes and guest interaction proxies.

To enable guest OS processing and start configuring it (accounts and guest interaction proxies):

1. Select **Enable application-aware processing**.

   When you select this option, Veeam Backup & Replication enables application-aware processing with the default settings for all VMs. You can further disable application-aware processing for individual VMs and reconfigure the default settings.

2. If you have added Microsoft Windows VMs to be processed, specify which guest interaction proxy Veeam Backup & Replication can use to perform different guest processing tasks:

   - If you want Veeam Backup & Replication to select the guest interaction proxy automatically, leave **Automatic selection** on the **Guest interaction proxy** field.
   - If you want to explicitly specify which servers will perform the guest interaction proxy role, click **Choose**. In the **Guest Interaction Proxy** window, click **Prefer the following guest interaction proxy server**, and select the necessary proxies.

   For more information on the guest interaction proxy, requirements and limitations for it, see Guest Interaction Proxies.

3. From the **Guest OS credentials** list, select a user account that will be used to connect to guest OSes and that has enough permissions. For more information on the permissions and requirements for the user account, see Permissions for Guest Processing.

   [For Microsoft Windows VMs] Veeam Backup & Replication will also use this account to deploy the non-persistent runtime components or use (if necessary, deploy) persistent agent. For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

   [For Linux VMs] If you installed persistent agent components for VMs running Linux or Unix operating systems, select *Use management agent credentials* from the list. For more information, see Persistent Agent Components.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. For more information on adding credentials, see the Credentials Manager section.

   > **NOTE**
   >
   > If you plan to use Kerberos authentication, check limitations and requirements listed in section Guest Processing.

4. To specify credentials for individual workloads, click **Credentials**. Then select the necessary workload and set user credentials for it.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose a VM whose settings you want to customize. Then select the VM in the list and define the necessary settings.

5. To check whether Veeam Backup & Replication can connect to VMs using the specified guest OS credentials and can deploy the non-persistent runtime components or connect to persistent agent components on the guest OSes, click **Test Now**.

> **TIP**
>
> You can specify guest processing settings for vApps and for specific VMs.

After you have enabled application-aware processing for all VMs and configured other settings required for guest processing, you can disable application-aware processing for individual VMs and change the default settings. For more information, see the following sections:

- Application-aware processing
- Microsoft SQL Server transaction log settings
- Oracle archived log settings
- PostgreSQL settings
- File exclusion settings
- Pre-freeze and post-thaw scripts

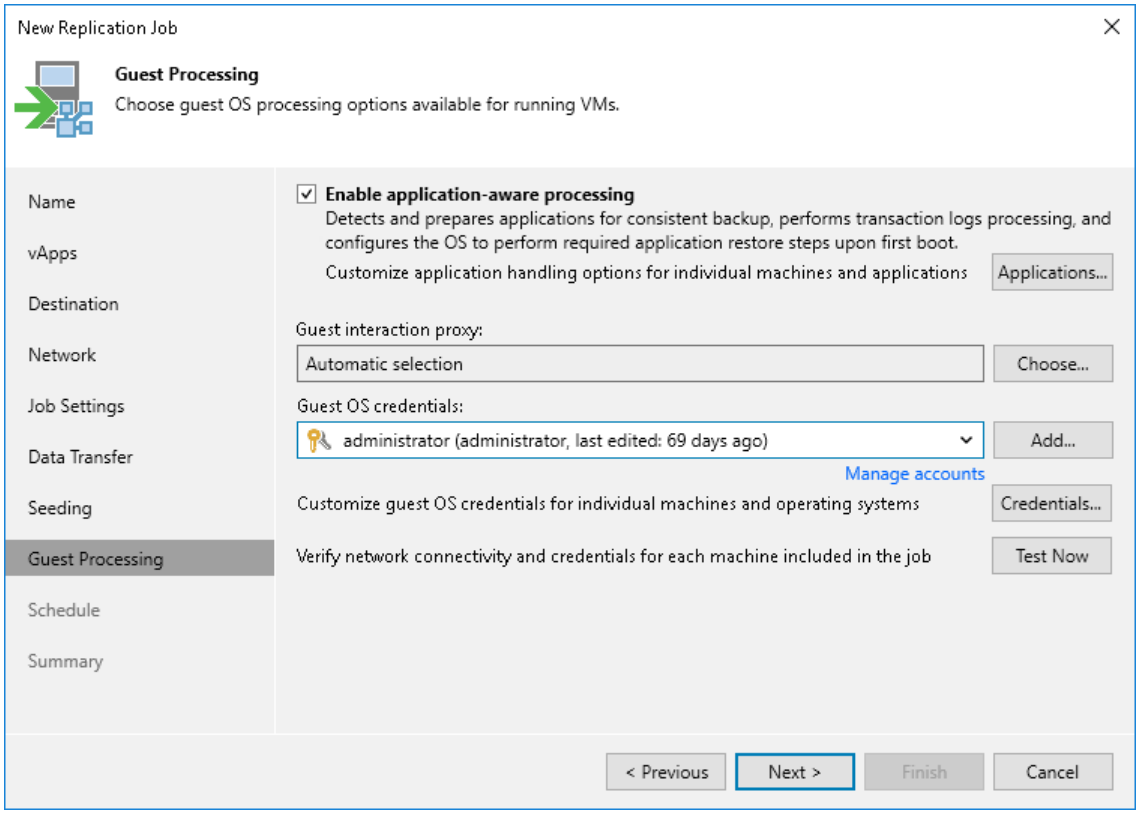


## Application-Aware Processing and Transaction Logs

Application-aware processing helps create transactionally consistent replicas. The transactionally consistent replicas guarantee proper recovery of applications without data loss. For more information on application-aware processing, see Application-Aware Processing.

To configure general application-aware processing settings and specify whether Veeam Backup & Replication processes transaction logs or creates copy-only replicas:

1. At the **Guest Processing** step of the wizard, check that you have selected the **Enable application-aware processing** check box and configured guest OS credentials.
2. At the **Guest Processing** step of the wizard, click **Applications**.
3. In the **Application-Aware Processing Options** window, select workloads for which you want to configure application-aware processing, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. On the **General** tab, in the **Applications** section, specify the behavior scenario for application-aware processing:
   - Select **Require successful processing** if you want Veeam Backup & Replication to stop the replication process if any error occurs during application-aware processing.
   - Select **Try application processing, but ignore failures** if you want to continue the replication process even if an error occurs during application-aware processing. This option guarantees that replication will continue working. However, the resulting replica will be crash consistent, not transactionally consistent.
   - Select **Disable application processing** if you want to disable application-aware processing for the workload.
5. [For Microsoft Exchange and Microsoft SQL Server] In the **VSS Settings** section, specify if Veeam Backup & Replication must process transaction logs or create copy-only replicas:
   a. Select **Process transaction logs with this job** if you want Veeam Backup & Replication to process transaction logs.

      [For Microsoft Exchange VMs] With this option selected, the non-persistent runtime components or persistent components running on the VM guest OS will wait for replication to complete successfully and then trigger truncation of transaction logs. If the replication job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

      [For Microsoft SQL Server] You will need to configure how to process transaction logs.

      > **TIP**
      >
      > To configure log processing for Oracle and PostgreSQL databases, switch to the Oracle and PostrgeSQL tabs.

   b. Select **Perform copy only** if you use another tool to perform guest level processing, and this tool maintains consistency of the database state. Veeam Backup & Replication will create a copy-only replica for the selected VMs. The copy only replica preserves the chain of full and differential files and transaction logs on the VM. For more information, see Microsoft Docs.
6. [For Microsoft Windows VMs] In the **Persistent guest agent** section, select the **Use persistent guest agent** check box to use for application-aware processing persistent guest agents on each protected VM.

   By default, Veeam Backup & Replication uses non-persistent runtime components. Veeam Backup & Replication deploys runtime components on each protected VM when the replication job starts, and removes the runtime components as soon as the replication job finishes.

   For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

[For Linux VMs] To use persistent guest agents, you must install Management Agent on protected VMs. For more information, see Persistent Agent Components.

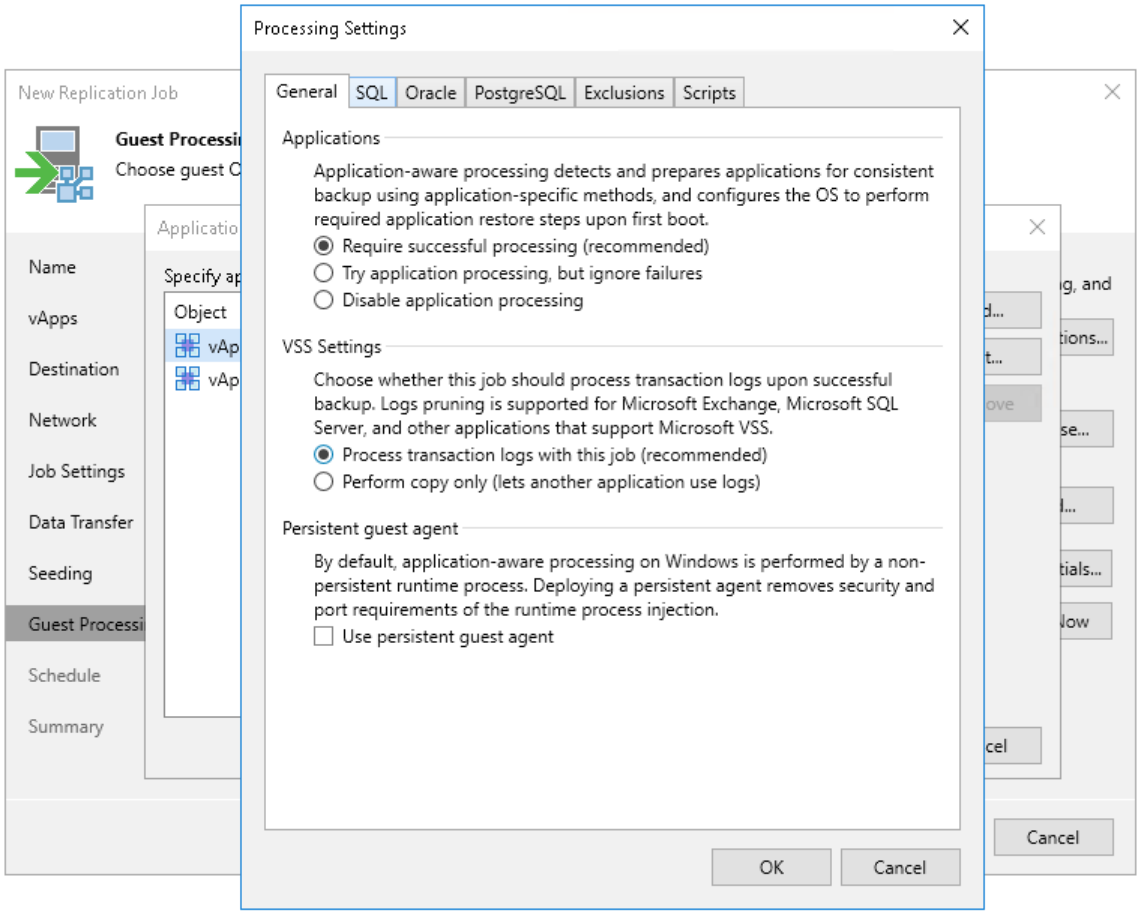


## Microsoft SQL Server Transaction Log Settings

The **SQL** tab is available for VMs that run Microsoft SQL Server and if you have selected **Process transaction logs with this job** when configuring application-aware processing.

To create transactionally consistent backups of an Microsoft SQL Servers, you must check that application-aware processing is enabled and then specify settings of transaction log processing.

# Enabling Application-Aware Processing

Before configuring transaction log processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Microsoft SQL Sever and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

# Specifying Transaction Log Settings

In the **Processing Settings** window, switch to the **SQL** tab and specify how transaction logs must be processed:

- If you want Veeam Backup & Replication to trigger truncation of transaction logs only after the job completes successfully, select **Truncate logs**.

  In this case, the non-persistent runtime components or persistent components will wait for replication to complete and then trigger truncation of transaction logs. If the replication job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

  - If you do not want Veeam Backup & Replication to truncate logs at all, select **Do not truncate logs**.

  This option is recommended if you are using another backup tool to perform VM guest-level backup or replication, and this tool maintains consistency of the database state. In such scenario, Veeam Backup & Replication will not trigger transaction log truncation. After you fail over to the necessary restore point of the VM replica, you will be able to apply transaction logs to get the database system to the necessary point in time between replication job sessions.

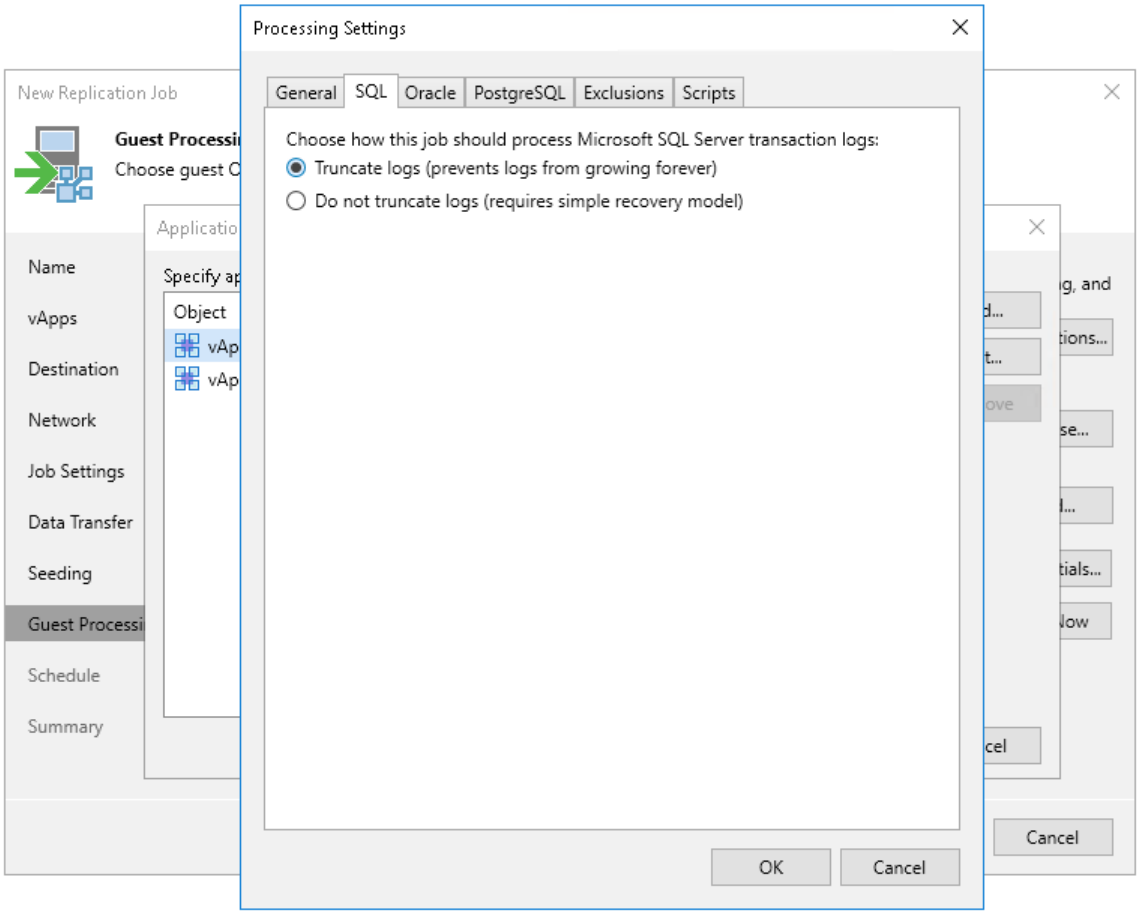


# Oracle Archived Log Settings

The **Oracle** tab applies to Linux VMs that run Oracle.

To create transactionally consistent backups of an Oracle server, you must check that application-aware processing is enabled and then specify settings of archive log processing.

# Enabling Application-Aware Processing

Before configuring archive log processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Oracle server and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

# Specifying Archive Log Settings

To configure how Veeam Backup & Replication must process archive logs of an Oracle server:

1. In the **Processing Settings** window, switch to the **Oracle** tab.
2. From the **Specify Oracle account with SYSDBA privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the Oracle databases. The account that you plan to use must have privileges described in section Permissions.

   You can select **Use guest credentials** from the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle databases.
3. In the **Archived logs** section, specify how to process archived logs:
   - If you want to preserve archived logs on the VM guest OS, select **Do not delete archived logs**. When the replication job completes, the non-persistent runtime components or persistent components will not truncate transaction logs.

     It is recommended that you select this option for databases where the ARCHIVELOG mode is turned off. If the ARCHIVELOG mode is turned on, archived logs may grow large and consume all disk space.
   - If you want to delete archived logs older than <N> hours, select **Delete logs older than <N> hours** and specify the number of hours.

- If you want to delete archived logs larger than <N> GB, select **Delete logs over <N> GB** and specify the size. The specified size refers to the log size of each database, not all databases on the selected Oracle server.

The non-persistent runtime components or persistent components running on the VM guest OS will wait for the replication job to complete successfully and then trigger transaction logs truncation using Oracle Call Interface (OCI). If the job does not manage to replicate the Oracle VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

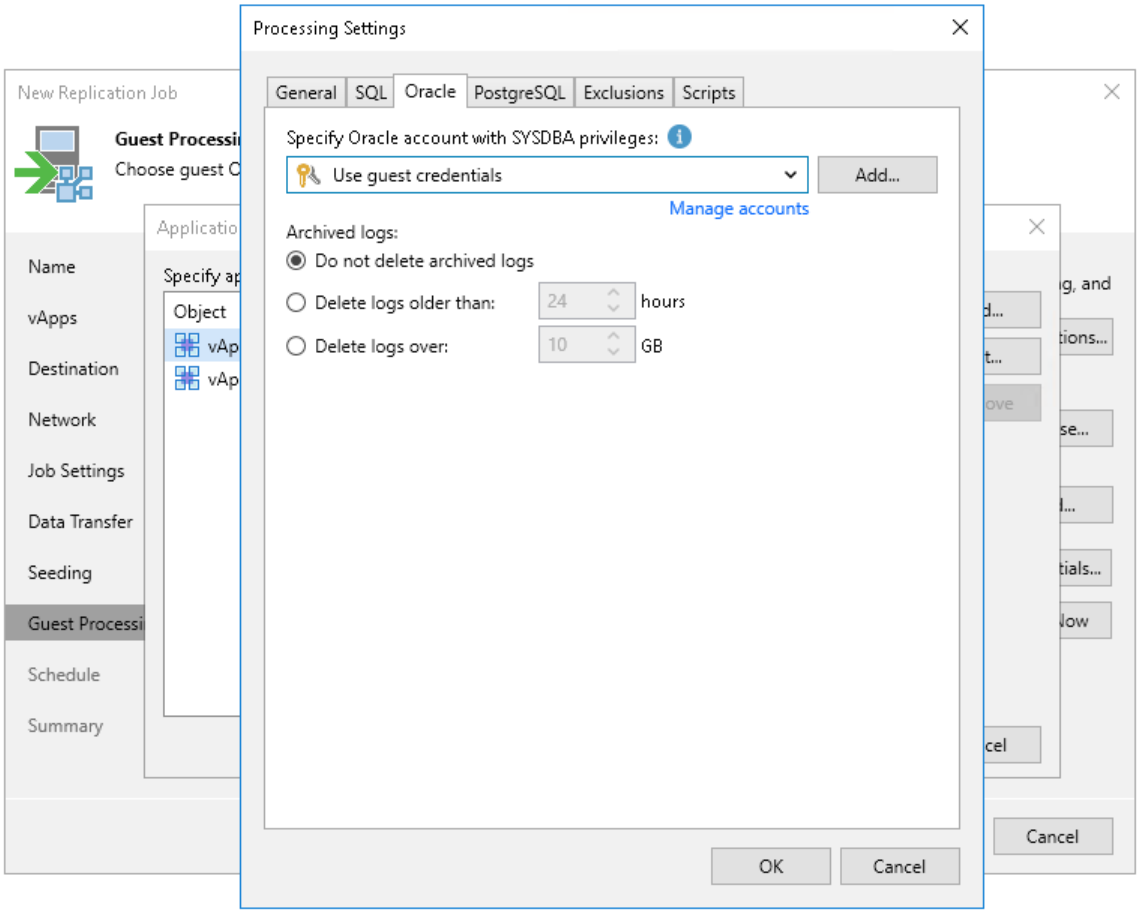


## PostgreSQL Settings

The **PostgreSQL** tab applies to Linux VMs that run PostgreSQL.

To create transactionally consistent backups of a PostgreSQL VM, you must check that application-aware processing is enabled and then specify settings of WAL files processing.

# Enabling Application-Aware Processing

Before configuring WAL files processing, check that application-aware processing is enabled:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.

3. In the displayed list, select the PostgreSQL VM and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.

4. In the **Processing Settings** window, on the **General** tab, check that **Require successful processing** or **Try application processing, but ignore failures** option is selected in the **Applications** area.

**NOTE**

By default, Veeam Backup & Replication recursively scans the `/etc/postgresql`, `/var/lib/postgresql` and `/var/lib/pgsql` directories for the configuration files of PostgreSQL instances. If you keep configuration files in custom directories, the *pgsqlagent agent* will use its own `VeeamPostgreSQLAgent.xml` configuration file that is located in the `/etc/veeam/` directory. The pgsqlagent agent configuration file must be a single line XML.

To explicitly include or exclude specific configuration files from rescan, you can add the following commands to the `VeeamPostgreSQLAgent.xml` file:

- `ExcludeConfigDirs` — use this element to exclude configuration files.
- `AddConfigDirs` — use this element to include configuration files.

For example: `<config AddConfigDirs="/opt/psql/" ExcludeConfigDirs="/var/lib/postgresql/13/main45/,/var/lib/postgresql/13/maindd/" />`

# Specifying WAL Files Settings

To define how Veeam Backup & Replication will process WAL files on this VM, do the following:

1. In the **Processing Settings** window, click the **PostgreSQL** tab.
2. From the **Specify PostgreSQL account with superuser privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the PostgreSQL instance. The account must have privileges described in section Permissions. You can select **Use guest credentials** from the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the PostgreSQL instance.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. Note that if you select the **System user without password file (peer)** option in the **The specified user is** area, you can add a user account without specifying a password.

3. In the **The specified user is** section, specify how the user selected in the **Specify PostgreSQL account with superuser privileges** drop-down list will authenticate against the PostgreSQL instance:
   - Select **Database user with password** if the account is a PostgreSQL account, and you entered the password for this account in the Credentials Manager.

> **NOTE**
>
> If you want Veeam Backup & Replication to use the user name map authentication, select **Database user with password** and leave the password field empty. Consider the following:
>
> - **Guest OS credentials** specified at the **Guest Processing** step of the wizard will be used as the System-Username.
> - PostgreSQL account specified at the step 2 will be used as the PG-Username.
>
> For more information about the user name maps, see PostgreSQL documentation.

- Select **Database user with password file (.pgpass)** if the password for the account is defined in the `.pgpass` configuration file on the PostgreSQL server. For more information about the password file, see PostgreSQL documentation.
- Select **System user without password file (peer)** if you want Veeam Backup & Replication to use the peer authentication method. In this case, Veeam Backup & Replication will apply the OS account as the PostgreSQL account.

> **IMPORTANT**
>
> If you have added a new PostgreSQL account and want to use it with the peer authentication method, make sure that you have added this account as a Linux user with sufficient permissions.

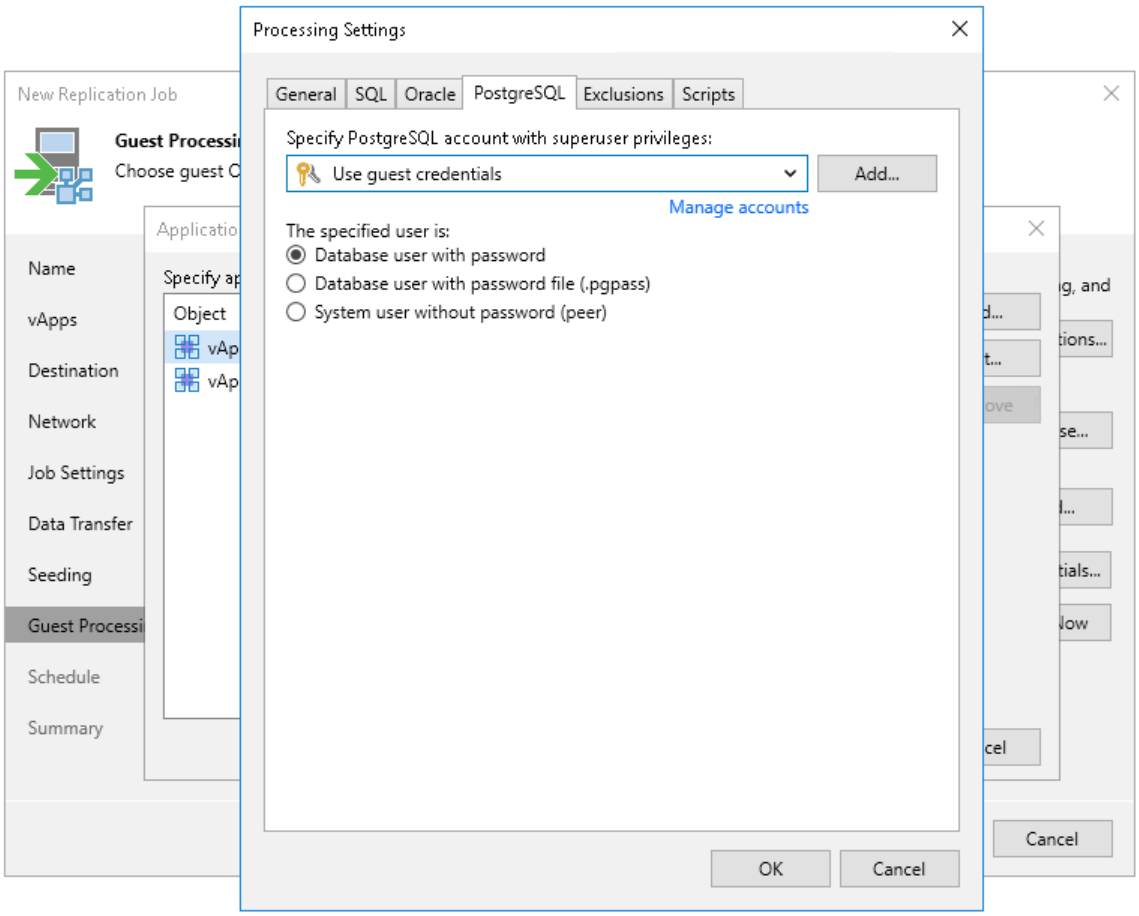


## Guest OS File Exclusion Settings

These settings apply only to Microsoft Windows workloads.

To exclude guest OS files and folders from being replicated:

1. At the **Guest Processing** step of the wizard, check that you have selected the **Enable application-aware processing** check box and configured guest OS credentials.
2. In the **Application-Aware Processing Options** list, select workloads for which you want to exclude files and folders and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
3. In the **Processing Settings** window, switch to the **Exclusions** tab and specify whether you want to exclude or include files and folders:
   - To remove individual files and folders from replicas, select **Exclude the following files and folders** and click **Add**.
   - To include only the specified files and folders in replicas, select **Include only the following files and folders** and click **Add**.
4. In the **Specify Folder** window, specify which files and folders you want to include or exclude. For the methods that you can use to specify the list of exclusions or inclusions, see VM Guest OS Files.

> **NOTE**
>
> When you select files to be included or excluded, consider requirements and limitations that are listed in section Requirements and Limitations for VM Guest OS File Exclusion.

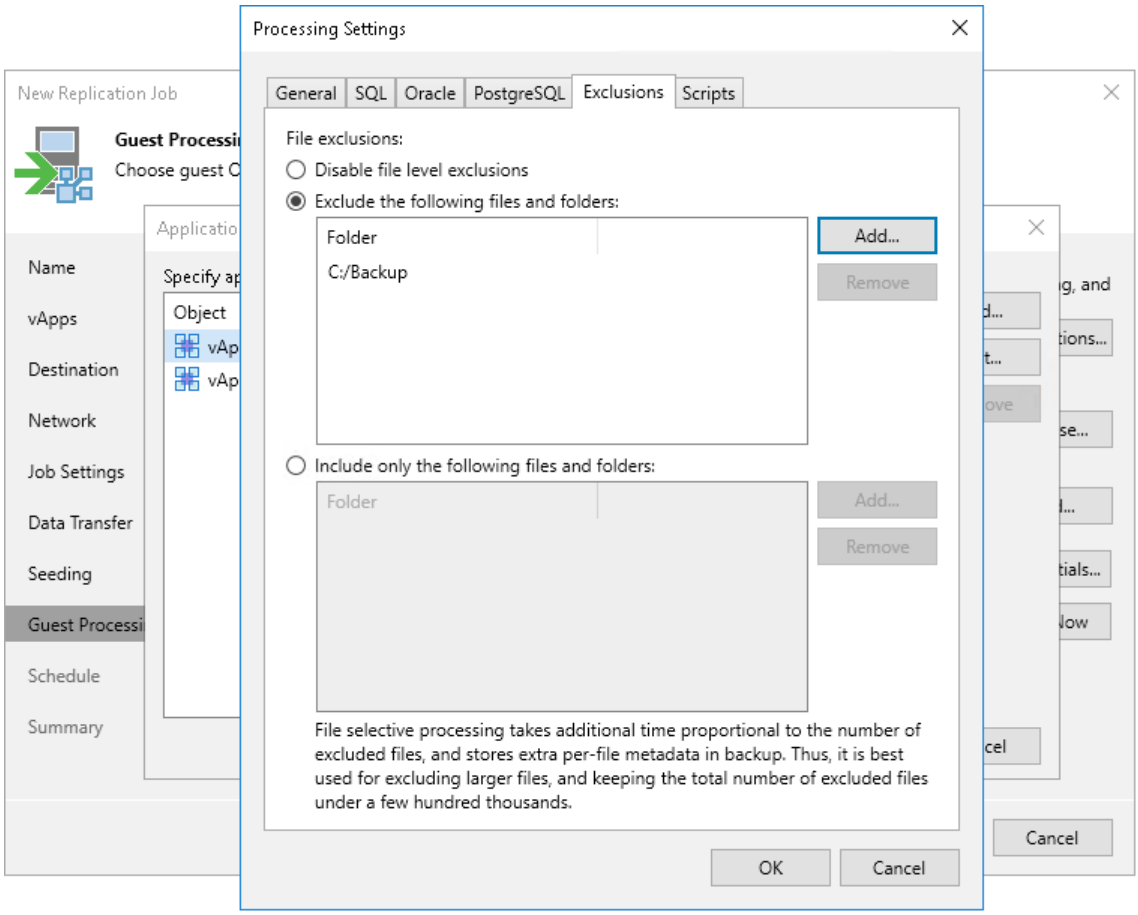

# Pre-Freeze and Post-Thaw Scripts

If you plan to replicate VMs running applications that do not support VSS, you can instruct Veeam Backup & Replication to run custom pre-freeze and post-thaw scripts for these VMs. The pre-freeze script quiesces the VM file system and application data to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script brings the VM and applications to their initial state.

To specify pre-freeze and post-thaw scripts:

1. At the **Guest Processing** step of the wizard, check that you have selected the **Enable application-aware processing** check box and configured guest OS credentials.
2. At the **Guest Processing** step of the wizard, click **Applications**.
3. In the **Application-Aware Processing Options** list, select workloads for which you want to configure scripts, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. Click the **Scripts** tab.
5. In the **Script processing mode** section, select a scenario for script execution:
   - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the replication process if scripts fail.
   - Select **Ignore script execution failures** if you want to continue the replication process even if script errors occur.
   - Select **Disable script execution** if you do not want to run scripts for the VM.
6. In the **Windows scripts** section, specify paths to scripts for Microsoft Windows VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.
7. In the **Linux scripts** section, specify paths to scripts for Linux VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

   If you plan to replicate a VM container with Microsoft Windows and Linux VMs, you can select to execute both Microsoft Windows and Linux scripts. When replication starts, Veeam Backup & Replication will automatically determine which OS type is installed on the VM and use the correct scripts for this VM.

> **TIP**
>
> Beside pre-freeze and post-thaw scripts for VM quiescence, you can instruct Veeam Backup & Replication to run custom scripts before the job starts and after the job completes. For more information, see Script Settings.

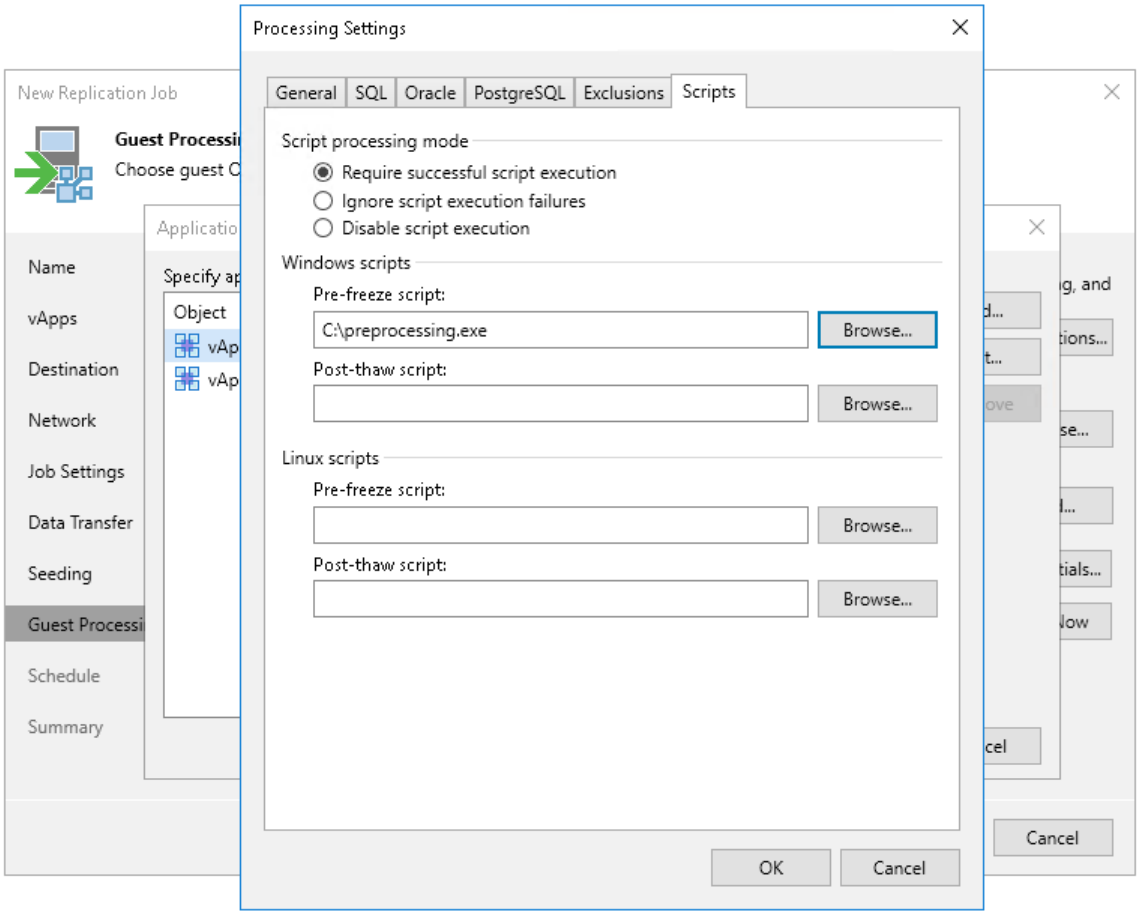
Processing Settings
General
SQL
Oracle
PostgreSQL
Exclusions
Scripts
Script processing mode
Require successful script execution
Ignore script execution failures
Disable script execution
Windows scripts
Pre-freeze script:
C:\preprocessing.exe
Browse...
Post-thaw script:
Browse...
Linux scripts
Pre-freeze script:
Browse...
Post-thaw script:
Browse...
OK
Cancel
New Replication Job
Name
vApps
Destination
Network
Job Settings
Data Transfer
Seeding
Guest Processi
Schedule
Summary
Object
Cancel

# Step 13. Define Job Schedule

At the **Schedule** step of the wizard, select to run the VMware Cloud Director replication job manually or schedule the job to run on a regular basis.

To define a job schedule:

1. Select the **Run the job automatically** check box. If you do not select this check box, you will have to start the job manually to perform VMware Cloud Director replication.
2. Define scheduling settings for the job:
   - To run the job at specific time daily, on defined week days or with specific periodicity, select **Daily at this time**. Use the fields on the right to configure the necessary schedule.
   - To run the job once a month on specific days, select **Monthly at this time**. Use the fields on the right to configure the necessary schedule.

     > **NOTE**
     >
     > When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a set time interval, select **Periodically every**. In the field on the right, select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** and use the time table to define the permitted time window for the job. In the **Start time within an hour** field, specify the exact time when the job must start.

     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run the job by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.
   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.
   - To chain jobs, use the **After this job** field. In the common practice, jobs start one after another: when job *A* finishes, job *B* starts and so on. If you want to create a chain of jobs, you must define the time schedule for the first job in the chain. For the rest of the jobs in the chain, select the **After this job option** and choose the preceding job from the list.

     > **NOTE**
     >
     > The **After this job** function will only start a job if the first job in the chain is started automatically by schedule. If the first job is started manually, jobs chained to it will not be started.

3. In the **Automatic retry** section, define whether Veeam Backup & Replication should attempt to run the job again if the job fails for some reason. During a job retry, Veeam Backup & Replication processes failed vApps only. Enter the number of attempts to run the job and define time spans between them. If you select continuous schedule for the job, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job sessions.

4. In the **Backup window** section, determine a time interval within which the job must be completed. The backup window prevents the job from overlapping with production hours and ensures it does not provide unwanted overhead on your production environment. To set up a backup window for the job:

   a. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**.

   b. In the **Time Periods** section, define the allowed hours and prohibited hours for vApp replication. If the job exceeds the allowed window, it will be automatically terminated.

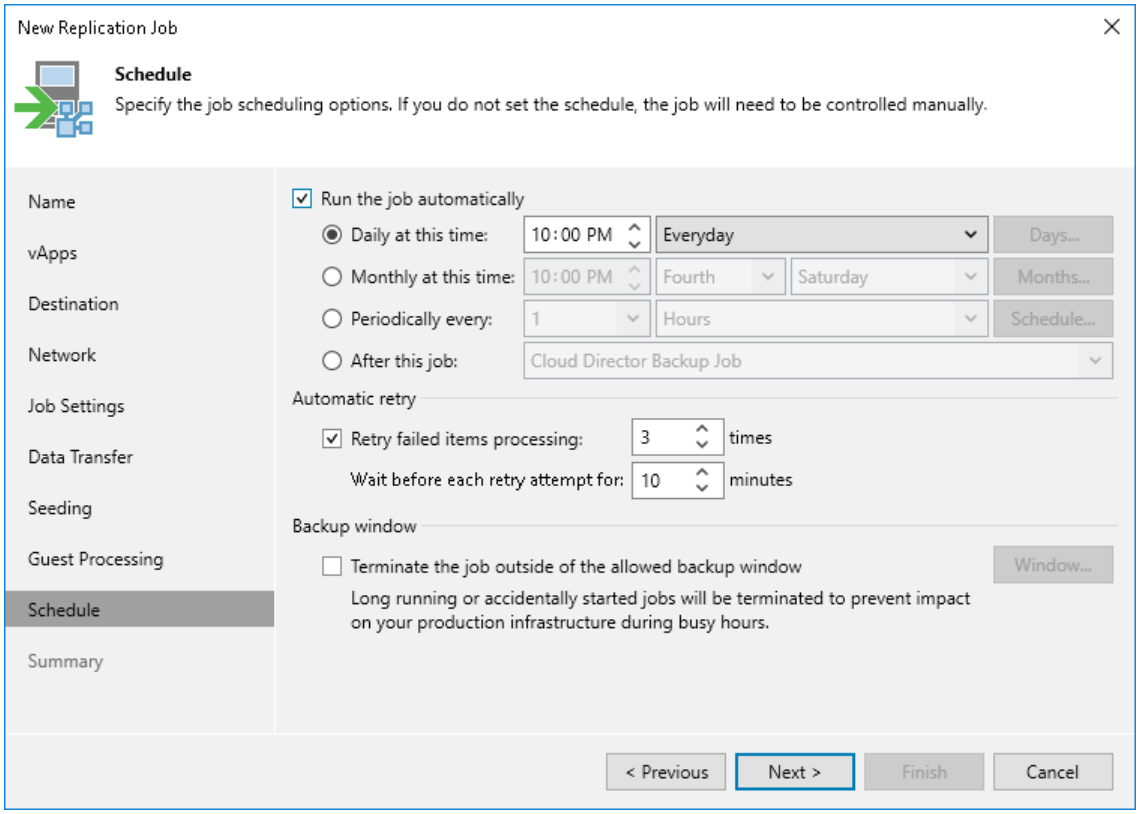

# Step 14. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the VMware Cloud Director replication job. If you want to start the job right after you close the wizard, leave the **Run the job when I click Finish** check box selected, otherwise clear the check box. Then click **Finish** to close the wizard.

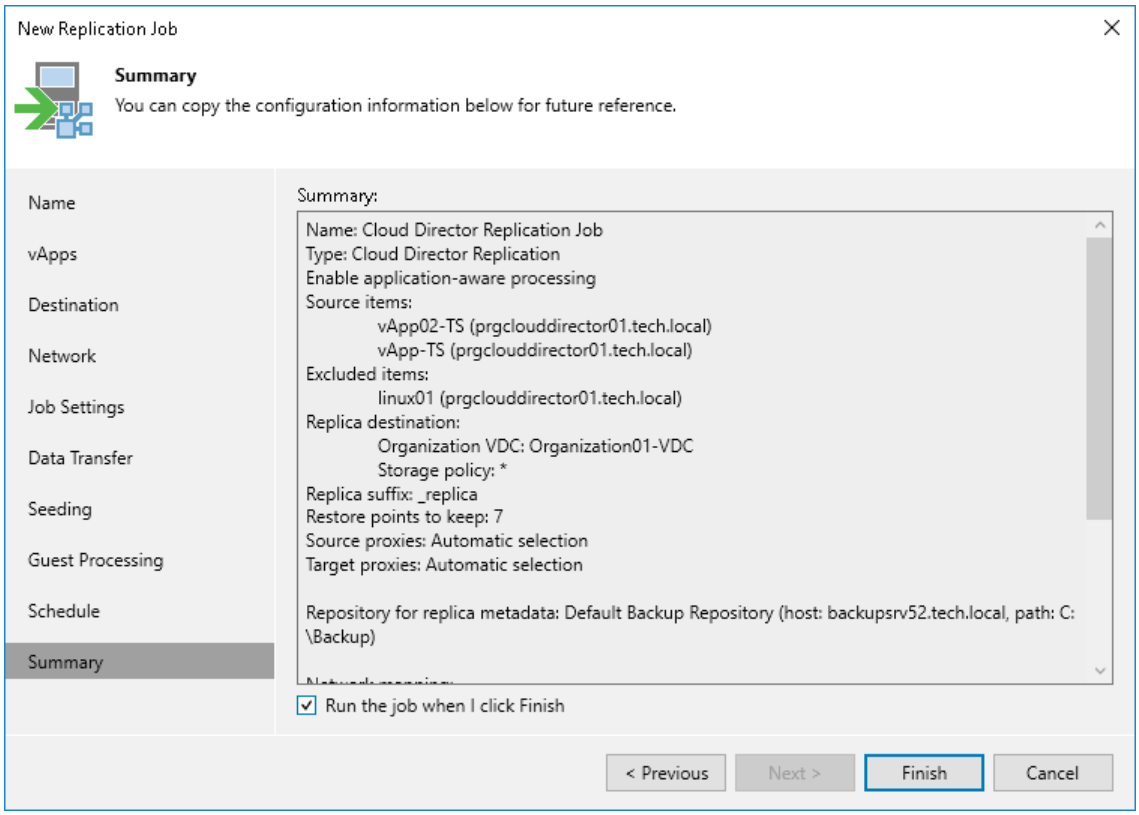

# Creating vApp Replica Seeds

To use seeding, you must have backups of replicated vApps in a backup repository in the disaster recovery (DR) site. These backups are known as replica seeds. For more information on seeding and when to use it, see vApp Seeding and Mapping.

If you do not have replica seeds in the DR site, do the following:

1. Create a backup of vApps that you plan to replicate as described in section Creating Backup Jobs. As the target repository for this job, select a backup repository in the production site. Then run the job.

   If you already have backups containing the necessary vApps, there is no need to configure and run a new backup job. For seeding, you can use any existing backups created by Veeam Backup & Replication. The backup must include VBK and VBM files. If you have a full backup and a chain of forward increments, you can use VIB files together with the VBK and VBM files. In this case, Veeam Backup & Replication will restore vApps from the seed to the latest available restore point.

2. Copy the backup from the backup repository in the production site to a backup repository in the DR site.

   You can move the backup using a file copy job or any other appropriate method, for example, copy the backup to a removable storage device, ship the device to the DR site and copy backups to the backup repository in the DR site.

   If you do not have a backup repository in the DR site, you need to create the repository as described in section Backup Repositories.

   > **IMPORTANT**
   >
   > You cannot copy backups to a scale-out backup repository in the DR site.

3. After the backup is copied to the backup repository in the DR site, perform rescan of this backup repository as described in section Rescanning Backup Repositories. Otherwise, Veeam Backup & Replication will not be able to detect the copied backup.

# Managing Cloud Director Replication Jobs

After you create VMware Cloud Director replication jobs, you can edit, disable and delete them.

To view all created VMware Cloud Director replication jobs, open the **Home** view and navigate to the **Jobs** > **Replication** node. The working area displays the full list of the created replication jobs. Here, you can manage the jobs.

You can manage VMware Cloud Director replication jobs with the following options:

- Retrying VMware Cloud Director Replication Jobs
- Editing VMware Cloud Director Replication Jobs
- Disabling VMware Cloud Director Replication Jobs
- Deleting VMware Cloud Director Replication Jobs

## Retrying Cloud Director Replication Jobs

The retry option is necessary when a VMware Cloud Director job fails and you want to retry this operation again. When you perform a retry, Veeam Backup & Replication restarts the same procedure only for the failed VMs added to the vApp and will not process VMs that have been processed successfully. As a result, the replication job will take less time and will not consume as many resources as when processing a whole vApp.

To a retry a failed VMware Cloud Director replication job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary VMware Cloud Director replication job and select **Retry** on the ribbon. Alternatively, you can right-click the necessary VMware Cloud Director replication job and select **Retry**.

## Editing Cloud Director Replication Jobs

To edit a VMware Cloud Director replication job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary VMware Cloud Director replication job and select **Edit** on the ribbon. Alternatively, you can right-click the necessary VMware Cloud Director replication job and select **Edit**.
4. Follow the instructions provided in Creating Cloud Director Replication Job.

## Disabling Cloud Director Replication Jobs

After you disable the job, Veeam Backup & Replication will pause it and will not process this job according the specified schedule. You can enable a disabled job at any time.

To disable a VMware Cloud Director replication job:

1. Open the **Home** view.

2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary policy and select **Disable** on the ribbon. Alternatively, you can right-click the necessary policy and select **Disable**.

> **TIP**
>
> To enable a disabled VMware Cloud Director replication job, select it and click **Disable** once again.

# Deleting Cloud Director Replication Jobs

After you delete the job, the replicas created by this job are displayed under the **Replicas** node.
Veeam Backup & Replication allows you to delete only stopped VMware Cloud Director replication jobs. You can also permanently delete a job from Veeam Backup & Replication and from the configuration database.

Veeam Backup & Replication allows you to delete only disabled VMware Cloud Director replication jobs.

To delete a VMware Cloud Director replication job:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **Replication** node.
3. In the working area, select the necessary VMware Cloud Director replication job and select **Delete** on the ribbon. Alternatively, you can right-click the necessary VMware Cloud Director replication job and select **Delete**.

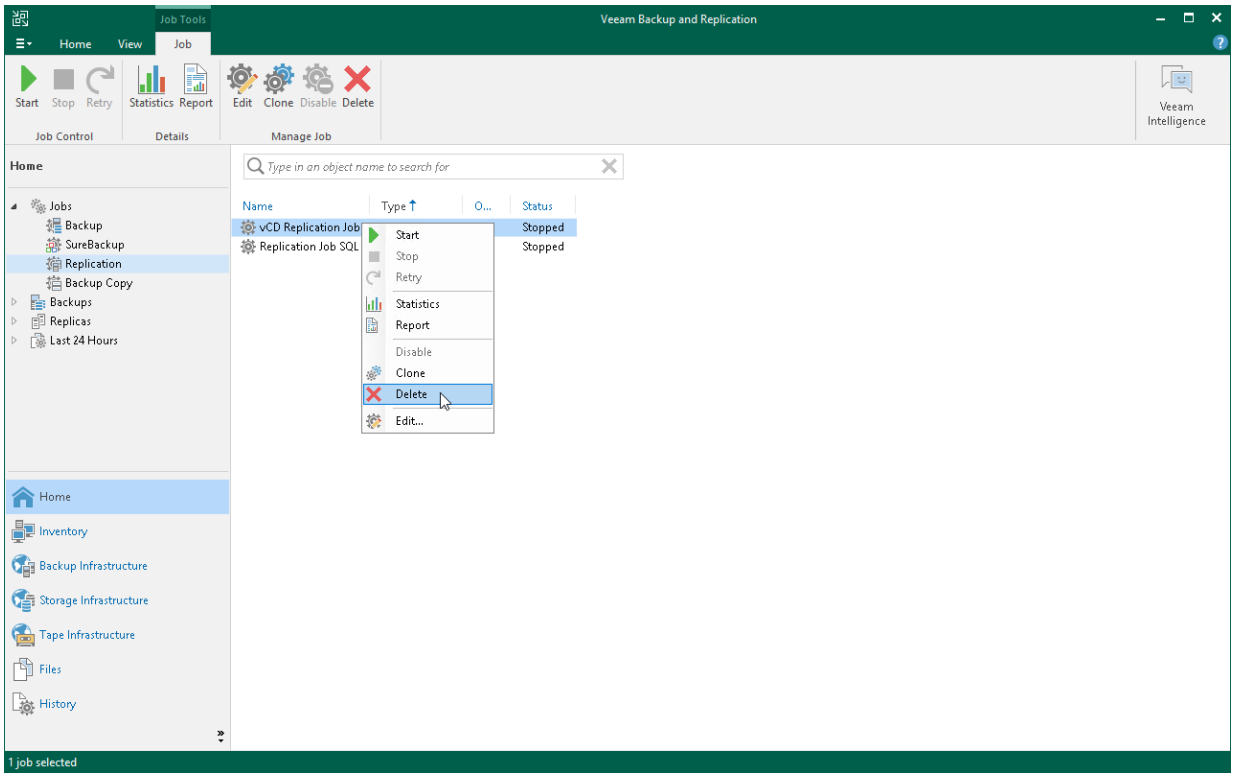

# Managing Cloud Director Replicas

To view all created replicas, open the **Home** view and navigate to the **Replicas** node. The working area displays the full list of the created replicas. Here, you can view replica properties and delete replicas from the configuration database or disk.

## Viewing Replica Properties

You can view replica properties that provide the following information:

- Available restore points
- Date of restore points creation
- Data size and replica status

To view replica properties:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, right-click the necessary replica and select **Properties**. Alternatively, select **Properties** on the ribbon.

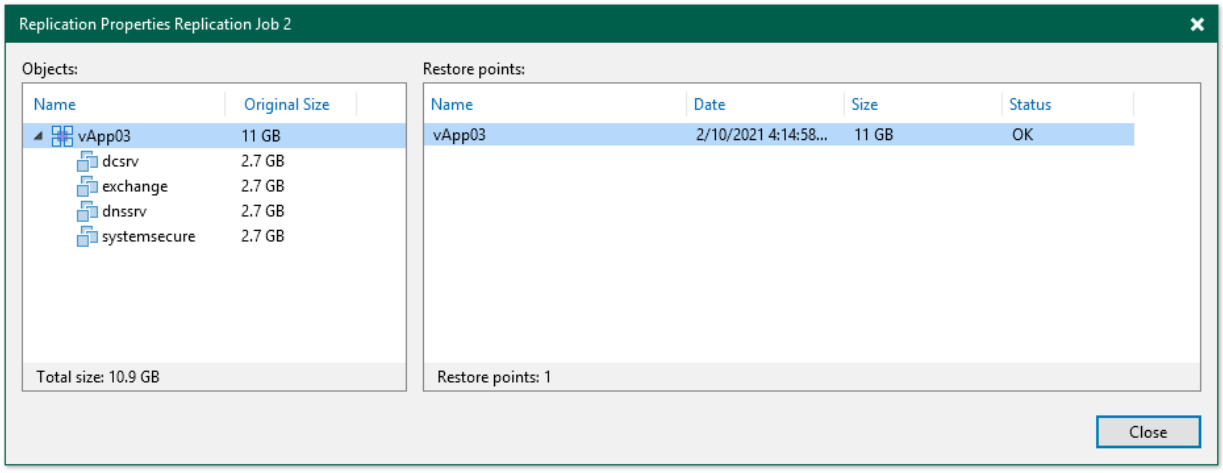


## Rescanning Replicas

Rescan is a process that allows Veeam Backup & Replication to verify that the following data written to configuration database is up to date: information about the disaster recovery (DR) site, its components and VM containers. The replica rescan process is performed the following way:

1. Veeam Backup & Replication gathers information on replicas that are currently available in the DR site.
2. Veeam Backup & Replication compares this information with information stored in the configuration database about replicas from this DR site.
3. If information about replicas from the DR site differs from information stored in the configuration database about these replicas, Veeam Backup & Replication updates the configuration database.

To rescan replicas, do the following:

1. Open the **Home** view.
2. In the inventory pane, right-click the **Replicas** node and select **Rescan replicas**.

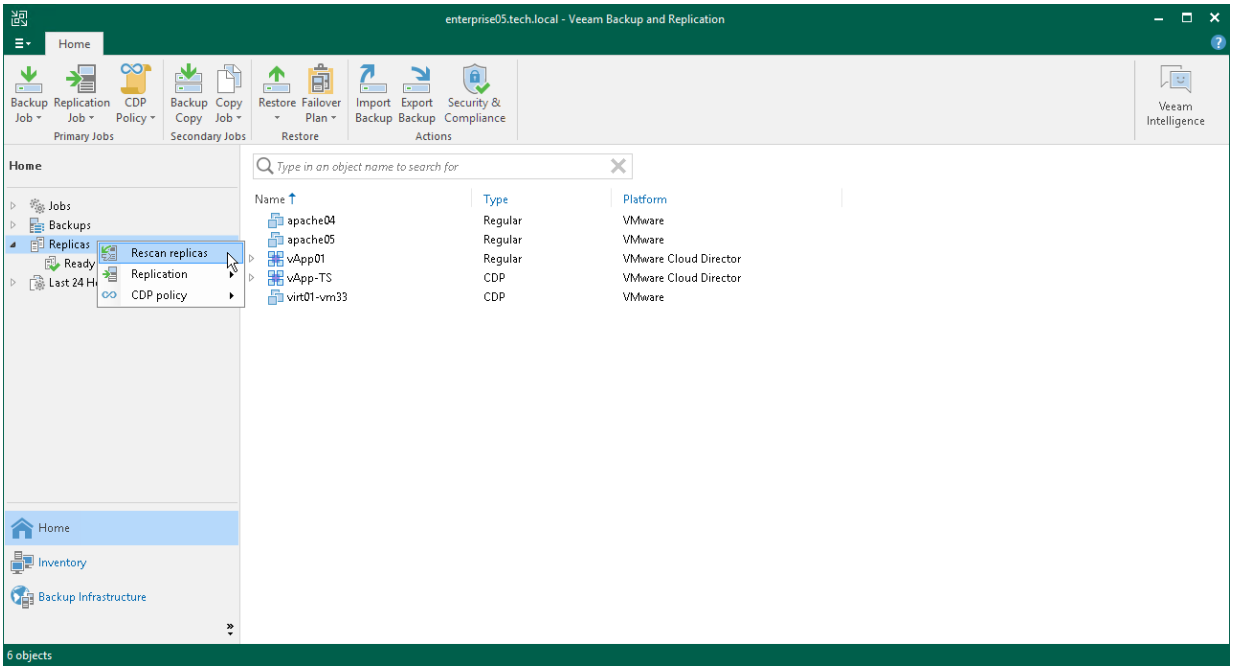


# Removing from Configuration

When you remove replicas from the configuration, Veeam Backup & Replication removes records about the replicas from the configuration database, stops showing the replicas in Veeam Backup & Replication console and also stops synchronizing their state with the state of the source VMs. However, the actual replicas remain on hosts.

To remove records about replicas from the Veeam Backup & Replication console and configuration database:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.
3. In the working area, select replicas in the *Ready* state and click **Remove from** > **Configuration** on the ribbon. Alternatively, right-click one of the selected replicas and select **Remove from configuration**.

**NOTE**

Consider the following:

- The **Remove from configuration** operation can be performed only for replicas in the *Ready* state. If the replica is in the *Failover* or *Failback* state, this option is disabled.
- When you perform the **Remove from configuration** operation for a vApp that is replicated as a standalone object, Veeam Backup & Replication removes this vApp from the initial replication job. When you perform the **Remove from configuration** operation for a vApp that is replicated as part of a VM container, Veeam Backup & Replication adds this vApp to the list of exclusions in the initial replication job. For more information, see Exclude Objects from Replication Job.

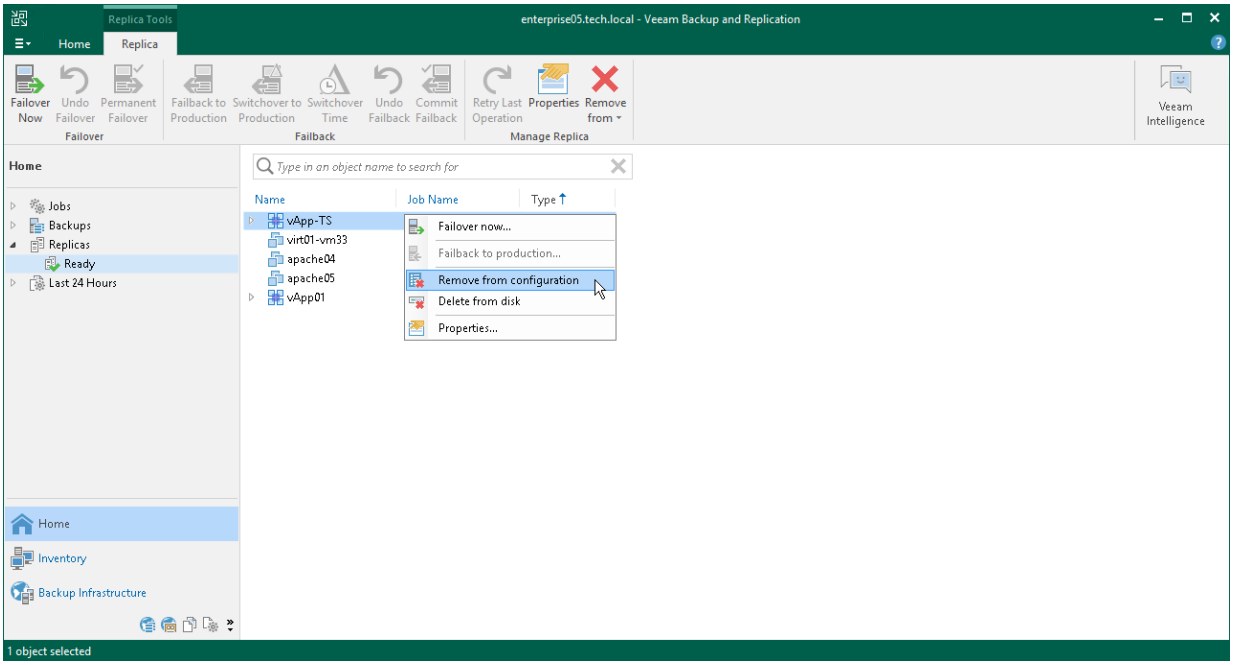


# Deleting from Disk

When you delete replicas from disks, Veeam Backup & Replication removes the replicas not only from the Veeam Backup & Replication console and configuration database, but also from host storage.

> **NOTE**
>
> Consider the following:
>
> - You can delete records only about replicas that are in the *Ready* state.
> - Do not delete replica files from the destination storage manually, use the **Delete from disk** option instead. If you delete replica files manually, subsequent replication sessions will fail.
> - Unlike the **Remove from configuration** operation, the **Delete from disk** operation does not remove the processed workload from the initial replication job. This means that the replication process will restart for this workload. To avoid this, you can exclude the workload from the replication job or disable the job.

To delete replica files from disks:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.

3. In the working area, select the necessary replica and click **Remove from** > **Disk** on the ribbon. As an alternative, right-click the replica and select **Delete from disk**.

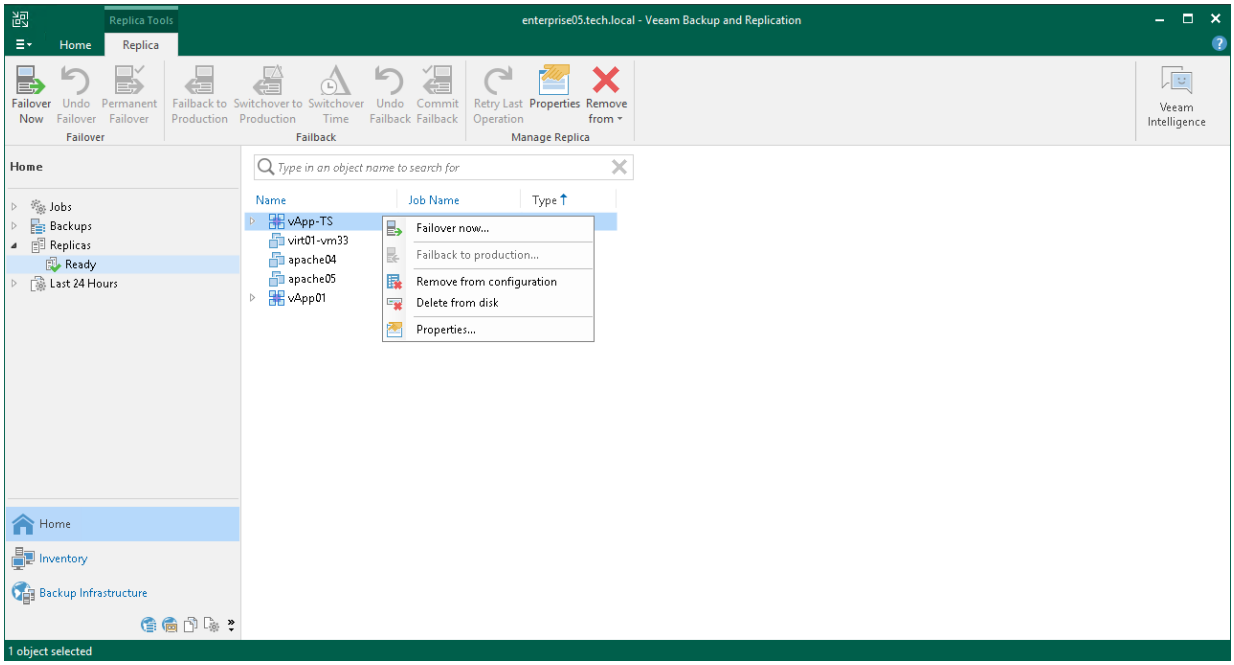

# Failover and Failback for Cloud Director

Failover and failback are operations that allow you to manage your production and disaster recovery (DR) sites if a disaster strikes. Failover is a process of switching from the vApp on the source organization VDC to its replica on a target organization VDC that is set up as the disaster recovery (DR) site. Failback is a process of returning from the replica to the source vApp or a new vApp.

Veeam Backup & Replication provides the following failover and failback operations:

- **Perform failover**

  When you perform failover, you shift all processes from the source vApp in the production organization VDC to the replica in the DR organization VDC. Failover is an intermediate step that needs to be finalized: you can perform permanent failover, perform failback or undo failover.

- **Perform permanent failover**

  When you perform permanent failover, you permanently switch from the source vApp to a replica and use this replica as the production vApp. The source vApp is excluded from VMware Cloud Director replica processing.

- **Undo failover**

  When you undo failover, you switch back to the source vApp and discard all changes made to the replica while it was running. For example, you can use the undo failover scenario if you have failed over to the replica for testing and troubleshooting purposes, and you do not need to synchronize the source vApp state with the current state of the replica.

- **Perform failback**

  When you perform failback, you switch back to the source vApp and send to the source vApp all changes that took place while the replica was running. If the source organization VDC is not available, you can recover a vApp with the same configuration as the source vApp and switch to it.

  When you perform failback, changes are only sent to the source or recovered vApp but not published. You must test whether the source or recovered vApp works with these changes. Depending on the test results, you can do the following:

  - **Commit failback**. When you commit failback, you confirm that the source or recovered vApp works as expected and you want to get back to it.
  - **Undo failback**. When you undo failback, you confirm that the source or recovered vApp is not working as expected and you want to get back to the replica.

The following scheme can help you decide which steps are preferable when you fail over to a replica.

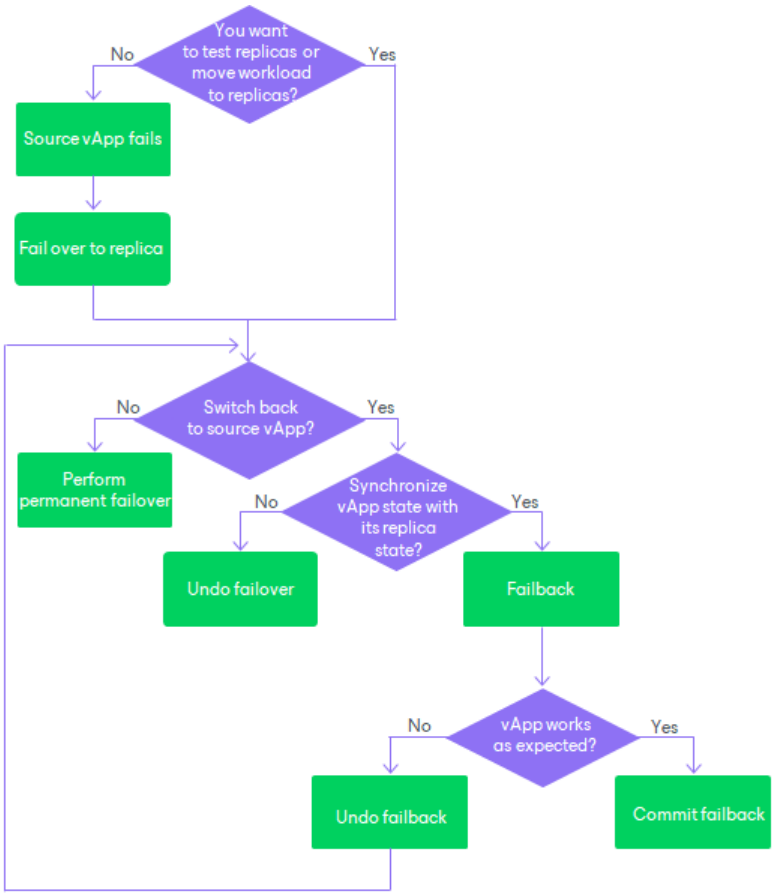


# Failover

Failover is a process when Veeam Backup & Replication switches processes from the source vApp in the production organization VDC to its replica in the disaster recovery organization VDC. During failover, Veeam Backup & Replication recovers the replica to the required restore point and shifts all I/O processes from the source vApp to its replica. As a result, you have a fully functional vApp within several minutes, and your users can access services and applications with minimum disruption.

You can fail over to replicas not only when a disaster strikes the production organization VDC, but also to test replicas for recoverability. If the source vApps and replicas are located in the same network, consider temporary disconnecting the source vApps from the network to avoid IP address or machine name conflicts.

# How Failover Works

The failover operation is performed in the following way:

1. Veeam Backup & Replication puts all replication activities on hold.
2. The state of the replica is changed from *Ready* to *Processing*.
3. Veeam Backup & Replication recovers a replica to the required restore point.
4. Veeam Backup & Replication powers on the replica.

   The source vApp still exists and failover does not change the vApp state: if the vApp is powered on when you perform failover, it remains powered on when failover completes; if the vApp is powered off, it remains in this state.

5. All changes made to the replica while it is running in the *Failover* state are written to the delta file and stored in the target host.
6. After failover completes successfully, the vApp state is changed from *Processing* to *Failover*.

# Finalizing Failover

Failover is an intermediate step that needs to be finalized. You can perform the following operations:

- Undo failover
- Perform permanent failover
- Perform failback

# In This Section

- Performing Failover
- Performing Failover Retry

## Performing Failover

For more information on failover, see Failover and Failback for Cloud Director and Failover.

To perform failover, do the following.

### Before You Begin

Before you fail over to a replica, check the following prerequisites:

- The failover operation can be performed for vApp that have been successfully replicated at least once.
- Replicas must be in the *Ready* state.
- The VMs added to the target vApp must be powered off.

## Step 1. Launch Failover Wizard

To launch the **Failover** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vCloud Director** > **Restore from replica** > **Entire vApp** > **Failover vApp to a replica**.
- Open the **Home** view, in the inventory pane select **Replicas**. In the working area, select the necessary replica, right-click one of the selected replica and click **Failover Now**. Alternatively, click **Failover Now** on the ribbon.

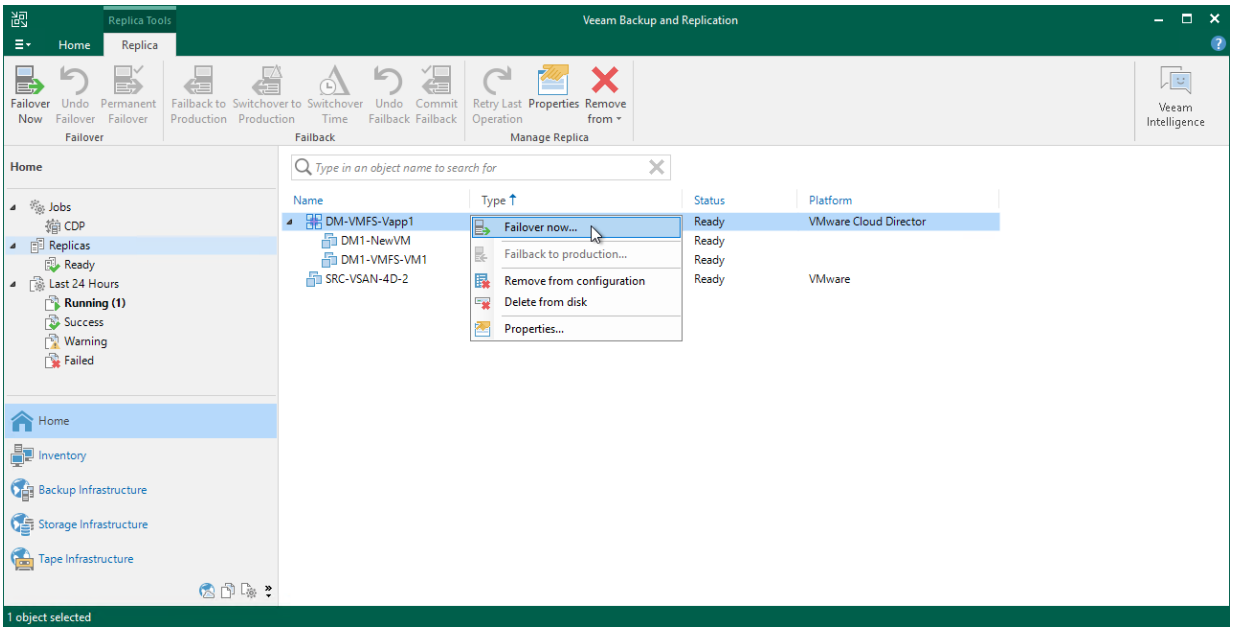

## Step 2. Select vApps

At the **vApp** step of the wizard, you can modify a list of vApps from which you fail over. To add vApps, click **Add** > **From infrastructure** if you want to add vApps from the virtual infrastructure, or **Add** > **From replicas** if you want to add vApps from existing replicas. Then select the necessary vApps. If you select organizations or organization VDCs, Veeam Backup & Replication will expand them to a vApp list.

> **NOTE**
>
> Make sure that vApps you select from the virtual environment have been successfully replicated at least once.

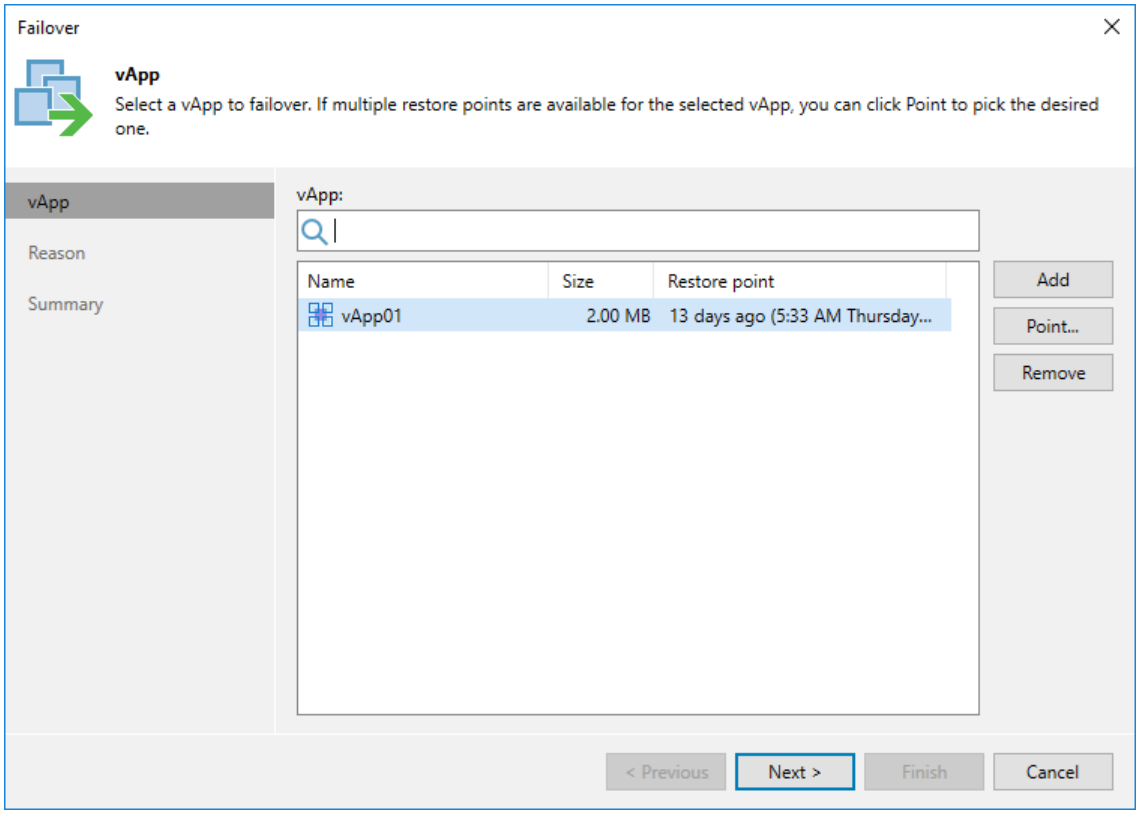

## Step 3. Select Restore Point

By default, Veeam Backup & Replication uses the latest valid restore point of the replica. However, you can fail over to an earlier state of the vApps. If you have chosen to perform failover for several vApps, you can select the necessary restore point for every vApp in the list.

To select a restore point for a vApp:

1. In the **vApp** list, select a vApp.
2. Click **Point** on the right.
3. In the **Restore Points** window, select a restore point to which you want to fail over.

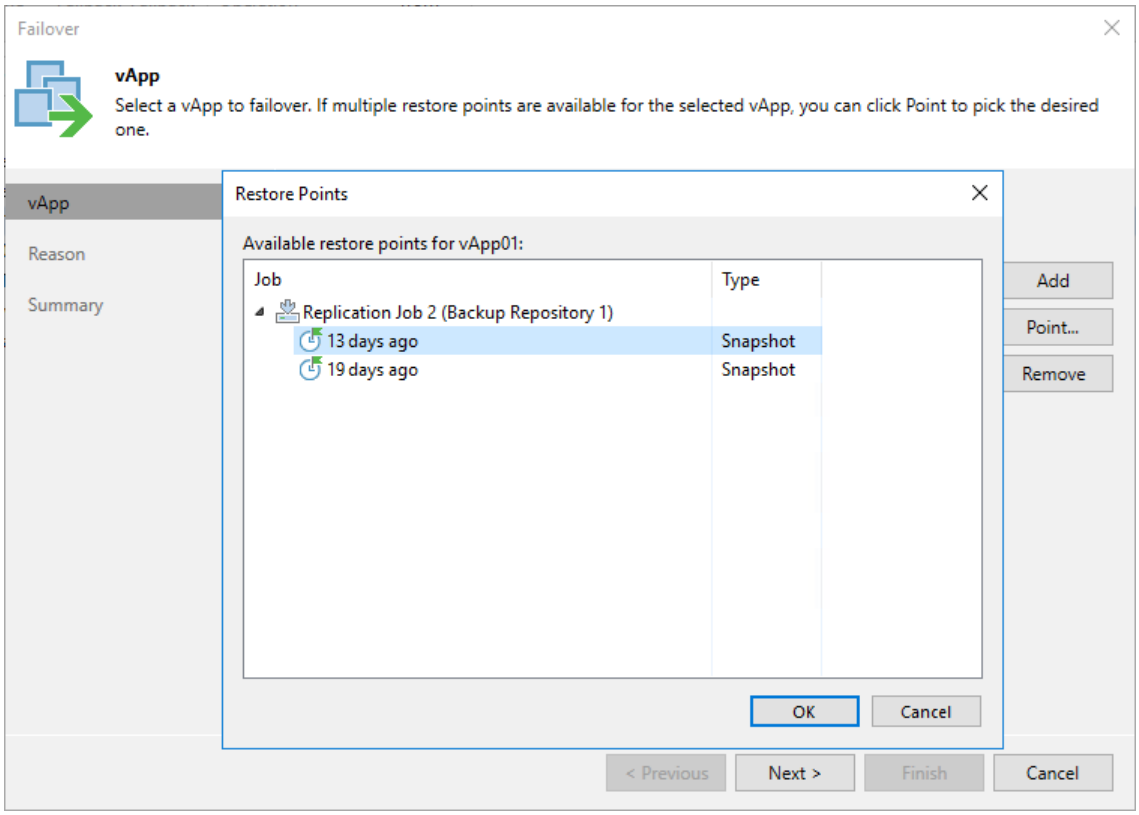

## Step 4. Specify Failover Reason

At the **Reason** step of the wizard, enter a reason for failing over to the replicas. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

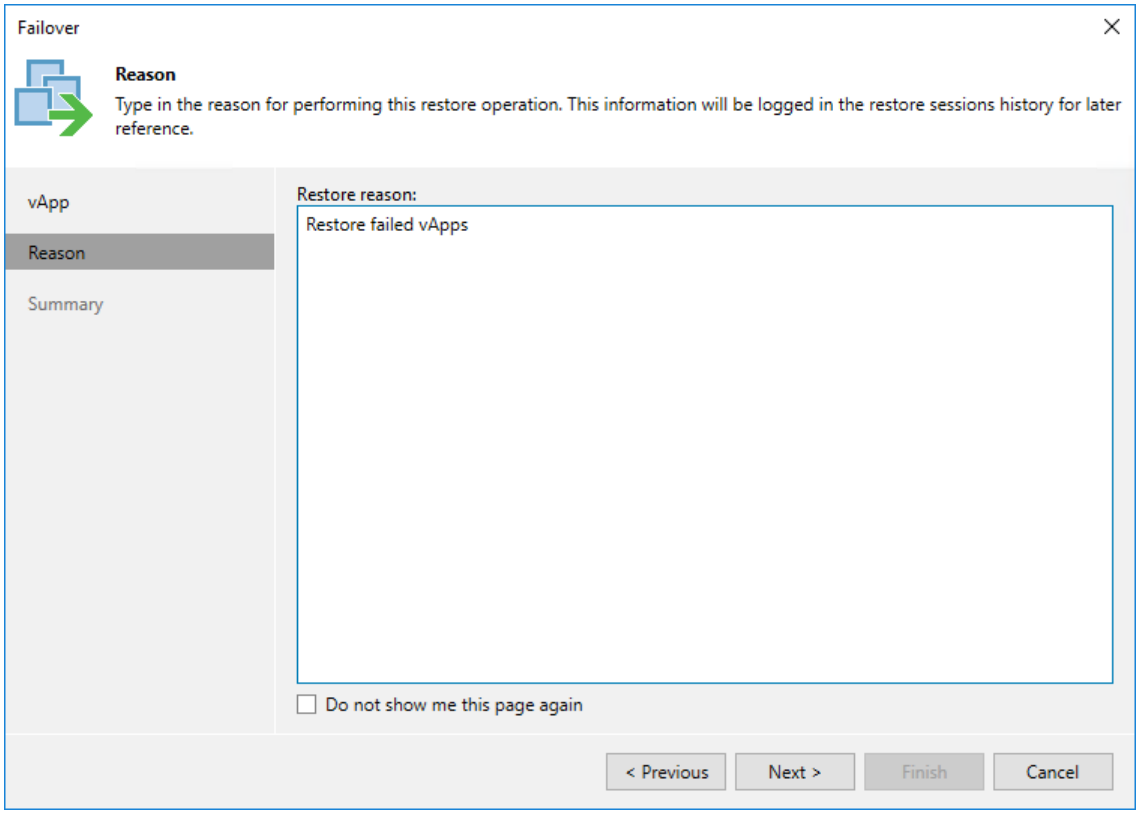

### Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the failover task. Then click **Finish** to start the failover process.

When the failover process is complete, the replicas will be started on the target hosts.

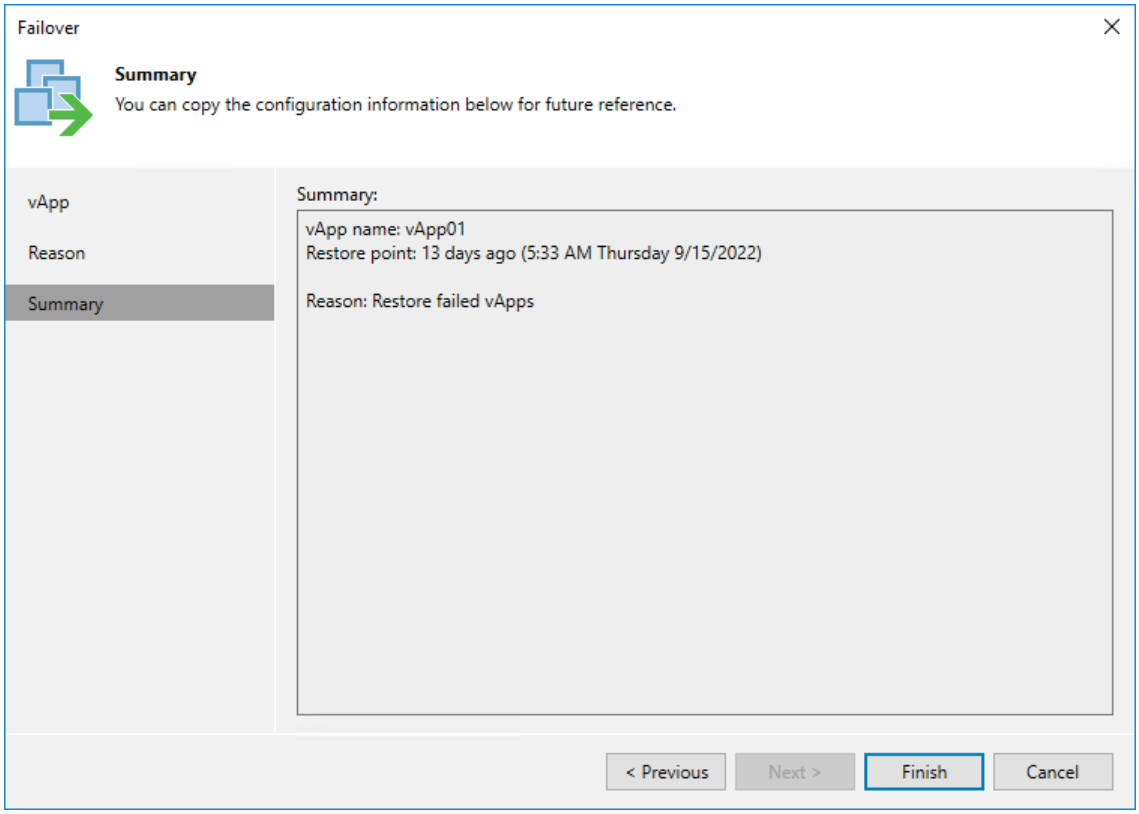


# What You Do Next

Failover is an intermediate step that needs to be finalized. You can finalize failover in the following ways:

- Perform permanent failover
- Undo failover
- Perform failback

## Performing Failover Retry

The failover retry option is necessary when failover of vApps fails with the *Incomplete* state. When you perform a retry, Veeam Backup & Replication restarts failover only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, failover takes less time and does not consume as many resources as when processing a whole vApp.

To retry failover:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary vApp and select **Retry Failover** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry failover**.

# Permanent Failover

Permanent failover is one of the ways to finalize failover. When you perform permanent failover, you permanently switch processes from the source vApp to its replica. As a result, the replica stops acting as a replica and starts acting as the production vApp.

The permanent failover operation is performed in the following way:

1. Veeam Backup & Replication removes restore points of the replica from the replication chain and deletes associated files from the datastore. Changes that were written to the delta disks are committed to the replica to bring the replica to the most recent state.
2. Veeam Backup & Replication removes the replica from the list of replicas in the Veeam Backup & Replication console and the replica becomes the production vApp.
3. To protect the replica from corruption after permanent failover is complete, Veeam Backup & Replication reconfigures the job or policy and adds the source vApp to the list of exclusions. When replication job starts, the source vApp is skipped from processing. As a result, no data is written to the working replica.

## Performing Permanent Failover

For more information on permanent failover, see Failover and Failback for Cloud Director and Permanent Failover.

To perform permanent failover, do one of the following:

- Open the **Home** view, in the inventory pane select **Replicas** > **Active**. In the working area, select the necessary vApp and click **Permanent Failover** on the ribbon.
- Open the **Home** view, in the inventory pane select **Replicas** > **Active**. In the working area, right-click the necessary vApp and select **Permanent Failover**.

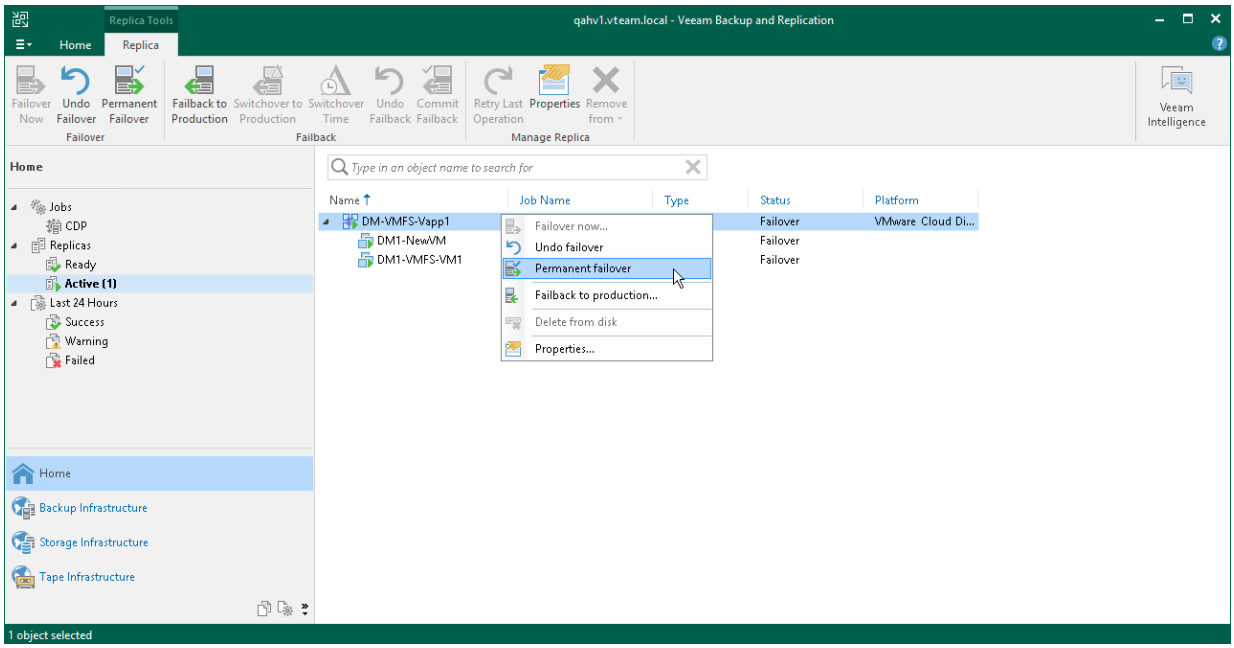

## Performing Permanent Failover Retry

If permanent failover failed, you can retry this operation. When you perform a retry, Veeam Backup & Replication restarts permanent failover only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, permanent failover takes less time and does not consume as many resources as when processing a whole vApp.

To perform a retry:

1. Open the **Home** view, in the inventory pane, navigate to the **Replicas** > **Active** node.
2. In the working area, select the necessary vApp and select **Retry Permanent Failover** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry permanent failover**.

# Failover Undo

Failover undo is one of the ways to finalize failover. When you undo failover, you switch back from a vApp replica to the original vApp. Veeam Backup & Replication discards all changes made to the vApp replica while it was in the *Failover* state.

The failover undo operation is performed in the following way:

1. Veeam Backup & Replication reverts the replica to its pre-failover state. To do this, Veeam Backup & Replication powers off the vApp replica and gets it back to the latest restore point in the replication chain.
2. The state of the replica gets back to *Ready*, and Veeam Backup & Replication resumes replication activities for the original vApp on the source host.

## Undoing Failover

For more information on failover undo, see Failover and Failback for Cloud Director and Failover Undo.

To undo failover:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, select the necessary replica and click **Undo Failover** on the ribbon. Alternatively, right-click the necessary replica and select **Undo Failover**.

4. In the displayed window, click **Yes** to confirm the operation.

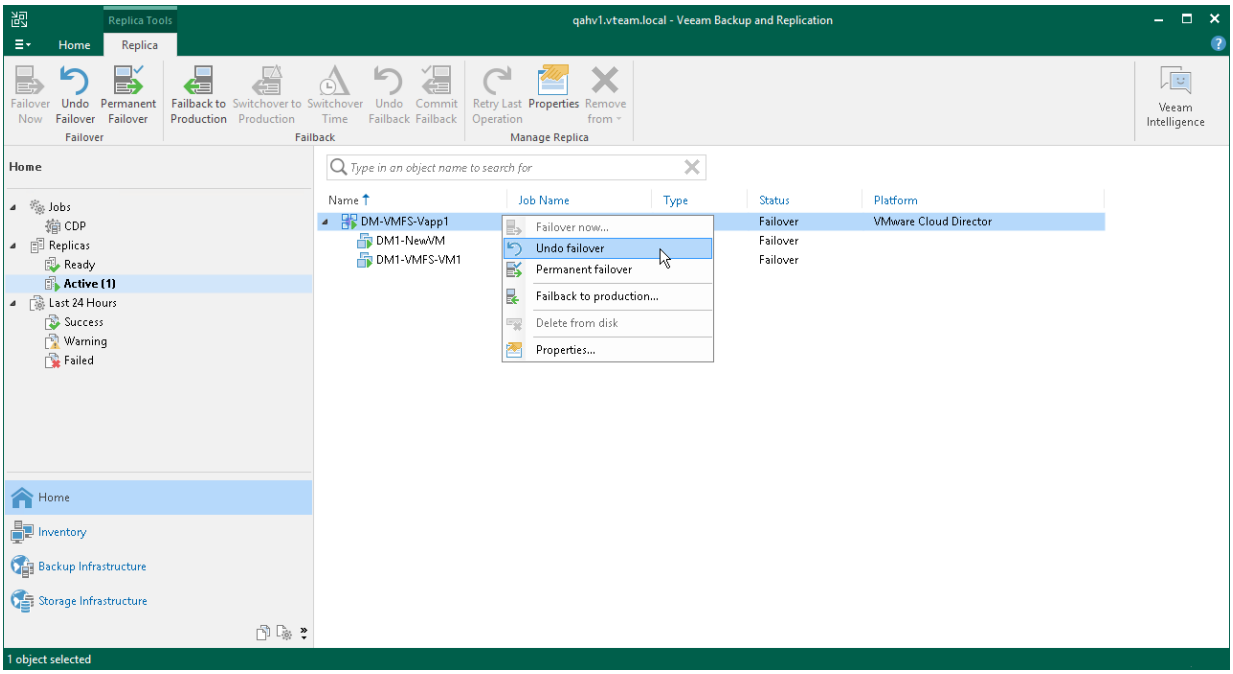


## Performing Failover Undo Retry

If the failover undo operation failed, you can retry this operation. When you perform a retry, Veeam Backup & Replication restarts the failover undo operation only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, the failover undo operation takes less time and does not consume as many resources as when processing the whole vApp.

To perform a retry:

1. Open the **Home** view, in the inventory pane, navigate to the **Replicas** > **Active** node.
2. In the working area, select the necessary vApp and select **Retry Undo Failover** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry undo failover**.

# Failback

Failback is an option that allows you to switch back from the replicated vApp on a disaster recovery organization VDC to the production vApp. When you perform failback, you switch back to the production VM from a VM replica, shift I/O processes from the disaster recovery site to the production site.

You can perform failback in the following ways:

- Fail back to the source vApp in the original location.
- Fail back to a vApp already recovered to a new location. This vApp must be recovered before you perform failback. For example, you can recover the VM from a backup.
- Fail back to a vApp recovered from a replica to a new location, or to any location but with different settings. The vApp will be recovered from the replica during the failback process.

The first two options help you decrease recovery time and the use of the network traffic because Veeam Backup & Replication needs to transfer only differences between the source or recovered vApp and replica. For the third option, Veeam Backup & Replication needs to transfer the whole vApp data, including its configuration and virtual disk content. Use the third option if there is no way to use the source vApp or restore it from a replica.

Veeam Backup & Replication performs failback in two phases:

- **First phase**: Veeam Backup & Replication synchronizes the state of the production vApp (the source vApp, an already recovered vApp or a vApp recovered from the replica) with the current state of the replica. This phase may take a lot of time especially if vApp is large. While Veeam Backup & Replication performs the first phase of failback, VMs from replicas are still up and running, users can access these VMs and perform daily routine tasks as normal. All changes made to vApps during the first phase of a failback a written to a delta file.
- **Second phase**: Veeam Backup & Replication transfers all changes made to the replica during the first phase of failback to the production vApp, switches all processes from the replica to the production vApp and turns off the replica.

  The time when the second phase starts depends on how you want to switch from the replica to the production vApp. You can switch to the production vApp automatically, at the scheduled time or manually. If you select to switch automatically, the second phase will start right after the first phase finishes. If you select to switch at the scheduled time or manually, the second phase will start at the time you want.

The process of failing back to the source vApp or a source vApp restored in a different location differs from the process of failing back to a specific location:

- How failback to the source vApp or a source vApp restored in a different location works
- How failback to a specific location works

# How Failback to Source vApp or Source vApp Restored in Different Location Works

When you fail back to the source vApp or an already recovered vApp, Veeam Backup & Replication performs the following operations during the first phase:

1. Veeam Backup & Replication calculates the difference between disks of the production vApp and disks of the replica in the *Failover* state. The calculation of the difference helps Veeam Backup & Replication understand what data needs to be transferred to the production vApp and to synchronize its state with the state of the replica.

   [For ESXi hosts prior to version 7.0] If you fail back to the source vApp in the source location and you have enabled the **Quick rollback** option, this calculation can be performed much faster. For more information on the **Quick rollback** option, see Quick Rollback.

2. Veeam Backup & Replication transfers the data that was detected as different to the production vApp. The transferred data is written to the production vApp.

3. Veeam Backup & Replication changes the state of the replica from *Failover* to *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the following operations:

1. The guest OS of the replica is shut down or the replica is powered off.

   If VMware Tools are installed on the VM added to the replica, Veeam Backup & Replication tries to shut down the replica guest OS. If nothing happens in 15 minutes, Veeam Backup & Replication powers off the vApp replica. If VMware Tools are not installed on the VM added to the replica or the vApp is suspended, Veeam Backup & Replication powers off the vApp. The replica remains powered off until you commit failback or undo failback.

2. Veeam Backup & Replication calculates the difference between disks of the production vApp and disks of the replica. The calculation of the difference helps Veeam Backup & Replication understand what data was changed while the replica was in the *Ready to switch* state.

3. Sends data changed on the replica while it was in the *Ready to switch* state to the production vApp.

4. The state of the replica is changed from *Ready to switch* to *Failback*.

5. [If you fail back to a recovered vApp] Veeam Backup & Replication updates the ID of the source vApp in the Veeam Backup & Replication configuration database. The ID of the source vApp is replaced with the ID of the recovered vApp.

6. If you have selected to power on the production vApp after failback, Veeam Backup & Replication powers on the production vApp on the host.

# How Failback to Specific Location Works

When you fail back to a vApp recovered from a replica, Veeam Backup & Replication performs the following operations during the first phase:

1. Veeam Backup & Replication requests VMware Cloud Director to create on the target organization VDC an empty vApp with the same configuration as the replica. VMware Cloud Director server registers the created production vApp.

2. Veeam Backup & Replication transfers data of the replica to the production vApp to update the production vApp state to the replica state.

3. Veeam Backup & Replication changes the state of the replica from *Failover* to the *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the same operations as described in section How Failback to Source vApp or Already Recovered vApp Works except for the step 2.

Failback is an intermediate step that needs to be finalized. If the production vApp works as expected and you want to get back to it, commit failback. If the vApp does not work as expected, undo failback.

## Quick Rollback

Quick rollback helps you significantly reduce the failback time. You can use quick rollback if you fail back from a replica to the source vApp in the original location.

During failback, Veeam Backup & Replication calculates differences between VM disks of the source vApp and disks of the replica. With the quick rollback option enabled, Veeam Backup & Replication compares only those disk sectors that have changed during the replica was in the *Failover* state instead of comparing entire disks. To get information about the changed disk sectors, Veeam Backup & Replication uses VMware vSphere Changed Block Tracking (CBT).

As a result of enabling quick rollback, difference calculation becomes much faster. After the differences are calculated, Veeam Backup & Replication performs failback in a regular way: transport changed blocks to the source vApp, powers off the replica and synchronizes the source vApp with the replica once again.

# Requirements for Quick Rollback

To perform quick rollback, make sure that the following requirements are met:

- You fail back to the source vApp in the original location.
- Do not use quick rollback if the problem occurred at the vApp hardware level, storage level or due to a power loss.

  Use quick rollback if you fail back to the source vApp that had a problem at the guest OS level — for example, there was an application error or a user accidentally deleted a file on the source VM guest OS.
- CBT must be enabled for the source vApp.

# Limitations for Quick Rollback

During the first replication job session after failback with quick rollback, CBT on the source vApp is reset. Due to that Veeam Backup & Replication will read data of the entire vApp.

## Performing Failback

For more information on failback, see Failover and Failback for Cloud Director and Failback.

To perform failback, do the following:

### Before You Begin

Before you perform failback, check the following prerequisites:

- vApps for which you plan to perform failback must be successfully replicated at least once.
- Replicas must be in the *Failover* state.

## Step 1. Launch Failback Wizard

To launch the **Failback** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware Cloud Director** > **Restore from replica** > **Entire vApp** > **Failback to production**.
- Open the **Home** view, in the inventory pane select **Replicas** > **Active**. In the working area, right-click the necessary replica and select **Failback to production**. Alternatively, click **Failback to Production** on the ribbon.

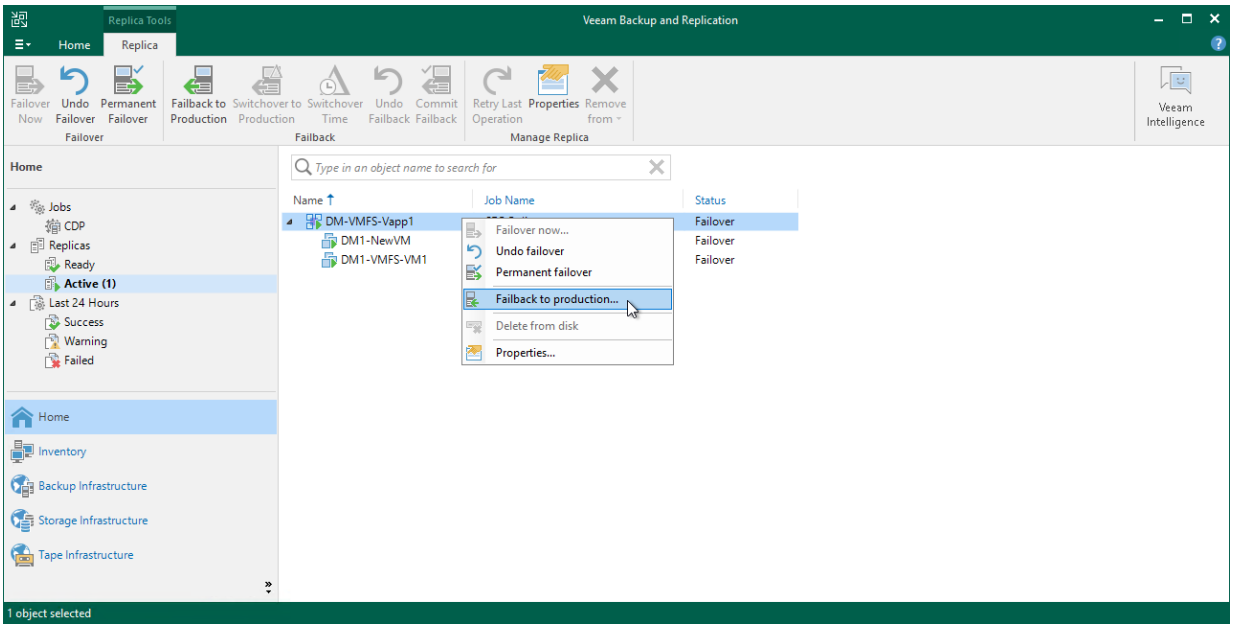

## Step 2. Select Replicas

At the **Replica** step of the wizard, select replicas from which you want to fail back.

To update the list of replicas that are ready for failback (replicas in the *Failover* state), click **Populate**.

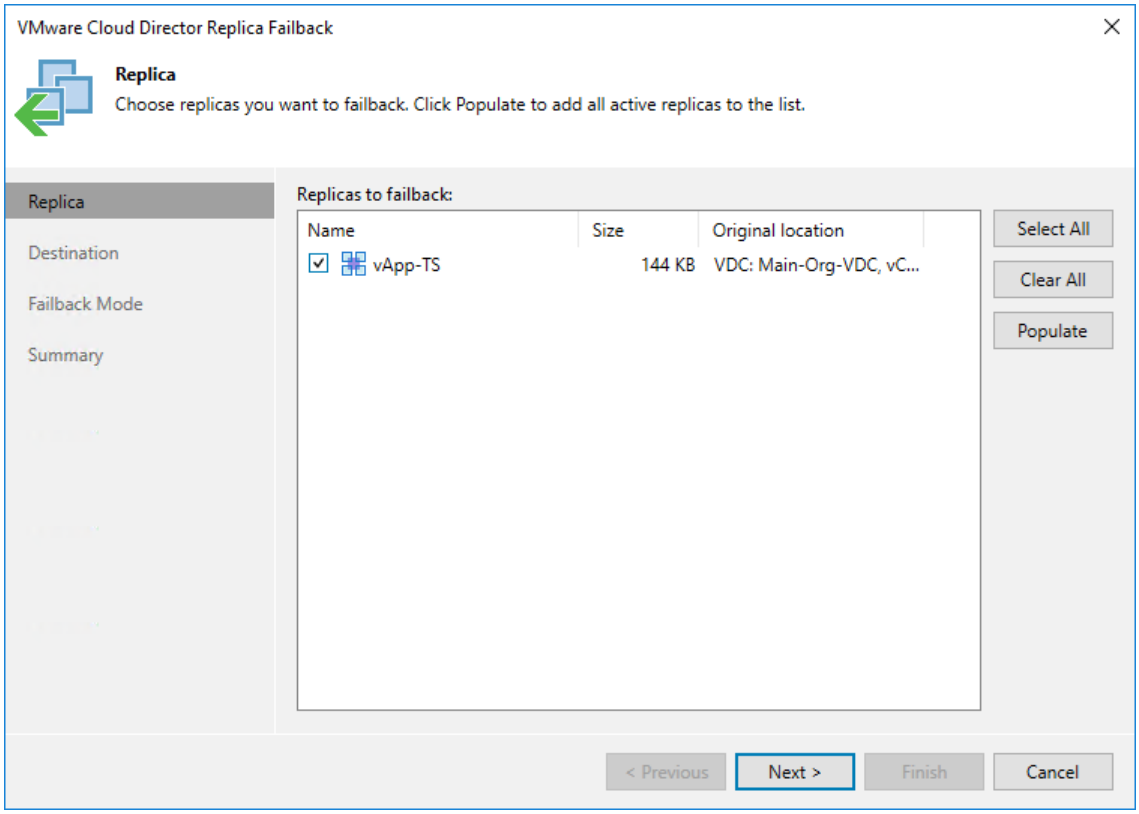

## Step 3. Select Failback Destination

At the **Destination** step of the wizard, select the failback destination and backup proxies for vApp data transport during failback:

1. Select a destination for failback. Veeam Backup & Replication supports the following options:
   - **Failback to the original vApp** – select this option if you want to fail back to the source vApps that reside on the source hosts. Veeam Backup & Replication will synchronize the state of the source vApps with the current state of their replicas to apply any changes that occurred to the replicas while running in the disaster recovery (DR) site.

     If this option is selected, you will proceed to the **Failback Mode** step of the wizard.
   - **Failback to the original vApp restored in a different location** – select this option if the source vApps have already been recovered to a new location, and you want to switch to the recovered vApps from their replicas. Veeam Backup & Replication will synchronize the state of the recovered vApps with the current state of the vApp replicas to apply any changes that occurred to the replicas while running in the DR site.

     If this option is selected, you will proceed to the **Target vApp** step of the wizard.

     > **TIP**
     >
     > You can restore to another VMware Cloud Director.
   - **Failback to the specified location** – select this option if you want to recover the source vApps from replicas. You can recover vApps to a new location, or to any location but with different settings (such as network settings, virtual disk type, configuration file path and so on). Select this option if there is no way to fail back to the source vApp or an already recovered vApp.

     If you select this option, the wizard will include additional steps.

   If you select one of the first two options, Veeam Backup & Replication will send to the source/recovered vApps only differences between the existing virtual disks of VMs included in the vApps.
   Veeam Backup & Replication will not send replica configuration changes such as different IP address or network settings (if replica Re-IP and network mapping were applied), new hardware or virtual disks added while the replicas were in the *Failover* state.

   If you select **Failback to the specified location**, Veeam Backup & Replication will send to the specified location whole replica data, including configurations and virtual disk content.

2. To select which backup proxies will be used for data transfer, click **Pick backup proxies for data transfer**.

   By default, Veeam Backup & Replication selects proxies automatically. Before processing a new vApp in the vApp list, Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication selects the most appropriate proxy basing on the following information: transport modes that the backup proxies can use and the current workload on the backup proxies.

   If want to select proxies manually and if vApps and their replicas reside in different sites, select at least one backup proxy in the production site and one proxy in the disaster recovery site. If vApps and replicas reside in the same site, you can use the same backup proxy as the source and target one.

   We recommend that you select at least two backup proxies in each site to ensure that failback will be performed in case one proxy fails or looses the network connection.

4. [For failback to the source vApps] If you want to fasten failback, and the source vApps had problems at the guest OS level, select the **Quick rollback** check box.

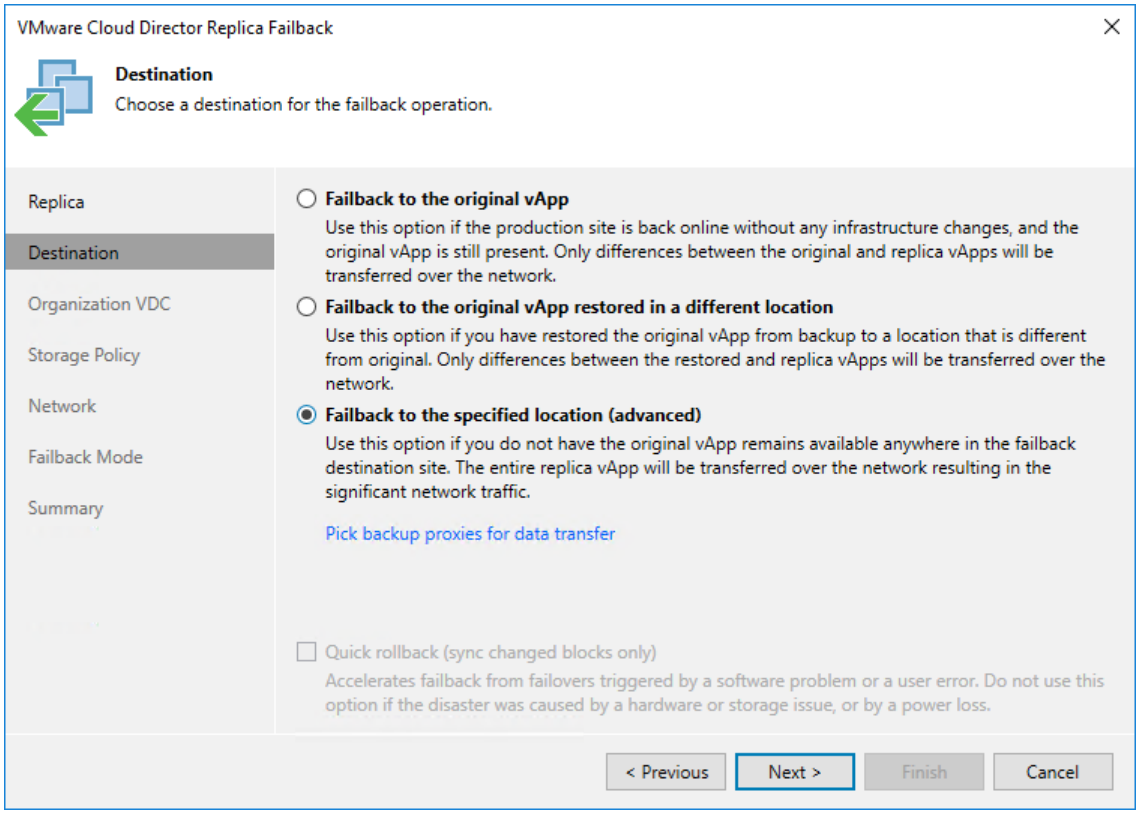


# Restoring Storage Policies

If the replicated vApp was associated with the storage policy, in the failback to original location scenario, Veeam Backup & Replication will associate the restored vApp with this storage policy.

When you click **Next**, Veeam Backup & Replication will check storage policies in the virtual environment and compare this information with the information about the replica storage policy. If the original storage policy has been changed or deleted, Veeam Backup & Replication will display a warning. You can select one of the following options:

- **Current** – the restored VM will be associated with the profile with which the source VM in the production environment is currently associated.
- **Default** – the restored VM will be associated with the profile that is set as default for the target datastore.
- **Stored** – the restored VM will be associated with the profile that was assigned to the source VM at the moment of replication.

For more information, see Storage Profiles.

## Step 4. Specify Organization VDCs

The **Organization VDC** step is available if you have selected the **Failback to the specified location** option at the **Destination** step.

At the **Organization VDC** step of the wizard, specify names for the restored vApps and the organization VDCs to which Veeam Backup & Replication will add restored vApps. To do this, select the necessary vApp and use the Name and VDC buttons.

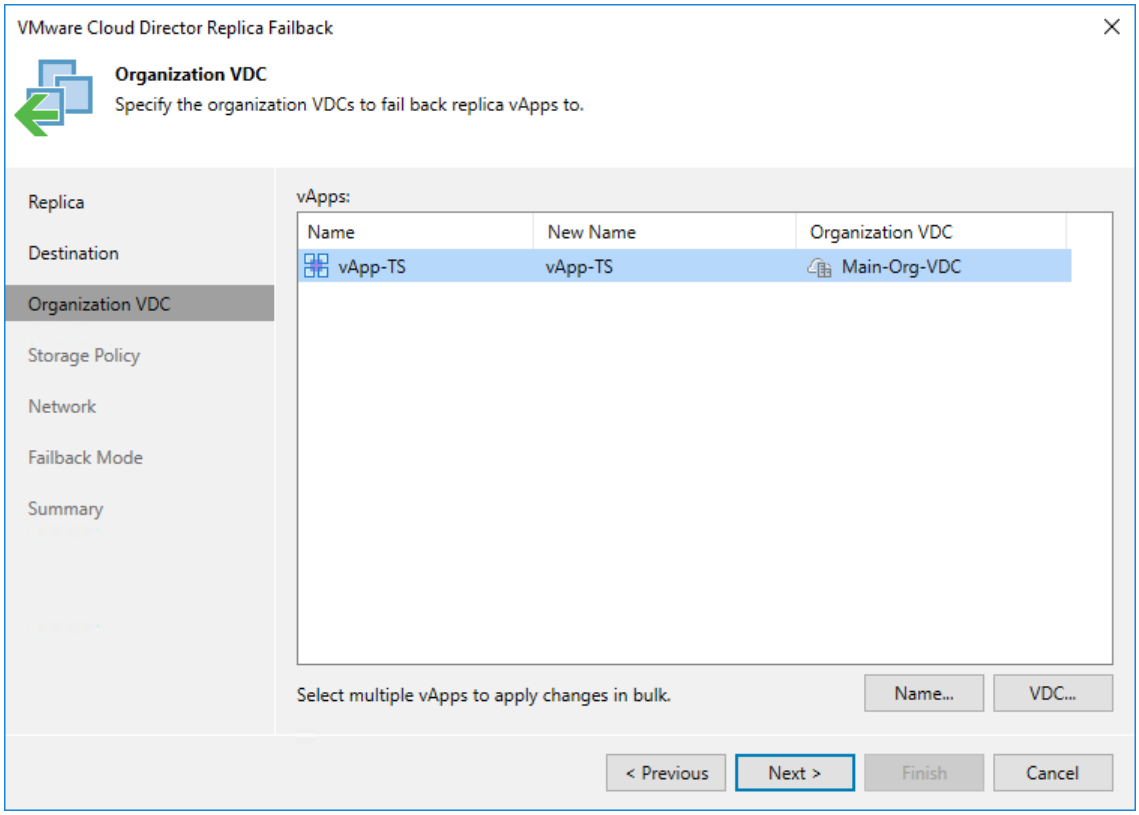

## Step 5. Specify Storage Policies

The **Storage Policy** step is available if you have selected the **Failback to the specified location** option at the **Destination** step.

At the **Storage Policy** step of the wizard, specify storage policies that Veeam Backup & Replication will apply to vApps that you want to restore:

1. In the **vApps** list, select vApps for which you want to change the policy and click **Policy**.
2. In the **Select storage policy** window select the policy that you want to apply.

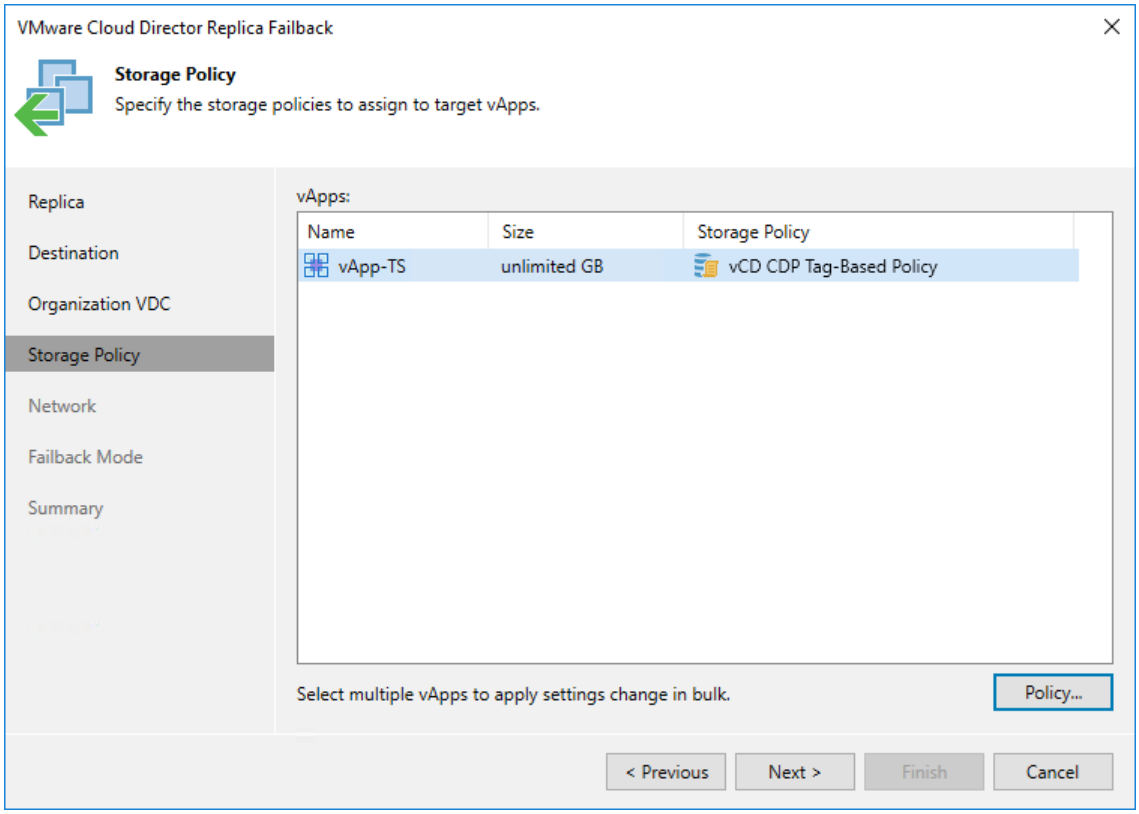

## Step 6. Configure Network Mapping

The **Network** step is available if you have selected the **Failback to the specified location** option at the **Destination** step.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the DR site to networks in the site where recovered vApps reside. Veeam Backup & Replication will use the network mapping table to update configuration files of VMs added to vApps on the fly, during the failback process.

To change networks to which the restored vApps will be connected:

1. In the **Network connections** list, select the necessary vApps and click **Network**.

   If vApps are connected to multiple networks, select the necessary network and click **Network**.

2. In the **Select network** window, select networks to which vApps must be connected after restore.

If you do not want to connect restored vApps to any virtual network, select the necessary vApps and click **Disconnect**.

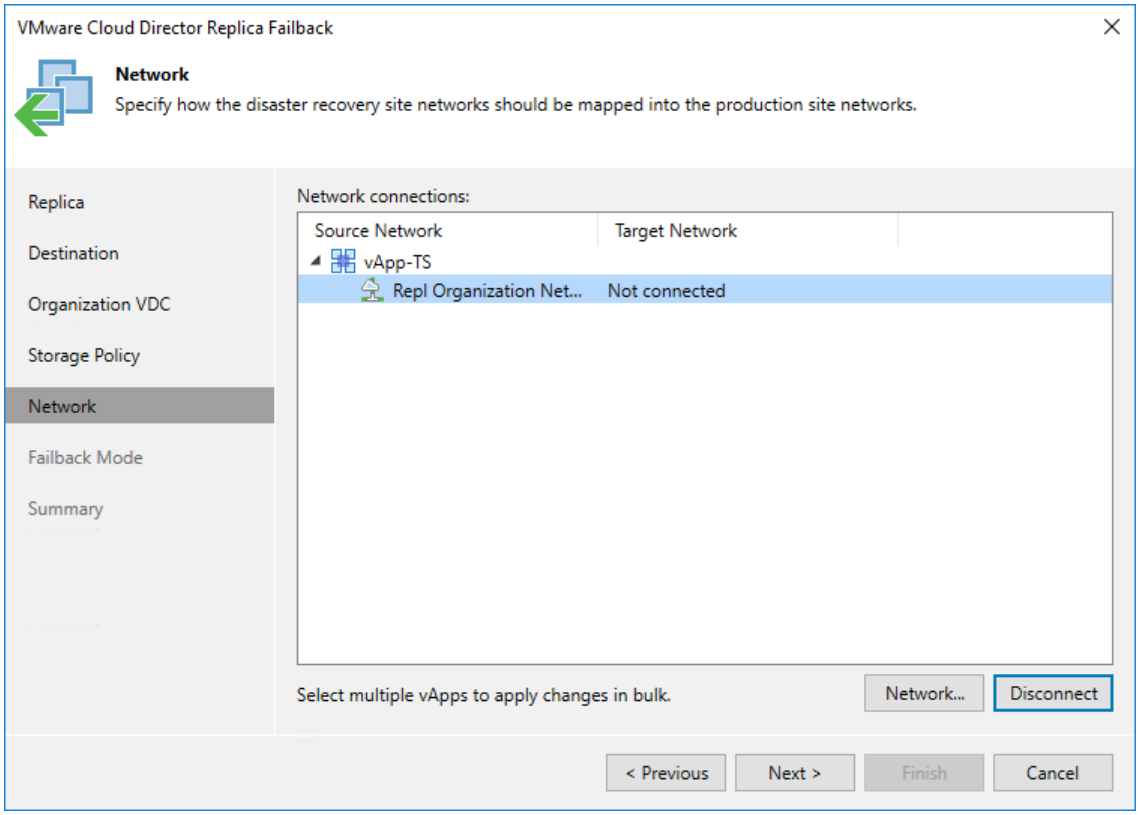

## Step 7. Select Target vApp

The **Target vApp** step is available if you have selected the **Failback to the original VM restored in a different location** option at the **Destination** step.

At the **Target vApp** step of the wizard, specify to which vApps you want to fail back from replicas. These vApps must be already restored from backups in the required location.

By default, Veeam Backup & Replication fails back the replica to the source vApps. If you want to specify the target vApp manually, perform the following steps:

1. Select a replica and click **Edit**.
2. In the **Selects Objects** window, select a vApp or vApp container to which you want to fail back.
3. Click **Add**.

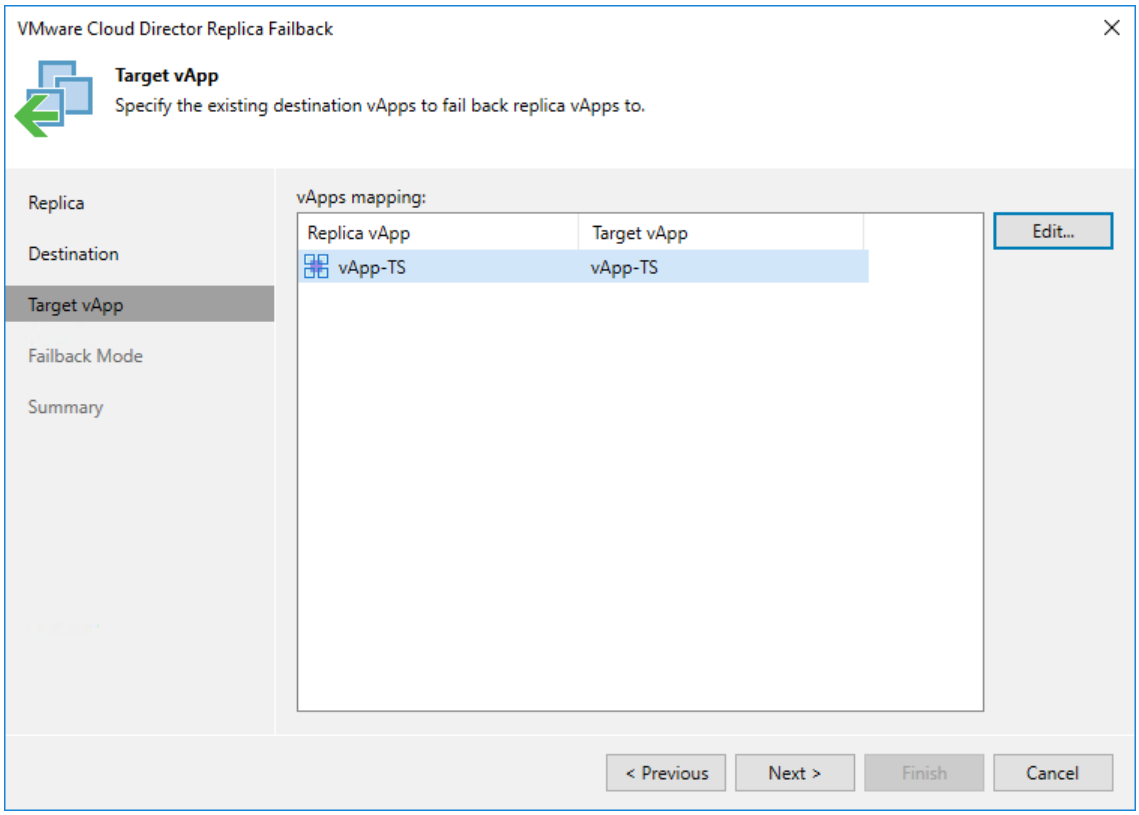

## Step 8. Schedule Switch to Production vApps

At the **Failback Mode** step of the wizard, specify when switch from replicas to production vApps must be performed:

- Select **Auto** if you want Veeam Backup & Replication to perform the switch automatically right after the state of the production vApps is synchronized with the state of their replicas.
- Select **Scheduled** if you want Veeam Backup & Replication to perform the switch at a specific time.
- Select **Manual** if you want to perform the switch manually.

If you select the **Scheduled** or **Manual** option, you can further reset/set the scheduled time or switch to the production vApps manually. For more information, see Changing Switching Time and Switching to Production vApps Manually.

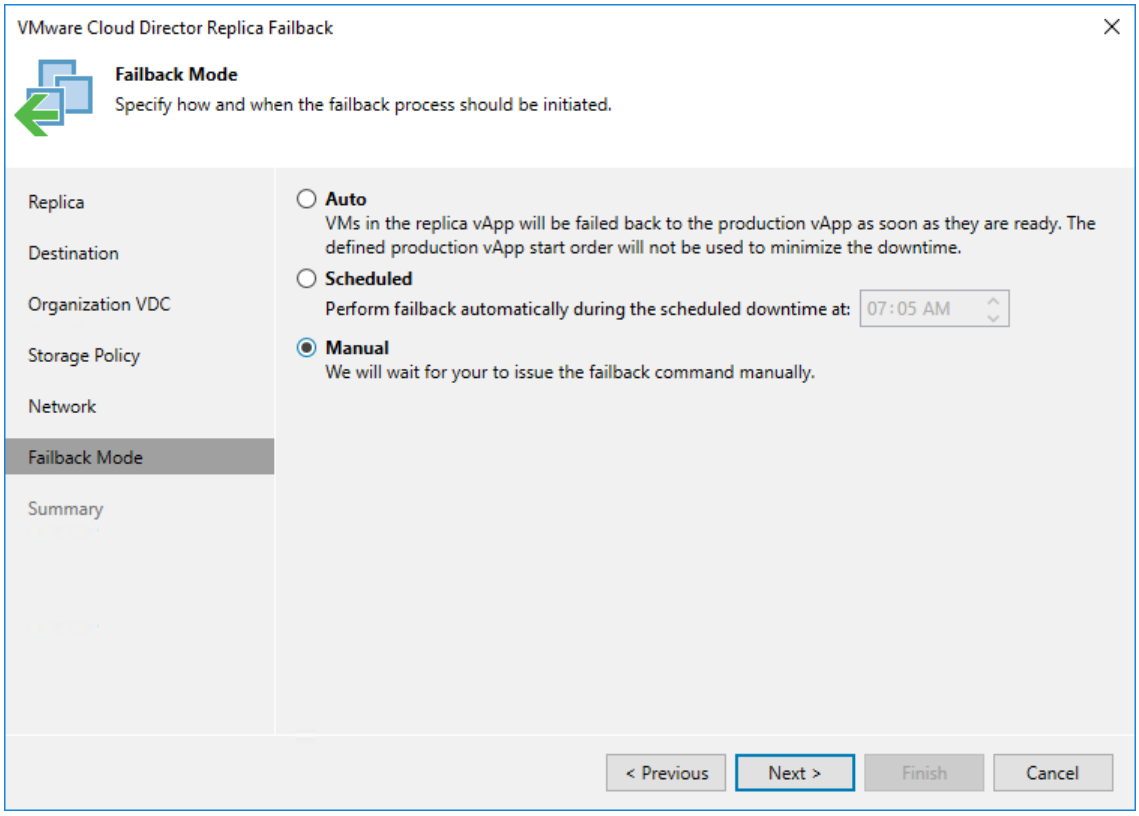

### Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured failback settings. If you want to power on the production VMs right after the switch to production operation is performed, select the **Power on target vApp after failback** check box. Then click **Finish**.

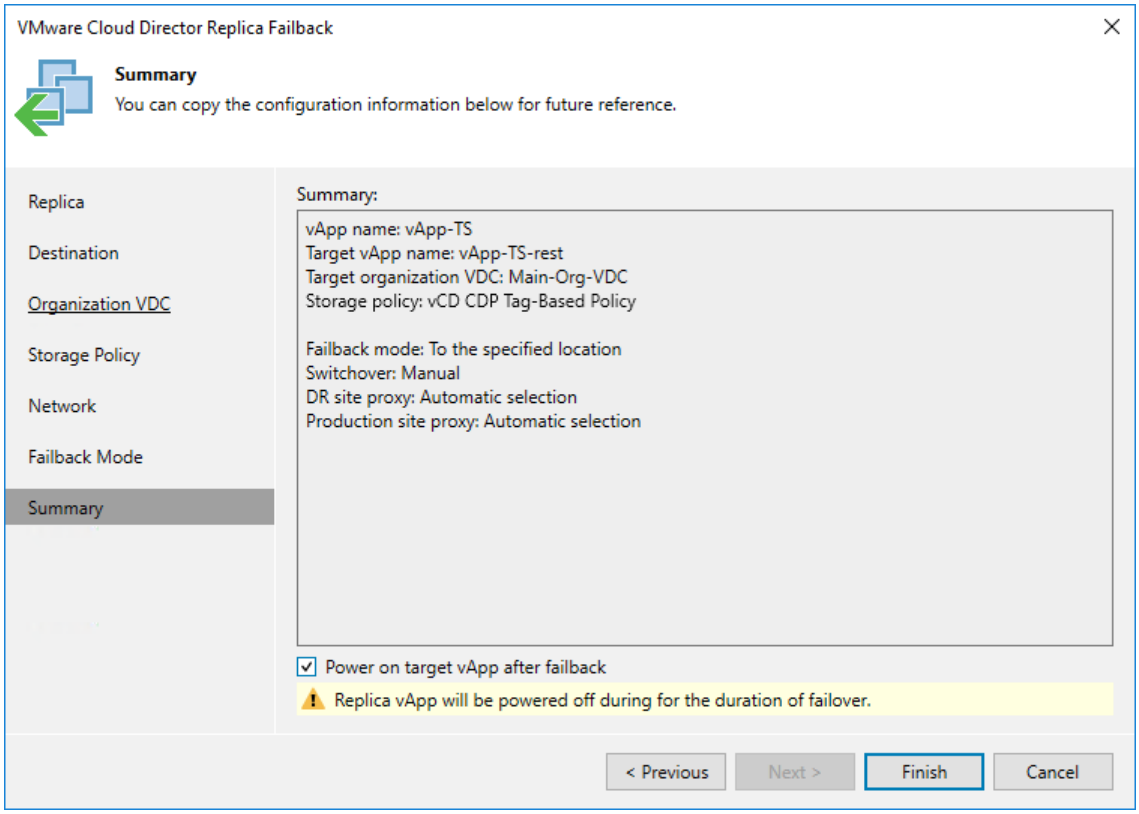


## What You Do Next

Failback is an intermediate step that needs to be finalized. You can finalize failback in the following ways:

- Commit failback
- Undo failback

## Changing Switching Time

The following instructions apply if you have selected to switch from replicas to production vApps manually or at the scheduled time at at the **Failback Mode** step of the **Failback** wizard.

To change the time when the switch from replicas to production vApps must be performed:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the vApp in the *Ready to switch* state and select **Switchover Time** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Change switchover time.**

If the switching time operation failed, you can retry this operation again. To perform a retry, in the working area, select the necessary vApp and select **Retry Switchover Time** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry switchover time.**

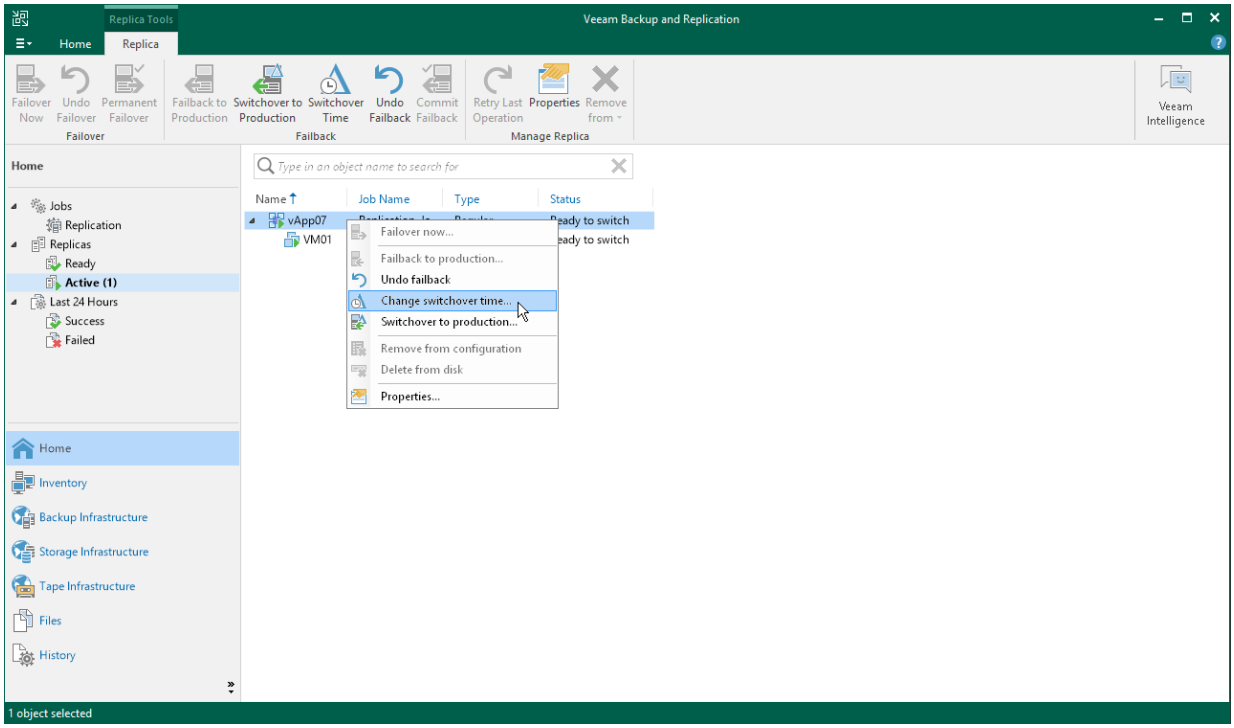


## Switching to Production vApps Manually

The following instructions apply if you have selected to switch from replicas to production vApps manually at the **Failback Mode** step of the **Failback** wizard.

To switch to a production vApp from its replica, do the following:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary vApp and select **Switchover to production** on the ribbon. As an alternative, you can right-click the necessary vApp and select **Switchover to production**.

If the switch to production operation failed, you can retry this operation again. To perform a retry, in the working area, select the necessary vApp and select **Retry Switchover to Production** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry switchover to production**.

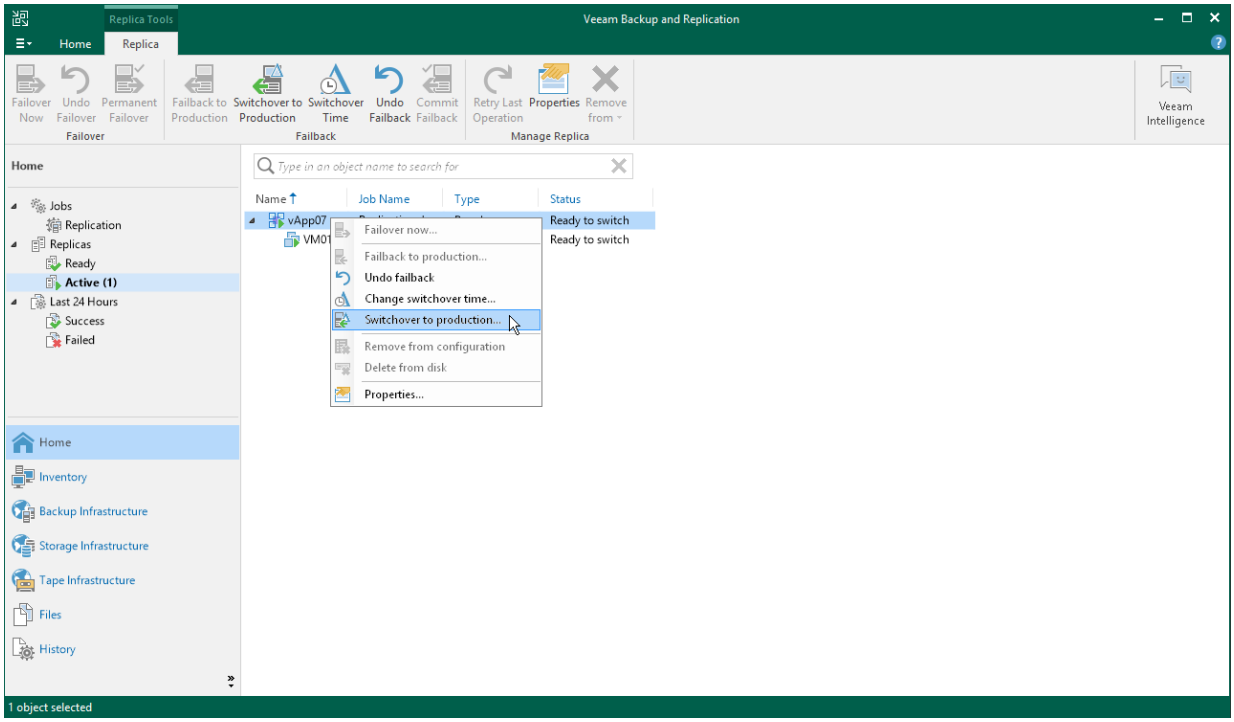


# What You Do Next

After you switch to the production VM, you must finalize failback. You can finalize failback in the following ways:

- Commit failback
- Undo failback

## Performing Failback Retry

The failback retry option is necessary when failback of vApps fails with the *Incomplete* state. When you perform a retry, Veeam Backup & Replication restarts failback only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, failback takes less time and does not consume as many resources as when processing the whole vApp.

To retry failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary vApp and select **Retry Failback** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry failback**.

# Failback Undo

Failback undo is one of the ways to finalize failback. You can use this option if the vApp to which you failed back (the production vApp) works in a wrong way and you want to get back to the replica.

The failback undo operation is performed in the following way:

1. Veeam Backup & Replication powers off the production vApp.
2. Veeam Backup & Replication reverts the replica to its pre-failback state.
3. Veeam Backup & Replication powers on the replica and changes the replica state from *Failback* or *Ready to Switch* to *Failover*.

## Undoing Failback

For more information on failback undo, see Failover and Failback for Cloud Director and Failback Undo.

To undo failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary replica and click **Undo Failback** on the ribbon. Alternatively, you can right-click the necessary replica and select **Undo Failback**.

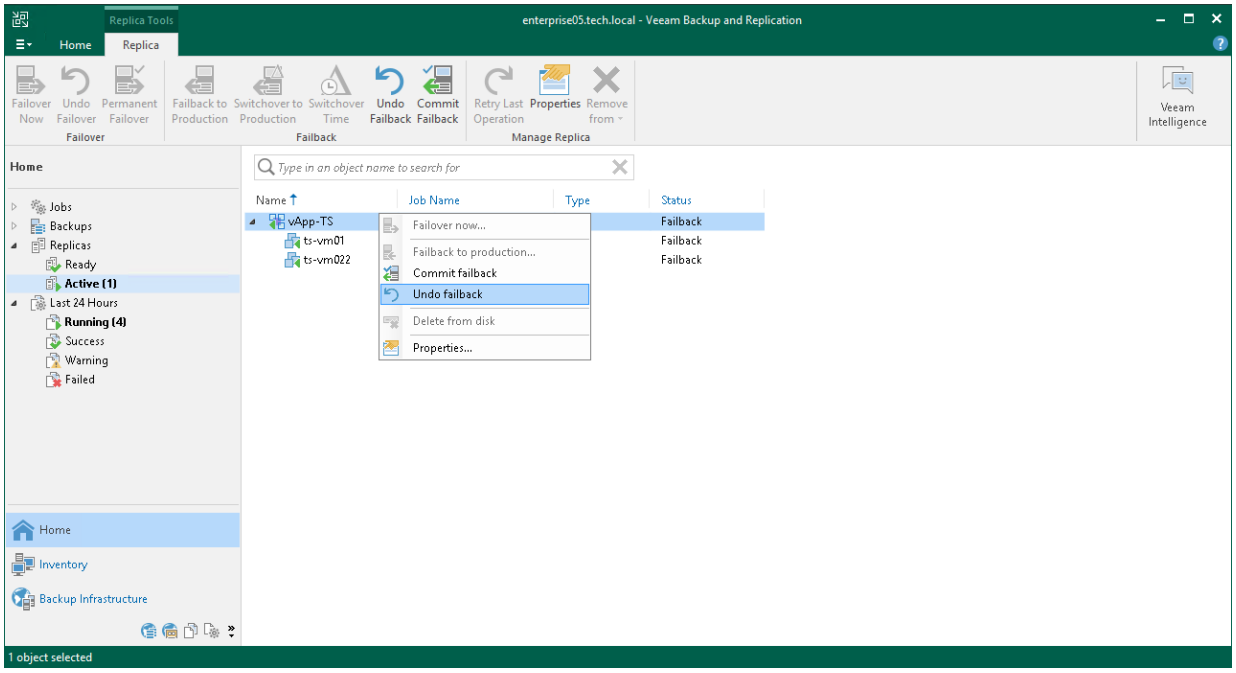


## Performing Failback Undo Retry

If the failback undo operation failed, you can retry this operation. When you perform a retry, Veeam Backup & Replication restarts the failback undo operation only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, the failback undo operation takes less time and does not consume as many resources as when processing the whole vApp.

To perform a retry:

1. Open the **Home** view, in the inventory pane, navigate to the **Replicas** > **Active** node.
2. In the working area, select the necessary vApp and select **Retry Undo Failover** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry undo failback**.

# Failback Commit

Failback commit is one of the ways to finalize failback. When you commit failback, you confirm that the vApp to which you failed back (the production vApp) works as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production vApp.

> **NOTE**
>
> If during failback, you have selected to switch to the production VM manually, you must first perform the switchover.

The failback commit operation is performed in the following way:

1. Depending on whether you have failed back to the source vApp or recovered vApp:
   - If you have failed back to a vApp recovered from a backup or replica, Veeam Backup & Replication reconfigures all existing jobs where the source vApp is present and adds the source vApp to the list of exclusions. The recovered vApp takes the role of the source vApp and is included into all jobs instead of the excluded vApp. When the VMware Cloud Director replication process starts, Veeam Backup & Replication processes the recovered vApp instead of the former source vApp.
   - If you have failed back to the source vApp, the replication job or policy is not reconfigured. When the replication process starts, Veeam Backup & Replication still processes the source vApp.
2. Veeam Backup & Replication changes the state of the replica from *Failback* to *Ready*.

During failback commit, the failback delta disk that saves the pre-failback state of a replica is not deleted. Veeam Backup & Replication uses this delta disk as an additional restore point for replica. With the pre-failback delta disk, Veeam Backup & Replication needs to transfer fewer changes and therefore puts less load on the network when replication activities are resumed.

## Committing Failback

For more information on failback commit, see Failover and Failback for Cloud Director and Failback Commit.

To commit failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. In the working area, select the necessary replica and click **Commit Failback** on the ribbon. As an alternative, you can right-click the replica and select **Commit failback**.

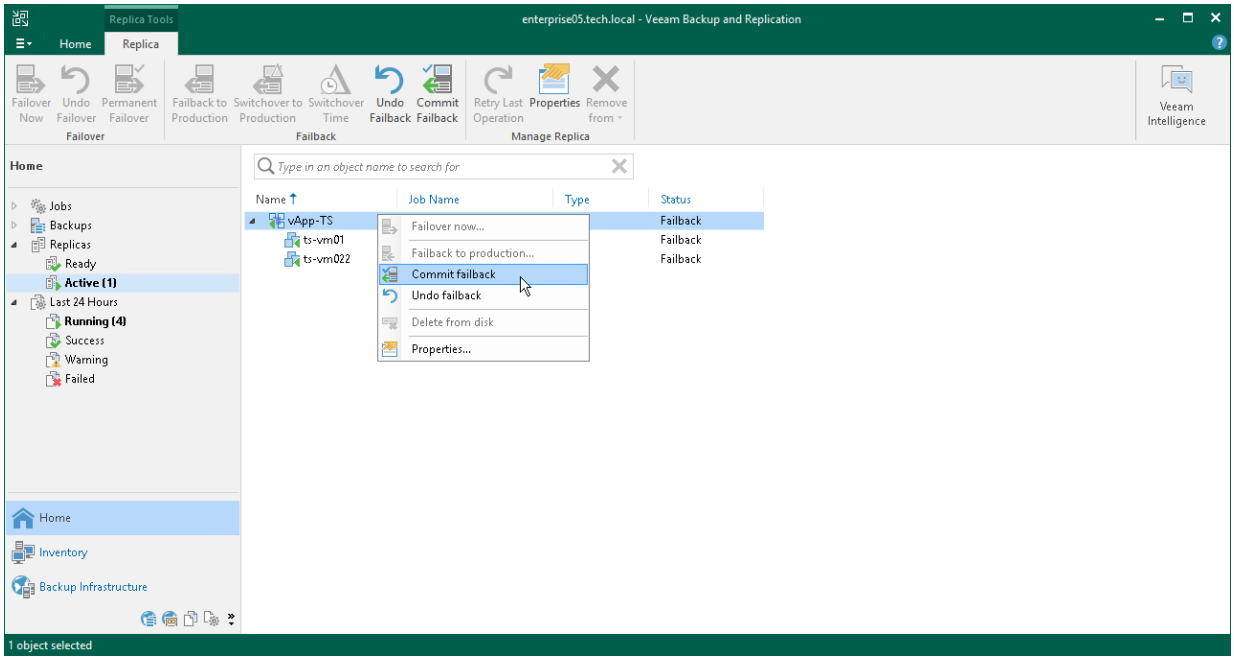


# Performing Failback Commit Retry

If the failback commit operation failed, you can retry this operation. When you perform a retry, Veeam Backup & Replication restarts the failback commit operation only for the failed VMs that are added to vApps. Veeam Backup & Replication does not process VMs that have been processed successfully. As a result, the failback commit operation takes less time and does not consume as many resources as when processing the whole vApp.

To perform a retry:

1. Open the **Home** view, in the inventory pane, navigate to the **Replicas** > **Active** node.
2. In the working area, select the necessary vApp and select **Retry Commit Failback** on the ribbon. Alternatively, you can right-click the necessary vApp and select **Retry commit failback**.

# Data Recovery for VMware Cloud Director

Veeam Backup & Replication enables full-fledged restore of VMs to VMware Cloud Director. You can restore separate VMs to vApps, as well as VM data.

For restore, Veeam Backup & Replication uses VM metadata saved to a backup file and restores specific VM attributes. As a result, you get a fully-functioning VM in VMware Cloud Director, do not need to import the restored VM to VMware Cloud Director and adjust the settings manually.

Backed-up objects can be restored to the same VMware Cloud Director hierarchy or to a different VMware Cloud Director environment. Restore options include:

- Instant Recovery
- Full restore for vApps and VMs
- Restore of VM disks
- Restore of VM files
- Guest OS file restore for VMs

# VM Recovery

VMware Cloud Director VM recovery includes the following methods:

- **Instant Recovery to Cloud Director vApp** — to instantly recover Cloud Director VMs directly from compressed and deduplicated backup files to vApps. Instant Recovery helps improve recovery time objectives (RTO), minimize disruption and downtime of production workloads. However, Instant Recovery provides for VMs "temporary spares" with limited I/O performance. To provide the recovered VMs full I/O performance, you must finalize Instant Recovery — migrate the recovered VMs to production environment. If you do not want to migrate the recovered VM, you can stop publishing it. This removes the recovered VM.

  For more information, see Performing Instant Recovery to Cloud Director vApp.

- **Instant Recovery to VMware vSphere** — to instantly recover Cloud Director VMs as regular VMware vSphere VMs. In this case, the VM will be restored at the level of the underlying vCenter Server, and the Instant Recovery process will be the same as for regular VMware vSphere VMs.

  For more information, see Performing Instant Recovery to VMware vSphere.

- **Restore of VMs to Cloud Director vApp** — to recover entire VMs to vApps. When you recover VMs, you extract VM images from backups to the production storage. This restore takes more resources and time to complete than Instant Recovery to Cloud Director vApp but recovers VMs with full I/O performance. You also do not need to perform additional steps to finalize the recovery process.

  For more information, see Restoring Entire VMs into Cloud Director vApp.

- **Entire VM restore to VMware vSphere** — to recover Cloud Director VMs as regular VMware vSphere VMs. In this case, the VM will be restored at the level of the underlying vCenter Server, and the entire VM restore process will be the same as for regular VMware vSphere VMs.

  For more information, see Restoring Entire VMs to VMware vSphere.

## How Restore of Regular and Standalone VMs to VMware Cloud Director Works

Veeam Backup & Replication lets you restore regular VMs that are part of vApps and standalone VMs that were created in the VMware Cloud Director tenant portal.

When you restore regular or standalone VMs back to the VMware Cloud Director hierarchy, the restore process includes the following steps:

1. Veeam Backup & Replication uses the captured vApp metadata to define the vApp settings and VM original location in the VMware Cloud Director hierarchy.
2. Veeam Backup & Replication restores VMs from the backup file to their original location or to a different location. Additionally, Veeam Backup & Replication restores all VM settings.

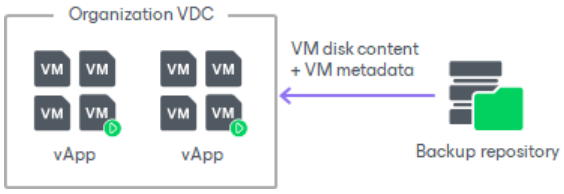

# How Restore of Linked Clone VMs to VMware Cloud Director Works

Veeam Backup & Replication lets you restore linked clone VMs – VMs that were deployed from a VM template using the fast provisioning technology. There are several mechanisms for processing linked clone VMs.

# Restore of Existing VMs

If you are restoring a VMware Cloud Director linked clone VM that exists in the VMware Cloud Director hierarchy, the restore process includes the following steps:

1. Veeam Backup & Replication uses the captured vApp metadata to define the initial settings of the VM.
2. Veeam Backup & Replication calculates a signature for the consolidated VM disk in the backup file (containing the VM template data and data of the delta file) and the signature for the VM existing in VMware Cloud Director. Veeam Backup & Replication then compares the disk signatures to define what data blocks have changed.
3. Veeam Backup & Replication restores only changed data blocks from the backup file and writes them to the user delta file.

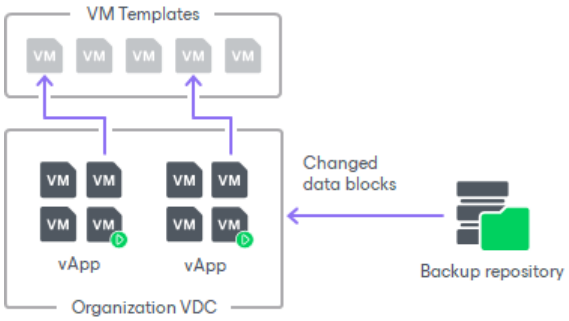


# Restore of Deleted VMs

If you are restoring a VM that no longer exists in VMware Cloud Director hierarchy, the restore process includes the following steps:

1. Veeam Backup & Replication uses VMware Cloud Director to create a new linked clone VM from the VM template that the user selects. The new VM has a blank user delta file.
2. Veeam Backup & Replication calculates a signature for the consolidated VM disk in the backup file (containing the VM template data and data of the delta file) and the signature for the created VM in VMware Cloud Director. Veeam Backup & Replication then compares the disk signatures to define what data blocks need to be restored.
3. Veeam Backup & Replication restores only those data blocks that need to be restored from the backup file and writes them to the blank user delta file.

By default, Veeam Backup & Replication links the VM to the same VM template that was used by the original VM. During restore, Veeam Backup & Replication checks the settings of the VM template to which the restored VM is linked: verifies connection settings, makes sure the disk size coincide and so on.

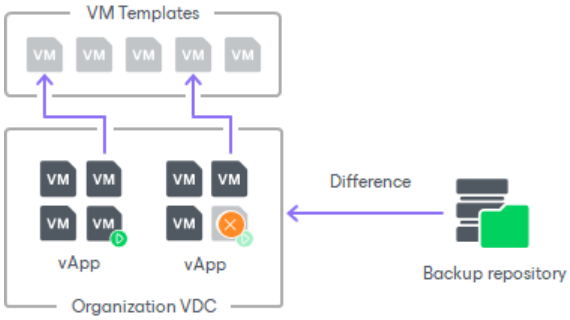


# Restore of Linked Clone VMs as Regular VMs

In some cases, Veeam Backup & Replication can restore a VM from a backup file as a regular VM. This type of restore is accomplished in the following situations:

- You have intentionally chosen to restore a linked clone VM as a regular VM.
- You are restoring a VM to the organization VDC which has the fast provisioning option disabled.
- A VM template to which the restored VM should be linked is not accessible in the location to which the VM is restored.

In this case, Veeam Backup & Replication uses the same algorithm as for restore of full VMs in the virtual environment. It retrieves the data of the consolidated VM disk from the backup file and restores the VM in the VMware Cloud Director hierarchy.

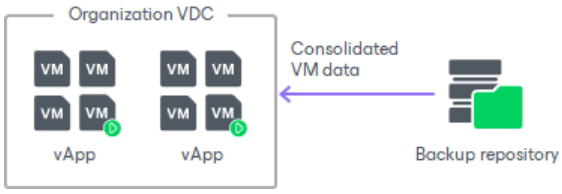


# Performing Instant Recovery to Cloud Director vApp

With Instant Recovery, you can immediately start a VM from a backup file stored in the backup repository. Instant Recovery accelerates the restore process, allows you to improve RTOs and decrease downtime of production VMs.

When you instantly recover a VM to VMware Cloud Director, Veeam Backup & Replication uses the vPower NFS datastore, just as with other VMware VMs. To import the VM to the vApp, Veeam Backup & Replication needs to associate the vPower NFS datastore with some storage policy. To do this, Veeam Backup & Replication creates for the underlying vCenter Server an auxiliary storage policy — *Veeam-InstantVMRecovery*, and displays it in VMware Cloud Director.

The created storage policy is added to the Provider VDC and organization VDC hosting the vApp to which the VM is restored. When the vPower NFS datastore is mounted to the ESXi host, the vPower NFS datastore is associated with the *Veeam-InstantVMRecovery* storage policy. After that, the VM is instantly restored in a regular manner and imported to the selected vApp.

When an Instant Recovery session is finished, the storage policy is not deleted from the Provider VDC, it remains on vCenter Server. This helps speed up all subsequent Instant Recovery operations. However, the storage policy is deleted from the organization VDC as organization VDC settings can be accessed only by organization administrators.

Before you start Instant Recovery, check prerequisites. Then use the **Instant Recovery to VMware Cloud Director** wizard to recover the necessary VM.

## Before You Begin

Before you perform Instant Recovery, check the following prerequisites:

- You can restore only those VMs whose placement policy is the same as the default placement policy of the target organization VDC (VDC where the vApp to which you restore VMs reside). For more information on placement policies, see VMware Docs.
- You can perform Instant Recovery for a VM that has been successfully backed up at least once.
- You must have at least 10 GB of free disk space on the datastore where write cache folder is located. This disk space is required to store virtual disk updates for the restored VM.

  By default, Veeam Backup & Replication writes virtual disk updates to the `IRCache` folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`. The write cache is not used when you select to redirect virtual disk updates to a VMware vSphere datastore when configuring the job.
- If you are recovering a VM to the production network, make sure that the original VM is powered off to avoid conflicts.
- If you want to scan VM data for viruses, check the secure restore requirements and limitations.
- Consider that during Instant Recovery, Veeam Backup & Replication restores standalone VMs as regular VMware Cloud Director VMs.
- [For migration of recovered VMs] You must disable the VM discovery option in VMware Cloud Director settings. For more information on where you can change the option, see VMware Docs.

## Step 1. Launch Instant Recovery to VMware Cloud Director Wizard

To launch the **Instant Recovery to VMware Cloud Director** wizard, do one of the following:

- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Select the machine that you want to restore and click **Instant Recovery** > **VMware Cloud Director** on the ribbon.
  - Right-click the machine that you want to restore and select **Instant recovery** > **VMware Cloud Director**.
- On the **Home** tab, click **Restore** > **VMware Cloud DIrector**. In the **Restore** window, select **Restore from backup** > **VM restore** > **Entire VM restore** > **Instant recovery** > **Instant recovery to VMware Cloud Director**.
- Open the **Inventory** view. On the **View** tab, click **Cloud Director View**. In the inventory pane, expand the **vCloud Director** hierarchy. In the working area, right-click the VM you want to restore and select **Restore** > **Instant recovery** > **VMware Cloud Director**.

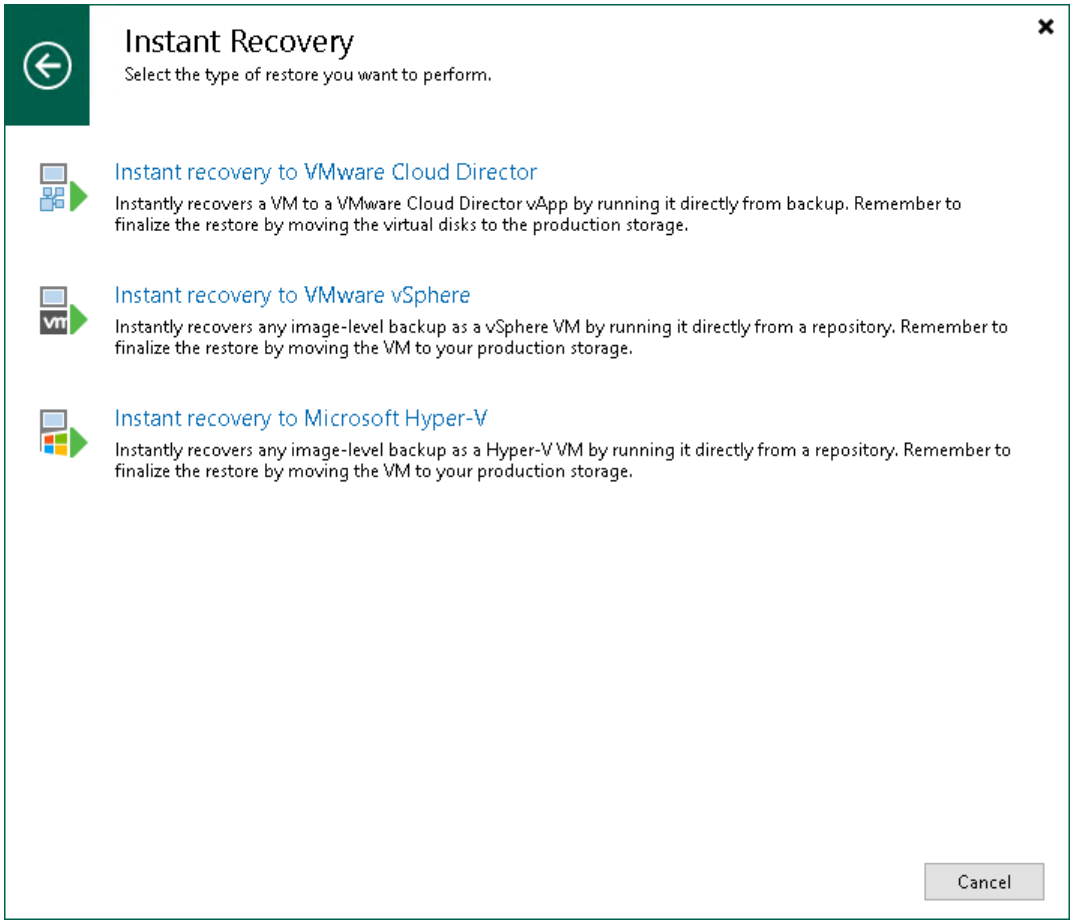

# Step 2. Select VMs

At the **Virtual Machine** step of the wizard, select a VM that you want to recover.

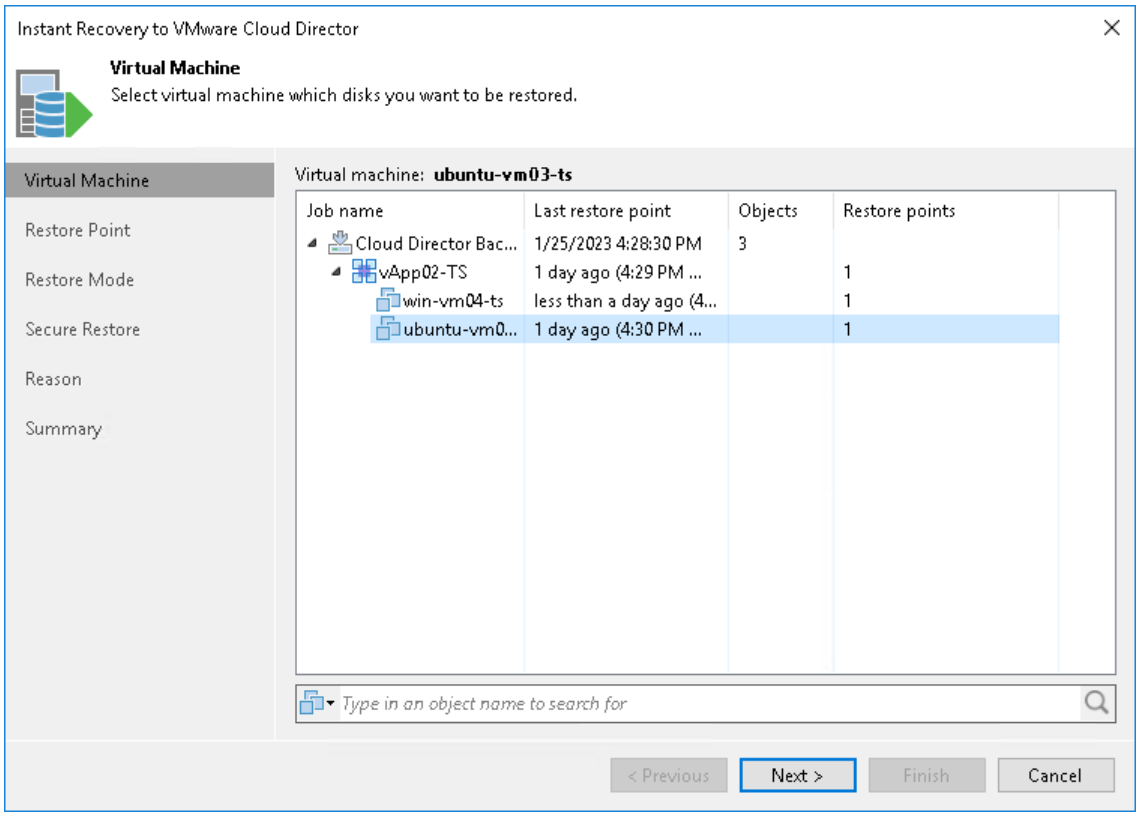

## Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select the restore point for the VM.

In the **Location** column, you can view a name of a backup repository where a restore point resides.

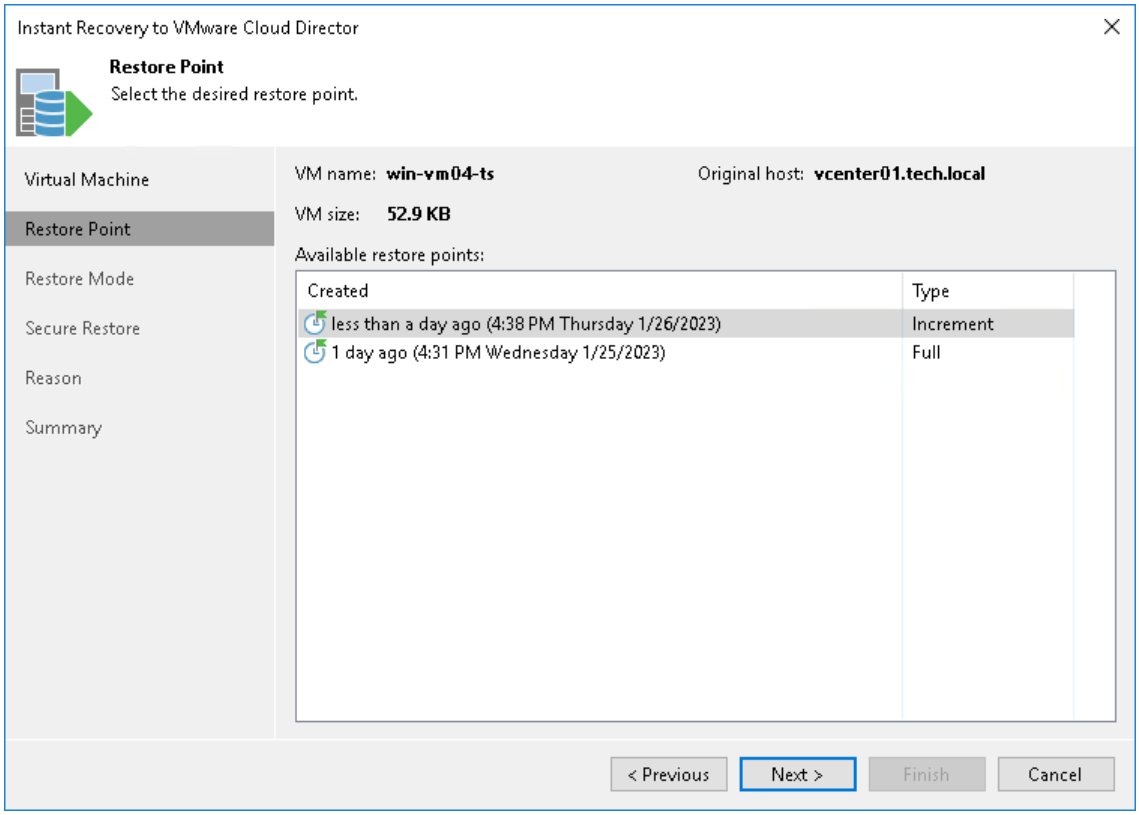

## Step 4. Select Restore Mode

At the **Restore Mode** of the wizard, choose the necessary restore mode.

1. Choose a restore mode:
   - Select **Restore to the original location** if you want to restore the VM with its initial settings to its original location. If this option is selected, you will pass directly to the Reason step of the wizard.
   - Select **Restore to a new location, or with different settings** if you want to restore the VM to a different location or with different settings (such as vApp, VM name, network settings and so on). If this option is selected, the **Instant Recovery** wizard will include additional steps for customizing VM settings.
2. Select the **Restore VM tags** check box if you want to restore tags that were assigned to the original VM, and assign them to the restored VM. Veeam Backup & Replication will restore the VM with original tags if the following conditions are met:
   - The VM is restored to its original location.
   - The original VM tag is still available on the source vCenter Server.

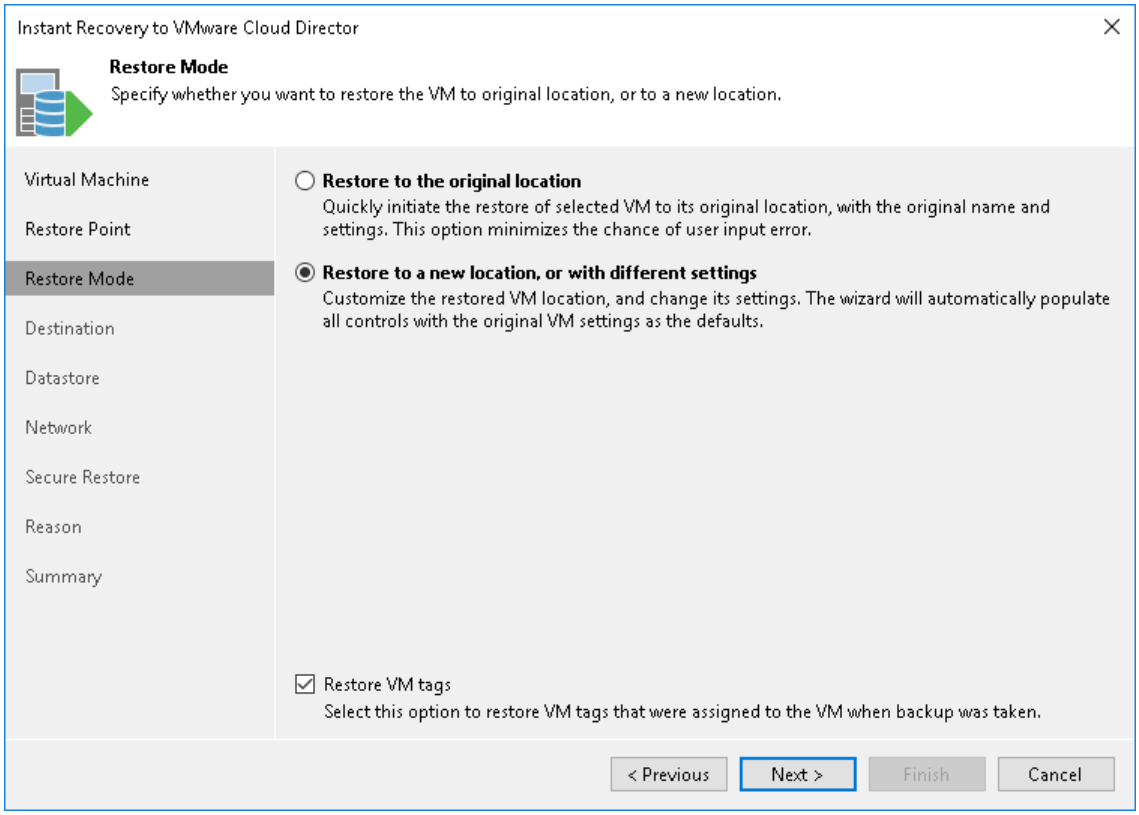

# Step 5. Select Destination for Restored VM

The **Destination** step of the wizard is available if you have chosen to change the location and settings of the restored VM.

Select a destination and specify a name for the restored VM:

1. In the **vApp** field, specify a vApp to which the VM must be restored. By default, Veeam Backup & Replication restores the VM to its original vApp.
2. In the **Restored VM name** field, enter a name under which the VM must be restored and registered. By default, Veeam Backup & Replication uses the original name of the VM. If you are restoring the VM to the same vApp where the original VM is registered and the original VM still resides there, it is recommended that you change the VM name to avoid conflicts.

> **NOTE**
>
> Veeam Backup & Replication checks the lease term for the vApp to which the VM is restored. In case the lease period has expired, the lease will be automatically updated.

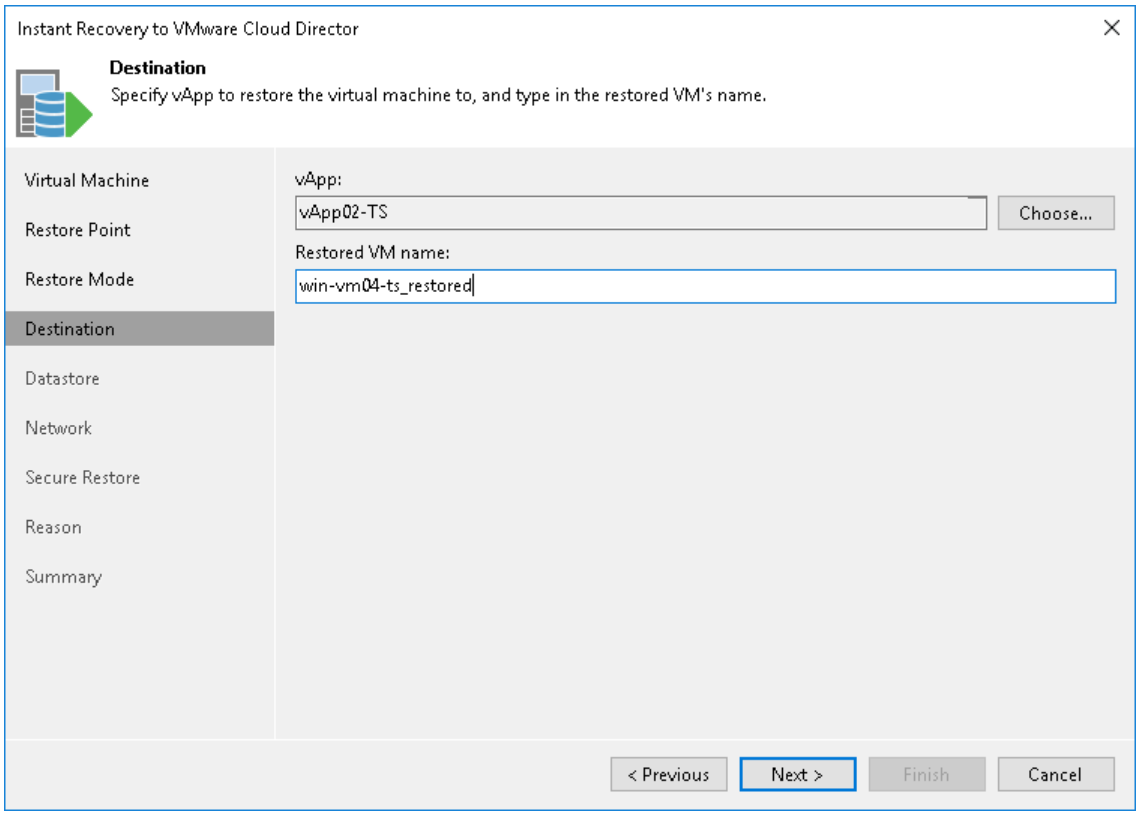

# Step 6. Select Destination for Virtual Disk Updates

The **Datastore** step of the wizard is available if you have chosen to change the location and settings of the restored VM.

Select the location for holding the VM disk changes when the VM is restored. By default, disk changes are stored directly on the vPower NFS server. However, you can store disk changes on any datastore in your virtual environment.

To select a datastore:

1. Select the **Redirect virtual disk updates** check box.
2. From the **Datastore** list, choose the necessary datastore. You can select only a datastore that is available in the organization VDC hosting the vApp to which the VM is restored.

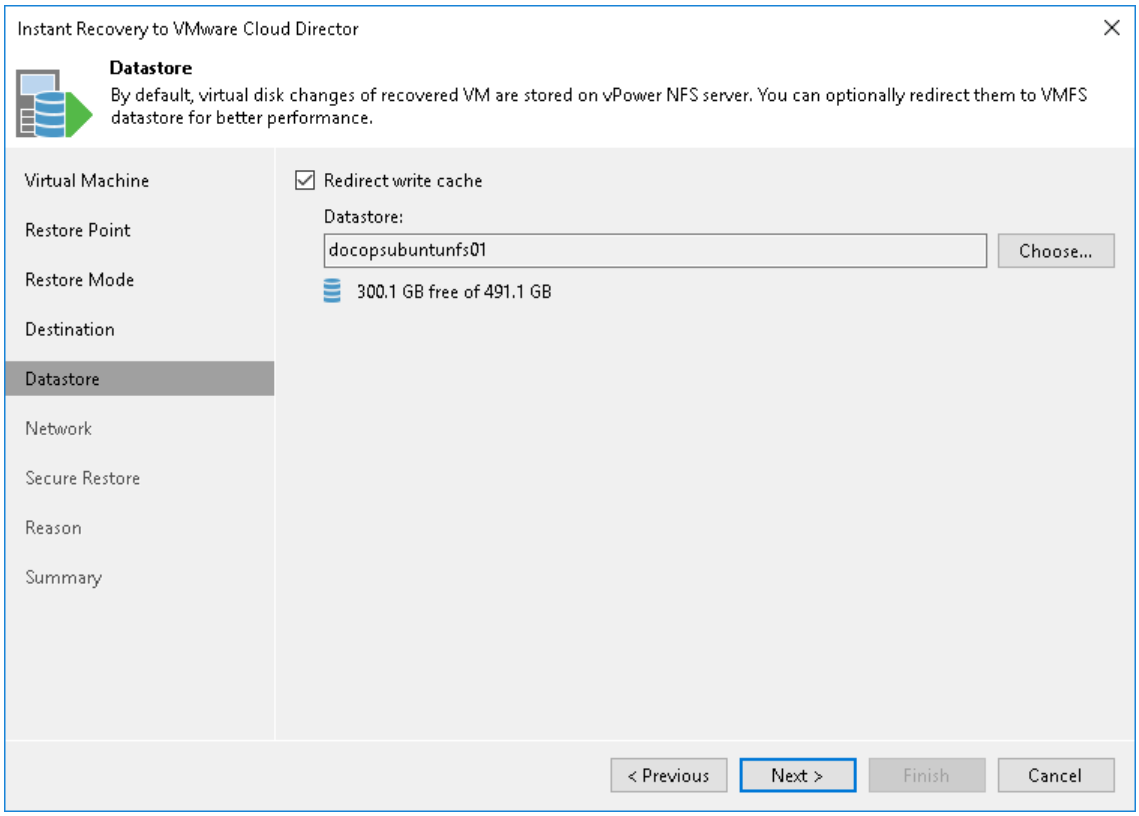

# Step 7. Select Destination Network

The **Network** step of the wizard is available if you have chosen to change the location and settings of the restored VM.

To select networks to which restored VMs must be connected:

1. Select a VM in the list and click **Network**.
2. The **Select Network** window displays all networks that are configured for the destination vApp. From the list of available networks, choose a network to which selected VM should have access upon restore.

   To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.
3. To prevent the restored VM from accessing any network, select it in the list and click **Disconnect**.

Veeam Backup & Replication maps the network settings you define and network settings of the original VM. If necessary, Veeam Backup & Replication makes changes to the network settings of the recovered VM. For example, if the original VM was connected to the network using the static IP mode and you have selected to connect a recovered VM to a network using the dynamic IP mode, Veeam Backup & Replication will change the network settings to the dynamic mode.

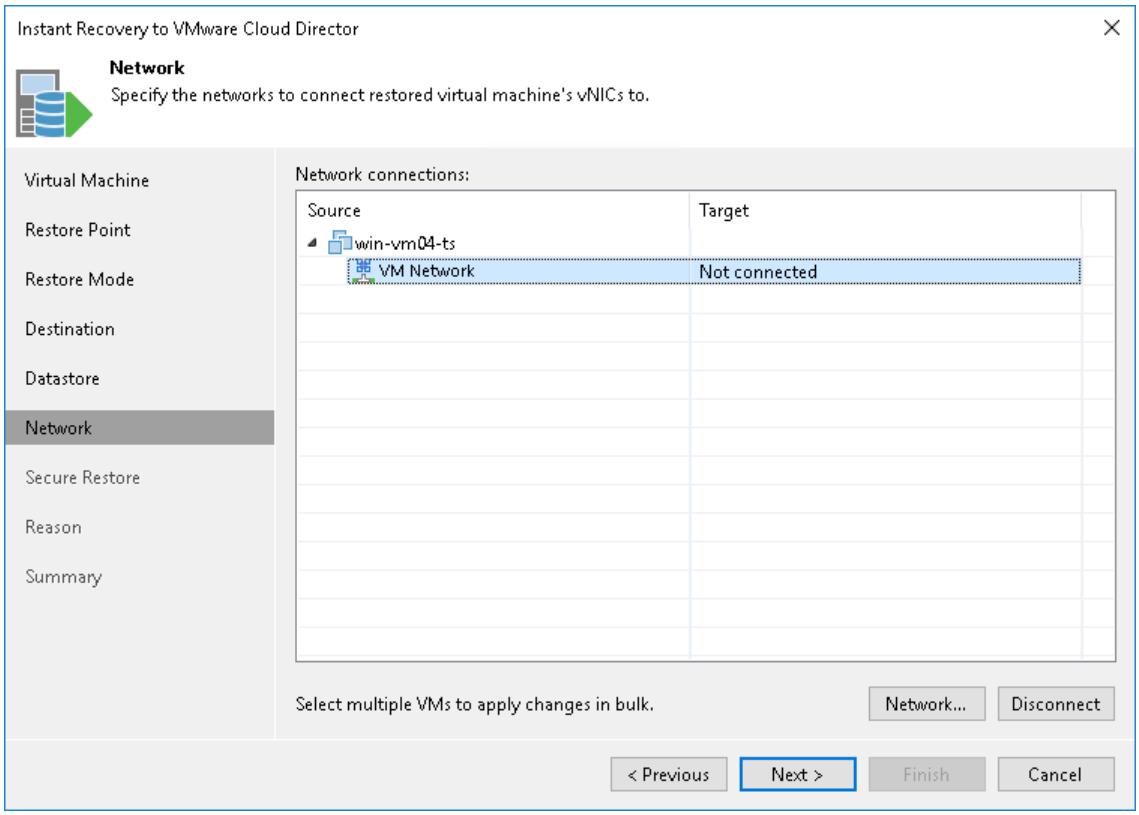

# Step 8. Specify Secure Restore Settings

This step is available if you restore Microsoft Windows VMs.

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Select which action Veeam Backup & Replication will take if scan finds a virus threat:
   - **Proceed with recovery but disable network adapters.** Select this action if you want to restore the VM with disabled network adapters (NICs).
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.

3. Select the **Continue scanning all remaining files after the first occurrence** check box if you want the antivirus software to continue the VM data scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Antivirus Scan Results.

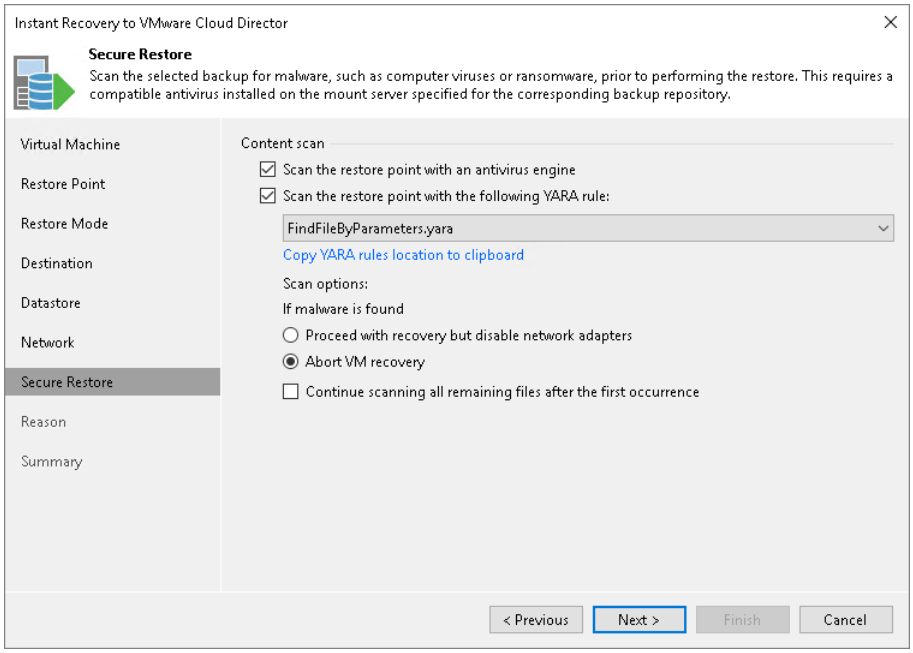

# Step 9. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for performing Instant Recovery of the VM. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

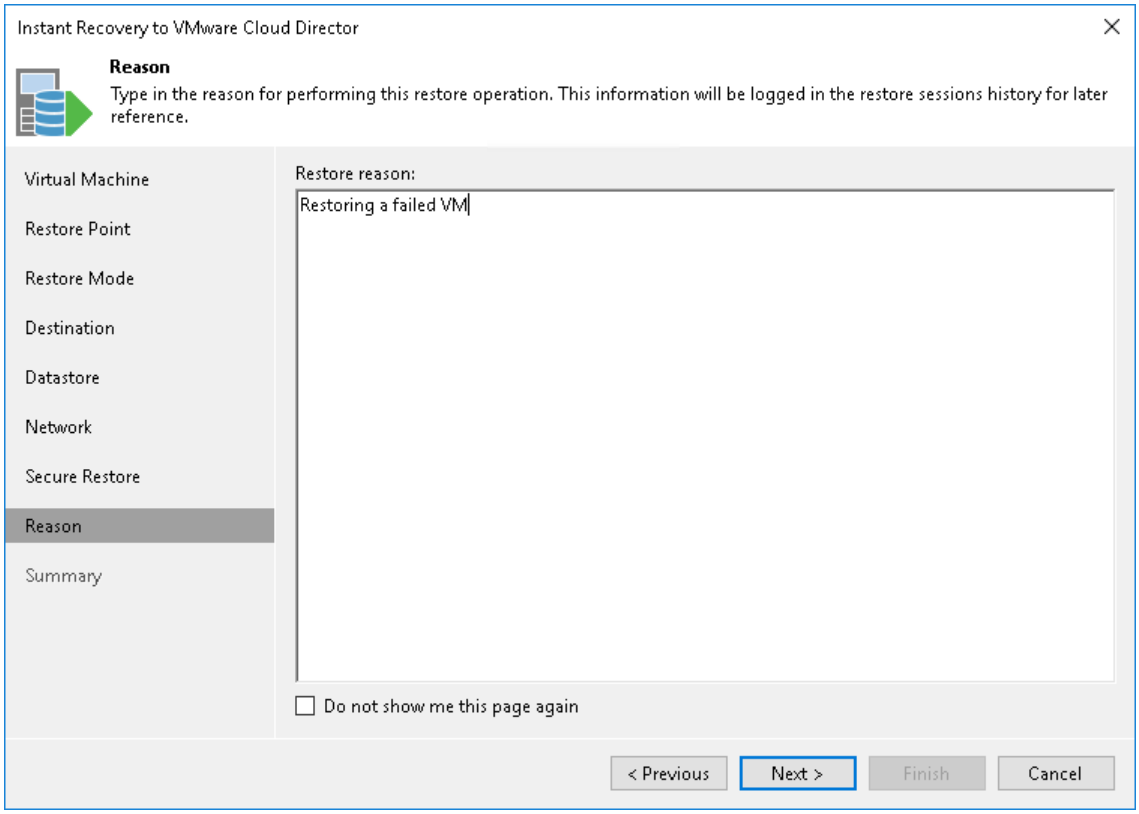

## Step 10. Verify Instant Recovery Settings

At the **Summary** step of the wizard, specify additional settings for Instant Recovery:

1. If you want to start the recovered VM, select the **Power on VM automatically** check box.
2. Check the specified settings of Instant Recovery and click **Finish**. Veeam Backup & Replication will recover the selected VM in the specified destination.

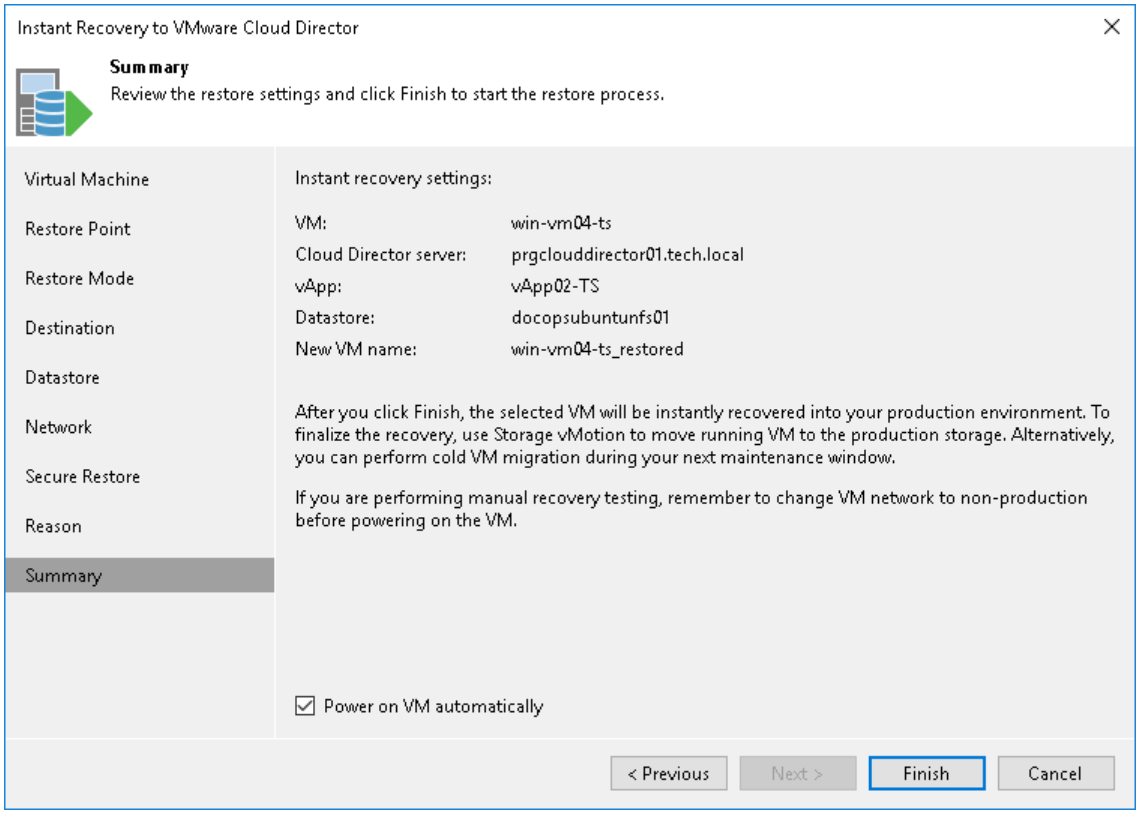


# What You Do Next

Finalizing Instant Recovery

# Performing Instant Recovery to VMware vSphere

The process of Instant Recovery for VMware Cloud Director VMs does not differ from the regular Instant Recovery process. For more information, see Performing Instant Recovery to VMware vSphere.

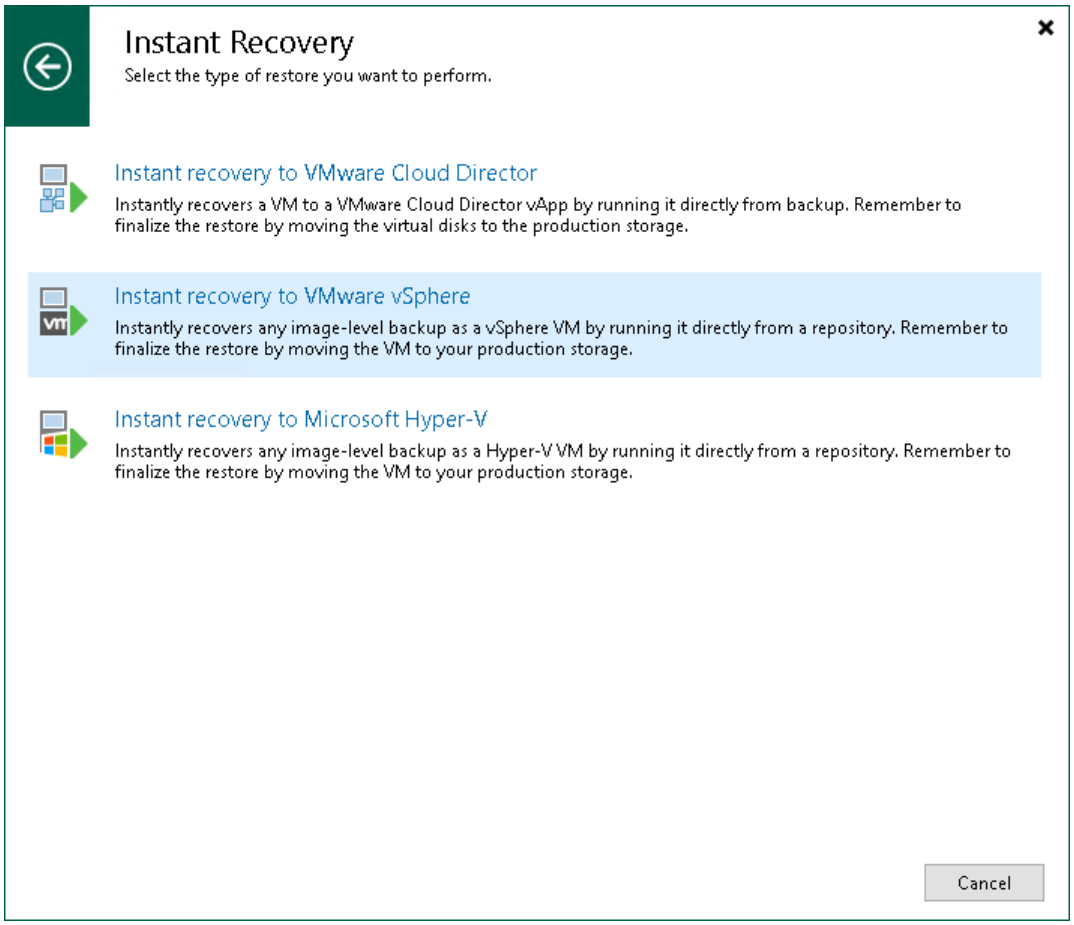


# Finalizing Instant Recovery

All VMs restored with Instant Recovery are displayed in the **Home** view, under the **Instant Recovery** node.

To check the progress of Instant Recovery and view session details:

1. Open the **Home** view.
2. In the inventory pane, click the **Instant Recovery** node.
3. Right-click the VM in the working area and select **Properties**.

Alternatively, you can open the **History** view, select the **Instant Recovery** node under **Restore** in the inventory pane and double-click the necessary instant restore session.

After the VMs have been successfully recovered, you must finalize the process. For this, test the recovered VMs and decide whether to migrate them to production environment or stop publishing.

- Testing recovered VMs
- Migrating recovered VMs
- Stop publishing recovered VMs

# Testing Recovered VMs

To test the recovered VMs before you migrate them to production, you can launch VMware Remote Console software from the Veeam Backup & Replication console.

> **IMPORTANT**
>
> Before you launch VMware Remote Console, make sure that this software is installed on the machine where the Veeam Backup & Replication console runs.

To open a VM console in Veeam Backup & Replication:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click a VM and select **Open VM console**.

# Migrating Recovered VM

To migrate the restored VM to production:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click the VM and select **Migrate to production**. Veeam Backup & Replication will launch the Quick Migration wizard.

   During migration, Veeam Backup & Replication will restore the VM from the backup file and additionally move all changes that were made while the VM was running from the backup in the Instant Recovery mode.

> **TIP**
>
> When you pass though the **Quick Migration** wizard, enable the **Delete source VM files upon successful migration** option. Veeam Backup & Replication will restore the VM to production and automatically stop the Instant Recovery session. If you do not enable this option, the Instant Recovery session will still be running, and you will need to unpublish the recovered VM manually.

# Stop Publishing Recovered VM

If you have disabled the **Delete source VM files upon successful migration** option in the Quick Migration settings, you must unpublish the VM manually. After you unpublish the VM, the Instant Recovery session will end and the recovered VM will be unmounted from the vPower NFS server. The migrated VM will remain on the production environment.

To unpublish a recovered VM:

1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click the VM and select **Stop publishing**.

**TIP**

After the VM has been published from the backup, you can open the VM console directly from Veeam Backup & Replication. To do this, in the working area right-click the VM and select **Open VM Console**.

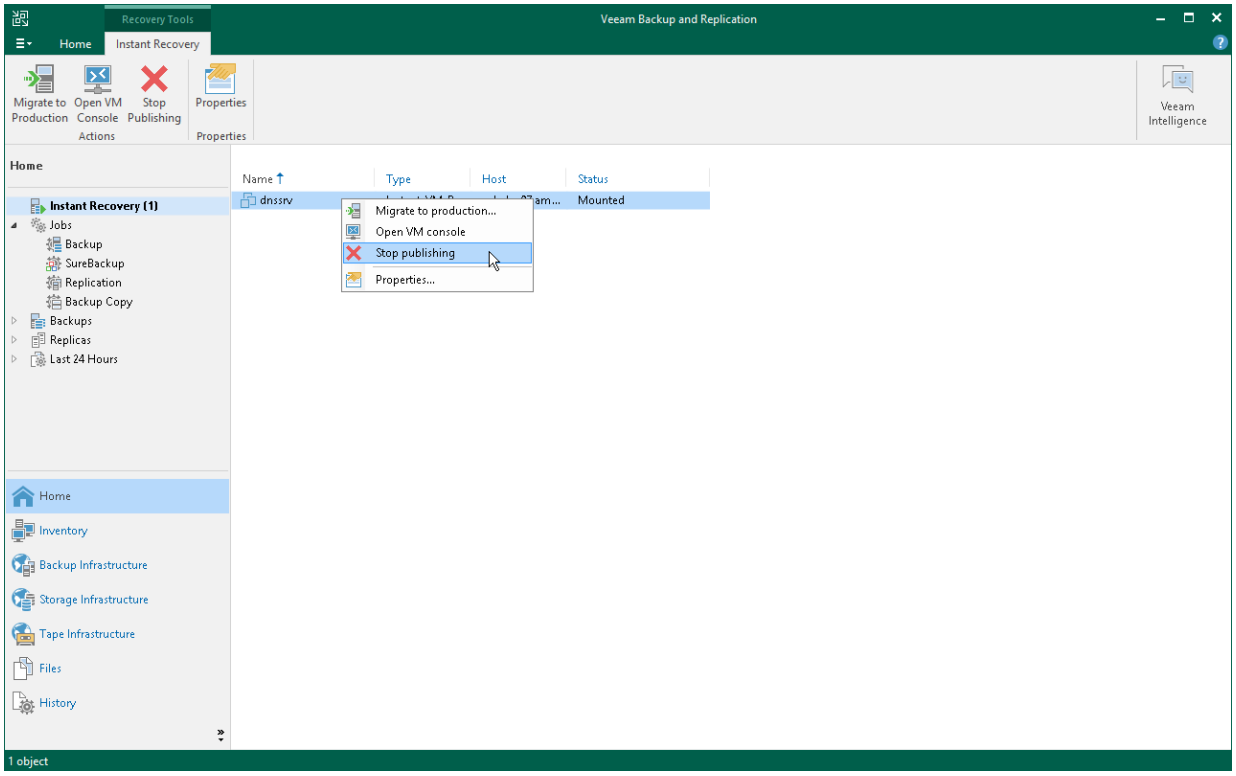


# Restoring Entire VMs to Cloud Director vApp

You can restore one or several VMs from VMware Cloud Director backups back to VMware Cloud Director.

The VMware Cloud Director VM can be restored to its original location — to a vApp in which the VM is already registered, or to a different location. You can restore a VM that already exists, for example, if the original VM is corrupted or you want to revert to an earlier state of the VM, or a VM that no longer exists, for example, if the VM was deleted by mistake. If you restore a VM that already exists, the original VM is overwritten with that from the VMware Cloud Director backup.

When restoring VMs to the VMware Cloud Director hierarchy, make sure that you select the **Restore into vCloud vApp** option. If you select the **Restore into vSphere infrastructure** option, the VM will be restored at the level of the underlying vCenter Server. To get a fully functional VM managed by VMware Cloud Director, you will need to manually import the restored VM to the VMware Cloud Director hierarchy.

Before you restore a VM to the VMware Cloud Director hierarchy, check prerequisites. Then use the **VMware Cloud Director Entire VM Restore** wizard to restore the necessary VM.

## Before You Begin

Before you restore VMware Cloud Director VMs to a vApp, consider the following:

- You can restore only those VMs whose placement policy is the same as the default placement policy of the target organization VDC (VDC where the vApp to which you restore VMs reside). For more information on placement policies, see VMware Docs.
- Veeam Backup & Replication restores standalone VMs as regular VMware Cloud Director VMs in the following cases:
  - If you selected to restore a standalone VM to a different location.
  - If the original standalone VM no longer exists in the VMware Cloud Director hierarchy.
- It is not recommended to restore VMware Cloud Director VMs to a vApp that already contains a standalone VM. After the restore process is complete, the vApp may not work as expected.
- If you restore linked clone VMs to a different location, make sure that fast provisioning is enabled at the level of the target organization VDC. Otherwise, Veeam Backup & Replication will restore the linked clone VM to a selected vApp as a regular VM.
- If you want to scan VM data for viruses, check the secure restore requirements and limitations.

## Step 1. Launch VMware Cloud Director Entire VM Restore Wizard

To launch the **VMware Cloud Director Entire VM Restore** wizard, do one of the following:

- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup. Select a VM you want to restore and click **Entire VM** > **VMware Cloud Director** on the ribbon.
- Open the Home view. In the inventory pane, select Backups. In the working area, expand the necessary backup. Right-click the VM you want to restore and select **Restore entire VM** > **VMware Cloud Director**.
- On the **Home** tab, click **Restore** and select **VMware Cloud Director**. In the **Restore** window, select **Restore from backup** > **VM restore** > **Entire VM restore** > **Entire VM restore** > **Entire VM restore to VMware Cloud Director**.
- Open the **Inventory** view. On the **View** tab, click **Cloud Director View**. In the inventory pane, expand the VMware Cloud Director hierarchy. In the working area, right-click the VM you want to restore and select **Restore** > **Restore entire VM** > **VMware Cloud Director**.

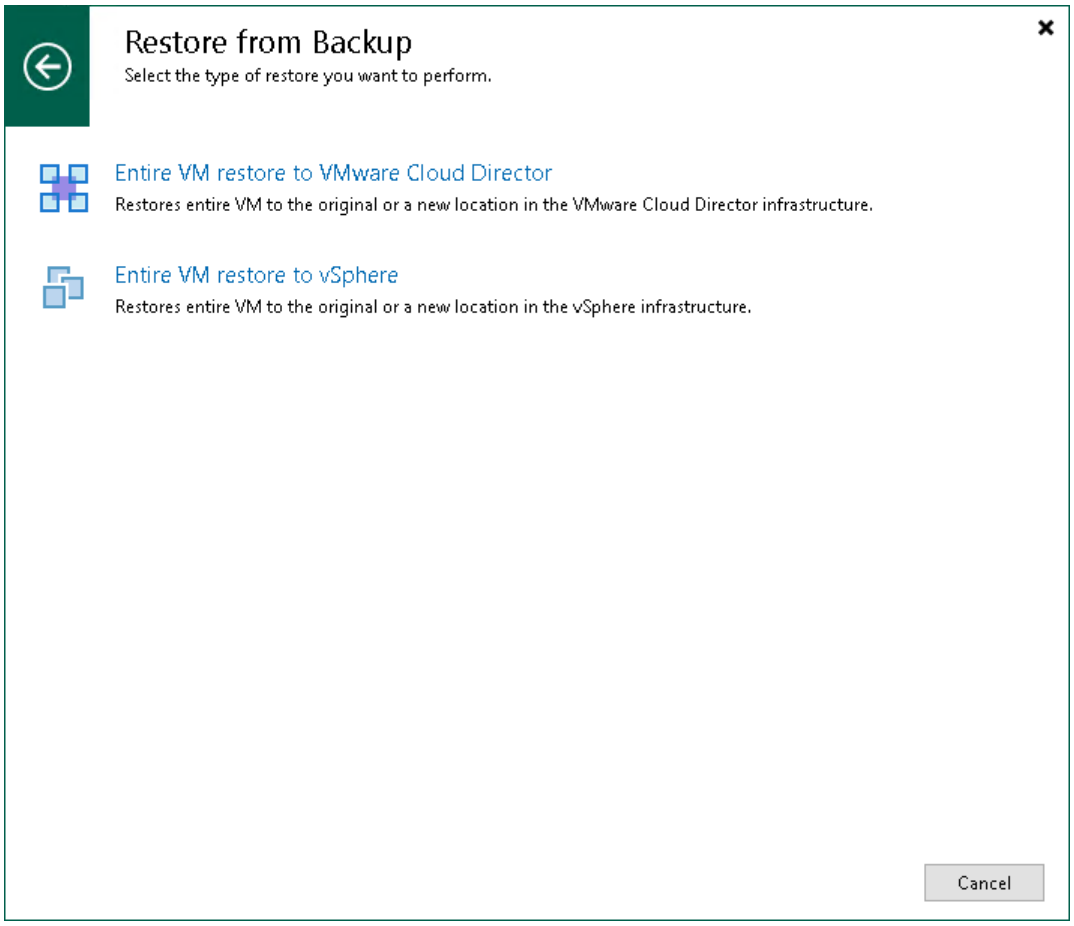

# Step 2. Select VMs to Restore

At the **Objects to Restore** step of the wizard, select one or several VMs to restore.

To add a VM, click **Add VM** and select where to browse for VMs:

- **From infrastructure** – browse the VMware Cloud Director hierarchy and select VMs to restore. Note that the VM you select from the VMware Cloud Director hierarchy must be successfully backed up at least once.
- **From backup** – browse existing backups and select VMs under backup jobs.

To facilitate selection, use the search field at the bottom of the **Select VMs** window: enter an object's name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

To add VMs to the list, you can also use the search field at the top of the window:

1. Enter a VM name or a part of it in the search field and Veeam Backup & Replication will search existing backups for the specified VM and display matching results.
2. To add the VM to the list, double-click it in the list of search results.
3. If the necessary VM is not found, click the **Show more** link to browse existing backups and choose the necessary VM.

To remove a VM from the list, select it and click **Remove** on the right.

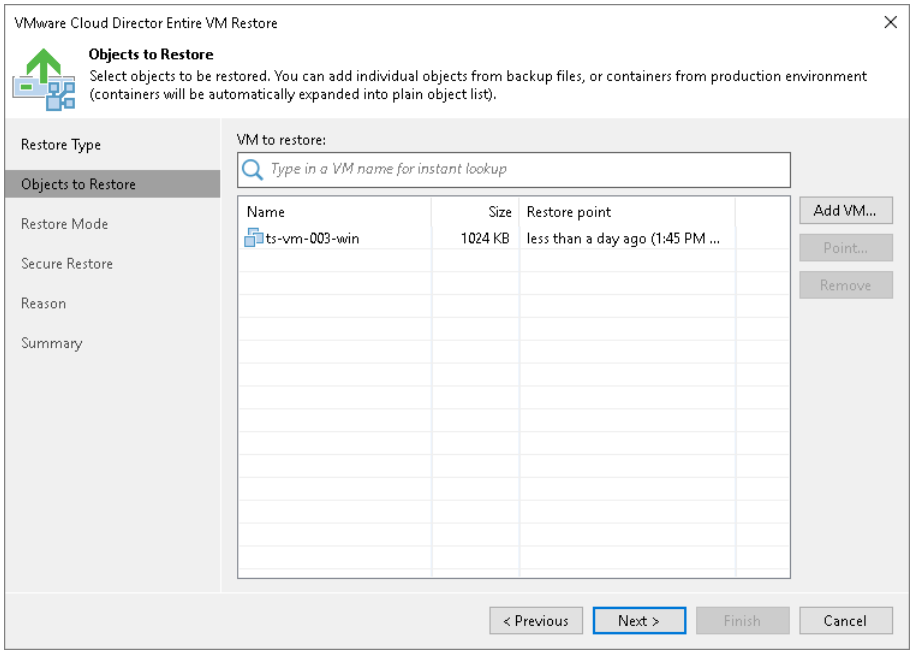

# Step 3. Select Restore Point

You can select the restore point for the VM.

By default, Veeam Backup & Replication uses the latest valid restore point to recover a VM. However, you can restore a VM to an earlier state. If you have chosen to restore multiple VMs, you can select a different restore point for every VM specifically.

To select a restore point for a VM:

1. Select a VM in the list and click **Point** on the right.
2. In the **Restore Points** window, select the restore point that must be used to recover the VM.

   In the **Location** column, you can view a name of a backup repository where a restore point resides.

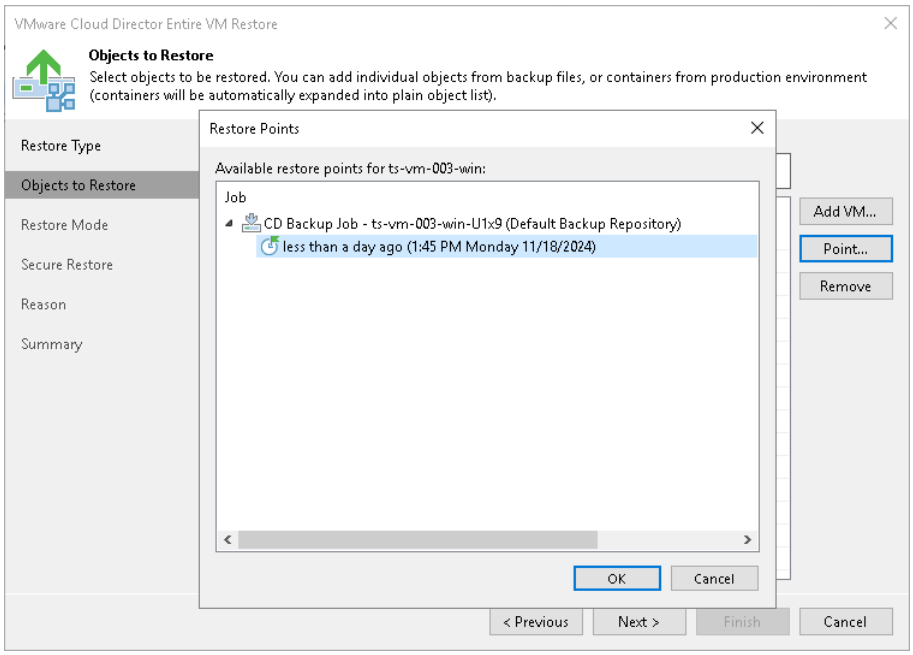

## Step 4. Select Restore Mode

At the **Restore Mode** step of the wizard, choose the necessary restore mode and backup proxy for VM data transport:

1. Choose a restore mode:
   - Select **Restore to original location** if you want to restore the VMs with initial settings and to the original location. If this option is selected, you will immediately pass to the Summary step of the wizard.
   - Select **Restore to a new location, or with different settings** if you want to restore the VMs to a different location or with different settings (such as VM location, network settings, fast provisioning settings and so on). If this option is selected, the **vCloud Full VM Restore** wizard will include additional steps for customizing VM settings.
2. [For VM restore to the original location] Select the **Quick rollback** check box if you want to use incremental restore for the VMs. Veeam Backup & Replication will use CBT to get data blocks that are necessary to revert the VMs to an earlier point in time, and will restore only these data blocks from the backup. Quick rollback significantly reduces the restore time and has little impact on the production environment.

   It is recommended that you enable this option if you restore the VMs after a problem that occurred at the level of the VM guest OS: for example, there has been an application error or a user has accidentally deleted a file on the VM guest OS. Do not enable this option if the problem has occurred at the VM hardware level, storage level or due to a power loss.
3. Click the **Pick proxy to use** link to select backup proxies over which VM data must be transported to the target datastore. You can assign backup proxies explicitly or instruct Veeam Backup & Replication to automatically select backup proxies.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the target datastore and will automatically assign optimal proxy resources for processing VM data.

     During the restore process, VMs are processed simultaneously. Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use for writing data to target, current workload on these backup proxies, and selects the most appropriate resources for VMs processing.

- If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that will be used for restore. It is recommended that you select at least two proxies to ensure that VMs are recovered should one of backup proxies fail or lose its connectivity to the target datastore during restore.

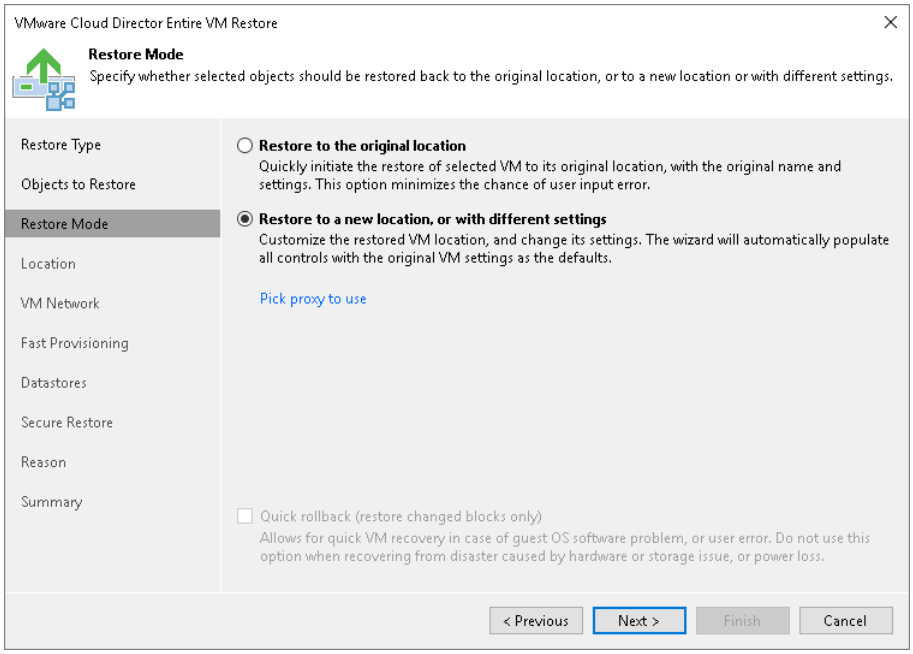

## Step 5. Select VM Location

The **Location** step of the wizard is available if you have chosen to change the location and settings for the restored VMs.

# Specifying Restore Location

By default, Veeam Backup & Replication restores the VM to its original location. To restore the VM to a different location:

1. Select the VM in the list and click **vApp**.
2. From the VMware Cloud Director hierarchy, choose a vApp in which the restored VM must be registered.

   To facilitate selection, use the search field at the bottom of the window: enter the vApp name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

# Changing Names

To change the VM name:

1. Select a VM in the list and click **Name**.
2. In the **Change Name** window, enter a new name explicitly or specify a change name rule by adding a prefix and suffix to the original VM name.
3. You can also change VM names directly in the list: select a VM, click the **New Name** field and enter the name to be assigned to the recovered VM.

# Restoring Tags

To restore tags that were assigned to the original VM and assign them to the restored VM, select the **Restore vSphere VM tags** check box. Veeam Backup & Replication will restore the VM with original tags if the following conditions are met:

- The VM is restored to its original location.

- The original VM tags are still available on the source vCenter Server.

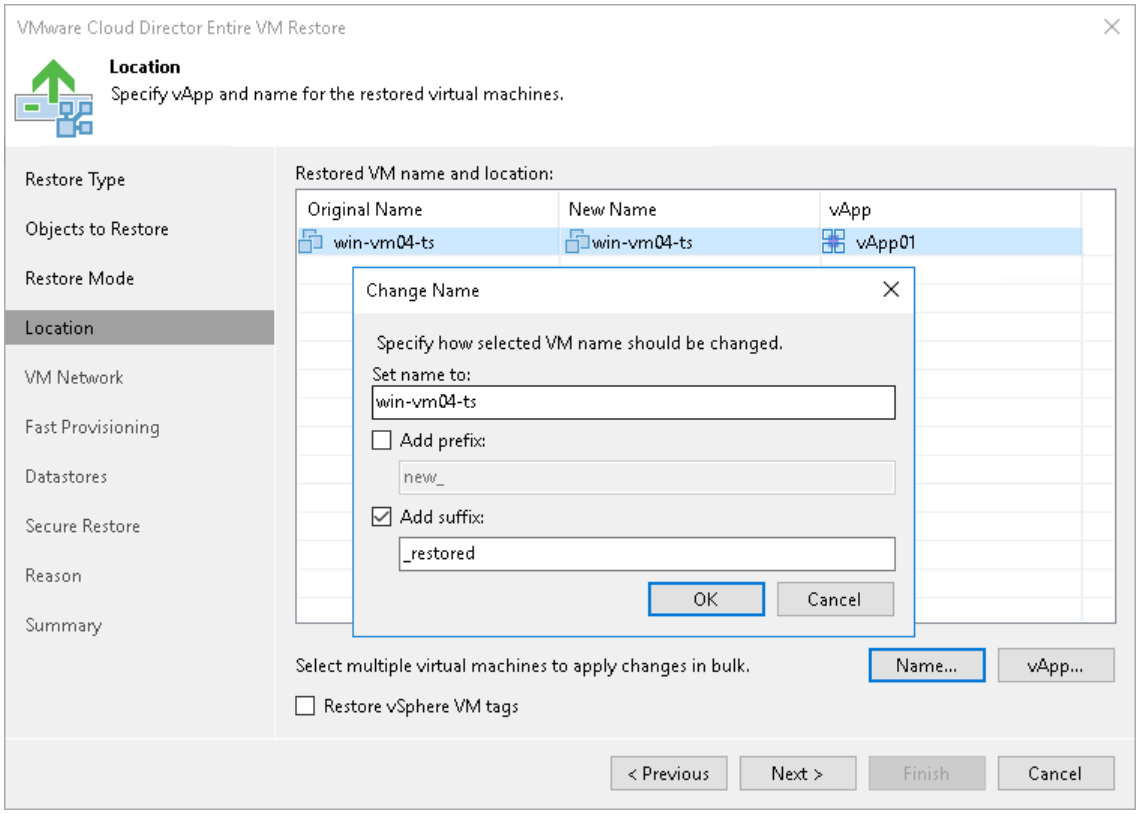

# Step 6. Select Destination Network

The **VM Network** step of the wizard is available if you have chosen to change the location and settings of the restored VMs.

To select networks to which the restored VM must be connected:

1. Select a VM in the list and click **Networks**.
2. The **Select Network** window displays all networks that are configured for the destination vApp. From the list of available networks, choose a network to which the restored VM must have access upon restore.

   To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

To prevent the restored VM from accessing any network, select it in the list and click **Disconnected**.

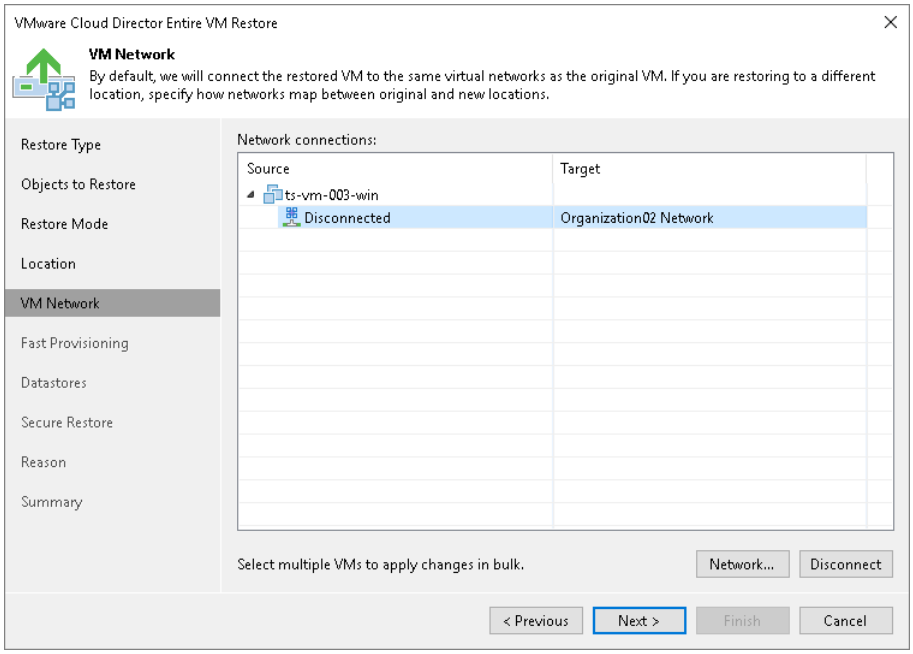

# Step 7. Select Template to Link

The **Fast Provisioning** step of the wizard is available if you have chosen to change the settings of the restored VMs.

To select a VM template:

1. Select a VM in the list and click **Set Template**.
2. From the VMware Cloud Director hierarchy, choose a template to which the restored VM must be linked.

   To facilitate selection, use the search field at the bottom of the window: enter a VM template name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

If you want to disable fast provisioning for the VM and restore it as a regular VM, select the VM in the list and click **Disable**.

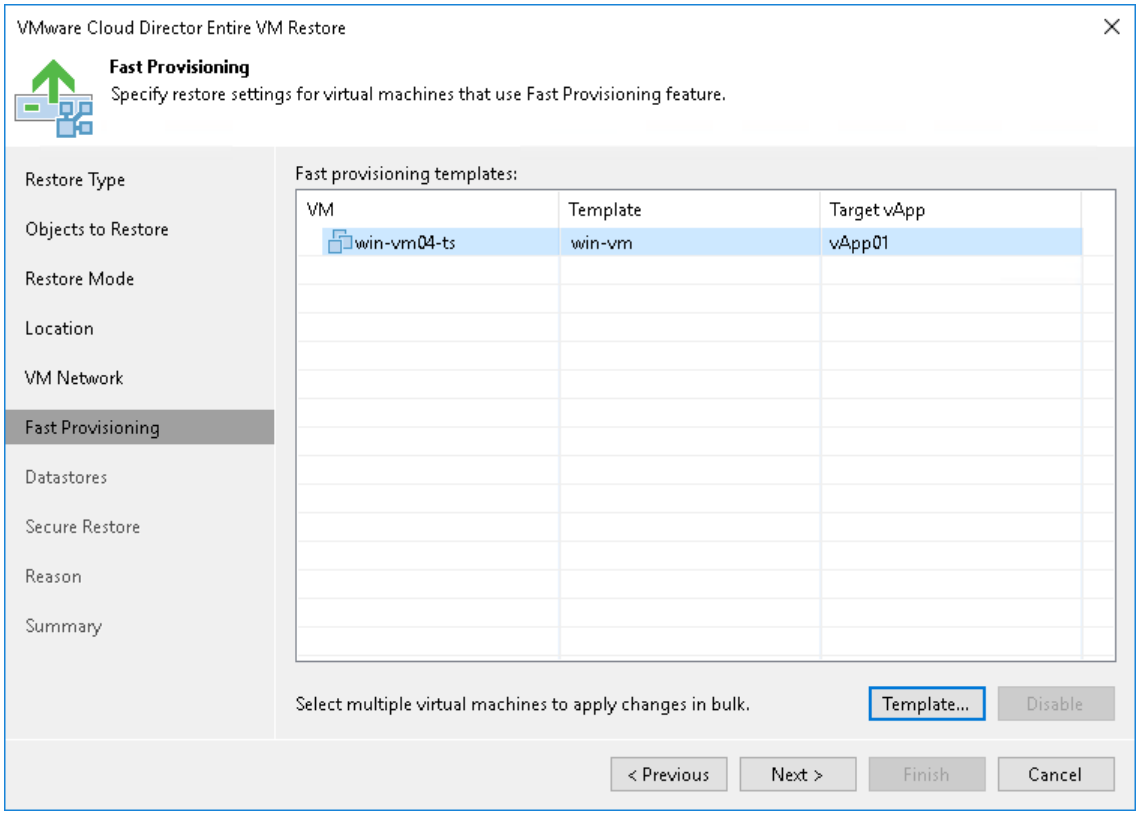

# Step 8. Select Storage Policy and Datastores

The **Datastores** step of the wizard is available if you have chosen to change the settings of the restored VMs.

To select a storage policy for the restored VM:

1. Select a VM in the list and click **Policy**.
2. In the displayed window, select the necessary policy for the VM.

If you have selected to disable fast provisioning at the previous step of the wizard, you must select a datastore on which disks of the restored VM will be placed.

1. Select a VM in the list and click **Datastore**.
2. In the displayed window, select the datastore on which the VM disks must be located.

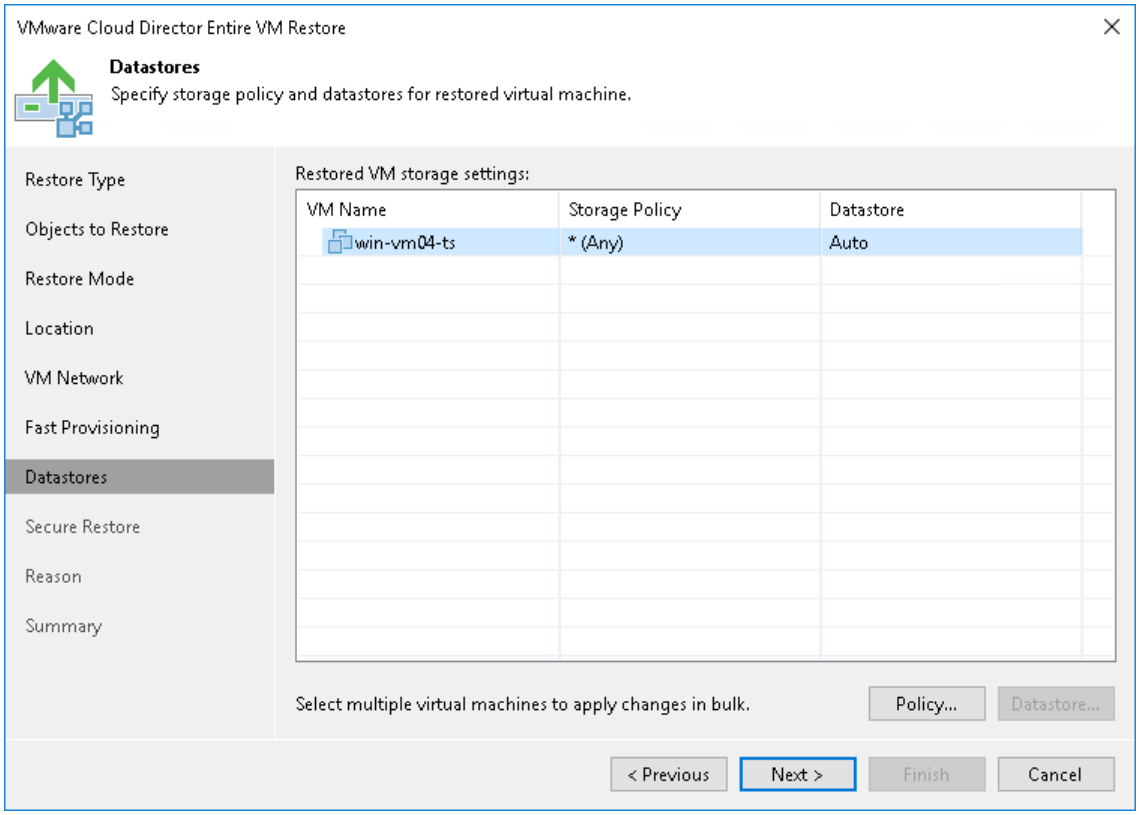

# Step 9. Specify Secure Restore Settings

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Select which action Veeam Backup & Replication will take if the scan finds a virus threat:
   - **Proceed with recovery but disable network adapters** — restore the VM with disabled network adapters (NICs).
   - **Abort VM recovery** — cancel the restore session.

3. Select the **Continue scanning all remaining files after the first occurrence** check box if you want the antivirus software to continue the VM data scan after the first malware is found. For information on how to view the results of the antivirus scan, see Viewing Antivirus Scan Results.

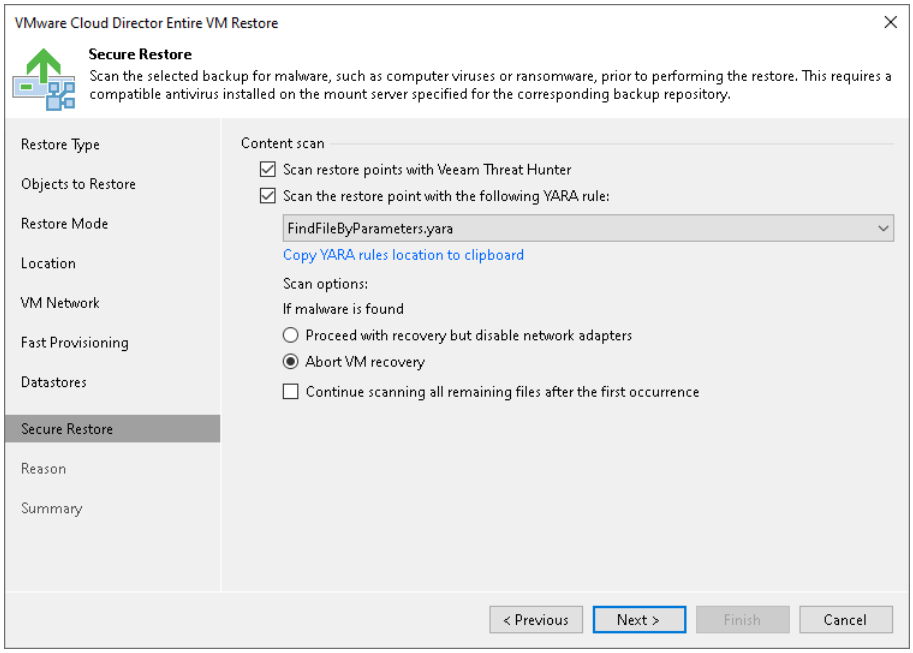

# Step 10. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the selected VMs. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

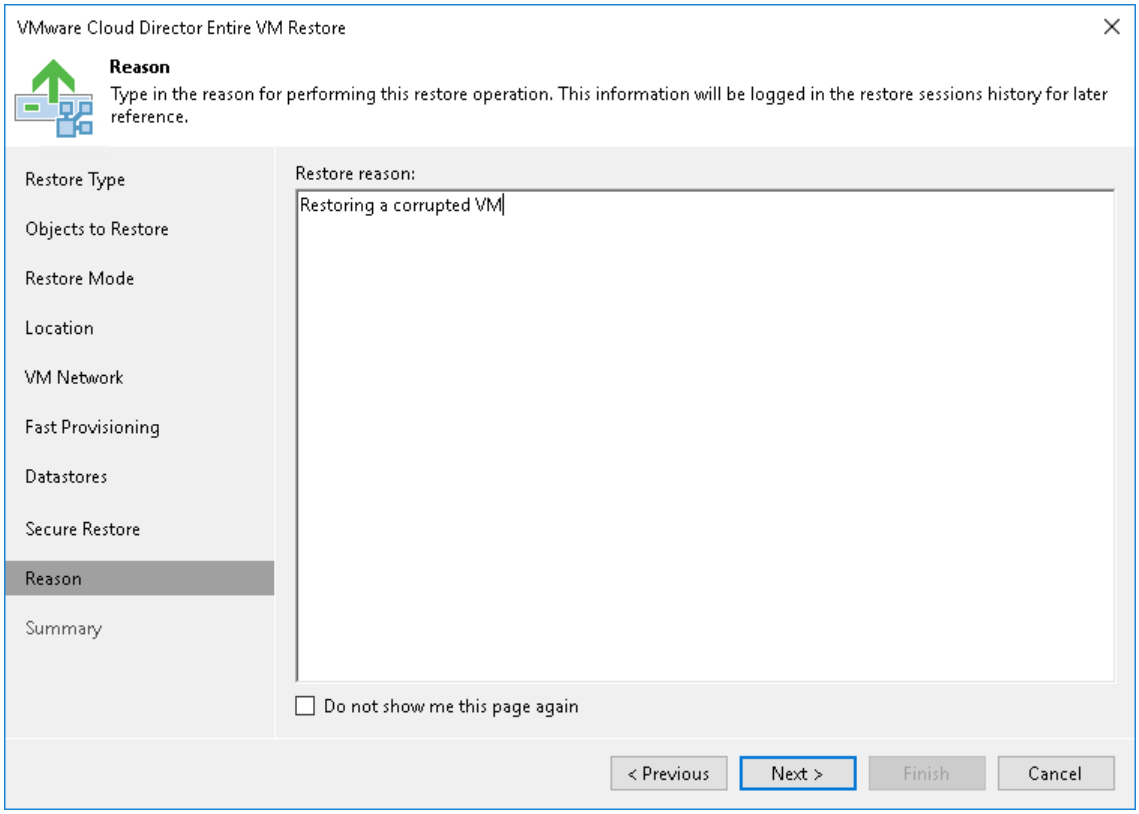

## Step 11. Verify Recovery Settings and Finish Working with Wizard

At the **Summary** step of the wizard, specify additional settings for VMs restore:

1. If you want to start the restored VMs, select the **Power on VM after restoring** check box.
2. Check the settings for VMs restore and click **Finish**. Veeam Backup & Replication will recover the VMs in the specified destination.

> **NOTE**
>
> Veeam Backup & Replication checks the lease term for the restored VMs. In case the lease period has expired, the lease will be automatically updated.

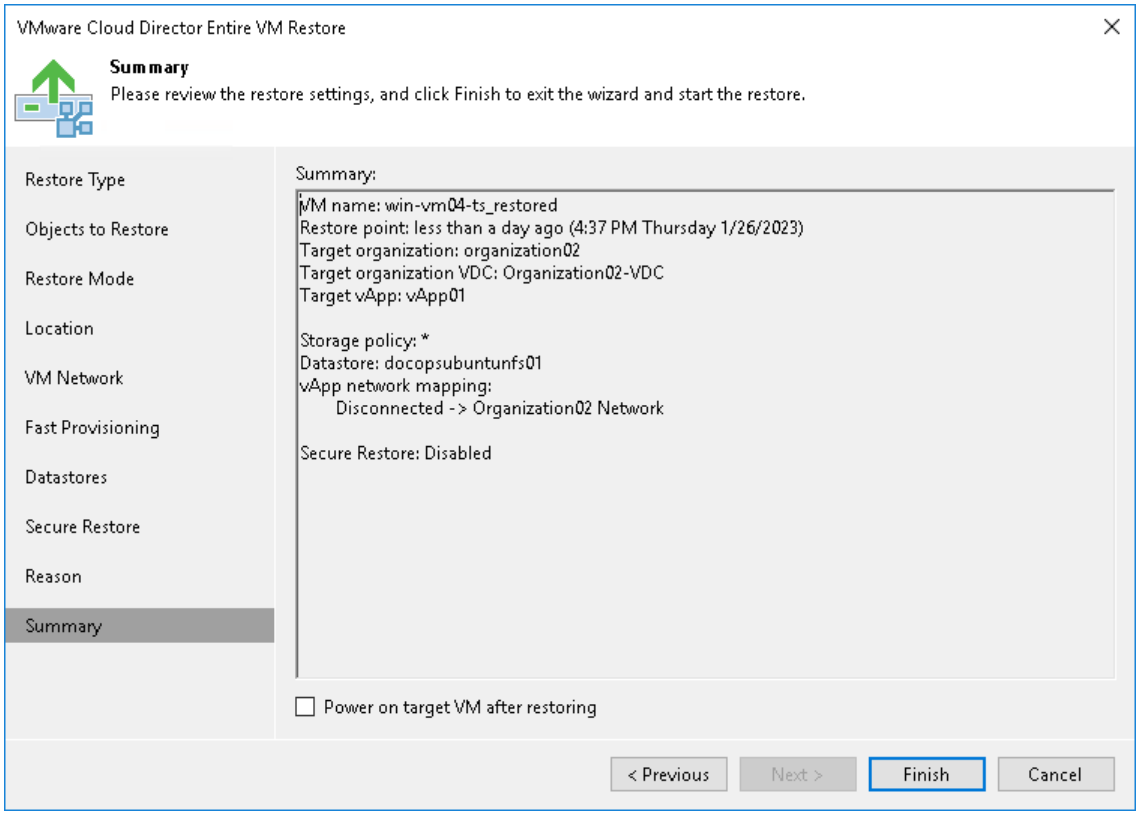


# Restoring Entire VMs to VMware vSphere

You can restore VMware Cloud Director VMs from the backup to the VMware vSphere infrastructure.

During restore, Veeam Backup & Replication neglects the vApp metadata saved to the backup file and performs a regular entire VM restore process. The VM is restored to the vCenter Server or ESXi host and is not registered in VMware Cloud Director. Cloud Director-specific features such as fast provisioning are not supported for such type of restore.

To launch the **Full VM Restore** wizard, do one of the following:

- Open the **Home** view, in the inventory pane select **Backups**. In the working area, expand the necessary backup, select the VMs you want to restore and click **Entire VM** > **VMware vSphere** on the ribbon.
- Open the **Inventory** view. In the inventory pane, expand the VMware Cloud Director hierarchy and select the vCenter Server. In the working area, right-click the VM you want to restore and select **Restore entire VM** > **VMware vSphere**.

Entire VM restore of VMware Cloud Director VMs does not differ from entire VM restore of regular VMware VMs. For more information, see Performing Entire VM Restore.

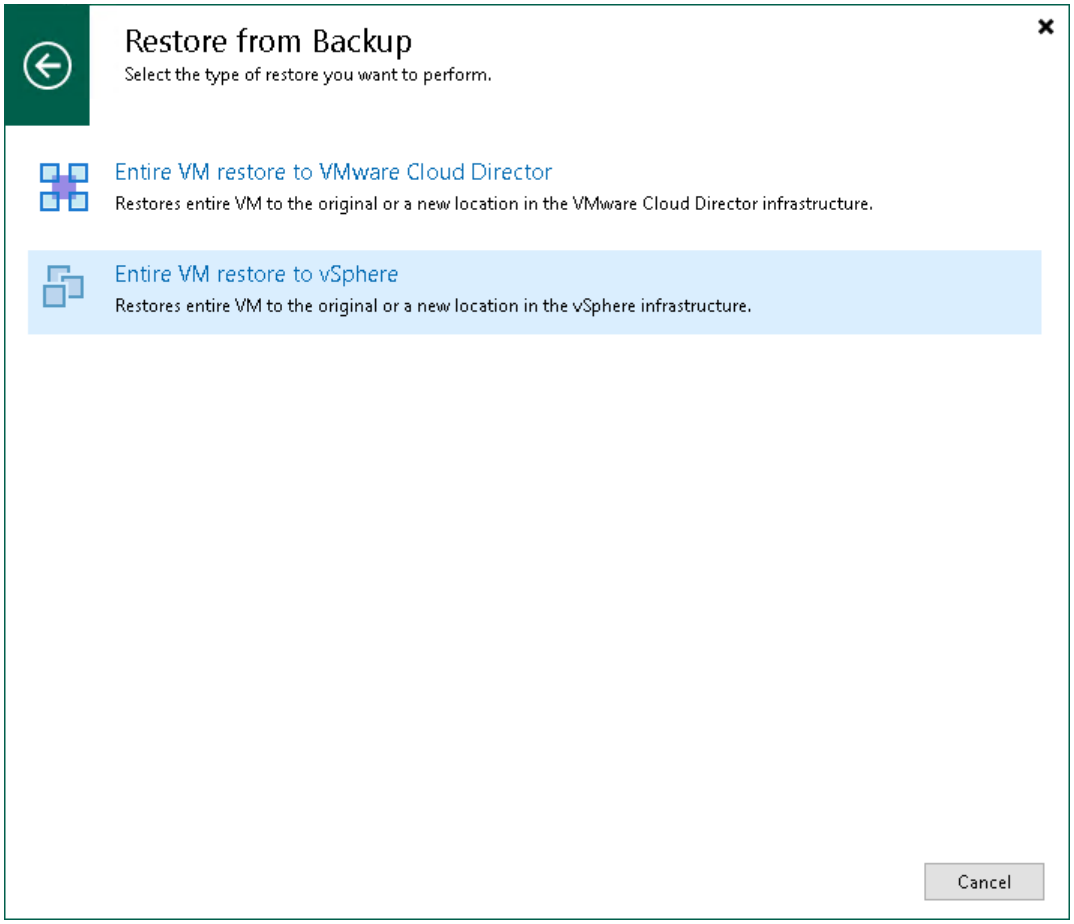

# vApp Recovery

With vApp recovery, you can restore the whole vApp from a backup to VMware Cloud Director.

vApps can be restored to their organization VDC or to any other organization VDC. You can restore the vApp that already exists, for example, in case the vApp is corrupted or you want to revert to an earlier state of the vApp, or the vApp that no longer exists, for example, if it was deleted by mistake. If you restore a vApp that already exists, the vApp is overwritten with that from the VMware Cloud Director backup.

During vApp recovery, Veeam Backup & Replication recovers VMs within one vApp one by one consequently. vApps are recovered in parallel.

> **IMPORTANT**
>
> You can restore only those VMs whose placement policy is the same as the default placement policy of the target organization VDC (VDC where the vApp to which you restore VMs reside). For more information on placement policies, see VMware Docs.

## Restoring Cloud Director vApps

To restore a vApp to VMware Cloud Director, use the **vCloud Full vApp Restore** wizard.

## Step 1. Launch Full vApp Restore Wizard

To launch the **Full vApp Restore** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware Cloud Director**. In the **Restore** window, select **Restore from backup** > **Entire vApp restore**.
- Open the **Home** view. In the inventory pane, select **Backups**. In the working area, expand the necessary backup and do one of the following:
  - Select the vApp and click **Restore vApp** on the ribbon.
  - Right-click the vApp and select **Restore VMware Cloud Director vApp**.

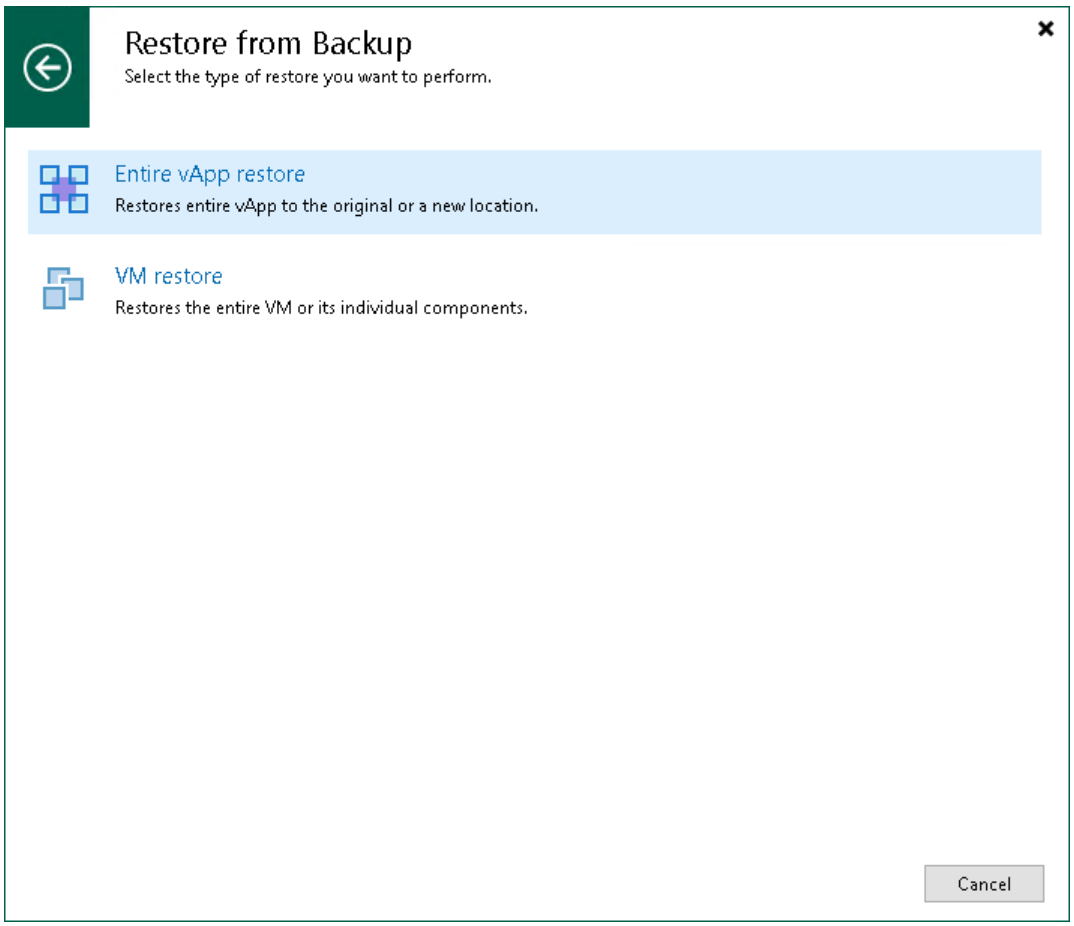

## Step 2. Select vApp to Restore

At the **Objects to Restore** step of the wizard, select the vApp you want to restore.

To add a vApp, click **Add vApp** and select where to browse for vApps:

- **From infrastructure** – browse the VMware Cloud Director hierarchy and select a vApp to restore. Note that the vApp you select from the VMware Cloud Director hierarchy must be successfully backed up at least once.
- **From backup** – browse existing backups and select the vApp under backup jobs.

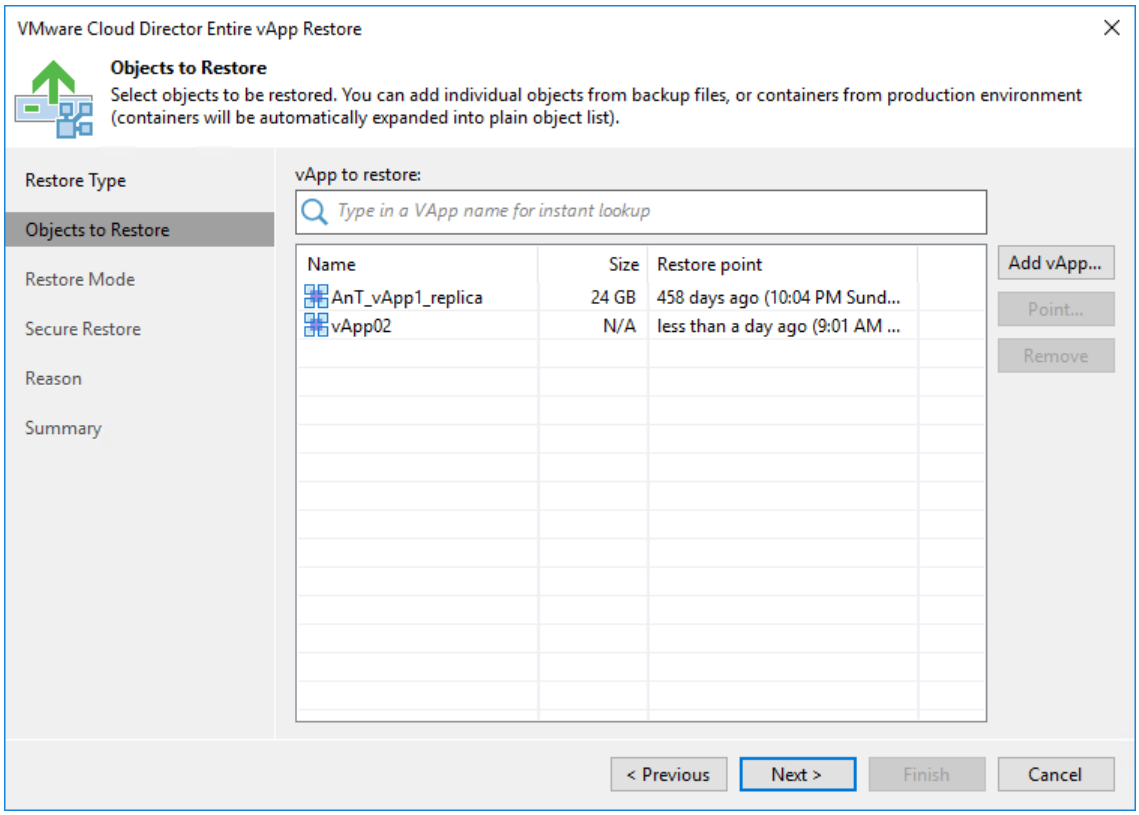

# Step 3. Select Restore Point

You can select the restore point for the vApp.

By default, Veeam Backup & Replication uses the latest valid restore point to recover the vApp. However, you can restore the vApp to an earlier state.

To select a restore point for the vApp:

1. Select the vApp in the list and click **Point** on the right.
2. In the **Restore Points** window, select a restore point that must be used to recover the vApp.

   In the **Location** column, you can view a name of a backup repository where a restore point resides.

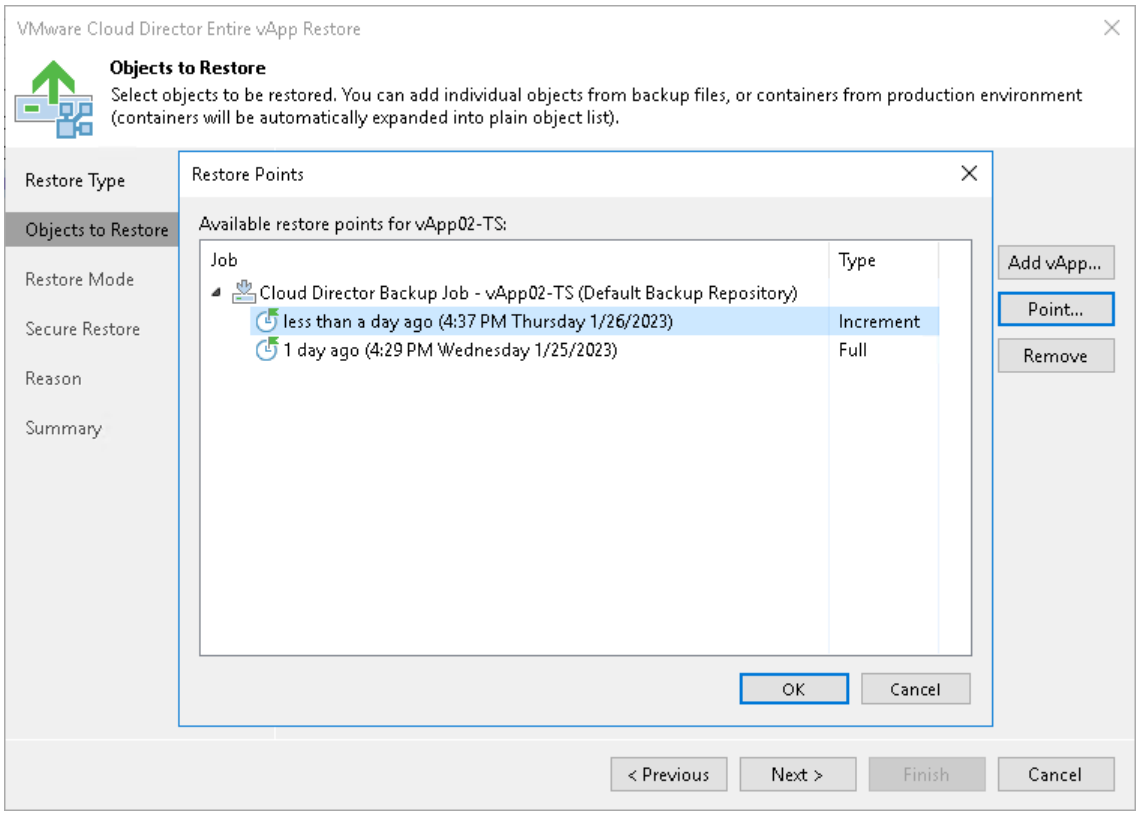

## Step 4. Select Restore Mode

At the **Restore Mode** step of the wizard, choose the necessary restore mode and backup proxy for VM data transport:

1. Choose a restore mode:
   - Select **Restore to original location** if you want to restore the vApp with its initial settings and to its original location. If this option is selected, you will immediately pass to the Summary step of the wizard.
   - Select **Restore to a new location, or with different settings** if you want to restore the vApp to a different location or with different settings (such as organization VDC, network settings, fast provisioning settings and so on). If this option is selected, the **vCloud Full vApp Restore** wizard will include additional steps for customizing vApp settings.
2. [For vApp restore to the original location] Select the **Quick rollback** check box if you want to use incremental restore for the vApp. Veeam Backup & Replication will query CBT to get data blocks that are necessary to revert the vApp to an earlier point in time, and will restore only these data blocks from the backup. Quick rollback significantly reduces the restore time and has little impact on the production environment.

   It is recommended that you enable this option if you restore VMs in the vApp after a problem that occurred at the level of the VM guest OS: for example, there has been an application error or a user has accidentally deleted a file on the VM guest OS. Do not enable this option if the problem has occurred at the VM hardware level, storage level or due to a power loss.
3. Click the **Pick proxy to use** link to select backup proxies over which vApp data must be transported to the target datastore. You can assign backup proxies explicitly or instruct Veeam Backup & Replication to automatically select backup proxies.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that are connected to the target datastore and will automatically assign optimal proxy resources for processing vApp data.

     During the restore process, VMs in the vApp are processed simultaneously.
     Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use for writing data to target, current workload on these backup proxies, and selects the most appropriate resources for VMs processing.

- If you choose **Use the selected backup proxy serves only**, you can explicitly select backup proxies that will be used for restore. It is recommended that you select at least two proxies to ensure that VMs are recovered should one of backup proxies fail or lose its connectivity to the target datastore during restore.

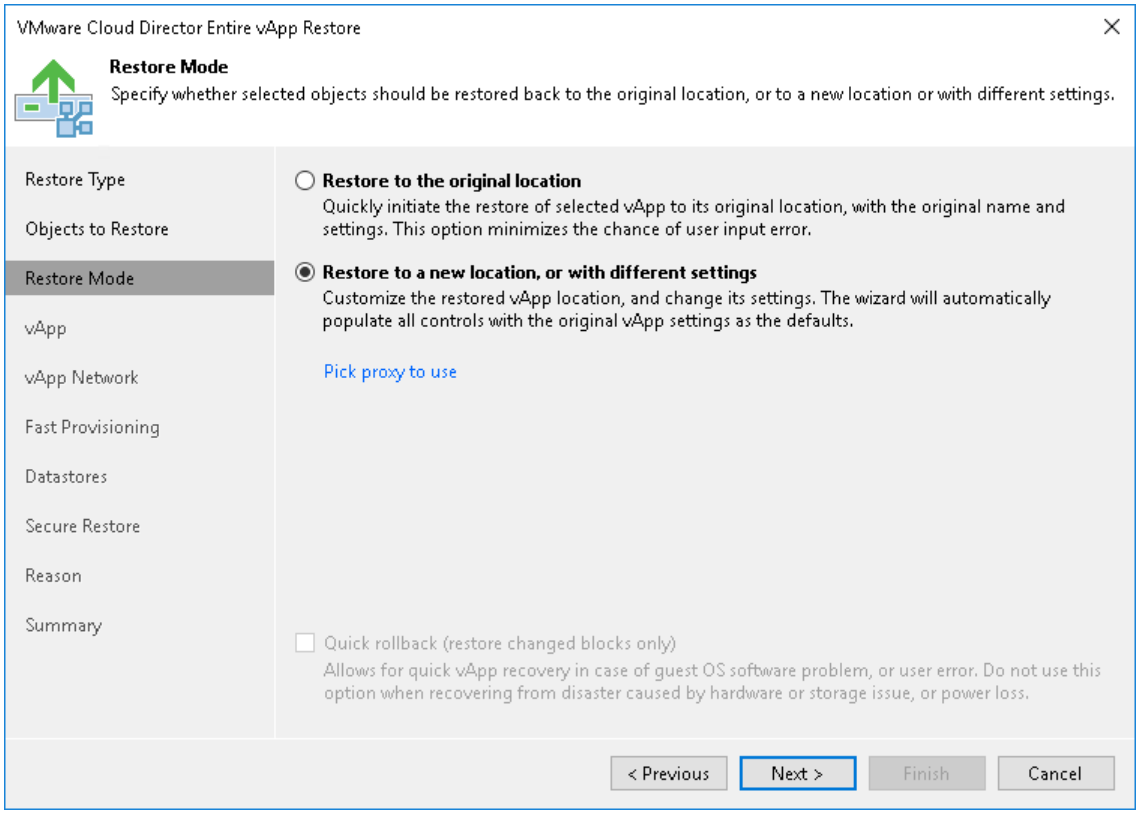

# Step 5. Select vApp Location

The **vApp** step of the wizard is available if you have chosen to change the location and settings of the restored vApp.

By default, Veeam Backup & Replication restores the vApp to its original location with its original name.

To restore the vApp to a different location:

1. Select the App in the list and click **_restoredDC**.
2. From the VMware Cloud Director hierarchy, choose an organization VDC where the selected vApp must be registered.

   To facilitate selection, use the search field at the bottom of the window: enter an object's name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

> **NOTE**
> If a vApp contains a standalone VM, Veeam Backup & Replication restores the standalone VM in such vApp as a regular VMware Cloud Director VM.

To change the vApp name:

1. Select the vApp in the list and click **Name**.
2. In the **Change Name** window, enter a new name explicitly or specify a change name rule by adding a prefix and suffix to the original vApp name.
3. You can also change the vApp name directly in the list: select a vApp, click the **New Name** field and enter the name to be assigned to the recovered vApp.

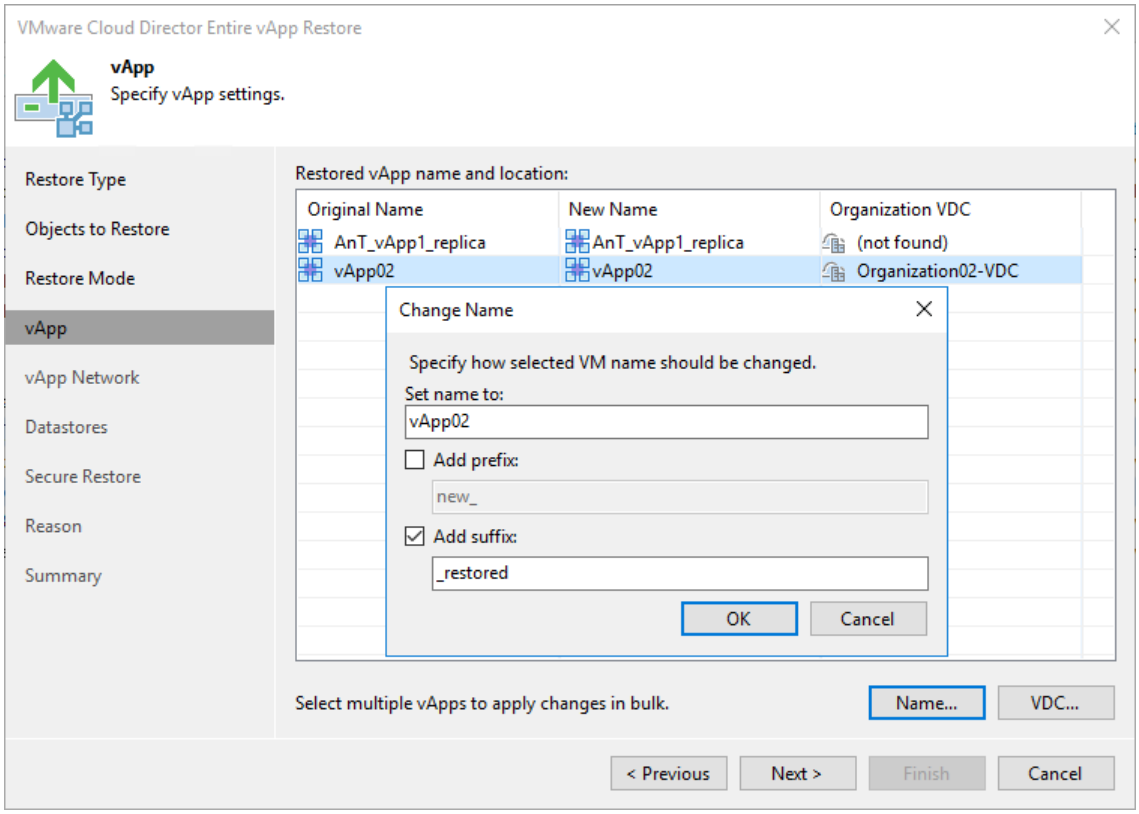

# Step 6. Select Destination Network

The **vApp Network** step of the wizard is available if you have chosen to change the location and settings of the restored vApp.

To select networks to which the restored vApp must connected:

1. Select the vApp in the list and click **Network**.
2. The **Select Network** window displays all networks that are configured for the destination organization VDC. From the list of available networks, choose a network to which selected vApp must have access upon restore.

   To facilitate selection, use the search field at the bottom of the window: enter a network name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

To prevent the restored vApp from accessing any network, select it in the list and click **Disconnect**.

> **NOTE**
>
> When Veeam Backup & Replication backs up a vApp, along with vApp networks data it saves information about organization networks to which the vApp is connected. If you restore the vApp to the original VDC organization and do not change the organization network settings, Veeam Backup & Replication attempts to connect the vApp to all source organization networks to which the vApp was connected at the moment of backup. If one or several source organization networks are not detected, for example, if they have been changed or removed by the time of restore, Veeam Backup & Replication will not be able to restore the vApp to its original organization. In this case, you will need to change the network settings for the restored vApp: map the vApp to other organization networks or disconnect it from organization networks at all.

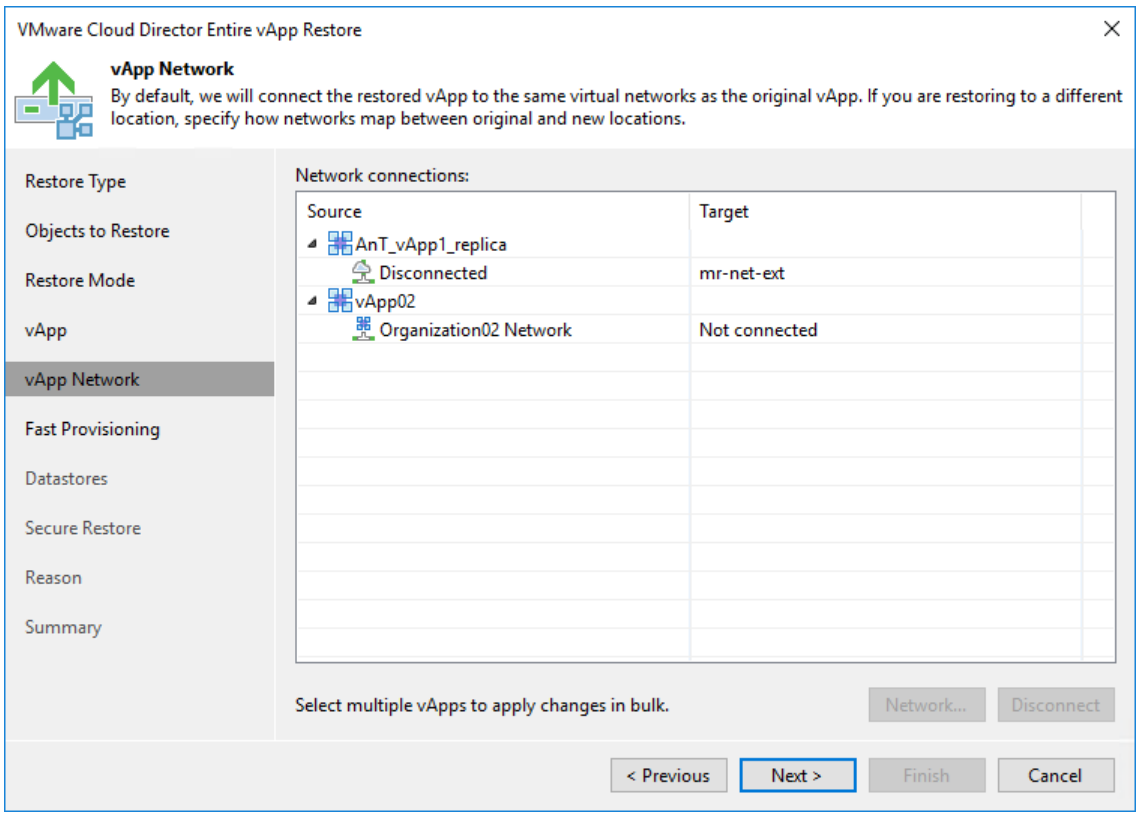

# Step 7. Select Template to Link

The **Fast Provisioning** step of the wizard is available if you have chosen to change settings of the restored vApp, for example, its name or location.

To select a VM template:

1. Select a VM in the list and click **Template**.
2. From the VMware Cloud Director hierarchy, choose a template to which the VMs from the restored vApp must be linked.

   To facilitate selection, use the search field at the bottom of the window: enter a VM template name or a part of it and click the **Start search** button on the right or press [Enter] on the keyboard.

If you want to disable fast provisioning for the VM and restore it as a regular VM, select the VM in the list and click **Disable**.

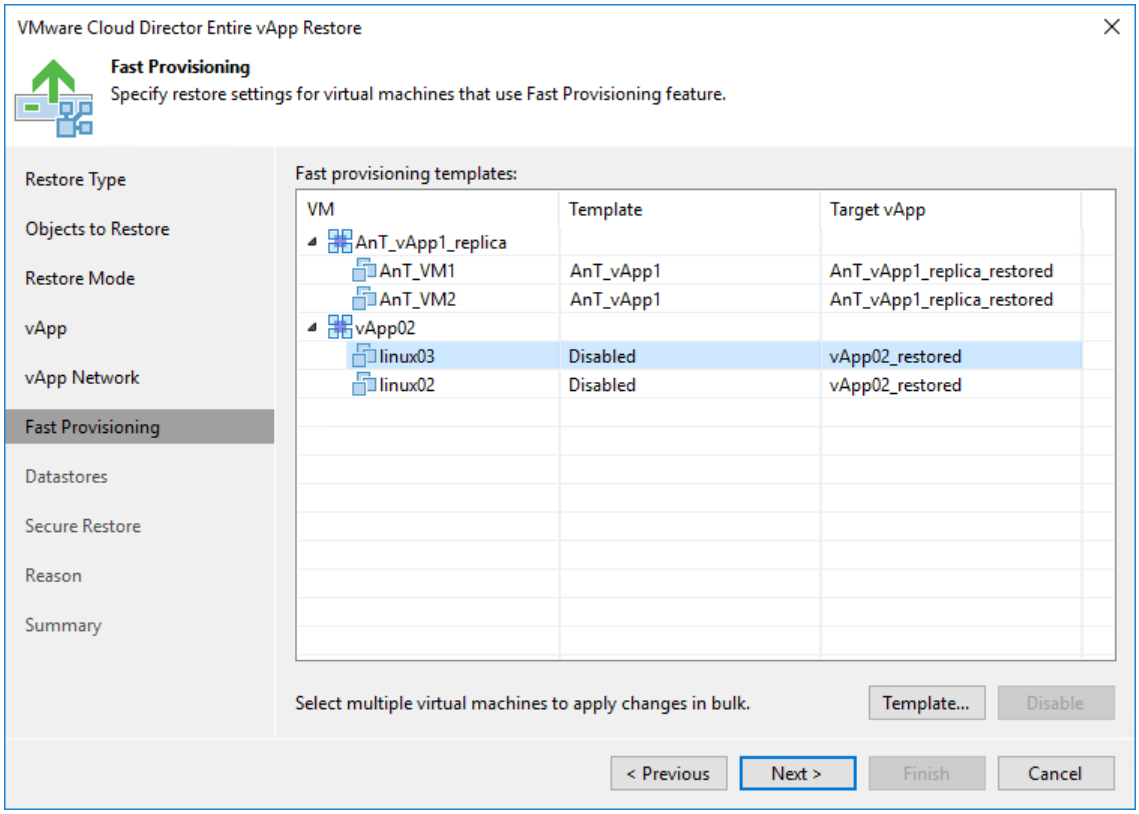

# Step 8. Select Storage Policy and Datastores

The **Datastores** step of the wizard is available if you have chosen to change settings of the restored vApp, for example, its name or location.

To select a storage policy for the vApp:

1. Select the vApp in the list and click **Policy**.
2. In the displayed window, select the necessary policy for the vApp.

If you have selected to disable fast provisioning at the previous step of the wizard, you must select a datastore on which the disks of restored VMs will be placed. To do this:

1. Select VM or vApp in the list and click **Datastore**.
2. In the displayed window, select the datastore on which the disks of the VM must be placed.

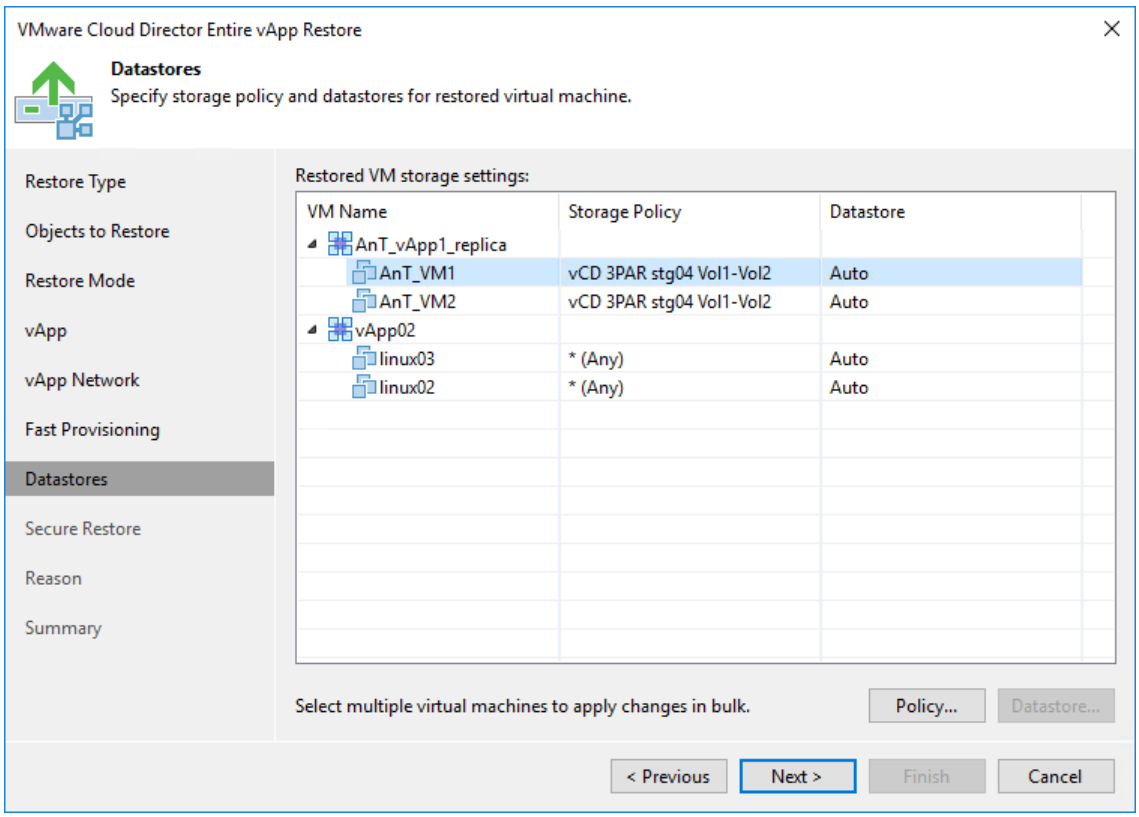

# Step 9. Specify Secure Restore Settings

At the **Secure Restore** step of the wizard, you can instruct Veeam Backup & Replication to scan machine data for malware activity before restoring the machine to the production environment. To scan the data, Veeam Backup & Replication uses the following engines:

- Veeam Threat Hunter — a signature-based scan engine. For more information, see Veeam Threat Hunter for Secure Restore.
- Third-party antivirus software. For more information, see Antivirus Scan for Secure Restore.
- YARA — a rule-based scan engine. For more information, see YARA Scan for Secure Restore.

To specify secure restore settings, do the following:

1. Select the method you want to use for data scan:
   - Select the **Scan restore points with Veeam Threat Hunter** option to use Veeam Threat Hunter.

     This option is available if you configured Veeam Threat Hunter as detection engine in Signature Detection settings.
   - Select the **Scan restore points with your existing antivirus software** to use third-party antivirus software.

     This option is available if you configured third-party antivirus in the Signature Detection settings.

     > **TIP**
     >
     > Click **Change** to open **Malware Detection Settings** window where you can change your scan engine to Veeam Threat Hunter.
   - Select **Scan the restore point with the following YARA rule** to use YARA as a scan engine.

     For a YARA rule to appear in the drop-down list, it must be placed in the `YaraRules` folder in the Veeam Backup & Replication product folder.

     > **TIP**
     >
     > To copy the path to the folder with YARA rules, click Copy YARA rules location to clipboard.
2. Select which action Veeam Backup & Replication will take if scan finds a virus threat:
   - **Proceed with recovery but disable network adapters.** Select this action if you want to restore the vApp VMs with disabled network adapters (NICs).
   - **Abort VM recovery**. Select this action if you want to cancel the restore session.

3. Select the **Scan the entire image** check box if you want to continue the vApp data scan after the first malware is found. For information on how to view results of the malware scan, see Viewing Antivirus Scan Results.

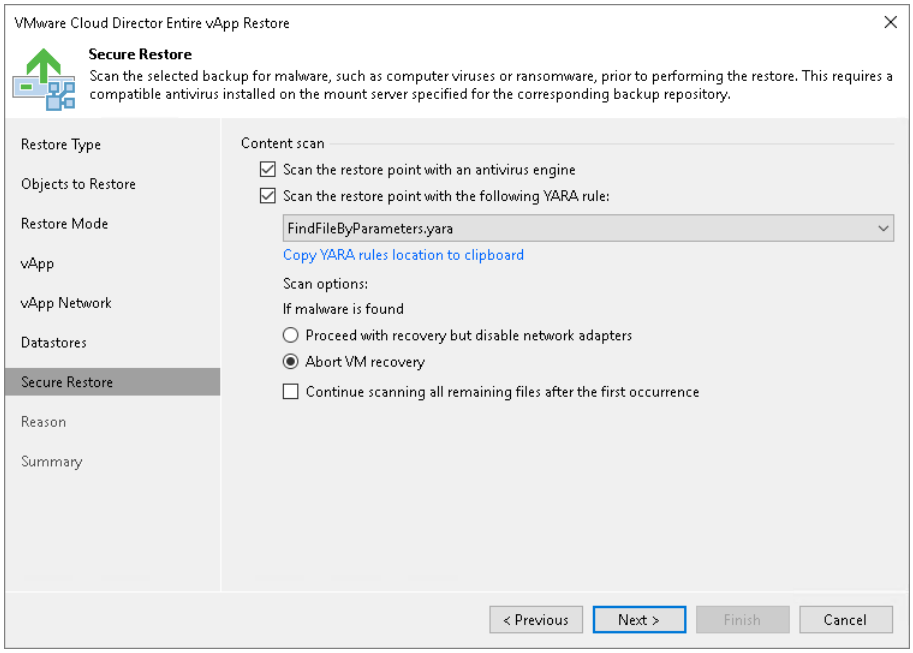

# Step 10. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the selected vApp. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to show this page, select the **Do not show me this page again** check box. If you want to unhide this page, follow the instructions described in this Veeam KB article.

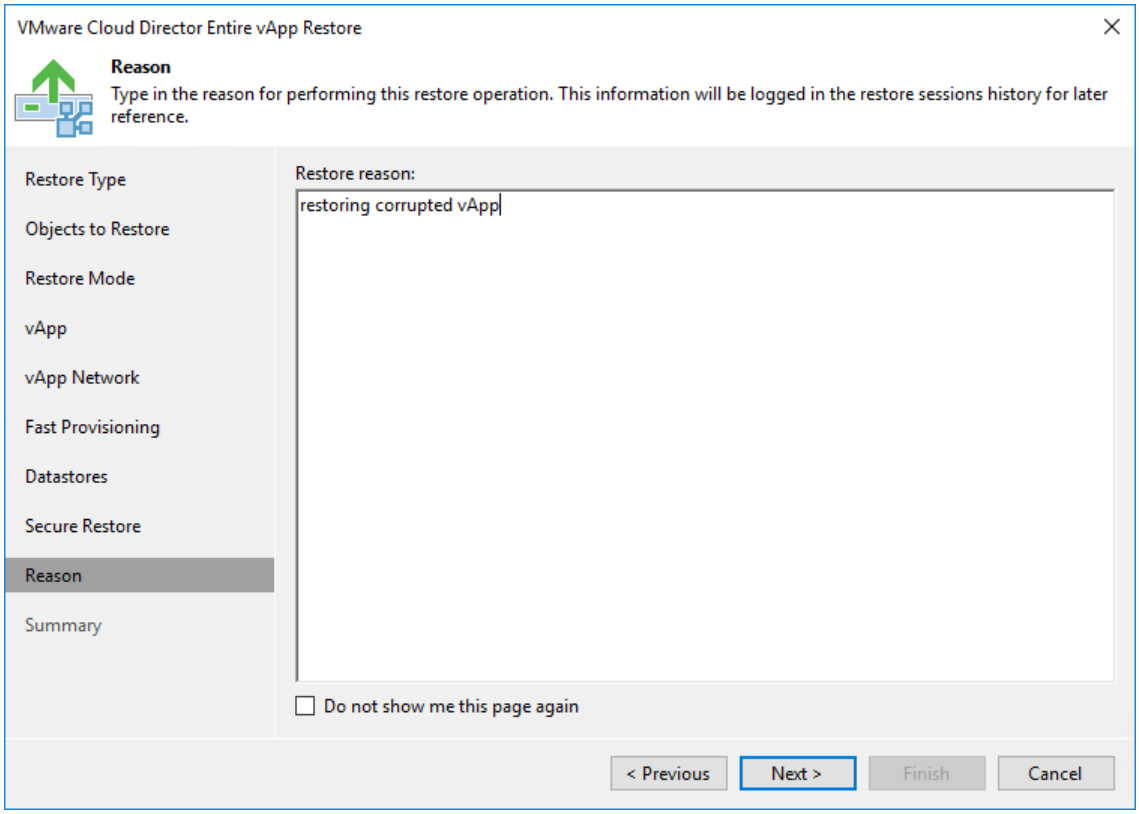

# Step 11. Verify Recovery Settings and Finish Working with Wizard

At the **Summary** step of the wizard, specify additional settings for vApp restore:

1. If you want to start VMs in the restored vApp, select the **Power on VM after restoring** check box.
2. Check the settings for vApp restore and click **Finish**. Veeam Backup & Replication will recover the vApp in the specified destination.

> **NOTE**
>
> Veeam Backup & Replication checks the lease term for the restored vApp. If the lease period has expired, the lease will be automatically updated.

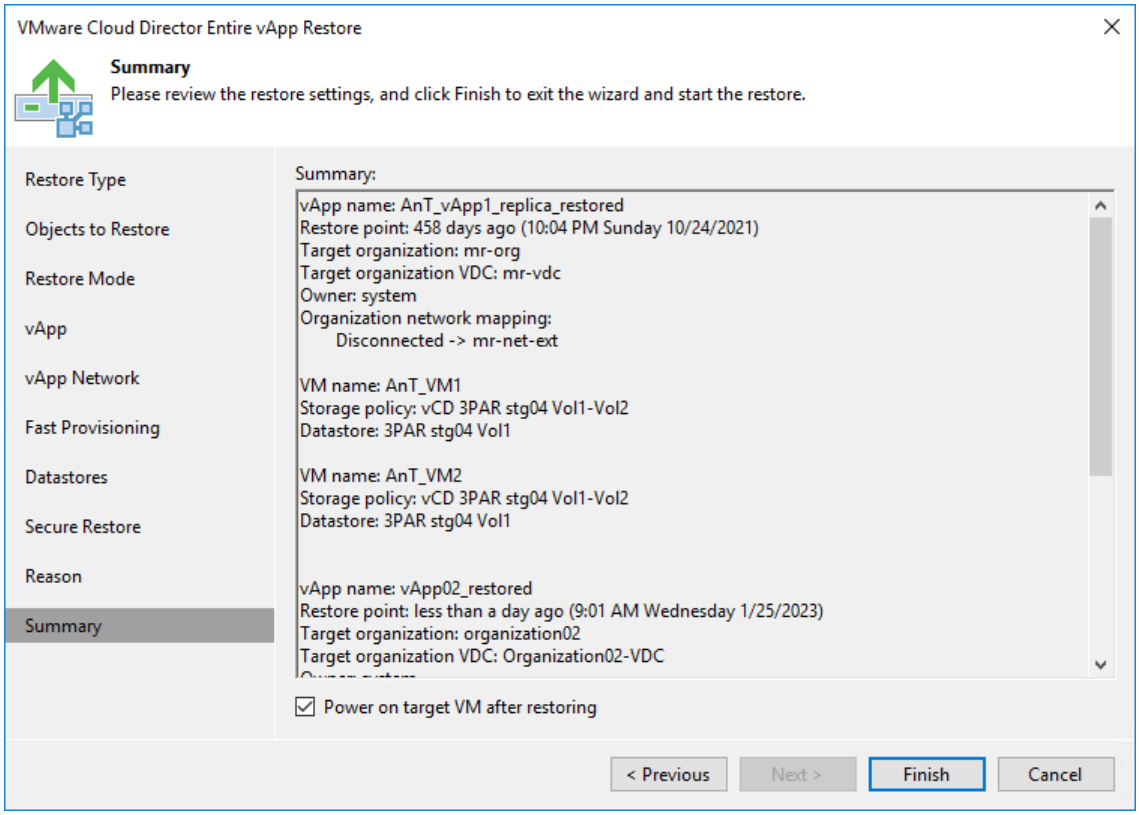

# Item Recovery

Item recovery includes the following methods:

- **VM file restore** — to restore VM files (.VMX, .VMXF and so on) without restoring the entire VM.

  The process of VM files restore for VMware Cloud Director VMs does not differ from that for regular VMware vSphere VMs. For more information, see Restoring VM Files.

- **Restoring VM hard disks** — to restore VM disks. When you restore disks, you extract them from backups to the production storage.

  The process of VM hard disks restore for VMware Cloud Director VMs does not differ from that for regular VMware vSphere VMs. For more information, see Restoring Virtual Disks.

- **Restoring VM guest OS files** — to recover individual guest OS files from Windows, Linux, Mac and other guest OS file systems.

  The process of VM guest files restore for VMware Cloud Director VMs does not differ from that for regular VMware vSphere VMs. For more information, see Restoring VM Guest OS Files (FAT, NTFS or ReFS) and Restoring VM Guest OS Files (Multi-OS).

# VMware Cloud Support

Veeam Backup & Replication supports the following VMware vSphere platforms:

- **VMware Cloud on AWS**. For the information on VMware Cloud on AWS support, see this Veeam KB article.
- **Microsoft Azure VMware Solution (AVS)**. For the information on AVS support, see this Veeam KB article.
- **Google Cloud VMware Engine**. For the information on GCVE support, see this Veeam KB article.
- **IBM Cloud for VMware Solutions**. For the information on IBM Cloud for VMware Solutions support, see this Veeam KB article.
- **Oracle Cloud VMware Solution**. For the information on Oracle Cloud VMware Solutions support, see this Veeam KB article.
- **VMware Cloud on Dell**. For the information on VMware Cloud on Dell support, see this Veeam KB article.
- **VMware Cloud Foundation (VCF)**. This platform is supported as individual VMware software components. VMware components listed on this page can be part of VCF. For the information on the correspondence of VMware components to the VCF version, see this VMware KB article.

# Tape Devices Support

Veeam provides native tape support that is fully integrated into Veeam Backup & Replication. You can administer all operations on tapes from your Veeam console.

For more information, see the Tape Device Support Guide.

# Storage System Snapshot Integration

To build the data protection and disaster recovery strategy, you can use capabilities of storage systems that host VM disks. Veeam Backup & Replication integrates with storage systems and offers advanced functionality that helps you decrease impact from backup and replication operations on the production environment and significantly improve RPOs.

For more information, see the Storage System Snapshot Integration Guide.

# Integration with Veeam Backup for AWS

Veeam Backup & Replication allows you to create and manage data protection and restore tasks for Amazon Elastic Compute Cloud (EC2), Amazon Relational Database Service (RDS), Amazon DynamoDB and Amazon Elastic File System (EFS) environments. For this, Veeam Backup & Replication uses a AWS Plug-in for Veeam Backup & Replication component. This component extends Veeam Backup & Replication functionality and provides access to Veeam Backup for AWS from the Veeam Backup & Replication console. For more information, see the Veeam Backup for AWS User Guide.

# Integration with Veeam Backup for Microsoft Azure

Veeam Backup & Replication allows you to create and manage data protection and restore tasks for Microsoft Azure environments. For this, Veeam Backup & Replication uses a Microsoft Azure Plug-in for Veeam Backup & Replication component. This component extends Veeam Backup & Replication functionality and provides access to Veeam Backup for Microsoft Azure from the Veeam Backup & Replication console. For more information, see the Veeam Backup for Microsoft Azure User Guide.

# Integration with Veeam Backup for Microsoft Entra ID

Veeam Backup & Replication allows you to create and manage data protection and restore tasks for the Microsoft Entra ID environment. For this, Veeam Backup & Replication uses an additional component — Veeam Backup for Microsoft Entra ID. This component extends Veeam Backup & Replication functionality and allows you to back up and restore Microsoft Entra ID tenant data including users, roles, administrative units and others. For more information, see the User Guide for Microsoft Entra ID.

# Integration with Veeam Backup for Google Cloud

Veeam Backup & Replication allows you to create and manage data protection and restore tasks for Google Cloud solution. For this, Veeam Backup & Replication uses a Google Cloud Plug-in for Veeam Backup & Replication component. This component extends Veeam Backup & Replication functionality and provides access to Veeam Backup for Google Cloud from the Veeam Backup & Replication console. For more information, see the Veeam Backup for Google Cloud User Guide.

# Integration with Veeam Backup for Nutanix AHV

Veeam Backup & Replication allows you to manage data protection and restore tasks for Nutanix AHV environments. For this, Veeam Backup & Replication uses an additional component: Nutanix AHV Plug-in for Veeam Backup & Replication. This component extends Veeam Backup & Replication functionality and allows you to back up and restore Nutanix AHV VMs. For more information, see the Veeam Backup for Nutanix AHV User Guide.

# Integration with Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization

Veeam Backup & Replication allows you to manage data protection and restore tasks for oVirt resources. For this, Veeam Backup & Replication uses an additional component: oVirt KVM Plug-in for Veeam Backup & Replication. This component extends Veeam Backup & Replication functionality and allows you to back up and restore oVirt VMs. For more information, see the Veeam Backup for OLVM and RHV User Guide.

# Integration with Veeam Backup for Proxmox VE

Veeam Backup & Replication allows you to manage data protection and restore tasks for Proxmox Virtual Environment. For this, Veeam Backup & Replication uses an additional component: Proxmox Virtual Environment Plug-in for Veeam Backup & Replication. This component extends Veeam Backup & Replication functionality and allows you to back up and restore Proxmox VE VMs. For more information, see the Veeam Backup for Proxmox VE User Guide.

# Integration with Kasten

Veeam Backup & Replication allows you to manage data protection and restore tasks for backups exported with Kasten policies. For this, Veeam Backup & Replication uses the Veeam Kasten Plug-in for Veeam Backup & Replication component.

This component extends Veeam Backup & Replication functionality and allows you to export backups created by Kasten policies to backup repositories. For more information, see the Veeam Kasten Integration User Guide.

# Veeam Agent Management

To back up physical machines running Windows, Linux, Unix or macOS operating systems, Veeam Backup & Replication uses backup agents installed on each computer. Veeam Backup & Replication operates as a centralized control center for deploying and managing Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for IBM AIX, Veeam Agent for Oracle Solaris and Veeam Agent for Mac (Veeam Agents).

For more information about Veeam Agents, see the Veeam Agent Management Guide.

# Veeam Cloud Connect

If you want to store your data in the cloud, you can connect to the service provider and write VM backups to cloud repositories or create VM replicas on cloud hosts.

For more information about Veeam Cloud Connect, see the Veeam Cloud Connect Guide.

# Advanced VMware vSphere Features

Veeam Backup & Replication lets you leverage the following VMware vSphere features and functionality during data protection and disaster recovery operations:

- VM tags
- Encrypted VMs
- Storage policies

# VM Tags

If you use vCenter Server tags to categorize objects in the virtual infrastructure, you can filter objects that you add to data protection and disaster recovery jobs and tasks by these tags. Use of tags facilitates object management. You can quickly configure jobs and tasks for VMs that belong to a specific category, for example, a certain department or SLA level.

To add objects by tags, switch to the **VMs and Tags** or **Tag combination** view in the **Add Objects** window. Veeam Backup & Replication will display objects categorized by tags.

In the **VMs and Tags** view, you can select only one tag and those objects that have this tag will be processed by a job. In the **Tags combination** view, you can select multiple tags and those object that have all the selected tags will be processed.

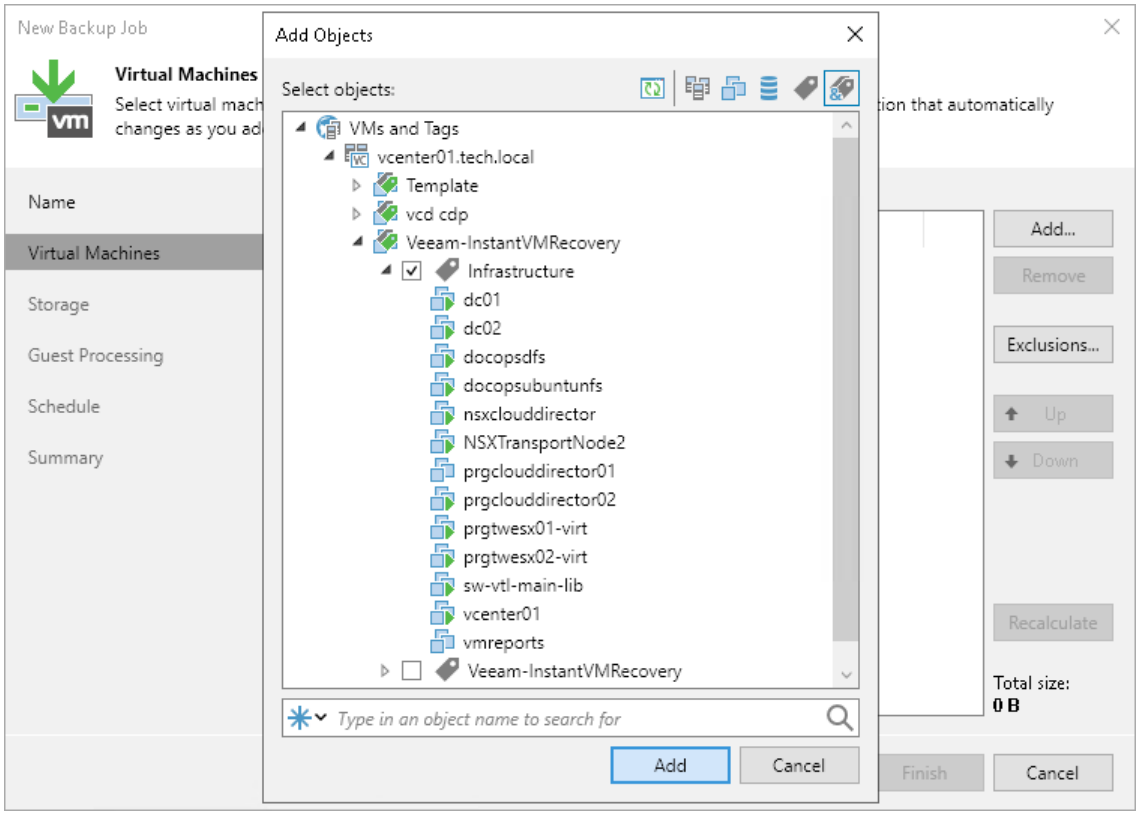


## Requirements and Recommendations for VM Tags

When you work with tags in Veeam Backup & Replication, consider the following requirements and recommendations:

- The *VirtualCenter.FQDN* parameter in the **Advanced Settings** of vCenter Server must contain the real fully qualified domain name of the vCenter Server host.
- A certificate installed on vCenter Server must contain the real fully qualified domain name of the vCenter Server host.
- The fully qualified domain name of the vCenter Server host must be accessible and resolved to its IP (and vice versa) from machines on which Veeam Backup & Replication services are installed (at least the Veeam Backup Service and Veeam Broker Service).

- A user account used for specific data protection and disaster recovery operations must have sufficient permissions on the vCenter Server. For more information, see Full VM Restore, Replica Failback and Cumulative Permissions sections in the Permissions Reference.
- If VM tags are not displayed in the Veeam Backup & Replication console for some reason, try restarting VMware vSphere services that are responsible for the tags functionality. In VMware vSphere earlier than 6.5, you must restart the vCenter Inventory Service. Starting from VMware vSphere 6.5, vCenter Inventory Services functionality is replaced by the vCenter Content Library and other services that are part of vCenter Server 6.5.

  When you upgrade to VMware vSphere 6.5, data from vCenter Inventory Service is migrated to the new database support services in vCenter Server 6.5. The vCenter Inventory Service may remain in the list of services, however, it is no longer used.
- If you are using the vSphere Fault Tolerance feature, assign VM tags in one of the following ways to provide proper VM backup processing by Veeam Backup & Replication:
  - Assign the required VM tag to the primary VM. In case of the failover to the secondary VM, the VM tag is copied to the secondary VM and this VM is used as a source for continuing Veeam Backup & Replication jobs.
  - Assign the required VM tag to the VM container that comprises the primary VM in the virtual infrastructure inventory. The primary VM inherits this tag. In case of the failover to the secondary VM, the VM tag is copied to the secondary VM and this VM is used as a source for continuing Veeam Backup & Replication jobs.

  Thus, there is no need to assign the VM tag to the secondary VM. Moreover, if you manually assign the same tag both to the primary VM and the secondary VM and use this tag to add objects to a VM backup job, this will result in the backup failure.

# Limitations for VM Tags

The VM tags support functionality has the following limitation:

Veeam Backup & Replication ignores the cardinality setting for VM tag categories. For example, you create a tag category *Priority* and set cardinality to **One tag per object**. In the tag category, you create two tags: *Normal* and *High*. You assign the Normal tag to a VM folder and the *High* tag to a VM in this folder. If you now configure a job that will process objects with the *Normal* tag, the VM with the *High* tag will also be added to this job (since Veeam Backup & Replication regards that VMs and templates in the VM container inherit the tag assigned to the container).

To overcome this situation, you can add to the list of exclusions the tag assigned to objects that you do not want to process.

# Encrypted VMs

> **NOTE**
>
> All limitations and considerations below also apply to a VM with a Virtual Trusted Platform Module (vTPM) as vTPM requires VM encryption to be enabled.

Veeam Backup & Replication provides support for VMware vSphere encrypted VMs.

- Backup of encrypted VMs
- Restore of encrypted VMs
- Replication of encrypted VMs
- Failback of encrypted VM replicas

> **NOTE**
>
> CDP meets specific requirements for VMs encrypted on VMware side. For more information, see the Virtual Machines section.

## Backup of Encrypted VMs

To back up VMware encrypted VMs, the backup infrastructure must meet the following requirements:

- VM encryption instances must be preconfigured in the virtual infrastructure: you must set up the Key Management Server (KMS), create the VM encryption policy, and assign it to VMs in advance.
- The backup proxy must be working in the **Virtual appliance** or **Network** transport mode:
  - The backup proxy working in the **Virtual appliance** transport mode must be deployed on an encrypted VM.
  - The backup proxy working in the **Network** transport mode uses the NBD protocol by default. If you want to use NBDSSL, select the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box in the **Transport Mode** window. Note that traffic encryption puts more stress on the CPU of an ESXi host and can decrease performance.

## Restore of Encrypted VMs

Veeam Backup & Replication supports the following restore options:

- Restore encrypted VM as encrypted or unencrypted.
- Restore unencrypted VM as encrypted.

The backup infrastructure must meet the following requirements:

- VM encryption instances must be preconfigured in the virtual infrastructure: you must set up the Key Management Server (KMS), create the VM encryption storage policy, and assign it to VMs in advance.
- The backup proxy must be working in the **Virtual appliance** or **Network** transport mode:
  - The backup proxy working in the **Virtual appliance** transport mode must be deployed on an encrypted VM.

- The backup proxy working in the **Network** transport mode uses the NBD protocol by default. If you want to use NBDSSL, select the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box in the **Transport Mode** window. Note that traffic encryption puts more stress on the CPU of an ESXi host and can decrease performance.

If you restore a VM as an encrypted one to the specified location, ensure that the target datastore is under the **VM Encryption Policy** node.

If a VM has several disks, you can optionally restore some disks as encrypted and some disks as unencrypted. Keep in mind, that even if one disk is restored as encrypted, the VM configuration file must also be placed on a datastore under the **VM Encryption Policy** node.

## Replication of Encrypted VMs

To replicate VMware encrypted VMs, the backup infrastructure must meet the following requirements:

- VM encryption instances must be preconfigured in the virtual infrastructure: you must set up the Key Management Server (KMS), create the VM encryption policy, and assign it to VMs in advance. Ensure that you use a common KMS or the KMS clusters at both sites use common encryption keys.

  > **NOTE**
  >
  > If you do not set up KMS, the replication job will not fail, but replicated VMs will not be encrypted in this case.

- Source and target backup proxies must be working in the **Virtual appliance** or **Network** transport mode:
  - The backup proxy working in the **Virtual appliance** transport mode must be deployed on an encrypted VM.
  - The backup proxy working in the **Network** transport mode uses the NBD protocol by default. If you want to use NBDSSL, select the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box in the **Transport Mode** window. Note that traffic encryption puts more stress on the CPU of an ESXi host and can decrease performance.

To replicate a VM as an encrypted one, place disks and the configuration file of the VM replica on datastores compatible with the VM encryption policy:

- At the **Destination** step of the wizard, click **Choose** near the **Datastore** field.
- In the **Select Datastore** window, select a datastore under the **VM Encryption Policy** node.

> **NOTE**
>
> Multi-OS guest OS file restore for encrypted VM replicas is not supported.

## Failback of Encrypted VM Replicas

To fail back VMware encrypted VMs replicas, the backup infrastructure must meet the following requirements:

- VM encryption instances must be preconfigured in the virtual infrastructure: you must set up the Key Management Server (KMS), create the VM encryption policy and assign it to VMs in advance. Ensure that you use a common KMS or the KMS clusters at both sites use common encryption keys.

- Source and target backup proxies must be working in the **Virtual appliance** or **Network** transport mode:
    - The backup proxy working in the **Virtual appliance** transport mode must be deployed on an encrypted VM.
    - The backup proxy working in the **Network** transport mode uses the NBD protocol by default. If you want to use NBDSSL, select the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box in the **Transport Mode** window. Note that traffic encryption puts more stress on the CPU of an ESXi host and can decrease performance.

If you fail back an encrypted VM replica to the specified location, ensure that the target datastore is under the **VM Encryption Policy** node.

# Storage Profiles

During backup, Veeam Backup & Replication preserves information about the storage policy associated with the VM, and stores this information to the backup file or replica metadata. When you restore the VM to its original location, Veeam Backup & Replication also restores information about the VM storage policy. The restored VM gets automatically associated with the original storage policy.

Veeam Backup & Replication restores the storage policy when you perform the following operations:

- Entire VM restore
- VM failback

Veeam Backup & Replication restores the storage policy only if you restore the VM to the original location. If you restore the VM to a new location, Veeam Backup & Replication does not preserve the storage policy for the VM.

In some cases, the original storage policy may be changed or missing by the time when you restore the VM. For example, the storage policy may be deleted, the original VM in the production environment may be associated with another storage policy and so on. In such situation, Veeam Backup & Replication displays a warning and lets you choose one of the following scenarios:

- Associate the VM with the current storage policy — the restored VM will be associated with the profile with which the original VM in the production environment is currently associated.
- Associate the VM with the default storage policy — the restored VM will be associated with the profile that is set as default for the target datastore.
- Associate the VM with the profile stored in the backup file — the restored VM will be associated with the profile that was assigned to the original VM at the moment of backup, and whose information is stored in the backup file.

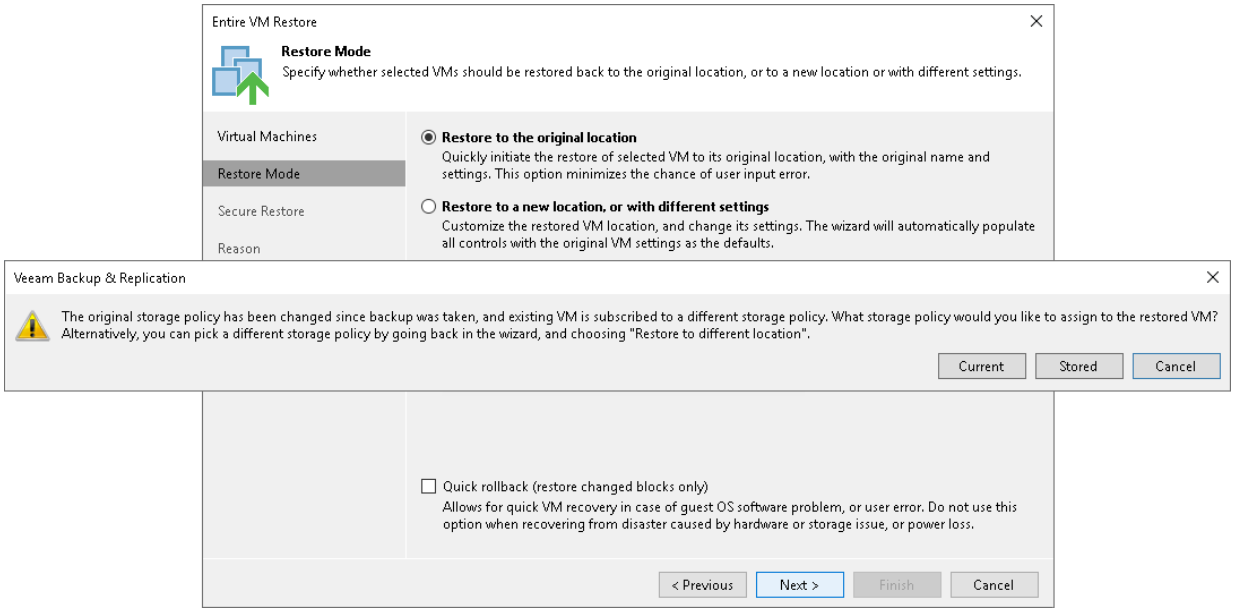

# Veeam Backup & Replication Utilities

You can use the following Veeam Backup & Replication utilities to perform advanced administration tasks in your backup infrastructure:

- Extract Utility
- Veeam VM Migrator Utility
- Veeam Configuration Database Connection Utility
- Veeam Backup Validator
- Veeam Backup Configuration Tool

# Extract Utility

Veeam Backup & Replication comes with an extract utility that can be used to recover machines from backup files. The extract utility does not require any interaction with Veeam Backup & Replication and can be used as an independent tool on Linux and Microsoft Windows machines.

The extract utility can be helpful, for example, if it is written to the tape next to machine backup files. In this case, you get a possibility to recover machines from backups at any moment of time even if backups are removed from Veeam Backup & Replication or Veeam Backup & Replication is not installed.

> **IMPORTANT**
>
> If you want to use the extract utility to work with backup files located on any of the extents of your scale-out backup repository, make sure that incremental and full backup files are located on the same extent.

The extract utility can be used in two interfaces:

- Graphic user interface (GUI)
- Command-line interface working in the interactive and regular mode

The extract utility is located in the installation folder of Veeam Backup & Replication, by default: `%PROGRAMFILES%\Veeam\Backup and Replication\Backup`. The folder contains three files for the extract utility:

- `Veeam.Backup.Extractor.exe` — utility working in GUI (can be used on Microsoft Windows machines only)
- `extract.exe` — utility working in the command-line interface, a version for Microsoft Windows
- `extract` — utility working in the command-line interface, a version for Linux

# Using Extract Utility in GUI

To restore machine data in the extract utility GUI:

1. Run the `Veeam.Backup.Extractor.exe` file from the installation folder of Veeam Backup & Replication.
2. In the **Backup file** field, specify a path to the backup file from which you want to restore machine data.
3. If the backup file is encrypted, the extract utility will require you to provide a password to unlock the backup file.
4. In the **Target folder** field, specify a path to the destination folder where machine data must be restored.
5. From the **Machines** list, select machines whose data you want to restore.
6. Click **Extract**. Machine data will be restored to the specified folder.

**IMPORTANT**

If you restore machine data in the extract utility GUI, consider the following:

- The extract utility can be started on Microsoft Windows machines only.
- If you plan to start the extract utility on the machine other than the backup server, make sure that you copy the `Veeam.Backup.Extractor.exe` file together with the extract.exe file from the product installation folder and store these files to the same folder on the destination machine. In the opposite case, the extract utility will fail to start.

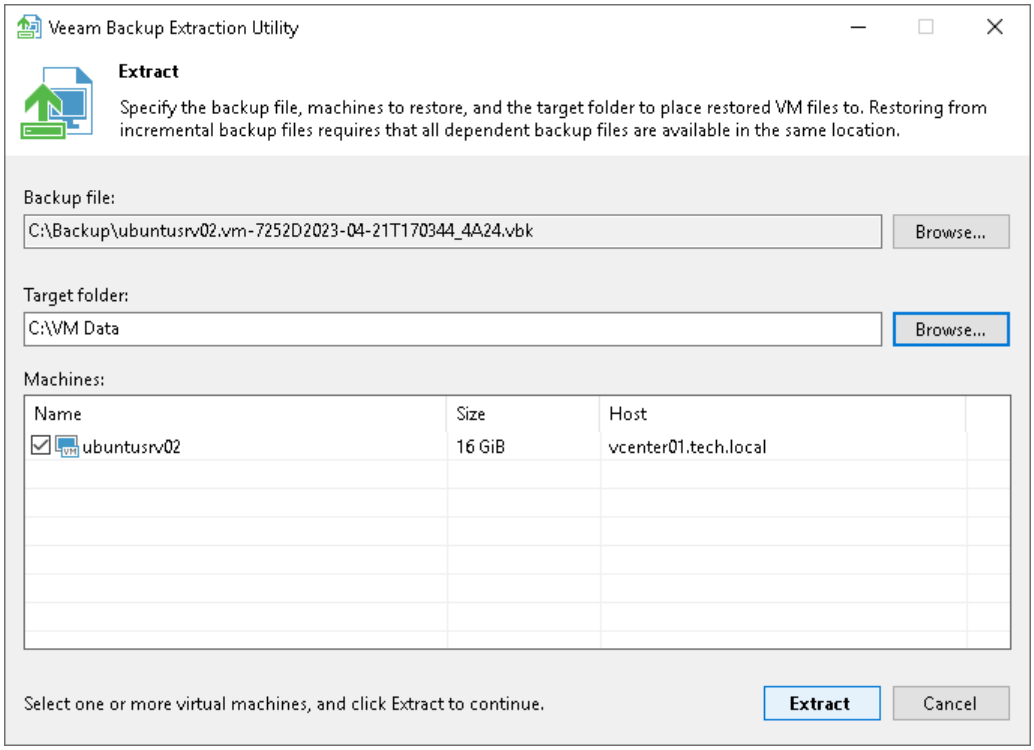

# Using Extract Utility in Interactive Mode

To start the extract utility in the interactive mode, run the `extract.exe` file from the product installation folder(in case of a Linux machine, run the `extract` file).

You will have to sequentially enter the following arguments:

1. A path to the backup file from which the machine must be restored. After you enter the path, the extract utility will display a list of all machines included in the backup and their description.
2. A name of the machine that you want to restore. If there is more than one machine with the specified name in the backup, you will be asked to specify the host on which the backed-up machine resides. If you want to restore all machines from the backup, press [Enter] on the keyboard.
3. If the backup was encrypted, password that was used to encrypt the backup file.
4. An output directory to which machines must be restored. If you want to restore machines to the current directory, press [Enter] on the keyboard.
5. The operation confirmation. Press [Y] on the keyboard to restore a machine to the directory you specified. If you want to abort the operation, press [Enter] on the keyboard.

# Using Extract Utility from Command Line

To run the extract utility from the command line, do one of the following:

- In the command line, change the current directory to the directory where the extract utility is located.
- Add the directory where the extract utility is located to the PATH variable.

Then you can perform the following actions:

- Run the extract utility in the interactive mode
- Display help information for the utility usage
- Display the list of all VMs in the backup file
- Getting encryption status of a backup file
- Restore all or selected VMs from the backup

## Running Extract Utility in Interactive Mode

This command runs the extract utility in the interactive mode.

## Syntax

```
extract.exe [-password backupkey] [pathtobackup]
```

## Parameters

| Parameter | Description | Required/Optional |
|---|---|---|
| **password** | Password for the encrypted backup file. | Required for encrypted backup files |
| **pathtobackup** | Path to the backup file from which machines must be restored. | Optional |

## Example

This command validates an encrypted backup file.

```
extract.exe -password "standard 1" "C:/Backup/Single/Backup Job Single Storage
D2022-10-03T132735_1E50.vbk"
```

# Displaying Help Information for Utility Usage

This command prints all variants of the extract utility usage along with required and optional parameters.

# Syntax

```
extract.exe -help
```

# Displaying List of Machines in Backup

This command displays the list of all machines in the backup file from which you want to perform restore.

# Syntax

```
extract.exe -dir [-vm vmname] [-host hostname] [-password backupkey] pathtoback
up
```

# Parameters

| Parameter | Description | Required/Optional |
|---|---|---|
| **vm** | Name of the machine that you want to restore. Use this parameter to filter machines in the backup. | Optional |
| **host** | Name of the host on which the initial machine resides. Specify this parameter to filter machines that have the same name but reside on different hosts.<br>**Note**: This parameter must be specified if the vm parameter is used. | Optional |
| **password** | Password for the encrypted backup file. | Required for encrypted backup files |
| **pathtobackup** | Path to the backup file from which the machine must be restored. | Required |

# Getting Encryption Status of Backup File

This command gets the encryption status of the backup file: encrypted or not encrypted.

## Syntax

```
extract.exe -getEncryptionStatus pathtobackup
```

## Parameters

| Parameter | Description | Required/Optional |
|---|---|---|
| **pathtobackup** | Path to the backup file from which the machine must be restored. | Required |

# Restoring VMs from Backup

This command restores data for all machines or for the selected machine from the backup file.

## Syntax

```
extract.exe -restore [-vm vmname] [-host hostname] [-password backupkey] pathto
backup [outputdir] [-log]
```

## Parameters

| Parameter | Description | Required/Optional |
|---|---|---|
| **vm** | Name of the machine that you want to restore. Use this parameter to filter machines in the backup. If you want to restore all machines from the backup file, do not specify this parameter. | Optional |
| **host** | Name of the host on which the initial machine resides. Specify this parameter to filter machines that have the same name but reside on different hosts.<br>**Note**: This parameter must be specified if the vm parameter is used. | Optional |
| **pathtobackup** | Path to the backup file from which the machine must be restored. | Required |
| **password** | Password for the encrypted backup file. | Required for encrypted backup files |

| Parameter | Description | Required/Optional |
| --- | --- | --- |
| **outputdir** | Path to the directory to which machine data must be restored. If this parameter is not specified, the machine will be restored to the current directory. | Optional |
| **log** | SwitchParameter. Enables log creation. The log file will be created in the current directory. | Optional |

# Example

This command restores `winsrv29` machine from an encrypted backup file to the `C:/Backup` directory.

```
extract.exe -restore -vm winsrv29 -password "standard 1" "C:/Backup/Single/Bac
kup Job Single StorageD2022-10-03T132735_1E50.vbk" C:/backup
```

# Veeam VM Migrator Utility

Veeam Backup & Replication tracks VMs in jobs using Managed Object Reference IDs (MORef-IDs), which change after migration or recreation of vCenter, causing MORef-ID misalignment.

Veeam VM Migrator utility is integrated into Veeam Backup PowerShell module and it allows you to resolve MORef-ID misalignment. As a result, your backup incremental chains will remain intact after an inventory change in vCenter.

The utility consists of the following cmdlets:

- `Set-VBRVmBiosUuid` — this cmdlet updates the BIOS UUIDs of existing VM entries within the Veeam Backup & Replication configuration database based on information from the old vCenter.
- `Set-VBRVCenterName` — this cmdlet modifies vCenter name by adding the `_old` suffix to its name.
- `Generate-VBRViMigrationSpecificationFile` — this cmdlet generates a migration task file which contains the list of mapping tasks.
- `Start-VBRViVMMigration` — this cmdlet starts MORef-IDs update.

For more information on the utility cmdlets, see Veeam PowerShell Reference guide.

## How VM Migrator Utility Works

Veeam VM Migrator utility performs migration the following way:

1. [Optional] Updates BIOS UUIDs of VMs located on the old vCenter to ensure future mapping accuracy.
2. [Optional] Renames the old vCenter name by adding the `_old` suffix to its name.
3. Scans old and new vCenters for VM references, and creates mapping tasks.
4. Generates a migration task file which contains mapping tasks.
5. Uses previously generated migration task file to perform VMs MORef-IDs update.

# Using Veeam VM Migrator Utility

Each procedure step includes running utility cmdlets in Veeam Backup PowerShell module.

To use Veeam VM Migrator utility, check the prerequisites and start Veeam PowerShell session.

## Before You Begin

Before using Veeam VM Migrator, consider the following:

- You must create a configuration backup of Veeam Backup & Replication database before using Veeam VM Migrator.
- Do not run backup jobs for VMs from the new vCenter until you finish working with Veeam VM Migrator. Otherwise, if there are any VMs backed up from the new vCenter, it will not be possible to migrate these VMs because new VM objects will already exist in the database.
- If you use the replication feature:
  - Veeam VM Migrator supports source vCenter migration only. If replication target vCenter is changed, you will need to manually map source VMs to their replicas.
  - Replica VMs must be mapped manually using the replica mapping feature in Veeam Backup & Replication console.

    For more information on replica mapping, see Replica Seeding and Mapping.
- Virtual labs should be recreated manually after migration.
- The VM containers will not be updated, you will need to manually re-add all folders, tags, datastores, clusters and resource pools in the list of VMs in backup and/or replica job settings.
- You can not resolve MORef-ID misalignment for jobs that target a Cloud Connect repository.

# Step 1. Start Veeam PowerShell Session

To use Veeam VM Migrator Utility, start Veeam PowerShell session in one of the following ways:

- **Using the Veeam Backup & Replication UI**: in the main menu of Veeam Backup & Replication, select **Console** > **PowerShell**.
- **From Windows PowerShell console**: start the console on your machine. For more information on starting the console, see Microsoft Docs.

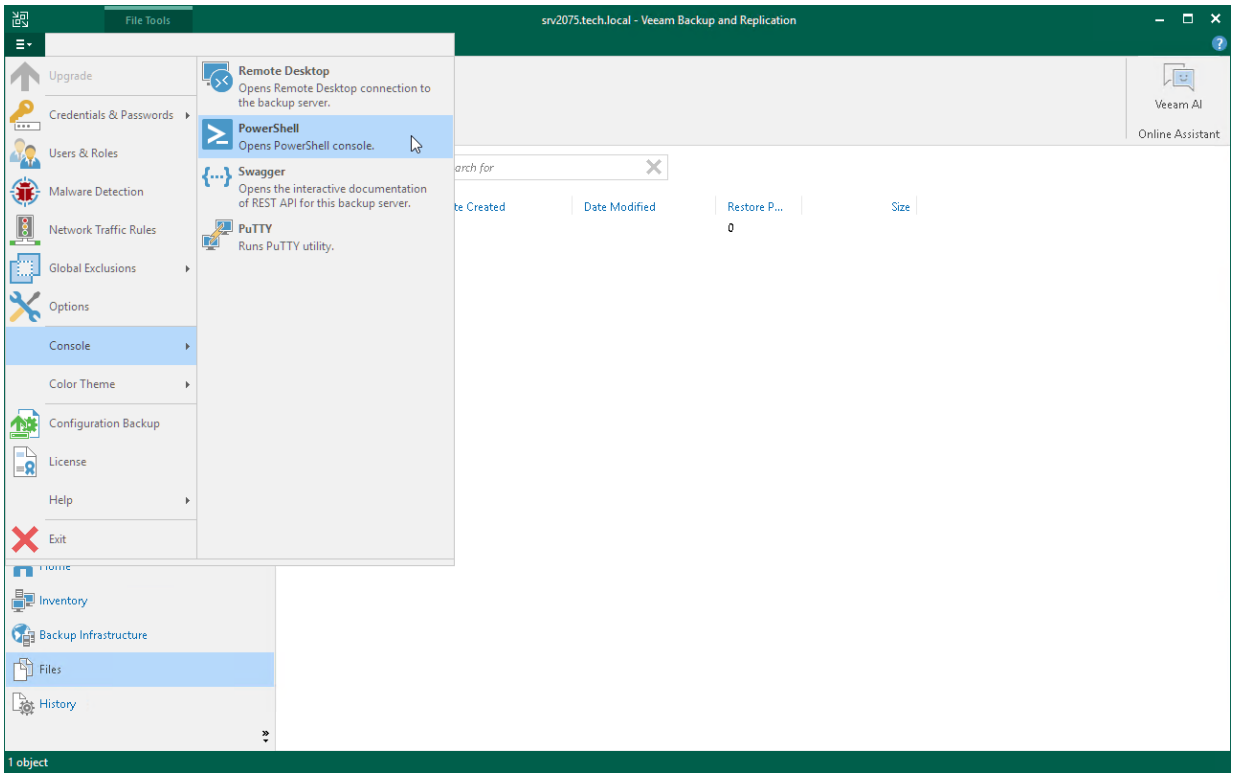

# Step 2. Update Existing VMs BIOS UUIDs

This is an optional step that can be performed to update the BIOS UUIDs of existing VM entries within the Veeam Backup & Replication configuration database based on information from the old vCenter, ensuring future matching accuracy. Skip this step if the old vCenter is no longer operational or if you have already migrated VMs to the new vCenter.

To update BIOS UUIDs of VMs, run the `Set-VBRVmBiosUuid` cmdlet and specify the `VCenterName` parameter value.

This example shows how to update BIOS UUIDs of VMs located on the old vCenter named `vcenter70`.

```
Set-VBRVmBiosUuid -VCenterName "vcenter70"
```

# Step 3. Modify Old VMware vCenter Name

This is an optional step that modifies the old vCenter name by adding the `_old` suffix to its name. Perform this step if the new vCenter has not been added yet, and it will have the same name as the old vCenter.

To rename the old vCenter, run the `Set-VBRVCenterName` cmdlet and specify the `OldVCenterName` parameter value.

This example shows how to rename the vCenter named `vcenter70`.

```
Set-VBRVCenterName -OldVCenterName "vcenter70"
```

# Step 4. Add New vCenter

This is an optional step which is required if the new vCenter has not been added to the backup infrastructure.

To add the new vCenter, see Adding VMware vSphere Servers section of the guide and complete the **New VMware Server** wizard.

After that, ensure that you have migrated all VMs for which you want to resolve MORef-ID misalignment. For more information, see Quick Migration.

# Step 5. Generate Migration Task File

To generate the migration task file, run the `Generate-VBRViMigrationSpecificationFile` cmdlet and specify the `ExportPath`, `NewVCenterName`,and `OldVCenterName` parameter values.

This example shows how to generate the migration task file for the old vCenter named `vcenter70_old` and the new vCenter named `vcenter70`. The migration task file will be generated at the following path: `C:\Folder`.

```
Generate-VBRViMigrationSpecificationFile -ExportPath C:\Folder -NewVCenterName
vcenter70 -OldVCenterName vcenter70_old
```

# Step 6. Review Migration Task File

After you generate the migration task file, you must carefully review and edit it. The file contains several sections with mapping tasks. Each line contains source VM MORef-ID, source VM name and BIOS UUID field values. These field values are followed by target VM MORef-ID and target VM name.

# Migration Task File Sections

| Section | Description |
|---|---|
| **hosts names** | This section lists vCenter server names (old and new). |
| **solid** | This section lists VMs which exist in the old vCenter and in the new vCenter with the same names and BIOS UUIDs. |
| **probable** | This section lists the following:<br>• VMs with the same BIOS UUID but with different names.<br>• VMs with the same name but different BIOS UUIDs.<br>• VMs which have more than one instance of the same VM name and BIOS UUID present in the database. |
| **old duplicates** | For entries having several database records with the same name and BIOS UUID, mapping is defined only for the VMs which last restore point is the newest restore point. The rest of such objects fall into the **old duplicates** section. Mapping for these objects is not defined. |
| **already migrated** | This section lists VMs which are already linked to the new vCenter, with the names and BIOS UUIDs matching the current values on the new vCenter. Mapping for these objects is not defined. |
| **not found** | This section lists all VMs that did not fall into other sections but previously existed in the old vCenter. Mapping for these objects is not defined. |

When editing migration task file, consider the following:

- In the migration task file field values must be separated by a tab.
- If you need to comment out a line, use a tab followed by double slash character "//".
- If task file contains multiple mappings for the same VM, they will be commented out. Uncomment one of them if you want to apply it.

# Step 7. Start Migration

The utility performs migration according to the migration task file. To start the migration, run the `Start-VBRViVMMigration` cmdlet and specify the `File` parameter value.

```
Start-VBRViVMMigration -File C:\Folder\vcenter70_old_to_vcenter70_migration_tas
k
```

This operation will update Veeam Backup & Replication database. If you encounter issues after migration, perform the restore of Veeam Backup & Replication configuration database. For more information, see Restoring Configuration Database.

# Step 8. Finish Migration

To finish the migration process, do the following:

- If you have replication jobs, map VM replicas manually to each replication job.
- Re-add VM containers to the list of source VMs in the jobs settings. Click **Recalculate** to make sure that all source objects were found.
- Recreate virtual labs.
- If you use exclusions to remove objects from processing in your jobs, verify and update the exclusion lists for each job.

# Veeam Configuration Database Connection Utility

Veeam Backup & Replication comes with the configuration database connection utility that allows you to manage connection settings for Veeam Backup & Replication, Veeam Backup Enterprise Manager and Microsoft Entra ID. Using this utility, you can:

- Connect to a different database on the same or another Microsoft SQL Server or PostgreSQL instance. If you specify a database that does not exist yet, it will be created on the selected server.
- Change the authentication method for the database connection. Possible methods are Microsoft Windows authentication and Microsoft SQL Server or PostgreSQL native authentication.

This section of the guide explains how to manage database connection settings for Veeam Backup & Replication. If you want to manage database connection settings for Veeam Backup Enterprise Manager and Microsoft Entra ID, see Veeam Backup Enterprise Manager Guide and User Guide for Microsoft Entra ID.

> **NOTE**
>
> Consider the following:
>
> - The configuration database connection utility supports only connection to configuration databases of the current version.
> - The configuration database connection utility is shared between Veeam Backup & Replication and Veeam Backup Enterprise Manager. If they are installed on the same machine, make sure these products are of the same version.

# Using Veeam Configuration Database Connection Utility

You can launch the configuration database connection utility the following ways:

- From the **Start** menu by clicking **Configuration Database Connection Settings**.
- By using the `Veeam.Backup.DBConfig.exe` file located in the installation folder. The default path to the folder is the following: `%PROGRAMFILES%\Common Files\Veeam\Backup and Replication\DBConfig`.
- By using the `Veeam.Backup.DBConfig.exe` file located in the ISO file. The path to the file is the following: `%ISO%:\Tools\DBConfig`.

To run the utility, you must have administrative rights on the local machine, as long as the utility makes changes to the registry. If prompted at the launch, choose **Run as administrator**.

To manage connection settings for Veeam Backup & Replication, check prerequisites and use the launched **Veeam Backup & Replication Configuration Database Connection Settings** wizard.

## Before You Begin

Before you configure database connection settings, consider the following:

- If you change the database to which Veeam Backup & Replication must be connected, make sure that the database to which Veeam Backup & Replication is currently connected is available. If not, you must stop the Veeam Backup Service on the machine where Veeam Backup & Replication is installed.
- If you change the authentication method for the database connection, make sure that user account has access rights to the following services:
  - Veeam Broker Service
  - Veeam CDP Coordinator Service
  - Veeam Cloud Connect Service
  - Veeam Backup Service
  - Veeam Data Analyzer Service
- Do not connect Veeam Backup & Replication to the database which is being used by another Veeam Backup & Replication. In that case, Veeam Backup & Replication will not be able to utilize the encrypted data stored in this database. If you want to migrate the configuration database, see Migrating Configuration Database to Another SQL Server or Migrating Configuration Database to PostgreSQL Server.
- When you migrate the configuration database to another server, you must use the Microsoft SQL Server credentials that have `CREATE ANY DATABASE` permission on the target Microsoft SQL Server. For details, see Microsoft Docs.

  After database creation this account automatically gets a `db_owner` role and can perform all operations with the database. If the current account does not have this permission, a Database Administrator may create an empty database in advance and grant the `db_owner` role to the account that will be used for migration of the configuration database.

# Step 1. Select Product

At the **Product** step of the wizard, select the database which settings you want to configure.

The utility detects what servers are installed on the local machine and displays available products for your choice. If you want to manage connection settings for Veeam Backup Enterprise Manager or Microsoft Entra ID backup repository, see Veeam Backup Enterprise Manager Guide or User Guide for Microsoft Entra ID.

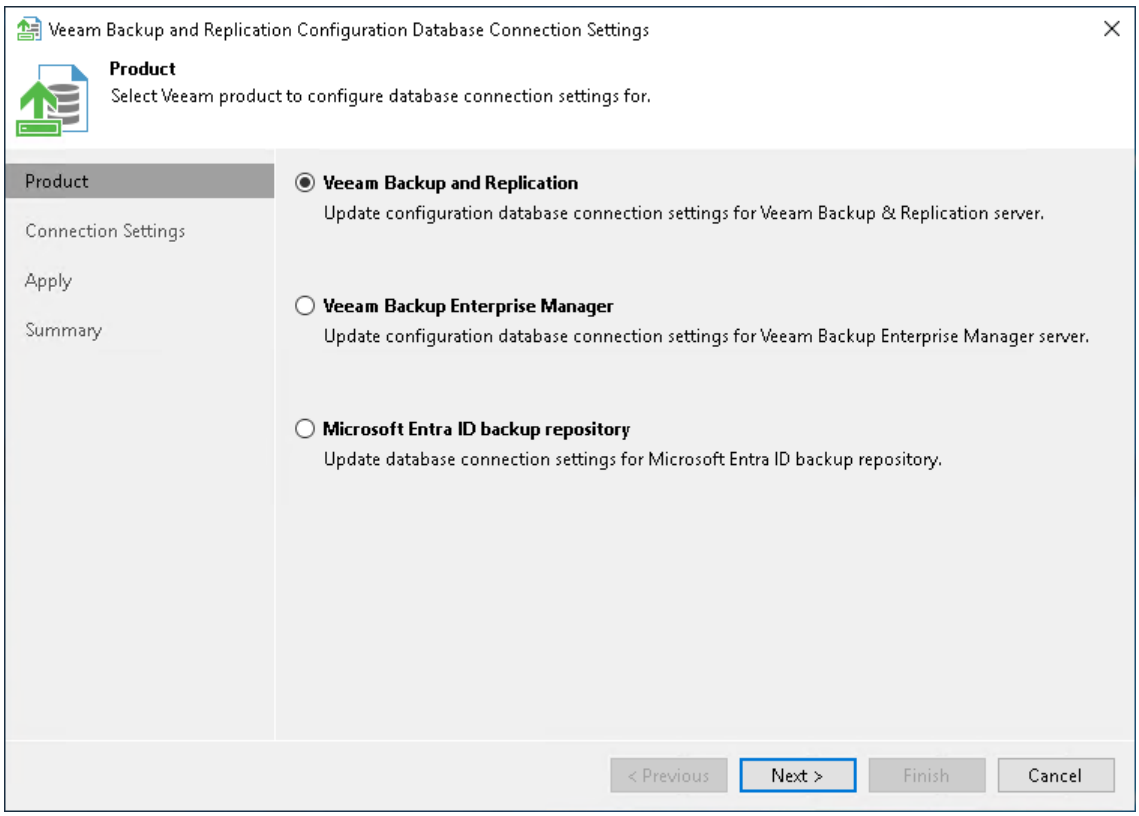

# Step 2. Specify Connection Settings

At the **Connection Settings** step of the wizard, provide the connection settings for the selected database.

Providing connection settings for Microsoft SQL Server:

1. Specify the Microsoft SQL Server database engine.
2. Specify an instance and a database name to which you want the Veeam Backup & Replication installation to connect in `localhost\instanceName` format.

   Both local and remote Microsoft SQL Server instances are supported. Instances available on the network are shown in the Instance name list. If necessary, click **Refresh** to get the latest information.

   If a database with the specified name does not exist on the selected instance, it will be created anew.
3. Select the authentication method that will be used for database connection.
   - If you plan to use the Microsoft Windows authentication, consider that the current service account will be used (that is, the account under which the Veeam Backup Service is running).
   - If you plan to use the Microsoft SQL Server native authentication, provide a login name and a password. To view the entered password, click and hold the eye icon on the right of the **Password** field.

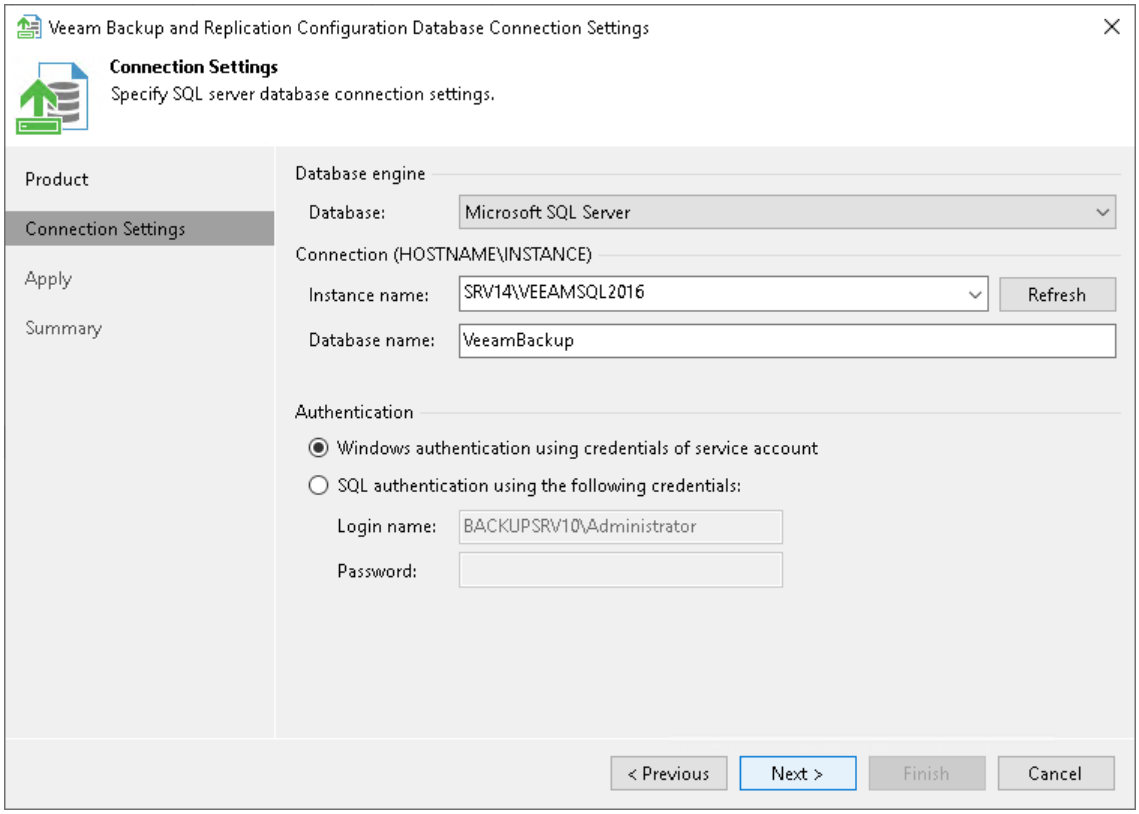


Providing connection settings for PostgreSQL:

1. Specify the PostgreSQL database engine.
2. Specify an instance and a database name to which you want the Veeam Backup & Replication installation to connect in `localhost:instancePort` format.

   Both local and remote PostgreSQL instances are supported. If a database with the specified name does not exist on the selected PostgreSQL instance, it will be created anew.

3. Select the authentication method that will be used for database connection:
   - If you plan to use the Microsoft Windows authentication, consider that the current service account will be used (that is, the account under which the Veeam Backup Service is running).
   - If you plan to use the PostgreSQL native authentication, provide a login name and a password. To view the entered password, click and hold the eye icon on the right of the **Password** field.

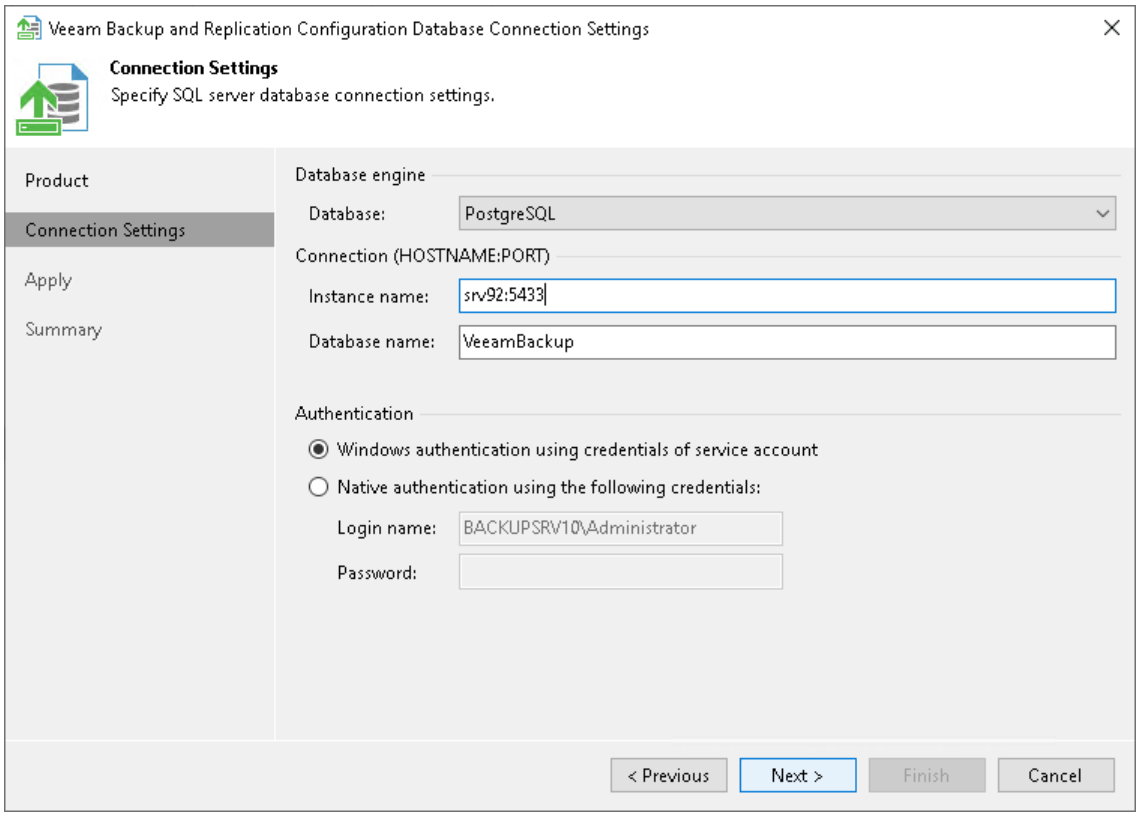

# Step 3. Apply Connection Settings

Before proceeding, the utility validates the specified settings to make sure that the user account has enough privileges to access the database.

- If you have selected the Microsoft Windows authentication method, the utility will check the privileges of the current user account (that is, the account under which the utility is running) to connect to specified Microsoft SQL Server.
- If you have selected the Microsoft SQL Server or the PostgreSQL native authentication method, the utility will check the privileges of the account you have specified.

To ensure that these accounts (as well as the account under which the Veeam Backup Service is running) have sufficient privileges for database access, you can contact your database administrator. Refer to the list of required permissions for Veeam Backup & Replication for detailed information.

For the new settings to be applied, the utility needs to stop Veeam Backup & Replication services and jobs that are currently running. Before proceeding to the **Apply** step, you must confirm the operation. After you confirm the operation by clicking **Yes**, Veeam Backup & Replication will force services and jobs to stop, and will apply database connection settings. For example, if you are configuring Veeam Backup & Replication database settings, the following prompt will be displayed.

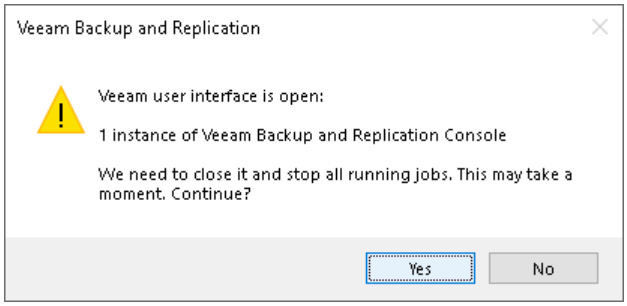

Wait for the operation to complete and click **Next** to proceed to the **Summary** step of the wizard. Previously stopped services will be started again at this moment.

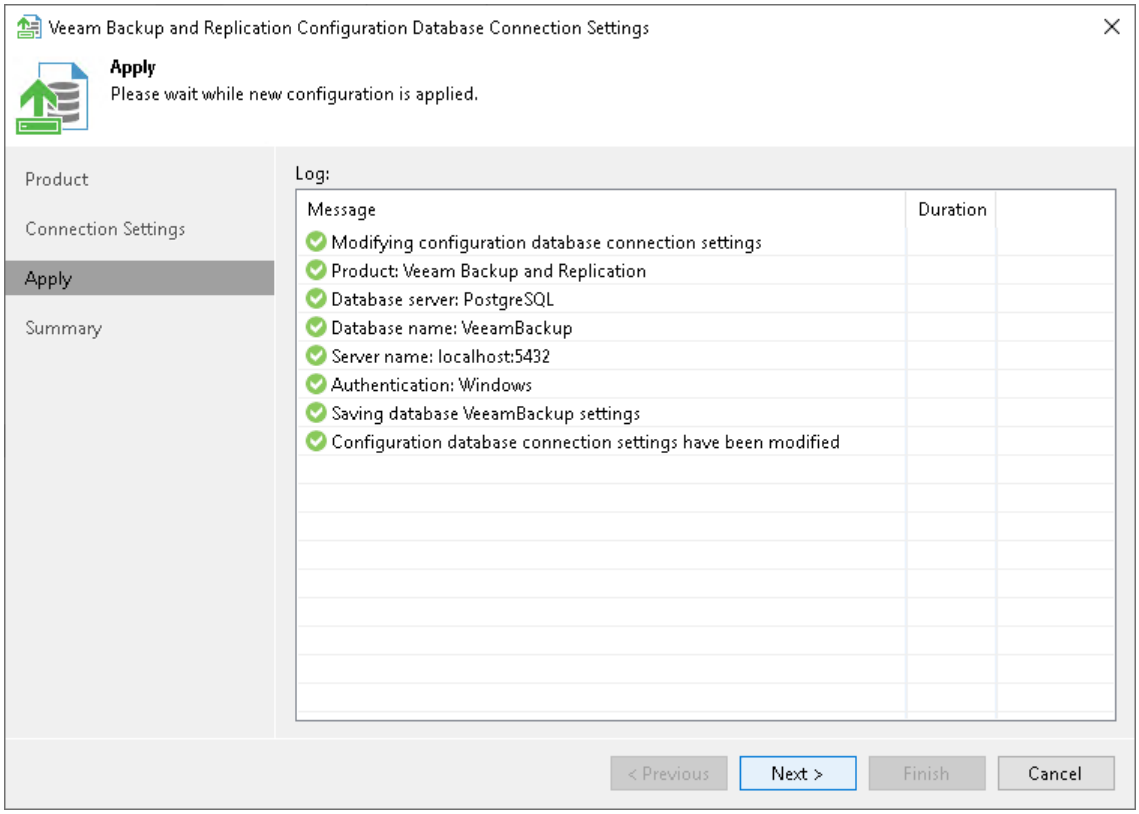

# Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, view the information about the changes in database connection settings. If you were configuring Veeam Backup & Replication database settings and you want the Veeam backup management console to be opened automatically after you finish working with the wizard, select the **Start the product automatically** check box.

> **NOTE**
>
> The **Start the product automatically** option is not available for Veeam Backup Enterprise Manager.

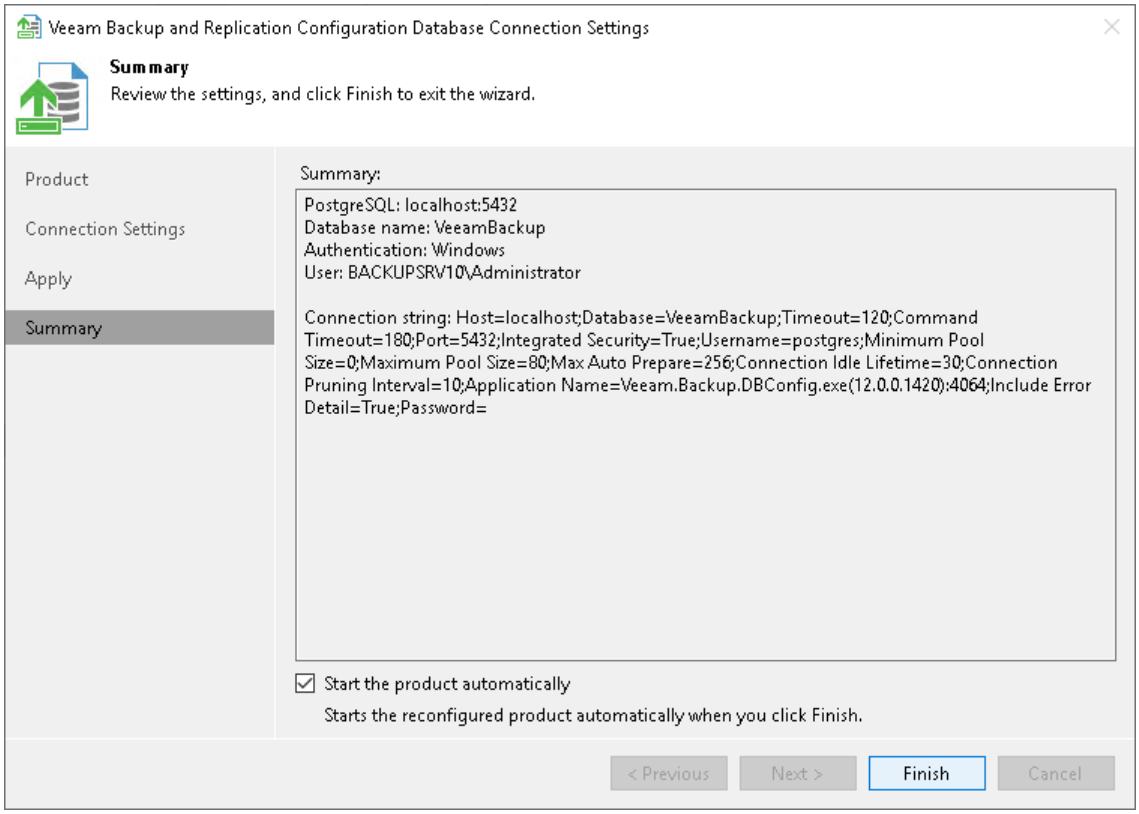

# Veeam Backup Validator

Veeam Backup Validator is an utility that verifies the integrity of a backup file without extracting VM data. Veeam Backup Validator is a command-prompt CRC check utility that tests a backup at the file level. You may need this utility to check whether backup files were damaged: for example, after hardware failures occurred in a backup storage side or if backup files were transferred over network.

For integrity validation, Veeam Backup Validator uses the checksum algorithm. When Veeam Backup & Replication creates a backup of a VM, it calculates a checksum for every data block in the backup file and attaches these checksums to data blocks. Veeam Backup Validator recalculates checksums for data blocks and compares them against the initial checksum values. If the results match, the backup file is viable.

# Using Veeam Backup Validator

Veeam Backup Validator is located on the backup server in the installation folder of Veeam Backup & Replication — by default, `%ProgramFiles%\Veeam\Backup and Replication\Backup\Veeam.Backup.Validator.exe`. If the default path was changed, you can find the actual path in the following registry value: `[HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication] CorePath`.

To run the utility, open the command prompt or the PowerShell console on the backup server and change the current folder to the folder where Veeam Backup Validator is located.

**IMPORTANT**

Consider the following:

- To run Veeam Backup Validator, you must use an account with administrative rights on the local machine. Also, you must disable multi-factor authentication (MFA) for this account. For more information, see Disabling MFA for Service Accounts.
- You cannot use Veeam Backup Validator to verify integrity of unstructured data backups.
- You cannot use Veeam Backup Validator to verify integrity of backups created in Veeam Cloud Connect repositories. For more information on Veeam Cloud Connect repositories, see the Cloud Repository section in the Veeam Cloud Connect Guide.
- The cloud provider can verify integrity of tenant backups only by using the **/file** parameter and only if the backup is not encrypted. They cannot verify encrypted backups and they cannot verify backups by using other parameters.
- [For backups stored in scale-out backup repositories] Veeam Backup Validator can validate only backups stored in the performance tier which consists of backup repositories (except object storage repositories). Make sure that incremental and full backup files are located on the same extent.

## Syntax

Veeam Backup Validator provides parameter sets that allow you to:

- Display Veeam Backup Validator help information.

```
Veeam.Backup.Validator /?
```

- Validate integrity of the content of all VMs or selected VMs in the specified backup.

```
Veeam.Backup.Validator.exe /backup:backupname|backupid [/vmname:vmname] [/
point:pointid] [/date:pointdate] [/time:pointtime] [/silence] [/skip] [/re
port:reportpath [/format:xml|html]]
```

- Validate integrity of the VM content in the specified backup file.

```
Veeam.Backup.Validator.exe /file:backupfile{1..*} [/username:username /pas
sword:password] [/vmname:vmname]
[/silence] [/skip] [/report:reportpath [/format:xml|html]]
```

# Parameters

| Parameter | Description | Required/Optional | Parameter Type | Notes |
|---|---|---|---|---|
| **/backup: backupname\| backupID** | Specifies a name or an ID* of a backup or backup copy job that you want to validate. | Required | String | Consider the following:<br>• Per-machine backup with separate metadata files is a default format of backup files. To validate these backup files, you must specify the child backup name — backup of a specific VM. To get the child backup name, use the following PowerShell commands:<br>`$backup = Get-VBRBackup -Name "<backup_name>"`<br>`$child_backups = $backup.FindChildBackups()`<br>`$first_child_backup_name = $child_backups[0].Name`<br>• For a backup copy job in the immediate mode, you must specify the name of its child job — task that copies backup job added as a source to the backup copy job. For example, if the `My Copy` backup copy job copies `Daily Backup` backup job, then you must set the parameter value to `My Copy\Daily Backup`.<br>If you want to validate the whole backup copy job, you must run the utility for each child job. |

| Parameter | Description | Required/Optional | Parameter Type | Notes |
|---|---|---|---|---|
| **/file:backupfile{1..*}** | Specifies a backup file (VBM, VBK, VIB, VRB) to be validated. | Required | String | Consider the following:<br>• To validate an incremental backup file, you must specify the whole backup chain starting from the VBK file up to the required incremental file. Each backup file must be specified using a separate `/file` parameter.<br>• If the file is located on a network share, make sure you specify a full path, for example: `\\172.16.16.198\TestShare\Empty VM encryptedD2017-09-22T172639.vbk`.<br>• Mapped network drives are not supported. For example, you cannot specify `z:\\172.16.16.198\TestShare`. |
| **/username: username /password: password** | To access files on network shares. Specifies account credentials that the utility uses to access network shares. | Required for network share | String | If you want to validate files located on different shares, make sure this account has access rights to all these shares. |
| **/vmname:vmname** | Specifies a name of the VM in the backup file to be validated. | Optional | String | — |

| Parameter | Description | Required/Optional | Parameter Type | Notes |
|---|---|---|---|---|
| **/point:pointID** | Specifies an ID* of the restore point to be validated.<br>**Note:** You must provide this parameter after the `/backup` parameter.<br>To get the restore point ID, run the Get-VBRRestorePoint cmdlet and retrieve the `PointID` property. | Optional | String | If not specified, Veeam Backup Validator will verify the latest restore point, that is, all backup files the restore point consists of. |
| **/date:pointdate** | Specifies the date when the validated restore point was created. | Optional | Date | Make sure to specify the date in the same format as used on the Veeam Backup server. For example:<br>• For the `mm/dd/yyyy` format, specify 08.30.2012.<br>• For the `dd/mm/yyyy` format, specify 30.08.2012. |
| **/time:pointtime** | Specifies the approximate time when the validated restore point was created. | Optional | Time | — |
| **/silence** | Defines whether to run validation in the silence mode. | Optional | Boolean | — |

| Parameter | Description | Required/Optional | Parameter Type | Notes |
|---|---|---|---|---|
| **/skip** | Defines whether to skip from processing VMs listed in the `vmname` parameter. | Optional | Boolean | In the `vmname` parameter, list all VMs that you want to skip. |
| **/report: reportpath [/format:xml\|html]** | Specifies a full path of a file where you want to store a report on validation results. The utility will generate a report on validation results and store it at the specified path. | Optional | String | Consider the following:<br>• You must specify the full file path, for example, `"C:\temp\validator-report.html"`.<br>• Supported report formats are HTML and XML. |

* You can get IDs of backup jobs and restore points from the Veeam Backup & Replication database using scripts or Management Studio.

# Examples

## Example 1. Validating Specific VM in Incremental Backup File

This example shows how to validate the `winsrv29` VM in the incremental backup file and write the result in the `report.html` file.

```
Veeam.Backup.Validator.exe /file:"C:\Backup\Backup Job Single Storage\Bac
kup Job Single StorageD2022-10-03T132735_1E50.vbk" /file:"C:\Backup\Backu
p Job Single Storage\Backup Job Single StorageD2022-10-28T122338_3EE4.vib
" /vmname:winsrv29 /report:"C:\report.html"
```

### Example 2. Validating Specific VM Backup

This example shows how to validate the `srv506` VM in the `Exchange Backup Job` backup file.

```
$backup = Get-VBRBackup -Name "Exchange Backup Job"
$child_backups = $backup.FindChildBackups()
$first_child_backup_name = $child_backups[0].Name
.\Veeam.Backup.Validator /backup:$first_child_backup_name /vmname:
srv506
```

### Example 3. Validating VM Backup Created After Specific Time and Date

This example shows how to validate the `srv506` VM in the `Exchange Backup Job` backup file created after June 2, 2024, 9:00 PM.

```
$backup = Get-VBRBackup -Name "Exchange Backup Job"
$child_backups = $backup.FindChildBackups()
$first_child_backup_name = $child_backups[0].Name
.\Veeam.Backup.Validator /backup:$first_child_backup_name /date:02.06.202
4 /time:21:00 /vmname:srv506
```

# Veeam Backup Configuration Tool

Veeam Backup & Replication comes with the `Veeam.Backup.Configuration.Tool.exe` utility that allows you to manage BCO files. BCO files are backup files that contain backups of configuration databases. Veeam Backup & Replication creates these files when it performs configuration backup. For more information on configuration backup, see Managing Configuration Database.

You can use the Veeam Backup Configuration tool in the following scenarios:

- You do not have information on the BCO file version, parameters and attributes. In this case, you can use the tool to get details on the BCO file version.
- Veeam Backup & Replication is not able to restore a configuration database using a specific BCO file. In this case, you can use the tool to check if the BCO file is corrupted. The Veeam Backup Configuration tool will perform the cyclic redundancy check (CRC) and will verify that the encrypted file is decrypted properly.
- You do not have information on the specific configuration database. In this case, you can get details on its version and also use the tool to back up this database.
- Your backup server is no longer available, but your configuration database is still up and running. You can use the tool to create a backup of the configuration database and migrate this database to another backup server. In this case, the tool produces the same configuration backup, as Veeam Backup & Replication creates when you run the configuration backup job.

> **NOTE**
>
> If you back up the configuration database using the Veeam Backup Configuration tool, you will not be able to choose the backup repository in which the configuration backup must be stored and the necessary retention settings. Veeam Backup & Replication will keep last 10 restore points of the configuration backup in the default backup repository. If you want to change these setting, see the Scheduling Configuration Backups section.

# Using Veeam Backup Configuration Tool

The Veeam Backup Configuration tool is located on the backup server in the installation folder of Veeam Backup & Replication. The default path is `%ProgramFiles%\Veeam\Backup and Replication\Backup\Veeam.Backup.Configuration.Tool.exe`. If the default path was changed, you can find the actual path in the following registry value: `[HKEY_LOCAL_MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication] CorePath`.

To run the tool, open the command prompt on the backup server and change the current folder to the folder where Veeam Backup Configuration tool is located.

> **IMPORTANT**
>
> To run the Veeam Backup Configuration tool, you must use an account with administrative rights on the local machine.

## Syntax

The Veeam Backup Configuration tool provides parameter sets that allow you to:

- Display help information for the Veeam Backup Configuration tool.

```
Veeam.Backup.Configuration.Tool /?
```

- Analyze a configuration backup file.

```
Veeam.Backup.Configuration.Tool /file:value /analyzefile
```

- Check whether a configuration backup is not corrupted.

```
Veeam.Backup.Configuration.Tool /file:value /checkfile
```

- Analyze a configuration database.

```
Veeam.Backup.Configuration.Tool /analyzedatabase [/servername:value] [/ins
tancename:value] [/serverport:value] [/initialcatalog:value] [/login:value
][/password:value]
```

- Back up a configuration database.

```
Veeam.Backup.Configuration.Tool /file:value /backupdatabase [/servername:v
alue][/instancename:value] [/serverport:value] [/initialcatalog:value] [/l
ogin:value] [/password:value] [/cryptfile]
```

# Parameters

| Parameter | Description |
|---|---|
| **/file:value** | Specifies a path to the configuration backup file. |
| **/analyzefile** | Analyzes a configuration backup file. |
| **/checkfile** | Checks whether the configuration backup is not corrupted. |
| **/analyzedatabase** | Analyzes a configuration database. |
| **/backupdatabase** | Backs up a configuration database. |
| **/databaseengine** | Specifies a database engine |
| **/servername:value** | Specifies a name of a SQL server. |
| **/instancename:value** | Specifies a name of a SQL instance. |
| **/serverport:value** | Specifies a port number of a SQL server. The tool will use this port to access the SQL server. |
| **/initialcatalog:value** | Specifies a name of a SQL database. |
| **/login:value** | Specifies a username that the tool will use to authenticate against a SQL server. |
| **/password:value** | Specifies a password that the tool will use to authenticate against a SQL server. |
| **/cryptfile** | Defines that the tool will encrypt a configuration backup file. |
| **/verbose** | Enables verbose output mode. |

# Examples

## Example 1

This example shows how to analyze the `193022052014.bco` configuration backup file. After you start the command, the command prompt will return the output.

```
Veeam.Backup.Configuration.Tool.exe /file:"c:\my files\193022052014.bco"
/analyzefile /verbose
```

## Example 2

This example shows how to analyze the `193022052014.bco` configuration backup file and back up the configuration database. The command will contain the following settings of the configuration database:

- The configuration database is located at the `WIN2008R2` SQL server.
- The name of the SQL instance is `VeeamSql2008`.
- The name of the SQL database is `VeeamBackup`.

```
Veeam.Backup.Configuration.Tool.exe /file:c:\backups\091323052014.bco
/backupdatabase /servername:WIN2008R2 /instancename:VeeamSql2008
/initialcatalog:VeeamBackup
```

# Veeam Backup & Replication Events

For more information about events written by Veeam Backup & Replication, see Event Reference.