



Veeam Agent for Linux

Version 6

User Guide

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	8
ABOUT THIS DOCUMENT	9
OVERVIEW	10
Solution Architecture	11
Standalone and Managed Operation Modes	12
Data Backup	14
Backup Types	15
How Backup Works	18
Backup Job	21
Backup Repository	25
Backup Chain	27
Data Compression	33
Data Encryption	35
Backup of Database Systems	39
Backup to Object Storage	46
Backup to Deduplicating Storage Appliances	53
Backup to Veeam Cloud Connect Repository	55
Data Restore	56
Volume-Level Restore	57
File-Level Restore	58
Veeam Recovery Media	59
Veeam Recovery Media Versions	60
Drivers in Veeam Recovery Media	61
Integration with Veeam Backup & Replication	62
Managing Veeam Agent in Veeam Backup & Replication	64
PLANNING AND PREPARATION	65
System Requirements	66
System Requirements for Nosnap Veeam Agent for Linux	74
System Requirements for Nosnap Veeam Agent for Linux on Power	80
RAM Requirements for Backup of Large Number of Files	85
Permissions	88
Ports	90
Live Patching Support	94
INSTALLATION AND CONFIGURATION	95
Before You Begin	96
Installing Veeam Agent for Linux	97
Connecting to Veeam Software Repository	98

Installing Veeam Agent for Linux with Kernel Module	99
Installing Nosnap Veeam Agent for Linux	102
Installing Nosnap Veeam Agent for Linux on Power	103
Installing Veeam Agent for Linux in Offline Mode	104
Installing Veeam Agent for Linux with Kernel Module in Offline Mode	105
Installing Nosnap Veeam Agent for Linux in Offline Mode	112
Installing Nosnap Veeam Agent for Linux on Power in Offline Mode	116
Configuring UEFI Secure Boot.....	118
Upgrading Veeam Agent for Linux	125
Upgrading Veeam Agent for Linux with Kernel Module	126
Upgrading Nosnap Veeam Agent for Linux	130
Upgrading Nosnap Veeam Agent for Linux on Power	131
Granting Permissions to Users	132
Performing Initial Setup	133
Step 1. Accept License Agreements	134
Step 2. Create Custom Veeam Recovery Media	135
Step 3. Install Product License	138
Configuring Advanced Settings.....	140
Managing Veeam Agent Operation Mode	142
Viewing Operation Mode	143
Resetting to Standalone Operation Mode	144
Connecting to Veeam Backup & Replication.....	145
Synchronizing with Veeam Backup Server	146
Exporting Logs to Veeam Backup Server	147
Uninstalling Veeam Agent for Linux	148
GETTING STARTED	150
GETTING TO KNOW USER INTERFACE	151
Veeam Agent for Linux Control Panel	152
Command Line Interface	154
Viewing Help	156
LICENSING	157
Product Editions	158
License Agreement	159
Installing License	160
Viewing License Information	162
Removing License	163
License Expiration.....	164
Managing License with Command Line Interface	165
Accepting License Agreements	166
Installing License	167

Viewing License Information	168
Removing License	169
PERFORMING BACKUP	170
Creating Custom Veeam Recovery Media	171
Creating Custom Veeam Recovery Media with Control Panel	173
Creating Custom Veeam Recovery Media with Command Line Interface	176
Creating Backup Jobs	178
Before You Begin	179
Creating Backup Job with Backup Job Wizard	181
Creating Backup Job with Command Line Interface	235
Starting and Stopping Backup Jobs	280
Starting Backup Job from Control Panel	281
Starting Backup Job from Command Line Interface	284
Creating Active Full Backups in Command Line Interface	286
Stopping Backup Job	287
Managing Backup Jobs	289
Viewing List of Backup Jobs	290
Viewing Backup Job Settings	291
Editing Backup Job Settings	293
Deleting Backup Job	299
Managing Backup Repositories	301
Creating Backup Repository	302
Viewing List of Backup Repositories	310
Editing Backup Repository Settings	311
Rescanning Veeam Backup Repository	313
Deleting Backup Repository	314
Managing Veeam Backup & Replication Servers	315
Connecting to Veeam Backup Server	316
Viewing List of Veeam Backup Servers	318
Viewing Backup Server Details	319
Editing Connection to Veeam Backup Server	320
Updating List of Veeam Backup Repositories	322
Deleting Connection to Veeam Backup Server	323
Managing Service Providers	324
Connecting to Service Provider	325
Viewing List of Service Providers	326
Editing Connection to Service Provider	327
Updating List of Cloud Repositories	330
Deleting Connection to Service Provider	331
Managing Backups	332

Viewing Backups	333
Viewing Backup Details	335
Viewing Restore Points in Backup	337
Importing Backups	339
Deleting Backups	342
PERFORMING RESTORE	343
Restoring from Veeam Recovery Media	344
Restoring Volumes	345
Restoring Files and Folders	394
Restoring Volumes with Command Line Interface	434
Before You Begin	435
Restoring from Backup	436
Restoring from Restore Point	442
Restoring Files and Folders with Recovery Wizard	449
Before You Begin	450
Step 1. Launch File Level Restore Wizard	451
Step 2. Select Backup and Restore Point	452
Step 3. Save Restored Files	456
Step 4. Stop Backup Mount Session	457
Restoring Files and Folders with Command Line Interface	458
Before You Begin	459
Restoring from Backup	460
Restoring from Restore Point	469
Exporting Backup to Virtual Disk	478
Exporting Backups	479
Exporting Restore Points	481
Restoring Data from Encrypted Backups	483
REPORTING	486
Viewing Job Session Progress	487
Viewing Real-Time Job Session Statistics	488
Viewing Job Session Result	490
Viewing Session Status	492
Viewing Session Logs	494
EXPORTING PRODUCT LOGS	495
Exporting Logs with Control Panel	496
Exporting Logs with Command Line Interface	498
GETTING SUPPORT	499
USING WITH VEEAM BACKUP & REPLICATION	500
Setting Up User Permissions on Backup Repositories	502
Managing License	505

Managing Instance Consumption by Veeam Agents	506
Assigning License to Veeam Agent	507
Viewing Licensed Veeam Agents and Revoking License	508
Performing Data Protection Tasks.....	510
Backing Up to Backup Repositories	511
Backing Up to Cloud Repositories	512
Performing Backup Copy for Veeam Agent Backups	514
Using SureBackup	515
Archiving Veeam Agent Backups to Tape	517
Restoring Data from Veeam Agent Backups.....	518
Restoring Veeam Agent Backup to vSphere VM	519
Restoring Veeam Agent Backup to Hyper-V VM	521
Restoring Veeam Agent Backup to Nutanix VM	523
Restoring Veeam Agent Backup to Proxmox VM	525
Restoring Veeam Agent Backup to oVirt KVM VM	527
Restoring Disk from Veeam Agent Backup to oVirt KVM	529
Restoring to Microsoft Azure	531
Restoring to Amazon EC2	532
Restoring to Google Compute Engine	534
Restoring Files and Folders.....	536
Restoring Application Items	537
Exporting Disks.....	538
Publishing Disks	547
Exporting Restore Point to Full Backup File	559
Performing Administration Tasks	560
Importing Veeam Agent Backups	561
Enabling and Disabling Veeam Agent Backup Jobs	563
Viewing Veeam Agent Backup Job Statistics	564
Deleting Veeam Agent Backup Jobs	565
Viewing Veeam Agent Backup Properties	566
Creating Recovery Token	567
Removing Veeam Agent Backups	569
Deleting Veeam Agent Backups from Disk	570
Configuring Global Settings	571
Assigning Roles to Users	572

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 Veeam Agent for Linux version 6.3.2.

Intended Audience

The user guide is intended for anyone who wants to use Veeam Agent for Linux to protect their computer.

Overview

Veeam Agent for Linux is a data protection and disaster recovery solution for physical endpoints and virtual machines running Linux-based operating systems.

Veeam Agent can be used by IT administrators who run Linux infrastructure to protect different types of computers and devices: servers, desktops and laptops. The solution runs inside the guest OS and does not need access to virtualization infrastructure components. Thus, Veeam Agent can be used to protect Linux server instances deployed in the public cloud, for example, in Microsoft Azure environment.

NOTE

Veeam Agent can operate in either standalone or managed mode. Depending on the mode, Veeam Agent provides different features and limitations. To learn more, see [Standalone and Managed Operation Modes](#).

Veeam Agent offers a variety of features to protect your data. You can create an entire system image backup, back up specific machine volumes or individual directories and files. Backups can be stored on a local hard drive, on an external hard drive, in a network shared folder, object storage repository or Veeam backup repository.

In case of a disaster, you can perform the following restore operations:

- Start the OS from the Veeam Recovery Media and use standard Linux command line tools to diagnose and fix problems.
- Perform bare metal restore.
- Restore necessary data from backups to its original location or a new location.


Veeam Agent integrates with Veeam Backup & Replication. Backup administrators who work with Veeam Backup & Replication can perform advanced tasks with Veeam Agent backups: restore files and disks from backups, manage backup jobs configured in Veeam Agent or backups created with these jobs.

Solution Architecture

Veeam Agent for Linux is set up on a Linux-based physical endpoint or virtual machine whose data you want to protect.

When you install the product, Veeam Agent deploys the following components:

- *Veeam Agent for Linux Service* is a service responsible for managing all tasks and resources in Veeam Agent. The *veeamservice* component is registered as a daemon in the Linux OS upon the product installation. The service is started automatically when you start the OS and runs in the background.
- *Veeam Agent for Linux Job Manager* is a process started by *Veeam Agent for Linux Service* for every backup job session.
- *Veeam Agent* that communicates with the *Veeam Agent for Linux Service* and *Veeam Agent for Linux Job Manager*. *Veeam Agent* is started by *Veeam Agent for Linux Manager* to perform data transfer operations of any kind: copy data from the backed-up volume to the backup location during backup, from the backup location to the target volume during restore, perform data compression, and so on.
- *Veeam Agent for Linux Driver* is a Veeam driver (Linux kernel module) used to create volume snapshots in the Linux OS and keep track of changed data blocks.
- To store its configuration data, Veeam Agent uses the SQLite database engine. SQLite requires only few files to install and takes little resources to run on a Linux OS.



Standalone and Managed Operation Modes

Veeam Agent can operate in two modes: *standalone mode* and *managed mode*. The current User Guide covers subjects related to Veeam Agent operating in the standalone mode only. Depending on the operation mode, Veeam Agent has different functionality and limitations.

Standalone Mode

In this mode, Veeam Agent operates as a standalone product. To use Veeam Agent operating in the standalone mode, you must manually install the product directly on the computer whose data you want to protect.

For Veeam Agent operating in the standalone mode, data protection, disaster recovery and administration tasks are performed by the user. You can also use Veeam Agent operating in the standalone mode with Veeam Backup & Replication. In this scenario, you can use backup repositories managed by Veeam Backup & Replication as a target location for Veeam Agent backups and use the Veeam Backup & Replication console to perform a number of tasks with Veeam Agent backup jobs and backups. To learn more, see [Integration with Veeam Backup & Replication](#).

You can also use Veeam Backup & Replication as a gateway for creating backups targeted at the following types of repositories:

- Veeam Cloud Connect repository. To learn more, see [Backup to Veeam Cloud Connect](#).
- Object storage repository.

With Veeam Agent operating in the standalone mode, you can also back up data directly to an object storage repository. To learn more about both options, see [Backup to Object Storage](#).

Managed Mode

In this mode, Veeam Agent operates under control from one of the following Veeam products:

- **Veeam Backup & Replication**

You can automate management of Veeam Agents on multiple computers in your infrastructure in the Veeam Backup & Replication console. You can configure Veeam Agent backup policies and perform other data protection and administration tasks on remote computers.

To use Veeam Agent operating in the managed mode, you must deploy the product in one of the following ways:

- From Veeam Backup & Replication
- Manually using external tools

To learn more about managed Veeam Agent deployment, see the [Protected Computers Discovery and Veeam Agent Deployment](#) section in the Veeam Agent Management User Guide.

For Veeam Agent managed by Veeam Backup & Replication, data protection, data restore and administration tasks are performed by a backup administrator in the Veeam Backup & Replication console. To learn about managing Veeam Agent in Veeam Backup & Replication, see the [Veeam Agent Management Guide](#).

- **Veeam Service Provider Console**

You can use Veeam Service Provider Console to manage Veeam Agents on multiple computers in your infrastructure. When Veeam Agent is managed by Veeam Service Provider Console, you can configure backup job settings, start and stop backup, change global settings, update and uninstall Veeam Agent and collect Veeam Agent data for monitoring and billing.

To manage Veeam Agent from Veeam Service Provider Console, you must install Veeam Service Provider Console management agent and Veeam Agent on the computer whose data you want to protect. After that, in Veeam Service Provider Console, you must activate Veeam Agent on the protected computer to set it into the managed operation mode.

For Veeam Agent managed by Veeam Service Provider Console, data protection, data restore and administration tasks are performed by a backup administrator in Veeam Service Provider Console.

Backup administrator can enable a read-only access mode for Veeam Agent installed on the protected computer. When you work directly with Veeam Agent operating in the read-only access mode, you can perform a limited set of operations, including:

- Running the backup job manually.
- Viewing backup session statistics.
- Restoring individual files.

To learn about deploying and managing Veeam Agent with Veeam Service Provider Console, see [Veeam Service Provider Console User Guides](#). Select the guide that suits your user role.

Data Backup

It is recommended that you regularly back up data stored on your machine. Backup creates a safety copy of your data. If any kind of disaster strikes, you can restore your data from the backup and be sure that you will not lose the necessary information.

You can set up Veeam Agent to perform automatic scheduled backups (triggered at specific time of the day), or you can choose to back up data manually when needed. You can back up the entire computer image, specific computer volumes or individual directories and files.

You can set up Veeam Agent to create multiple backups — with individual backup scope, upon individual schedule or in different locations. This functionality is available if Veeam Agent operates in the Server edition. To learn more about editions, see [Product Editions](#).

Backups created with Veeam Agent can be saved to the following locations:

- Removable storage device
- Local computer drive
- NFS or SMB (CIFS) network shared folder
- Veeam backup repository managed by a Veeam backup server
- Veeam Cloud Connect repository
- Object storage repository

Backup Types

Veeam Agent for Linux lets you create the following backup types:

- [Volume-level backup](#)
- [File-level backup](#)

Volume-Level Backup


You can set up Veeam Agent for Linux to create volume-level backup. The volume-level backup captures the whole image of a data volume on your computer. You can use the volume-level backup to restore a computer volume, specific files and folders on the volume or perform bare metal recovery.

Veeam Agent for Linux supports backup of the following types of computer volumes:

- Simple volumes
- LVM logical volumes
- BTRFS subvolumes


You can back up all computer volumes or specific computer volumes.

- When you back up the entire computer image, Veeam Agent captures the content of all volumes on your computer. The resulting backup file contains all volume data and Linux OS system data: system partition, partition table and bootloader.



- When you back up a specific computer volume, Veeam Agent captures only the data that resides on this specific volume: files, folder, application data and so on.

If you choose to back up the system volume (volume to which the root file system is mounted), Veeam Agent automatically includes the bootloader into the backup scope.




File-Level Backup


You can set up Veeam Agent for Linux to create file-level backup. The file-level backup captures only data of individual directories and files on the computer. You can use the file-level backup to restore files and directories that you have added to the backup scope.

With Veeam Agent for Linux, you can specify which files and directories to back up:


- You can include individual directories in the backup. When you include a directory in the backup, its subdirectories are automatically included in the backup too. When you recover from such backup, you will be able to restore directories that you have selected to back up, all subdirectories of these directories and files in these directories.



- You can exclude from the backup some subdirectories of the directories that are included in the backup. When you recover from such backup, you will be able to restore directories that you have selected to back up, specific subdirectories of these directories and files in these directories.



- You can include or exclude files of a specific type in/from the backup. You can specify file names explicitly or use UNIX wildcard characters to define include and exclude file name masks. When you recover from such backup, you will be able to restore directories that you have selected to back up with files whose names match the specified include masks.



Snapshot-Less File-Level Backup

You can set up Veeam Agent for Linux to create file-level backup in the snapshot-less mode. This allows you to back up data that resides in any file system mounted to the root file system of the Veeam Agent computer. For example, you can use the snapshot-less mode to back up data that resides in a file system that is not supported for snapshot-based backup with Veeam Agent, such as UFS, ZFS, GFS, GFS2, OCFS2, NILFS2 or bcache fs. You can also use it to back up data that resides in an NFS or SMB (CIFS) network shared folder.

To create backups in the snapshot-less mode, you must enable this mode in the properties of the file-level backup job. To learn more, see [Creating Backup Jobs](#).

In the snapshot-less mode, Veeam Agent does not create a snapshot of the backed-up volume. Instead, when the backup process starts, Veeam Agent reads files and directories that you selected to back up, and copies backed-up data to the target location.

IMPORTANT

During backup in the snapshot-less mode, Veeam Agent does not track whether files and directories have changed in their original location since the time when the backup process started. To make sure that data in the backup is in the consistent state, you must not perform write operations in the file system that contains the backed-up data until the backup process completes.

How Backup Works

Veeam Agent for Linux performs backup differently depending on the backup type:

- [Volume-level backup](#)
- [File-level backup](#)

How Volume-Level Backup Works

During volume-level backup, Veeam Agent performs the following operations for every [backup job session](#):

1. When a new job session starts, Veeam Agent creates a backup file in the target location.
2. In the backup file, Veeam Agent creates a disk for each backed-up disk. In disks, Veeam Agent creates blank partitions that have the same size and location as partitions in backed-up disks.
3. Veeam Agent creates a snapshot of the volume whose data you want to back up. The snapshot is created on the volume that has enough free disk space to contain the snapshot data. To create a snapshot, Veeam Agent uses the *Veeam Agent for Linux Driver*.

The snapshot helps make sure that the data on the volume is consistent and does not change at the moment of backup. If a data block is about to change on disk during backup, Veeam Agent will copy this block to the snapshot. After the data block is overwritten on the source location, its original copy will remain intact in the snapshot.

NOTE

Consider the following:

- If you instruct Veeam Agent to back up a database system, Veeam Agent prepares databases for backup before creating a snapshot of the volume. To learn more, see [Backup of Database Systems](#).
- During backup of data that resides in the BTRFS file system, Veeam Agent does not use its driver to create a snapshot. Instead, Veeam Agent leverages BTRFS capabilities to create a BTRFS snapshot.

4. [For incremental backup] Veeam Agent uses the *Veeam Agent for Linux Driver* to detect what blocks have changed on the volume since the previous job session. The driver keeps this information as a changed block tracking map in the RAM of your computer.


Mind that every time the driver is unloaded or the Veeam Agent computer is rebooted, the changed block tracking map is reset as well. In such case, to detect what data blocks have changed since the previous job session, Veeam Agent rescans the entire data added to the backup scope and creates a new changed block tracking map. In this case, backup requires greater time.

To learn about full and incremental backup, see [Backup Chain](#).

5. Veeam Agent copies the partition table and bootloader located on the hard disk to the backup file in the target location.
6. [For incremental backup] Veeam Agent calculates checksums for each data block and compares them with checksums from the backup file created during the previous job session. If checksums do not match, Veeam Agent will copy the data block to the target location during the next backup process step.
7. Veeam Agent copies data from the following sources:
 - Data that did not change on disk during backup is transferred from the source volume.

- Data that changed on disk during backup is transferred from the snapshot.

After all the data is transferred, Veeam Agent removes the snapshot.



How File-Level Backup Works

During file-level backup, Veeam Agent performs the following operations for every [backup job session](#):

1. When a new job session starts, Veeam Agent creates a backup file in the target location.
2. In the backup file, Veeam Agent creates a disk. The disk contains a volume with the ext4 file system.
3. Veeam Agent creates a snapshot of the volume which data you want to back up. The snapshot is created on the volume that has enough free disk space to contain the snapshot data. To create a snapshot, Veeam Agent uses the *Veeam Agent for Linux Driver*.

The snapshot helps make sure that the data on the volume is consistent and does not change at the moment of backup. If a data block is about to change on disk during backup, Veeam Agent will copy this block to the snapshot. After the data block is overwritten on the source location, its original copy will remain intact in the snapshot.

TIP

Consider the following:

- You can also set up Veeam Agent to create a file-level backup in the snapshot-less mode. This mode allows you to back up data that resides in any file system mounted to the root file system of the Veeam Agent computer. However, Veeam Agent does not track whether source files have changed since the backup process start. To learn more, see [Snapshot-Less File-Level Backup](#).
- Compared to the volume-level backup, the file-level backup, Veeam Agent does not provide changed block tracking mechanism and does not split source files into data blocks. As a result, if you plan to back up a significant amount of data, the file-level backup will require greater time, and created backup files will have greater size.


For example, you have a 1 GB file, and since the previous backup session only one data block of this file has changed. In case of the file-level backup, Veeam Agent will send the whole 1 GB file to the target again.

4. [For incremental backup] To detect files that changed on the Veeam Agent computer since the previous backup session, Veeam Agent reads file metadata and compares last modification time of files in the original location and files in the backup created during the previous job session. If the file has modification time later than the previous job session start time, Veeam Agent considers the file as changed.

To learn about full and incremental backup, see [Backup Chain](#).

5. Veeam Agent copies data that you selected for backup to the target location. As part of this process, Veeam Agent performs the following operations:
 - a. Enumerates all files in the source location.
 - b. For each enumerated file, creates a target file in the volume inside the backup file.
 - c. Opens the source and the target files.
 - d. Copies file data to the target location from the following sources:
 - Data blocks that did not change on disk during backup are transferred from the source volume.
 - Data blocks that changed on disk during backup are transferred from the snapshot.
 - e. Closes the source and target files.

After all backed-up files and directories are transferred, Veeam Agent removes the snapshot.



Backup Job

To back up your data, you must configure a backup job. The backup job settings define what data you want to back up, what the target location and retention policy for created backups are and how to back up your data. If necessary, you can re-configure the backup job and change its settings at any time.

NOTE

You cannot change the backup job type from volume-level to file-level, and vice versa.

In Veeam Agent for Linux, you can configure several backup jobs with different settings. For example, you can configure one backup job to create volume-level backup and another backup job to create file-level backup. You can configure backup jobs targeted at different backup locations to keep several copies of your backed-up data. You can also configure several backup jobs with individual schedule to fine-tune automatic backup creation process.

NOTE

You can create more than one backup job only if Veeam Agent operates in the Workstation or Server edition. To learn more, see [Product Editions](#).

Veeam Agent launches the backup job according to the schedule you define. You can schedule the job to start at specific time daily or on specific week days. You can also start a backup job manually to perform backup on demand when needed.

Backup job scheduling settings are configured globally for all accounts of the Linux OS. As a result, Veeam Agent can start a backup job automatically regardless of the currently running user session.

Backup Job Scripts

You can instruct Veeam Agent for Linux to run custom scripts within the backup job session:

- [Pre-job and post-job scripts](#) – Veeam Agent runs these scripts before the backup job starts and after the backup job completes. You can use pre-job and post-job scripts, for example, to configure email notifications about jobs performed by Veeam Agent.
- [Pre-freeze and post-thaw scripts](#) (in the Server edition only) – Veeam Agent runs these scripts before and after creating a snapshot. For example, the pre-freeze script may quiesce the file system and application data to bring the Linux OS to a consistent state before Veeam Agent creates a snapshot. After the snapshot is created, the post-thaw script may bring the file system and applications to their initial state.

Consider the following about using backup job scripts:

- Scripts must be created beforehand. You must specify paths to them in the job settings. Veeam Agent supports scripts in the SH file format.
- Scripts must have UNIX line endings (LF).
- Script settings are enabled at the job level. If Veeam Agent operates in the Server edition and you want to configure multiple backup jobs, you can specify individual scripts for each job.
- If you use relative paths in your scripts, during script execution such paths will refer to the root directory. For example, the script may have an output that must be saved to a new file. If you specify a relative path to that file or only a file name, the file will be created in the root directory. To specify a different location for a file, use a full absolute path.

Pre-Job and Post-Job Scripts

You can instruct Veeam Agent for Linux to run custom pre-job and post-job scripts. Veeam Agent executes the pre-job script directly before the backup job starts. After the backup job completes, Veeam Agent executes the post-job script.

Veeam Agent starts the backup job regardless of the pre-job script result. If the pre-job script fails to execute, Veeam Agent will always start the backup job. Then, after the backup job completes, Veeam Agent will execute the post-job script.

The script is considered to be executed successfully if a "0" is returned.

The default time period for script execution is 10 minutes. After this period expires, Veeam Agent stops executing the script and displays a warning message in the job session. You can change the script timeouts in the `/etc/veeam/veeam.ini` configuration file with the `timeoutPrePost` parameter where the time period is set in seconds.

Pre-Freeze and Post-Thaw Scripts

You can instruct Veeam Agent for Linux to run custom pre-freeze and post-thaw scripts. Veeam Agent executes the pre-freeze script before creating a snapshot. After the snapshot is created, Veeam Agent executes the post-thaw script.

NOTE

Veeam Agent does not execute pre-freeze and post-freeze scripts if during the backup job a snapshot is not created.

The script is considered to be executed successfully if a "0" is returned.

By default, if the pre-freeze or post-thaw script fails to execute, Veeam Agent does not start the backup job. However, you can instruct Veeam Agent to ignore errors that occur during the script execution process. To allow Veeam Agent to start backup jobs regardless of the script execution result, in the `/etc/veeam/veeam.ini` configuration file, uncomment the `ignoreFreezeThawFailures` parameter and set its value to `true`.

If Veeam Agent is set up to ignore script errors, and the pre-freeze or post-thaw script fails to execute, Veeam Agent will start the backup job. After the job successfully completes, Veeam Agent will display the *Warning* status for the job session.

The default time period for script execution is 10 minutes. After this period expires, Veeam Agent stops executing the script. You can change the script timeouts in the `/etc/veeam/veeam.ini` configuration file with the `timeoutFreezeThaw` parameter where the time period is set in seconds.

NOTE

You can specify pre-freeze and post-thaw scripts only if Veeam Agent for Linux operates in the Server edition. If these scripts were enabled for the job while Veeam Agent operated in the Server edition, and then Veeam Agent has switched to another edition (for example, to the Free edition after the license has expired), the backup job will fail. You will need to delete the existing job and create a new backup job without pre-freeze and post-thaw scripts enabled.

File System Indexing

You can instruct Veeam Agent for Linux to create an index of files and directories located on the Veeam Agent computer during backup. File indexing allows you to search for specific files inside Veeam Agent backups and perform 1-click restore in Veeam Backup Enterprise Manager.

File indexing is enabled at the job level. You can specify granular indexing settings for each job.

IMPORTANT

Indexing mechanism does not recognize file exclusion masks. If you specify masks to exclude certain files in a file-level backup job, Veeam Agent for Linux will nevertheless index all files located in the directories that have been selected for backup.

For example, you have included the `/home` directory into the backup and specified the `*.pdf` exclusion mask. The *Index everything* option is enabled for the backup job. In this case, when you browse the resulting backup in Veeam Backup Enterprise Manager, PDF files will be displayed in the `/home` directory as if they were backed up.

Considerations and Limitations

- Veeam Agent for Linux on Power does not support file system indexing.
- Veeam Agent for Linux must have either Workstation or Server license installed.
- The following utilities must be installed on the computer: `gzip` and `tar` (standard utilities for majority of Linux distributions). These utilities are provided along with the product in the product installation media.

NOTE

Consider the following:

- File system 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 with such backup job. For more information, see the [Preparing for File Browsing and Restore](#) section in the Veeam Backup Enterprise Manager User Guide.
- If SELinux is enabled in the Linux OS, file system indexing may fail.

Automatic Job Retries

Veeam Agent supports automatic job retries if a scheduled backup job fails for any reason — for example, if the backup repository is not available or connection to it is interrupted during the backup job execution.

If the backup job fails, Veeam Agent will automatically create a new session for this backup job with the *Pending* status. By default, Veeam Agent for Linux retries a failed job 3 times with an interval of 10 minutes.

Veeam Agent automatically restarts the backup job under the following conditions:

- If the backup job was launched automatically according to a schedule and failed for any reason. Veeam Agent will not perform a backup job retry if the backup job ended with the *Success* or *Warning* status.

- [For object storage targets] If during the job session, a backup health check detects corrupted data in the backup that resides in an object storage repository. By default, Veeam Agent will launch a job retry. Veeam Agent will retry the backup job up to 3 times if the backup job was run on a schedule. Veeam Agent will retry the backup job only once if the backup job was launched manually. For more information, see [Health Check for Object Storage](#).

Backup Repository

A backup job configured in Veeam Agent for Linux creates backup files in a backup repository. A backup repository is a directory on the storage where you want to keep backup files. You can use the following types of disk-based storage to create a backup repository:

- Local (internal) storage of the protected machine (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share.
- Veeam Backup & Replication 12.3 backup repository (including deduplication appliances).

NOTE

To benefit from all recent enhancements and improvements, make sure to install the latest available build of Veeam Backup & Replication.

- Veeam Cloud Connect 12.0 or later backup repository.
- Object storage repository, such as S3 Compatible storage, Amazon S3, Google Cloud or Microsoft Azure Blob.

IMPORTANT

Consider the following about backup repositories:

- [For local storage] A backup repository should be created on a separate volume from the volume that contains data you plan to back up.
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported.

Veeam Agent for Linux works with backup storage differently depending on the way you configure and start backup jobs – with the Veeam Agent control panel or command line interface.

Backup Location and Control Panel

If you use the Veeam Agent control panel to perform backup tasks, you do not have to deal with backup repositories. When you specify a target location for backup in the Backup Job wizard, Veeam Agent configures the backup repository automatically. Veeam Agent saves path to the specified backup location, assigns to this location a unique name and ID and saves this information in the database. The information is used by Veeam Agent and is not displayed in the control panel.

If you target a backup job at the network shared folder, every time the backup job starts, Veeam Agent will automatically mount the shared folder to the `/tmp/veeam` directory in the computer file system and create a backup file in this directory. After the backup job completes, Veeam Agent will automatically unmount the network shared folder.

You can target several backup jobs to individual backup locations or use the same target location for several backup jobs. This may be useful if you want to back up different types of data to separate locations or to keep all backed-up data at one place.

Backup Repository and Command Line Interface

If you work with Veeam Agent for Linux using the command line interface, you must deal with backup repositories depending on the target location selected for the backup job.

If you target a backup job at a local directory or network shared folder, you must create a repository before you configure a backup job:

- In case of a local directory, you specify a name for the repository and a local directory in which Veeam Agent will create backup files. To learn more, see [Creating Repository in Local Directory](#).
- In case of a network shared folder, you specify a name for the repository, a path to the network shared folder in which Veeam Agent will create backup files, a type of the network shared folder and additional mounting options.

Every time the backup job starts, Veeam Agent will automatically mount the shared folder to the `/tmp/veeam` directory in the computer file system and create a backup file in this directory. After the backup job completes, Veeam Agent will automatically unmount the network shared folder. To learn more, see [Creating Repository in NFS Share](#) and [Creating Repository in SMB Share](#).

If the directory to which the shared folder should be mounted resides on the backed-up volume, the backup job may fail.

- In case of object storage, you specify a name for the storage provider, a name for the repository and settings to access the storage account and bucket/container. To learn more, see [Creating Repository in Object Storage](#).


If you target a backup job at a Veeam backup repository or cloud repository, you do not need to create repositories. Before configuring the backup job, you must connect to the Veeam backup server or Veeam Cloud Connect service provider. To learn more, see [Connecting to Veeam Backup Server](#) and [Connecting to Service Provider](#).

You can configure several backup repositories and target different backup jobs at these repositories. This may be useful if you want to back up different types of data to separate locations or to keep several copies of your backed-up data.

Backup Chain


Every backup job session produces a new backup file in the target location. Backup files make up a backup chain. The backup chain can contain files of two types: full backups and incremental backups.

- During the first backup job session, Veeam Agent performs full backup. It copies all data that you have chosen to back up (entire volumes and folders) and stores the resulting full backup file (VBK) in the target location. The full backup takes significant time to complete and produces a large backup file: you have to copy the whole amount of data.
- During subsequent backup job sessions, Veeam Agent performs incremental backups. It copies only new or changed data relatively to the last backup job session and saves this data as an incremental backup file (VIB) in the target location. Incremental backups typically take less time than full backup: you have to copy only changes, not the whole amount of data.



After several backup cycles, you have a chain of backup files in the target location: the first full backup file and subsequent incremental backup files. Every backup file contains a restore point for backed-up data. A restore point is a "snapshot" of your data at a specific point in time. You can use restore points to roll back your data to the necessary state.

To recover data to a specific restore point, you need a chain of backup files: a full backup file plus a set of incremental backup files following this full backup file. If some file from the backup chain is missing, you will not be able to roll back to the necessary state. For this reason, we recommend that you do not delete separate backup files manually. To learn more, see [Deleting Backups](#).



Types of Backup Files

Veeam Agent produces backup files of the following types:

- VBK – full backup file.
- VIB – incremental backup file.
- VBM – backup metadata file. The backup metadata file is updated with every backup job session. It contains information about the computer on which the backup was created, every restore point in the backup chain, how restore points are linked to each other and so on. The backup metadata file is required for performing file-level and volume-level restore operations.

Short-Term Retention Policy


Restore points in the backup chain are not kept forever. They are removed according to the retention policy. The retention policy helps maintain the life cycle of restore points and make sure that backup files do not consume the whole disk space.

Veeam Agent for Linux retains the number of latest restore points defined by the user. During every backup job session, Veeam Agent for Linux checks if there is any obsolete restore point in the backup chain. If some restore point is obsolete, it is removed from the chain.


Removing Backups by Retention

When the obsolete restore points are removed by retention, Veeam Agent transforms the backup chain so it always contains a full backup file on which subsequent incremental backup files are dependent. To do so, Veeam Agent uses the following rotation scheme:


1. During every backup job session Veeam Agent adds a backup file to the backup chain and checks if there is an obsolete restore point.




2. If an obsolete restore point exists, Veeam Agent transforms the backup chain. As part of this process, it performs the following operations:
 - a. Veeam Agent rebuilds the full backup file to include in it data of the incremental backup file that follows the full backup file. To do this, Veeam Agent injects into the full backup file data blocks from the earliest incremental backup file in the chain. This way, a full backup 'moves' forward in the backup chain.



- b. Veeam Agent removes the earliest incremental backup file from the chain as redundant: its data has already been injected into the full backup file, and the full backup file includes data of this incremental backup file.



If the backup chain contains several obsolete restore points, the rebuild procedure is similar. Data from several restore points is injected to the rebuilt full backup file. This way, Veeam Agent makes sure that the backup chain is not broken, and you will be able to recover your data to any restore point.



Long-Term Retention Policy

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 Agent does not create any special new backup files – it uses backup files created while backup job runs and marks these backups with specific GFS flags.

To mark a backup file for long-term retention, Veeam Agent 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 Agent assigns depend on the configured [GFS retention policy settings](#).

NOTE

Consider the following:

- GFS flags can be assigned only to full backup files created during the time period specified in GFS policy settings.
- If you store your backups in an object storage repository managed by Veeam Backup & Replication and connection to this repository is set up through a gateway server, configuring active full backups is not required, Veeam Agent will create a full backup based on the last incremental backup and will assign a GFS flag to this full backup. If some data blocks required to create the full backup already reside in the object storage repository, the full backup will contain links to such data blocks. To avoid extra costs, Veeam Agent does not retrieve actual data blocks from the object storage repository.

If Veeam Agent assigns a GFS flag to a full backup file, this backup file can no longer be deleted or modified. Veeam Agent does not apply short-term retention policy settings to the full backup file. For example, Veeam Agent ignores the backup file when determining whether the number of allowed backup files is exceeded.

When the specified retention period ends, Veeam Agent 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.

Veeam Agent assigns GFS flags in the similar way as Veeam Backup & Replication does for VM backup files. To learn about logic behind GFS flags, see the [Assignment of GFS Flags](#) and [Removal of GFS Flags](#) sections in the Veeam Backup & Replication User Guide.

Limitations


When planning to use GFS retention policy, consider the following limitations:

- [Applicable to all backup targets except object storage] While applying the GFS retention policy, Veeam Agent does not create new full backup files. 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, the following applies:
 - Veeam Agent cannot transform the backup chain according to the short-term retention policy.
 - Veeam Agent is not able to merge data from incremental backup files into the full backup file.
- Veeam Agent 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.

Active Full Backup

When Veeam Agent performs active full backup, it produces a full backup file and adds this file to the backup chain.

The active full backup resets the backup chain. All incremental backup files use the latest active full backup file as a new starting point. A previously used full backup file and its subsequent incremental backup files remain on the disk. After the last incremental backup file created prior to the active full backup becomes outdated, Veeam Agent automatically deletes the previous backup chain. To learn more, see [Retention Policy for Active Full Backups](#).



You can create active full backups manually or schedule a backup job to create active full backups periodically. To do this, you can use the Veeam Agent for Linux control panel or command line interface.

- To learn how to configure active full backup schedule and create active full backups with the Veeam Agent for Linux control panel, see [Active Full Backup Settings](#) and [Starting Backup Job from Control Panel](#).
- To learn how to configure active full backup schedule and create active full backups with the Veeam Agent for Linux command line interface, see [Configuring Active Full Backup Schedule](#) and [Creating Active Full Backups](#).

Active Full Backup Schedule

You can schedule a backup job to create active full backups periodically. Active full backup schedule depends on the regular backup schedule.

- In case active full backup is scheduled on a week day, Veeam Agent modifies the regular schedule of the backup job.

For example, the regular backup schedule is set to Monday and Tuesday at 15:00. Active full backup schedule is set to Friday. In this case, the backup job schedule will contain information that the job must start on Monday, Tuesday and Friday at 15:00.


- In case active full backup is scheduled on a day of the month, Veeam Agent runs the backup job on this day at the same time as it must run upon the regular schedule.

Keep in mind that if the job is not scheduled to run automatically, Veeam Agent will not run active full backup. For more information on how to configure backup job schedule, see [Configuring Backup Schedule](#) and [Configuring Active Full Backup Schedule](#).


To be able to restore data from a Veeam Agent backup, you need to have a full backup file and a chain of subsequent incremental backup files on the disk. If you delete a full backup file, the whole chain of incremental backup files will become useless. In a similar manner, if you delete any incremental backup file before the point to which you want to roll back, you won't be able to restore data (since later incremental backup files depend on earlier incremental backup files).

For this reason, if you create an active full backup, in some days there will be more restore points on the disk than specified by retention policy settings. Veeam Agent 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. A full backup file is created on Sunday, incremental backup files are created Monday through Saturday, and an active full backup is scheduled on Thursday. Although the retention policy is already breached on Wednesday, the full backup is not deleted. Without the full backup, backup chain would be useless, leaving you without any restore point at all. Veeam Agent will wait for the next 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.



Keep in mind that if the backup job is set up to create periodic active full backups, Veeam Agent will never transform the backup chain. Instead, Veeam Agent will always wait for the next full backup file and the necessary number of incremental backup files to be created, and only then will delete the whole previous chain. In the example above, Veeam Agent will delete the previous chain every Saturday. As a result, although the retention policy is set to 3 restore points, the actual number of backup files on the disk will be greater most of the time.



In contrary, in a situation when you manually create a single active full backup, Veeam Agent will treat the active full backup in the same way as a regular full backup. If some restore point becomes obsolete, Veeam Agent will re-build the full backup file to include in it data of the incremental backup file that follows the full backup file. After that, Veeam Agent will remove the earliest incremental backup file from the chain as redundant.



Data Compression

Veeam Agent provides mechanisms of data compression. Data compression lets you decrease traffic going over the network and disk space required for storing backup files.

Data Compression

Data compression decreases the size of created backups but affects duration of the backup procedure. When you create a backup job in command line interface, Veeam Agent allows you to specify one of the following compression levels:

Compression Level	CLI Option	Compression Algorithm	Description
None	0	No compression	This compression level is recommended if you plan to store backup files on storage devices that support hardware compression and deduplication.
Dedupe-friendly	1	Rle	Optimized compression level for very low CPU usage. You can select this compression level if you want to decrease the load on the CPU of the Veeam Agent computer.
Optimal	2	Lz4	The default recommended compression level. It provides the best ratio between size of the backup file and time of the backup procedure.
High	3	Zstd 3	Provides up to 60% additional compression ratio over the Optimal level at the cost of 2x higher CPU usage and 2x slower restore.
Extreme	4	Zstd 9	Provides the smallest size of the backup file but reduces the backup performance. We recommend that you use the extreme compression level only on Veeam Agent computers with modern multi-core CPUs (6 cores recommended).

You can change data compression settings for existing backup jobs. New settings will not have any effect on previously created backup files in the backup chain. They will be applied to new backup files created after the settings were changed.

Compression settings are changed on the fly. You do not need to create a new full backup to use new settings — Veeam Agent will automatically apply the new compression level to newly created backup files.

Storage Optimization

Depending on the type of storage you select as a backup target, Veeam Agent uses data blocks of different size, which optimizes the size of a backup file and job performance. You can choose one of the following storage optimization options:

- **4 MB** – select this option for backup jobs that can produce very large full backup files (larger than 16 TB). With this option selected, Veeam Agent will use data block size of 4096 KB.
- **1 MB (default)** – select this option for backup to SAN, DAS or local storage. With this option selected, Veeam Agent will use data block size of 1024 KB.

The SAN identifies larger blocks of data and therefore can process large amounts of data at a time. This option provides the fastest backup job performance.

- **512 KB** – select this option for backup to NAS and onsite backup. With this option selected, Veeam Agent will use data block size of 512 KB. This option reduces the size of an incremental backup file because of reduced data block sizes.
- **256 KB** – select this option if you plan to use WAN for offsite backup. With this option selected, Veeam Agent will use data block size of 256 KB. This results in the smallest size of backup files, allowing you to reduce the amount of traffic over WAN.

Various factors can affect the overall size of the backup files. As a result, the smallest file size may not always be achieved.

NOTE

If you change storage optimization settings, the new settings will be applied only after an active full backup is created. Veeam Agent will use the new block size for the active full backup and subsequent backup files in the backup chain. For more information on scheduling active full backups, see [Active Full Backup Settings](#).

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 data to remote locations. To keep your data safe, you can use data encryption.

Data encryption transforms data to an unreadable, scrambled format with the help of a cryptographic algorithm and a secret key. 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 back to a readable format.

Encryption is performed on the trusted side depending on the backup target:

- Encryption is performed on the source side for all backup targets except the Veeam backup repository.
- Encryption is performed on the target side if you store backups in the Veeam backup repository.

Decryption is performed on the same side as encryption.

To create encrypted backups, you must enable the encryption option and specify a password that will be used for data encryption. To learn more, see [Data Encryption Settings](#).

NOTE

You cannot specify encryption options for the backup job if you have chosen to save Veeam Agent backups in a Veeam backup repository. For such jobs, encryption options are managed per repository by a backup administrator working with Veeam Backup & Replication. To learn more, see the [Encrypting Standalone Application Backups in Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

How Data Encryption Works


Data encryption is performed as part of the backup process. Encryption works at the source or target side depending on the backup target. As a result, encryption keys are not passed to the untrusted side, which helps to avoid data interception.

The source side encryption is performed at the [job](#) level. The target side encryption is performed at the [storage](#) level.

Job encryption

Encryption for the backup job is configured in the [data encryption settings](#). You should enable the encryption option and specify a password to protect data in backup files produced by the backup job.

When you enable encryption for a job, you must define a password to protect data processed by this job, and define a hint for the password. The password and the hint are saved in the job settings. Based on this password, Veeam Agent generates a user key.



The backup job processing with encryption enabled includes the following steps:

1. You enable encryption for the backup job and specify a password.
2. Veeam Agent generates the necessary keys to protect backup data.
3. Veeam Agent encrypts data blocks and transfers them to the target location already encrypted.
4. On the target storage, encrypted data blocks are stored in a resulting backup file.



To guarantee the flawless process of data encryption and decryption, consider the following advice:

Password

1. Use strong passwords that are hard to crack or guess. Consider the following recommendations:
 - a. The password must be at least 8 characters long.
 - b. The password must contain uppercase and lowercase characters.
 - c. The password must be a mixture of alphabetic, numeric and punctuation characters.
 - d. The password must significantly differ from the password you used previously.
 - e. 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.

2. Provide a meaningful hint for the password that will help you recall the password. The hint for the password must significantly differ from the password itself. The hint for the password is displayed when you select an encrypted backup server and attempt to unlock it.
3. Change passwords for encrypted jobs regularly. Use of different passwords helps increase the encryption security level.

Encryption for Existing Job

If you enable encryption, during the next job session Veeam Agent 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.

If you change the password for the already encrypted job, during the next job session Veeam Agent 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.

If you disable encryption, during the next job session Veeam Agent will automatically create a full backup file.

Storage encryption

You cannot specify encryption options for the backup job if you have chosen to save Veeam Agent backups in a Veeam backup repository. For such jobs, encryption options are managed per repository by a backup administrator working with Veeam Backup & Replication. To learn more, see the [Encrypting Standalone Application Backups in Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

How Data Decryption Works

When you restore data from an encrypted backup file, Veeam Agent performs data decryption automatically in the background or requires you to provide a password.

- If encryption keys required to unlock the backup file are available in the Veeam Agent database, you do not need to enter the password. Veeam Agent uses keys from the database to unlock the backup file. Data decryption is performed in the background, and data restore does not differ from that from an unencrypted one.

Automatic data decryption can be performed when you encrypt and decrypt the backup file on the same Veeam Agent computer using the same Veeam Agent database.

- If encryption keys are not available in the Veeam Agent database, you need to provide a password to unlock the encrypted file.

Data decryption is performed on the source or target side depending on the backup target. As a result, encryption keys are not passed to the untrusted side, which helps avoid data interception.

In Veeam Agent, the decryption process includes the following steps. Keep in mind that steps 1 and 2 are required only if you decrypt the file on the Veeam Agent computer other than the computer where the file was encrypted.

1. You select the backup from which you want to restore data. Veeam Agent notifies you that one or more files in the backup chain are encrypted and requires a password.
2. You specify a password for the imported file. If the password has changed once or several times, you need to specify the latest password. In Veeam Agent, you can use the latest password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

3. Veeam Agent reads the entered password and generates the user key based on this password . With the user key available, Veeam Agent performs decryption.

After the encrypted file is unlocked, you can work with it as usual.



Backup of Database Systems

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run one of the following database systems:

- [Oracle database system](#)
- [MySQL database system](#)
- [PostgreSQL database system](#)

To process database systems with Veeam Agent for Linux, enable application-aware processing for the backup job:

- If you work with Veeam Agent using the Veeam Agent control panel, configure application-specific settings at the **Advanced** step of the Backup Job wizard. To learn more, see [Specify Advanced Backup Settings](#).
- If you work with Veeam Agent using the command line interface, create the backup job, then specify application-specific settings for this job. To learn more, see [Creating Volume-Level Backup Job](#) and [Configuring Database Processing Settings](#).

Considerations and Limitations

When you back up database systems, consider the following:

- Nosnap Veeam Agent for Linux and nosnap Veeam Agent for Linux on Power do not support application-aware processing and cannot be used to back up database systems.
- You can specify settings for database system processing only if Veeam Agent for Linux operates in the Server edition.
- Veeam Agent supports processing of database systems for the volume-level backup only.
- If there are multiple database systems on the Veeam Agent computer, consider the following:
 - Veeam Agent supports processing of multiple PostgreSQL or Oracle database systems on one Veeam Agent computer.
 - Veeam Agent does not support processing of multiple MySQL database systems on one Veeam Agent computer.
 - Veeam Agent does not support processing of multiple database systems of different types on one Veeam Agent computer.
- Veeam Agent does not support 32-bit database systems installed on a 64-bit Linux OS.

Oracle Backup

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the Oracle database system.

NOTE

You can use Veeam Explorer for Oracle to restore Oracle databases from a Veeam Agent for Linux backup. For information about item-level recovery of Oracle systems, see the [Restoring Oracle Items](#) section of the Veeam Backup & Replication User Guide.

Requirements and Limitations for Oracle Processing

- Oracle Database versions 11g – 21c are supported for all operating systems supported by Veeam Agent for Linux. To learn more, see [System Requirements](#).
- Automatic Storage Management (ASM) is not supported.
- Oracle Real Application Clusters (RAC) are not supported.
- Oracle Grid Infrastructure is not supported.
- Oracle Database Express Edition (XE) is not supported.
- SAP on Oracle is not supported.
- Oracle Database architectures with Data Guard are not supported.

Authentication Methods

Veeam Agent for Linux can connect to the Oracle database system and perform Oracle archived logs backup and delete operations using one of the following account types:

- *System account* – Veeam Agent uses the account of the machine OS. To connect to the Oracle database system, the account must be a member of the group that owns configuration files for the Oracle database (for example, the oinstall group).
- *Oracle account* – Veeam Agent uses the Oracle account. To connect to the Oracle database system, the account must have SYSDBA rights.

How Oracle Processing Works

To ensure that the backed-up data is in the consistent state, Veeam Agent for Linux performs the Oracle database system processing using an internal component: *oralib*. To process the database system, Veeam Agent performs the following operations:

1. When the backup job starts, Veeam Agent obtains information about Oracle databases that run on the Veeam Agent machine.
2. Veeam Agent connects to the Oracle database and operates depending on the database state and mode:
 - [Shutdown state](#)
 - [Backup state](#)
 - [Running database in ARCHIVELOG mode](#)
 - [Running database in NOARCHIVELOG mode](#)

After Veeam Agent for Linux finishes database system processing, Veeam Agent proceeds to the next step of the backup process. To learn more, see [How Backup Works](#).

Processing of Database in Shutdown State

If the Oracle database is shut down, Veeam Agent skips it and tries to connect to the next Oracle database if there are multiple database instances on the machine. The skipped Oracle database will be included in the backup. You cannot restore such Oracle database as an independent item using Veeam Explorer for Oracle. To restore such database, you must restore the entire volume that contains the database. To learn more about restoring volumes, see [Volume-Level Restore](#).

Veeam Agent displays a warning message about the database that is shut down in the job session logs. The backup job does not fail.

Processing of Database in Backup State

If the database is in the backup state, depending on the selected Oracle Processing option, Veeam Agent performs application-aware processing differently:

- If the Oracle processing is set to **Require successful processing**, the backup job will fail.
- If the Oracle processing is set to **Try application processing, ignore failures**, Veeam Agent will skip the database that is in the backup state and if there are multiple databases in the system, will try to connect to the next database. The skipped database will not be included in the backup.

Processing of Running Database in ARCHIVELOG Mode

If the Oracle database is running in the ARCHIVELOG mode, the Oracle database system keeps archived logs that allow to recover all committed transactions of the database. To learn more, see [Oracle documentation](#).

If the database operates in the ARCHIVELOG mode, Veeam Agent performs the following operations:

1. Veeam Agent switches the database to the backup mode. Veeam Agent changes the database state using the Oracle functionality.
2. Veeam Agent creates a snapshot of the volume.
3. Veeam Agent returns the database to the initial state.

Processing of Running Database in NOARCHIVELOG Mode

If the Oracle database is running in the NOARCHIVELOG mode, the Oracle database does not create archived logs. Logs that are created before the database is switched to NOARCHIVELOG remain untouched. In this mode, you can restore the database only to the state in which the database is contained in the restore point. You cannot recover transactions subsequent to that full database backup.

If the database operates in the NOARCHIVELOG mode, Veeam Agent performs the following operations:

1. Shuts down the database using the Oracle functionality.
2. Creates a snapshot of the volume.
3. Returns the database to the initial state.

Archived Log Processing

In the ARCHIVELOG mode, the Oracle database system stores database archived logs to a certain location on the machine that runs the database system, as specified by the database administrator. Veeam Agent allows you to set up the following ways of archived logs processing:

- *Delete logs older than the specified time (in hours).* After the backup job completes, Veeam Agent deletes archived logs that are older than the specified time from the Veeam Agent machine. This helps make sure that logs do not overflow the storage space on the processed machine.
- *Delete oldest logs larger than the specified size (in GB).* After the backup job completes, Veeam Agent checks whether the total size of archived logs exceeds the specified size. After that, Veeam Agent deletes oldest archived logs that exceed the specified size from the processed machine. This helps make sure that logs do not overflow the storage space on the Veeam Agent machine.
- *Do not delete archived logs.* Log files remain untouched on the Veeam Agent machine.

Veeam Agent processes archive logs via Oracle Call Interface (OCI).

MySQL Backup

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the MySQL database system.

Requirements and Limitations of MySQL Processing

- Veeam Agent for Linux supports processing of MySQL database systems version 5.7 – 9.0.
- Configurations with multiple MySQL installations or instances on the same machine are not supported.
- MySQL Cluster versions are not supported.
- MySQL tables that use the MyISAM storage engine must be locked to keep them in consistent state while Veeam Agent is creating the system snapshot. To correctly process such tables, MySQL account must have the following instance-wide privileges:
 - `SELECT`. This privilege enables Veeam Agent to access tables' metadata and select for a lock the tables that use the MyISAM storage engine. Without this privilege, the processing of the MySQL database system will run successfully but MyISAM tables will not be locked, which may result in an inconsistent state of the backed up data.
 - `LOCK TABLES`. This privilege is required for locking the selected MyISAM tables. If some MyISAM tables are selected but the MySQL account does not have the `LOCK TABLES` privilege, the processing of the MySQL database system will fail.
 - `RELOAD` or `FLUSH TABLES`. If some MyISAM tables are selected but the MySQL account does not have either `RELOAD` or `FLUSH TABLES` privilege, the processing of the MySQL database system will fail.

To obtain information about the privileges that are assigned to an account, use MySQL functionality, for example, the `SHOW GRANTS` statement. To learn more, see [MySQL documentation](#).

Authentication Methods

Veeam Agent for Linux can connect to the MySQL database system using one of the following methods:

- *Password* – Veeam Agent uses the MySQL account credentials that you specify in the backup job settings.
- *Password file* – Veeam Agent uses the MySQL account credentials that are stored in the `.my.cnf` password file. To learn more about password file configuration, see [Preparing Password File for MySQL Processing](#).

How MySQL Processing Works

To ensure that the backed-up data is in the consistent state, Veeam Agent for Linux performs the MySQL database system processing. To process the database system, Veeam Agent performs the following operations:

1. When the backup job starts, Veeam Agent connects to the MySQL database system and obtains the list of tables.

2. Veeam Agent locks the base tables that use the MyISAM storage engine. Veeam Agent changes the table state using the MySQL functionality. Tables that use the InnoDB storage engine do not require locking.
Keep in mind that Veeam Agent supports processing of tables based on the MyISAM and InnoDB storage engines only. Veeam Agent does not support tables that use other storage engines.
3. Veeam Agent creates a snapshot of the volume.
4. Veeam Agent unlocks tables locked at Step 2.

After Veeam Agent unlocks tables, Veeam Agent proceeds to the next step of the backup process. To learn more, see [How Backup Works](#).

PostgreSQL Backup

You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the PostgreSQL database system.

NOTE

You can use Veeam Explorer for PostgreSQL to restore PostgreSQL instances from a Veeam Agent for Linux backup. For information about item-level recovery of PostgreSQL systems, see the [Restoring PostgreSQL Items](#) section of the Veeam Backup & Replication User Guide.

Requirements and Limitations for PostgreSQL Processing

- Veeam Agent supports processing of PostgreSQL database systems version 12 – 17.
- Veeam Agent does not support backup of PostgreSQL clusters.
- Veeam Agent does not support high availability cluster configurations and replication setups of PostgreSQL servers.

Authentication Methods

Veeam Agent for Linux can connect to the PostgreSQL database system using one of the following methods:

- *Database user with password* – Veeam Agent uses the PostgreSQL account credentials that you specify in the backup job settings.
- *Database user with password file* – Veeam Agent the PostgreSQL database system to use account credentials that are stored in the `.pgpass` password file. To learn more about password file configuration, see [Preparing Password File for PostgreSQL Processing](#).
- *System user without password* – Veeam Agent uses the peer authentication. In the peer authentication method, Veeam Agent for Linux uses the OS account as the PostgreSQL database user name.

How PostgreSQL Processing Works

After Veeam Agent for Linux finishes database system processing, Veeam Agent proceeds to the next step of the backup process. To learn more, see [How Backup Works](#).

To ensure that the backed-up data is in the consistent state, Veeam Agent performs the PostgreSQL database processing using an internal component: *pgsqlagent*. To process the database system, Veeam Agent performs the following operations:

1. When the backup job starts, Veeam Agent scans the Veeam Agent computer for PostgreSQL instances.

By default, Veeam Agent recursively scans the `/etc/postgresql`, `/var/lib/postgresql` and `/var/lib/pgsql` directories for the configuration files of PostgreSQL instances. If your instance is stored in a custom location, you must specify its location in the PostgreSQL configuration file – `VeeamPostgreSQLAgent.xml`. You must create this file in the `/etc/veeam/` directory. To explicitly include or exclude specific directories in/from processing, you can use the following elements in the configuration file:

- `AddConfigDirs` – use this element to specify paths to directories that you want Veeam Agent to scan.
- `ExcludeConfigDirs` – use this element to specify paths to directories that you do not want Veeam Agent to scan.

TIP

You can specify directories that you want to include and directories that you want to exclude in the same configuration file.

An example of the `VeeamPostgreSQLAgent.xml` file:

```
<config AddConfigDirs="/opt/psql/" ExcludeConfigDirs="/var/lib/postgresql/13/main45/,/var/lib/postgresql/13/maindd/" />
```

IMPORTANT

The configuration file must be formatted as a single line XML.

2. If a PostgreSQL instance is detected, Veeam Agent collects information about its state and configuration settings. The following Veeam Agent behavior depends on the collected information:
 - [Shutdown state](#)
 - [Backup state](#)
 - [Running instance with WAL level set as minimal](#)
 - [Running instance with WAL level set as archival, replica or logical](#)

To learn more about the WAL level setting, see [PostgreSQL documentation](#).

TIP

Veeam Agent stores all collected data about PostgreSQL instances in the `.VBM` file, which allows Veeam Agent to restore PostgreSQL instance as an application item. To learn more, see [Restoring PostgreSQL Items](#) in the Veeam Backup & Replication User Guide.

3. Veeam Agent creates a snapshot of the volume and proceeds to the next step of the backup process.

To learn more about backup process, see [How Backup Works](#).

Processing of Instance in Shutdown State

If the database instance is shut down, Veeam Agent skips it and tries to connect to the next instance. The skipped database instance will be included in the backup. You cannot restore such PostgreSQL instance as an independent item using Veeam Explorer for PostgreSQL. You can restore such database instance only using either volume-level or file-level restore. To learn more about restoring volumes and files, see [Data Restore](#).

Processing of Instance in Backup State

If the database instance is in the backup state, depending on the selected PostgreSQL Processing option, Veeam Agent performs the backup job differently:

- If the PostgreSQL processing is set to **Require successful processing**, the backup job will fail.
- If the PostgreSQL processing is set to **Try application processing, ignore failures**, Veeam Agent will skip the instance that is in the backup state and will try to connect to the next instance. The skipped database instance will not be included in the backup.

Processing of Instance with WAL Level Set as Minimal

If the database is running and the WAL level is set as minimal, Veeam Agent forces a WAL checkpoint. This command fastens the database system restore. To learn more, see [PostgreSQL documentation](#).

Keep in mind that the backup of a PostgreSQL instance with the minimal WAL level does not contain logs. As a result, you can restore your instance only to image-level backup state.

Processing of Instance with WAL Level Set as Archival, Replica or Logical

If the database instance is running and the WAL level is set as archival, replica or logical, Veeam Agent performs the following operations:

- Prepares the PostgreSQL instance and starts the on-line backup.
- Creates a snapshot of the instance.
- Stops the on-line backup.

Backup to Object Storage

If you want to store your data in a cloud-based or on-premises object storage, you can create backup jobs targeted at the object storage or another storage that uses object storage as a back end.

You can store Veeam Agent backups in the following object storage types:

- S3 compatible (including Wasabi Cloud Storage and IBM Cloud Object Storage)
- Amazon S3
- Google Cloud Storage
- Microsoft Azure Blob Storage
- Veeam Data Cloud Vault added as a Veeam backup repository or Veeam Cloud Connect repository. To learn more, see [Backup Destinations](#).

Depending on your backup infrastructure, object storage can be available in different configurations. To learn more, see the following subsections:

- [Backup Destinations](#)
- [Types of Connection to Object Storage in Veeam Backup & Replication](#)
- [Considerations and Limitations](#)

Backup Destinations

You can back up Veeam Agent computer data to object storage in the following ways:

- Directly to object storage. In this case, Veeam Agent connects to the object storage and creates a backup repository in this storage.

Keep in mind that to connect to object storage, you need to specify an account with access permissions to read and write data.

To learn more, see [Object Storage Settings](#).

- To object storage added as a Veeam backup repository. In this case, an object storage repository is created in Veeam Backup & Replication. Veeam Agent connects to Veeam Backup & Replication and uses this Veeam backup repository.

To learn more, see [Veeam Backup Repository Settings](#).

- To object storage added as a Veeam Cloud Connect repository. In this case, a Veeam Cloud Connect service provider creates an object storage repository in Veeam Backup & Replication and exposes it as a cloud repository to tenants. Veeam Agent connects to the service provider and uses the Veeam Cloud Connect repository.

To learn more, see [Veeam Cloud Connect Repository Settings](#).

Connection Types

If you back up data to object storage added as a Veeam backup repository or Veeam Cloud Connect repository, you must configure a repository beforehand on the Veeam Backup & Replication side. Depending on the repository configuration, Veeam Backup & Replication provides one of the following connection types to the repository in the object storage:

- Connection through a gateway server. With this connection type, Veeam Agent connects to the repository using a proxy component – a gateway server that is assigned in the Veeam Backup & Replication console. The backup data is transferred from the Veeam Agent computer to the gateway server, then it is transferred from the gateway server to the repository.
- Direct connection. With this connection type, Veeam Agent connects directly to the repository. The backup data is transferred from the Veeam Agent computer to the repository without proxy components. The access to this repository is managed by Application Programming Interface (API) that is provided by an external cloud service provider.

Considerations and Limitations

Before you configure a backup job to store backups in object storage, consider the following:

- Veeam Agent does not support backup to object storage for which lifecycle rules are enabled. Enabling lifecycle rules may result in backup and restore failures.
- You cannot back up data to the Microsoft Azure Blob Storage added in the direct connection mode under the general-purpose V1 storage account type.
- You cannot back up data to the Veeam Data Cloud Vault storage added in the direct connection mode.
- You can store backups only in those S3 compatible storage repositories that are accessible over the HTTPs protocol.
- [For object storage added as a Veeam backup repository or Veeam Cloud Connect repository] If you want to back up your data to the S3 compatible storage using a direct connection, you must additionally specify access permissions settings for the storage. To do so, enable the **Agents share credentials to object storage repository** or the **Provided by IAM/STS object storage capabilities** access control option. To learn more, see the [Managing Permissions for S3 Compatible Object Storage](#) section in the Veeam Backup & Replication User Guide.
- [For object storage added as a Veeam backup repository or Veeam Cloud Connect repository] Data recovery options are not available if you access the object storage repository using credentials with the read-only access permissions.
- [For object storage added as a Veeam backup repository] If you back up data to the S3 compatible object storage with multiple buckets, Veeam Agent will ignore the number of workloads set for one bucket and will store all backups, created under a specific account that has access to the repository, in a single child bucket. To learn more about multiple buckets, see the [Multiple Buckets for S3 Compatible Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.

Health Check for Object Storage

If you keep the backups of your Linux computer in an object storage repository, you can schedule regular health checks to validate integrity of the backups in the repository.

Consider the following about health check for object storage:

- Veeam Agent verifies metadata of the whole backup, not just the latest restore point.
- Veeam Agent does not read data from data blocks in the storage; Veeam Agent only lists data blocks to make sure all blocks in the storage are available for rebuilding every restore point in the active backup chain. This mechanism reduces the number of requests to the storage, which makes health check for object storage cost-efficient.

Configuring Health Check Schedule

If you want to run health checks for a backup that resides in an object storage repository, you must set a schedule according to which Veeam Agent will perform health checks. You can set the schedule [in the Backup Job wizard](#) or [in command line interface](#) to run health checks weekly or monthly on specific days.

When you configure a health check schedule, consider the following:

- Health check is run automatically during incremental backup job session on the days specified in the health check schedule. If the backup job runs several times on a specified day, health check is performed only with the first run of the backup job on that day.

Health check is not performed during the first full backup or subsequent active full backup job sessions.

- If Veeam Agent does not run any backup jobs on the day specified in the health check schedule, health check will be performed during the first backup job session following that day.

For example, you may have scheduled to run a health check every last day of a month, while the backup job is scheduled to run every day and to create an active full backup on Sundays. If the last day of a month falls on a Sunday, the health check will be performed on the following Monday with the first incremental backup job session on that day.

How Health Check Works

Veeam Agent performs a health check of a backup in the following way:

1. During the backup job session after a new incremental backup file is created, Veeam Agent starts the health check of the whole backup. Veeam Agent checks if the metadata of the backup is consistent, and no metadata is missing. Veeam Agent also checks if all data blocks for every restore point are available on the storage. Veeam Agent does not read data from data blocks.
2. If Veeam Agent does not find any corrupted data, the health check completes successfully. Otherwise, the health check completes with an error.

You can view the health check result in the session log. If during the health check, Veeam Agent finds corrupted data, it will also display information on where corrupt data has been detected – in metadata or blockstore, as well as list all restore points that share the corrupted data blocks.

Depending on the detected data inconsistency, Veeam Agent behaves in one of the following ways:

- If the health check detects corrupted metadata, Veeam Agent will mark the backup chain as corrupted in the Veeam Agent configuration database; the backup job session will fail. During the next scheduled or manual backup job session, Veeam Agent will create a full backup and will start a new backup chain. The corrupted backup chain will become orphaned and will remain in the repository – you can keep or delete it.

- If the health check detects corrupted data blocks in the latest restore point of the active backup chain, Veeam Agent launches a health check retry.

During the health check retry, Veeam Agent restarts the backup job to create a new restore point and transports data blocks from the Veeam Agent computer including the blocks that were corrupted in the object storage repository and the blocks that changed since the start of the backup job session that triggered the health check. Veeam Agent will not perform another health check after the job retry is finished successfully. The next health check will be run according to the defined schedule.

- If the health check detects corrupted data blocks in an inactive backup chain, Veeam Agent will not launch a health check retry. Veeam Agent will mark the backup and all related restore points as corrupted; the backup job session will end with a warning message.

NOTE

If you try to restore data from a corrupted backup, Veeam Agent will display a warning message informing you that the restore operation may fail or the restored data may be corrupted.

Backup Immutability

If you store your backup files in an object storage repository, Veeam Agent allows you to protect backup data from deletion or modification by making that data temporarily immutable. It is done for increased security: immutability protects data in your recent backups from loss as a result of attacks, malware activity or any other injurious actions.

IMPORTANT

Backup immutability uses native object storage capabilities. You may incur additional API and storage charges from the storage provider.

Supported Object Storage Types

Veeam Agent supports backup immutability for the following object storage types:

- Amazon S3
- S3 compatible storage that supports S3 Object Lock (including Wasabi)
- Microsoft Azure Blob Storage
- Veeam Data Cloud Vault

NOTE

Veeam Agent does not support backup immutability for the Google Cloud storage.

Before You Begin

Before you configure immutability for Veeam Agent backups, you must prepare the target storage account. Depending on the selected object storage type, perform the following actions:

- [S3 Compatible and Amazon S3 storage] When you create the S3 bucket, you must enable versioning and the S3 Object Lock feature for the bucket. For more information, see [AWS documentation](#).

- [S3 Compatible and Amazon S3 storage] After you create the S3 bucket with Object Lock enabled, make sure that the default retention is disabled to avoid unpredictable system behavior and data loss. To disable the default retention, edit the Object Lock retention settings as described in [AWS documentation](#).
- [Microsoft Azure Blob storage] You must enable blob versioning and version-level immutability support for the Azure container. For more information, see [Microsoft documentation](#).

Consider the following about backup immutability:

- The effective immutability period consists of the user-defined immutability period and the block generation period automatically appended by Veeam Agent. For more information, see [How Backup Immutability Works](#) and [Block Generation](#).
- [S3 Compatible and Amazon S3 storage] Veeam Agent will use the *compliance* retention mode for each uploaded object. For more information on retention modes of S3 Object Lock, see [AWS documentation](#).
- [Microsoft Azure Blob storage] Do not enable immutability for already existing containers in the Microsoft Azure Portal. Otherwise, Veeam Agent will not be able to process these containers properly and it may result in data loss.

Configuring Backup Immutability

Depending on how you create the backup job and configure [connection to an object storage repository](#), you can define backup immutability settings in one of the following ways:

- [Backup Job wizard] You must specify the immutability period at the Bucket step of the wizard. For more information, see [Object Storage Settings](#).
- [Command line interface] You must specify the immutability period in the advanced options of the command for creating the backup job. For more information, see [Creating Backup Job with Command Line Interface](#).

NOTE

If you want to create the backup job in command line interface, you must create the object storage repository first. For details, see [Creating Repository in Object Storage](#).

- If you create the backup job that is targeted at an object storage repository configured as a Veeam backup repository or Veeam Cloud Connect repository, the immutability period in the settings of the repository must be specified in Veeam Backup & Replication. For details, see the [Adding Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.

Backup Immutability and Retention Policy

Backup immutability operates with backup data and related metadata (checkpoints) on the object storage side. Retention policy operates with logical representation of the stored data, or restore points, on the Veeam Agent side. These two mechanisms act independently from each other.

Veeam Agent will remove the irrelevant restore points per the defined backup retention policy. If the data associated with the removed restore point is still immutable, such data will remain in the repository until expiration of the immutability period. After that it will be automatically removed from the storage.

Limitation of Backup Immutability

If you use Veeam Agent in the standalone mode, you can restore the immutable data that is associated with a restore point removed by retention policy only in Veeam Backup & Replication console. In Veeam Backup & Replication, you must perform the following actions:

1. Add the object storage repository that contains the necessary data to Veeam Backup & Replication. For more information, see the [Adding Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.
2. Roll back to the necessary checkpoint. For more information, see the [Immutability](#) section in the Veeam PowerShell Reference.
3. Remove the repository from the Veeam Backup & Replication infrastructure. For more information, see the [Removing Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

After that, you will be able to use Veeam Agent to restore data from the object repository in a regular manner.

How Backup Immutability Works

After you specify the immutability period for a backup and run the backup job for the first time, Veeam Agent will append an additional period to the specified immutability period depending on the type of the object storage repository:

- 30 days – for Amazon S3 object storage.
- 10 days – for other types of object storage repositories.

This additional period is called *block generation*. The resulting effective immutability period is the sum of the user-defined immutability period and the block generation period. All data blocks transferred to the target repository within the block generation period will have the same immutability expiration date. For example, data block *a* added on day 1 of the block generation period will have the same immutability expiration date as block *b* added on day 9. For more information, see [Block Generation](#).

During the effective immutability period, the following operations with backup data in the object storage repository will be prohibited:

- Manual removal of data from the backup repository.
- Removal of data by backup retention policy.
- Removal of data using any object storage provider tools.
- Removal of data by the technical support department of the object storage provider.

Extension of Effective Immutability Period

During each transfer of data to the object storage repository, Veeam Agent creates a new checkpoint file with metadata that describes the latest state of the backup in the storage. The immutable blocks of data from a previous checkpoint may be reused in the newly created checkpoint. Veeam Agent keeps reused, or dependent, blocks of data locked by continuously assigning them to new generations and extending their effective immutability period. This guarantees that the effective immutability period is no less than the immutability period defined by user.

During data transfer, the effective immutability period for the backup is set as follows:

- [For new data blocks in the checkpoint] Immutability is set anew. The user-defined immutability period is appended with a 10 or 30-day block generation period.

- [For data blocks reused from the previous checkpoint] Immutability is extended to the immutability expiration date set for the new blocks.
- [For data blocks that are not reused in the checkpoint] Immutability is not extended. Such data blocks will remain in the repository until their immutability period is over. After that Veeam Agent will automatically remove them from the repository.

Block Generation

When you specify an immutability period for the recent backups, Veeam Agent will automatically extend the immutability expiration date by 10 or 30 days depending on the type of the object storage repository:

- 30 days – for Amazon S3 object storage.
- 10 days – for other types of object storage repositories.

This period is called *block generation*. The block generation period serves to reduce the number of requests to the object storage repository, which results in lower traffic and reduced storage costs. You do not have to configure it, the block generation period is applied automatically.

When the block generation period is appended to the user-defined immutability period, it means there is no need to extend the immutability period for old data blocks when adding new data blocks to the backup during that block generation period.

Consider this example. When you create a full backup to start a backup chain, all data blocks transferred to the object storage repository are new. For these new blocks of data, Veeam Agent will add the block generation period to the specified immutability period. For example, if the immutability period is set by user to the default period of 30 days the effective immutability period with the added block generation period will become 40 days. The first full backup starts its generation that will last for 10 days. All new and reused data blocks within this block generation period will have the same immutability expiration date. For instance, a data block that was transferred to the target repository on day 9 will have the same immutability expiration date as a data block transferred on day 1. This mechanism guarantees that the effective immutability period for all the data blocks within a generation is no less than 30 days.

If a block generation period is over but data blocks from that generation are reused in the newly created checkpoint, their effective immutability period is automatically extended to ensure that the effective immutability period for all the data blocks in the new checkpoint is no less than the user-defined immutability period. For more information, see [How Backup Immutability Works](#).

Backup to Deduplicating Storage Appliances

You can store backups in the deduplicating storage appliances added as backup repositories. To learn the full lists of supported appliances and their requirements and limitations, see the [Deduplicating Storage Appliances](#) section in the Veeam Backup & Replication User Guide.

If you want to use immutability with the deduplicating storage appliances, consider the limitations listed in the following subsections:

- [HPE StoreOnce Immutability and Veeam Agents](#)
- [Dell Data Domain Immutability and Veeam Agents](#)

HPE StoreOnce Immutability and Veeam Agents

If you want to use immutability with HPE StoreOnce Catalyst repositories, make sure that the value for the Maximum ISV Controlled Data Retention setting in HPE StoreOnce exceeds or is equal to the value of each of the following settings configured in Veeam Backup & Replication:

- The immutability period specified in the backup repository settings.
- The long-term retention period configured in backup job settings.

If the value of at least one of these periods exceeds the maximum value set for immutability in HPE StoreOnce Catalyst Store, consider the following:

- If the immutability period specified in the backup repository settings exceeds the maximum value set for immutability in HPE StoreOnce Catalyst Store, Veeam Backup & Replication applies immutability to the created backups according to the maximum value set for immutability in HPE StoreOnce Catalyst Store.
- If the long-term retention period configured in backup job settings exceeds the maximum value set for immutability in HPE StoreOnce Catalyst Store, the immutability period for backups with GFS flags is set according to the maximum value set for immutability in HPE StoreOnce Catalyst Store. For example, if the backup is stored for 1 year, but the maximum value for immutability in HPE StoreOnce Catalyst Store is set to 6 months, the backups will be immutable for 6 months.

To learn more about immutability for the HPE StoreOnce Catalyst repositories, see the [HPE StoreOnce and Immutability](#) section in the Veeam Backup & Replication User Guide.

Dell Data Domain Immutability and Veeam Agents

If you want to use immutability with the Dell Data Domain backup repository, make sure that the values of the following settings configured in Veeam Backup & Replication lie in the range between the minimum and maximum retention periods configured in Dell Data Domain:

- The immutability period specified in the backup repository settings.
- The long-term retention period configured in backup job settings.

The minimum and maximum values are included into the range.

If the value of at least one of these periods lies outside the established range, consider the following:

- If the immutability period specified in the backup repository settings lies outside the range between the minimum and maximum retention periods configured in Dell Data Domain, the immutability for backups is set equal to the nearest range value.

- If the long-term retention period configured in backup job settings lies outside the range between the minimum and maximum retention periods configured in Dell Data Domain, the immutability period for backups with GFS flags is set equal to the nearest range value. For example, if the backup is stored for 1 year, but the range for immutability in Dell Data Domain is set to from 2 to 6 months, the backups will be immutable for 6 months.

To learn more about immutability for the Dell Data Domain backup repositories, see the [Retention Lock](#) section in the Veeam Backup & Replication User Guide.

Backup to Veeam Cloud Connect Repository

If you want to store your data in the cloud, you can connect to a Veeam Cloud Connect service provider (SP) and create Veeam Agent backups in a cloud repository. To do this, you must provide credentials of the tenant (or subtenant) account that you obtained from the SP and select a cloud repository as a target for backup files in the properties of the backup job. To learn more, see [Veeam Cloud Connect Repository Settings](#).

NOTE

Consider the following:

- You can create Veeam Agent backups in a cloud repository if the SP backup server runs Veeam Backup & Replication 12.0 or later.
- Backup to a cloud repository is available if Veeam Agent for Linux operates in the Workstation or Server edition.

Data Restore

Veeam Agent for Linux offers two data restore scenarios:

- You can perform volume-level restore to recover the entire system image of your computer or specific computer volumes. To learn more, see [Volume-Level Restore](#).
- You can perform file-level restore to recover individual files and directories. To learn more, see [File-Level Restore](#).

Volume-Level Restore

If data on a computer volume gets corrupted, you can restore this volume from the backup. For volume-level restore, you can use backups that were created at the volume level. File-level backups cannot be used for volume restore.

When you perform volume-level restore, Veeam Agent for Linux restores the entire content of the volume. It retrieves from the backup data blocks pertaining to a specific volume and copies them to the necessary location.

Keep in mind that you cannot browse the volume in the backup and select individual files and directories for restore. For granular file-level restore, you can use the [File-Level Restore](#) option.

A volume can be restored to its original location or new location. If you restore the volume to its original location, Veeam Agent for Linux overwrites data on the original volume. If you restore the volume to a new location, and the target disk contains any data, Veeam Agent for Linux overwrites data in the target location with data retrieved from the backup.

Limitations for Volume-Level Restore

Volume restore has the following limitations:

- You cannot restore the system volume to its original location.
- You cannot restore a volume to the volume on which the Linux swap space is hosted.
- You cannot restore a volume to the volume where the backup file used for restore is located.

To overcome the first two limitations, you can create a Veeam Recovery Media and use the **Volume Restore** wizard for volume-level restore. To learn more, see [Veeam Recovery Media](#).

File-Level Restore

If you have lost or modified files and directories on your computer by mistake, you can restore a copy of the necessary objects from the backup. For file-level restore, you can use a backup of any type:

- Volume-level backup
- File-level backup

Veeam Agent for Linux does not simply extract files and folders from the backup file. During file-level restore, Veeam Agent for Linux performs the following operations:

1. Veeam Agent for Linux associates the backup file with a loop device, for example, `/dev/loop0`, to make the backup file accessible as a block device.
2. Veeam Agent for Linux mounts the loop device to the mount point directory in the computer's file system.
 - For file-level restore with the Veeam Agent for Linux control panel or Veeam Recovery Media, Veeam Agent for Linux mounts the backup content to the `/mnt/backup` directory.
 - For file-level restore with the command line interface, you can specify a directory in which Veeam Agent for Linux should mount the backup content.

After the backup content is mounted, you can use Linux command line utilities or preferred file browser to work with restored files and directories. You can browse for files and directories in the mounted backup and copy them to their initial location or to a new location.

Veeam Recovery Media

Veeam Agent for Linux lets you use the Veeam Recovery Media – a recovery image of the Linux OS that provides an alternative way to boot your computer.

The recovery image includes a custom Linux OS with the limited functionality. It comprises Linux kernel and a set of GNU/Linux utilities necessary to boot the computer and perform basic administration tasks. If the OS installed on the computer fails to start for some reason, you can boot the recovery image OS. After booting, you can do the following:

- You can restore data from a backup to your computer. For this scenario, you must have a backup created with Veeam Agent for Linux.
- You can use Linux OS tools to diagnose problems and fix errors on your computer.

The recovery image can be helpful if one of the following errors occur:

- The OS on the computer fails to start.
- You want to perform bare metal restore from the backup on the computer without the OS and other software installed.
- You want to restore the system volume of the computer and so on.

Veeam Recovery Media is distributed as an ISO image. You can download the ISO image file from [this Veeam webpage](#): select an operating system to display the download links for the product and recovery ISO. You can burn the ISO image file to the following types of media:

- Removable storage devices such as USB drives or SD cards
- CD/DVD/BD

NOTE

Consider the following:

- You can also download the Veeam Recovery Media ISO image from the [Veeam software repository](#).
- For information about how to burn the ISO image to a removable storage device, as well as workaround for accessing the Veeam recovery UI, see [this Veeam KB article](#).

When you boot from the Veeam Recovery Media, you can use the recovery environment to fix the OS system errors on your computer or restore data from the backup. Veeam Agent for Linux offers a set of tools for the computer system image and data recovery:

- Restore volumes – the Veeam Recovery wizard to recover data on the original computer or perform bare metal recovery.
- Restore files – the File Level Restore wizard to restore files and folders to the original location or to a new location.
- Exit to shell – Linux shell prompt with standard utilities to diagnose problems and fix errors.

Veeam Recovery Media Versions

Veeam Agent for Linux offers Veeam Recovery Media for computers based on the x86 and x64 architecture with Linux kernel version 3.10 and later.

You cannot create custom Veeam Recovery Media for Veeam Agent computers that run Linux kernel version earlier than 3.10. For a workaround, see [this Veeam KB article](#).

You can download the recovery image from the following sources:

- [Veeam website](#)

Recovery image ISO files downloaded from the Veeam website have the following names:

- `veeam-recovery-media-6.3.2.1207_i386.iso` — for Veeam Agent computers based on the x86 architecture.
- `veeam-recovery-media-6.3.2.1207_x86_64.iso` — for Veeam Agent computers based on the x64 architecture.
- `veeam-recovery-media-6.3.2.1207_ppc64le.iso` — for Veeam Agent computers based on the IBM Power architecture.

- [Veeam software repository](#)

Recovery image ISO files downloaded from the Veeam software repository have the following names:

- `veeam-recovery-i386-6.0.0.iso` — for Veeam Agent computers based on the x86 architecture.
- `veeam-recovery-amd64-6.0.0.iso` — for Veeam Agent computers based on the x64 architecture.
- `veeam-recovery-ppc64le-6.0.0.iso` — for Veeam Agent computers based on the IBM Power architecture.

The size of the regular recovery image file depends on the Veeam Agent computer architecture: 571 MB for x86 computers and 663 MB for x64 computers and 656 MB for IBM Power machines.

Drivers in Veeam Recovery Media

The generic Veeam Recovery Media available for download from [the Veeam website](#) or [Veeam software repository](#) contains the following data:

1. Set of files required to start the recovery image OS from the recovery media.
2. Set of Veeam tools for the computer system image and data recovery.
3. Set of Linux command line tools to diagnose problems and fix errors on your computer. For the regular recovery image, in addition to the standard set of tools, you can install custom software from a software repository.
4. Drivers required to run hardware and devices on your computer in a regular way. The regular recovery image contains drivers included in the Linux kernel versions 6.1 and 6.11.

When you boot your computer from the Veeam Recovery Media, drivers from the Veeam Recovery Media are automatically loaded on the recovery image OS.

If your computer uses hardware that requires drivers not included in the generic Veeam Recovery Media, you can create a custom recovery image. Veeam Agent will copy the Linux kernel running on your computer with its currently loaded modules and include them into the custom Veeam Recovery Media. To learn more, see [Creating Custom Veeam Recovery Media](#).

Integration with Veeam Backup & Replication

IMPORTANT

To use Veeam Agent for Linux 6.3.2 with Veeam Backup & Replication, you must install Veeam Backup & Replication 12.3 on the Veeam backup server.

You can store backup files created with Veeam Agent for Linux on backup repositories managed by Veeam Backup & Replication. To do this, you must select a Veeam Backup & Replication backup repository as a target location in the properties of the backup job. To learn more about supported backup repositories, see the [Backup Repositories](#) and [Scale-Out Backup Repositories](#) sections in the Veeam Backup & Replication User Guide.

NOTE

Consider the following:

- The current guide covers subjects related to Veeam Agent for Linux operating in the standalone mode.
- You can also use Veeam Backup & Replication to manage Veeam Agent for Linux on computers in your infrastructure. As part of the Veeam Agent management scenario, you can remotely deploy Veeam Agent to your computers, as well as configure and manage Veeam Agent backup jobs in Veeam Backup & Replication. To learn more, see [Veeam Agent Management Guide](#).
- If you create a backup job with the Veeam Agent command line interface, you need to specify a Veeam backup repository in the backup job settings. Veeam backup repository appears in the list of backup repositories after you connect to a Veeam backup server. To learn more, see [Managing Veeam Backup & Replication Servers](#).

Veeam Agent for Linux works with the Veeam Backup & Replication backup repository as with any other backup repository. Backup files are stored to a separate folder; you can perform standard restore operations using these files.

Information about Veeam Agent backups stored on the Veeam Backup & Replication backup repositories, backup jobs and sessions becomes available in the Veeam Backup & Replication console:

- The Veeam Agent for Linux backup job is displayed in the list of jobs in Veeam Backup & Replication.
- Backup files created with Veeam Agent for Linux are displayed in the list of backups, under the **Agents** node.
- Performed job sessions are available in the **History** view of Veeam Backup & Replication.

Backup administrators working with Veeam Backup & Replication can perform a set of operations with Veeam Agent backups:

- Perform [data protection operations](#).
- Perform [restore operations](#).

- Perform [administrative tasks](#).

The screenshot shows the Veeam Backup and Replication console interface. The top navigation bar includes tabs for 'Home', 'Jobs', 'Inventory', 'Backup Infrastructure', and 'History'. The 'Home' tab is active, displaying a table of backup jobs. The table has columns for 'Job Name', 'Creation Time', 'Restore Points', 'Repository', and 'Platform'. The jobs listed are under the 'Agents' category.

Job Name	Creation Time	Restore Points	Repository	Platform
winsrv001 DailyBackup	1/15/2023 9:15 PM	3	Default Backup Repository	Windows
winsrv002 WeeklyBackup	1/15/2023 6:25 PM	4	Default Backup Repository	Windows
srv01 DailyBackup	1/12/2023 3:37 PM	3	Default Backup Repository	Linux
srv02 SystemBackup	12/6/2022 5:39 PM	4	Default Backup Repository	Linux

At the bottom of the console, a status bar indicates '1 backup selected'.

Managing Veeam Agent in Veeam Backup & Replication

Veeam Backup & Replication lets you automate management of Veeam Agent on multiple computers in your infrastructure. You can deploy Veeam Agent for Linux, configure Veeam Agent backup jobs and perform other data protection and administration tasks on remote computers. To use the Veeam Agent management functionality in Veeam Backup & Replication, you must install Veeam Backup & Replication on the Veeam backup server.

To learn more, see [Veeam Agent Management Guide](#).

Planning and Preparation

Before you install Veeam Agent for Linux, make sure that the target computer meets the system requirements, and all required ports are open.

System Requirements

The protected Linux computer must meet requirements listed in the table below.

NOTE

The following system requirements apply to the following Veeam Agent for Linux configuration:

- Veeam Agent for Linux version is 6.3.2.
- Veeam Agent for Linux is operating in the standalone mode.
To learn about system requirements for Veeam Agent managed by Veeam Backup & Replication, see the [System Requirements](#) section in the Veeam Agent Management Guide.
- Veeam Agent for Linux is installed with the `veeam-libs`, `veeam`, and Veeam kernel module packages.
To learn about system requirements for nosnap Veeam Agent for Linux, see [System Requirements for Nosnap Veeam Agent for Linux](#).
To learn about system requirements for nosnap Veeam Agent for Linux on Power, see [System Requirements for Nosnap Veeam Agent for Linux on Power](#).

Specification	Requirement
Hardware	Important! Check considerations and limitations that apply to the supported hardware. CPU: x86 or x64. Memory: 1 GB RAM or more. Memory consumption varies depending on the backup type and the total amount of backed-up data. For more information, see RAM Requirements for Backup of Large Number of Files. Disk Space: 100 MB free disk space for product installation. Network: 10 Mbps or faster network connection to a backup target. System firmware: BIOS or UEFI. Disk layout: MBR or GPT.

Specification	Requirement
OS	Important! Check considerations and limitations that apply to the list of supported OSes. Linux kernel version 2.6.32 to version 6.14 is supported. For the most up-to-date information on the Veeam Agent compatibility with Linux kernel versions of the supported distributions, see this Veeam KB article. Veeam Agent supports the 64-bit versions of the following Linux distributions: • Debian 10.13 – 12.11 • Ubuntu 16.04, 18.04, 20.04, 22.04, 22.10, 23.04, 23.10, 24.04, 24.10 and 25.04 • RHEL 6.4 – 9.6 and 10.0 • CentOS 7 • Oracle Linux 6 – 9.6 and 10.0 (RHCK) • Oracle Linux 6 (starting from UEK R2) – Oracle Linux 8 (up to UEK R6) • Oracle Linux 8 (UEK R7) – for information on installing Veeam Agent, see this Veeam KB article. • Oracle Linux 9 (UEK R7 up to 5.15.0-308.179.6.3.el9uek) • Oracle Linux 9 (UEK R8) – for information on installing Veeam Agent on Oracle Linux 9 with UEK R8, see this Veeam KB article. • SLES 12 SP4 and SP5, 15 SP1 – 15 SP6 • SLES for SAP 12 SP4 and SP5, 15 SP1 – 15 SP6 • Fedora 36, 37, 38 and 39 • openSUSE Leap 15.3 – 15.6 • openSUSE Tumbleweed has an experimental support status. • Rocky Linux 8.10, 9.3 – 9.6 and 10.0 • AlmaLinux 8.10, 9.3 – 9.6 and 10.0 Veeam Agent supports the 32-bit versions of the following distributions: • RHEL 6 • Oracle Linux 6 32-bit versions of other distributions are not supported.

Specification	Requirement
File System	Important! Check considerations and limitations that apply to the list of supported file systems. Veeam Agent for Linux supports consistent snapshot-based data backup for the following file systems: • BTRFS (for OSES that run Linux kernel 3.16 or later) • Ext 2/3/4 • F2FS • FAT16 • FAT32 • HFS • HFS+ • JFS • NTFS • ReiserFS • XFS The supported file system (except for BTRFS) can reside on a simple volume or LVM2 volume; volumes protected with encryption software such as dm-crypt are supported. BTRFS is supported only if it resides directly on a physical device with no additional abstraction layers (such as LVM, software RAID, dm-crypt and so on) below or above it. Other file systems, file systems that are not located on logical volumes, as well as network file systems like NFS or SMB shares can be backed up using the snapshot-less mode only. For details, see Snapshot-Less File-Level Backup.

Specification	Requirement
Software	Important! Check considerations and limitations that apply to the list of supported components. Protected computer must have the following components installed: • dkms • gcc • make • perl • linux-headers (for Debian-based systems) • kernel-headers (for RedHat-based systems) • kernel-devel (for RedHat-based systems) • kernel-uek-devel (for Oracle Linux with UEK) • libudev • libacl • libattr • lvm2 • libfuse2 (FUSE libraries for Debian-based and SLES-based systems) • fuse-libs (FUSE libraries for RedHat-based and Fedora systems) • libncurses5 • dmidecode • libmysqlclient • libpq5 • python3 • efibootmgr (for UEFI-based systems) • isolinux (for Debian-based systems) • syslinux (for RedHat-based systems) • btrfs-progs (for backup of BTRFS file system) • mksquashfs (for custom Veeam Recovery Media) • unsquashfs (for custom Veeam Recovery Media) • wget (for custom Veeam Recovery Media) • xorriso (for custom Veeam Recovery Media with EFI support) • tar (for file system indexing, log export and rotation) • gzip (for file system indexing, log export and rotation)

Considerations and Limitations

Hardware

- For virtual machines, only full virtualization type is supported. Oracle VM virtual machines are supported with [limitations](#). Virtual I/O (VirtIO) devices have [experimental support](#) status. Other containers and paravirtualized instances are not supported; backup of such devices may result in corruption of the source file system – for more information, see [this Veeam KB article](#).
- Devices managed by Veritas Volume Manager are not supported.

OS

- Linux kernel version 2.6.32 to version 6.14 is supported as long as you use kernels supplied by your distribution. Consider the following limitations of Veeam Agent compatibility with Linux kernel:
 - Fedora 38, 39 and openSUSE Tumbleweed are supported up to kernel 6.14.
 - Linux kernel 2.6.32-754.6.3 in CentOS / RHEL and Oracle Linux (RHCK) is not supported.

- Only GA versions of the [supported distributions](#) that have been released before the current version of Veeam Agent for Linux are supported.

If a new version of a supported Linux distribution is released after the release of the current version of Veeam Agent, Veeam Agent may require a patch to support this new OS version. For details on Veeam Agent compatibility with Linux OS versions, see [this Veeam KB article](#). Customers with a valid contract can request a patch from Veeam Support; for other customers, the support of the new Linux distribution will be provided with the next release of Veeam Agent.

- [For RHEL 6 and 7] Veeam Agent supports the following versions of RHEL with Extended Life-cycle Support Add-On: 6.10 and 7.9.
- [For RHEL 8 and 9] Veeam Agent supports the following versions of RHEL with Extended Update Support Add-On: 8.4, 8.6, 8.8, 8.10, 9.0, 9.2, 9.4 and 9.6.
- For the following distributions, we recommend installing Veeam kernel modules from pre-built binary packages provided by Veeam:
 - CentOS 7
 - RHEL 6.4 – 9.6 and 10.0
 - Rocky Linux 8.10, 9.3 – 9.6 and 10.0
 - AlmaLinux 8.10, 9.3 – 9.6 and 10.0
 - SLES 12 SP4 and SP5, 15 SP1 – 15 SP6
 - SLES for SAP 12 SP4 and 12 SP5, 15 SP1 – 15 SP6
 - openSUSE Leap 15.3 – 15.6

For other supported distributions, use the `dkms` packages instead of the pre-built binary packages with Veeam kernel modules.

Consider the following about Veeam kernel modules from pre-built binary packages:

- Pre-built binary `veeamsnap` kernel module packages require kernel 2.6.32-131.0.15 or later for RHEL 6 (excluding 2.6.32-279.el6.i686) and 3.10.0-123 or later for CentOS / RHEL 7.0 – 7.9.
- Pre-built binary `blksnap` kernel module packages require kernel 5.3.18 or later.

For details on installing Veeam Agent on every supported distribution, see [Installing Veeam Agent for Linux](#).

- To ensure proper functioning of the Veeam kernel module, verify that your system does not have any of the following modules installed: `hcdpdriver`, `snapapi26`, `snapapi`, `snapper`, `dattobd`, `dattobd-dkms`, `dkms-dattobd`, `cdr` or `cxbf`.
- The Linux OS must be set up to receive software updates from the default repositories enabled in the OS after installation.

- For cloud-based installations that use customized kernels (such as Linux distributions deployed from AWS Marketplace or Azure Marketplace), the Veeam kernel module has experimental support status. For details about experimental support, see [this Veeam KB article](#).
- RHEL, CentOS and Oracle Linux (RHCK) are supported up to certain kernel versions. For details, see [this Veeam KB article](#).
- Ubuntu with Linux kernel for KVM (Kernel-based Virtual Machine) is not supported. For the list of linux-kvm kernels for Ubuntu, see [Ubuntu documentation](#).
- AlmaLinux OS Kitten 10 is not supported.
- [Oracle Linux (UEK) 6.6 – 7.4] If the operating system has the FIPS mode enabled, you must sign the DKMS Veeam kernel module. For more information on automating the process of signing DKMS kernel modules, see [Linux documentation](#).

File System

- File-level backup has the following limitations:
 - Total size of all file systems must not exceed 216 TiB. This limitation applies to all file systems where files you plan to back up are located.
 - Size of a file included in a file-level backup must not exceed 16 TB.
 - Name of a file must not be larger than 254 bytes.

Keep in mind that characters that you can use in the file name may be encoded in 2 bytes or more.
- To store volume snapshots, the `blksnap` kernel module requires an Ext4, BTRFS or XFS file system.
- Veeam Agent supports backup of extended attributes with the following limitations:
 - Veeam Agent backs up extended attributes only with the following public namespaces: `system`, `security`, `trusted`, and `user`.
 - All extended attribute names and values of a file must not exceed 4096 bytes (size of a default ext4 file system block). Veeam Agent does not back up attributes exceeding the limit.

For the kernel version 4.13 or later, if a value of extended attribute exceeds the limit, Veeam Agent uses the `ea_inodes` feature. Backups created using the `ea_inodes` feature cannot be mounted on kernel versions up to 4.12.
- Each volume included in a backup must have a unique UUID.
- Veeam kernel module provides a RAM-based changed block tracking (CBT) mechanism. Every time the module is unloaded or Veeam Agent for Linux computer is rebooted, CBT data is reset. As a result, Veeam Agent reads the entire data added to the backup scope to detect what blocks have changed since the last job session, and incremental backup requires greater time.
- Consider the following about backup of computers that are used as cluster nodes:
 - To back up data on local disks, you can use [file-level](#) or [volume-level](#) backup.
 - Data on shared disks, clustered file systems or clustered LVM can be backed up only in the [snapshot-less file-level backup](#) mode.
 - Backup of clustered file systems using a native file system snapshot, or a snapshot created with the help of custom pre-job or post-job scripts, is not supported.

To protect data on computers that function as cluster nodes, you can also create backups using `nosnap` Veeam Agent for Linux. For more information, see [System Requirements for Nosnap Veeam Agent for Linux](#).

- Certain limitations for Dell PowerPath configuration apply. To learn more, see [this Veeam KB article](#).
- Backup of file and directory attributes (for example, a – append only, c – compressed, and so on) is not supported.
- Veeam Agent for Linux does not back up volumes that reside on USB devices and SD cards.
- Veeam Agent for Linux does not back up LVM snapshots.
- BFQ I/O scheduler is not supported.
- Sparse files are not supported. Veeam Agent backs up and restores sparse files as regular files.
- Backup of pseudo file systems, such as `/proc`, `/sys`, `tmpfs`, `devfs` and others, is not supported.
- During backup, network file systems are skipped unless explicitly included into the backup scope.
- Backup of BTRFS volumes and subvolumes with enabled file-system compression is not supported.

Software

IMPORTANT

Linux user account used to work with Veeam Agent for Linux must have `/bin/bash` set as the default shell.

- The following packages are not required for CentOS, RHEL and SLES distributions if a pre-built binary package with Veeam kernel module is to be installed.
 - `dkms`
 - `gcc`
 - `make`
 - `perl`
 - `kernel-headers` (for RedHat-based systems)
 - `kernel-devel` (for RedHat-based systems)

For details, see [Installing Veeam Agent for Linux](#).

- Version of the following packages varies according to the Linux kernel version that you use:
 - `linux-headers` (for Debian-based systems)
 - `kernel-headers` (for RedHat-based systems)
 - `kernel-devel` (for RedHat-based systems)
 - `kernel-uek-devel` (for Oracle Linux systems with UEK)
- For openSUSE and SLES distributions, either of the following packages is required: `libncurses5` or `libncurses6`.
- The `dmidecode` package is required for Veeam Agent management – a valid BIOS UUID must be obtainable either from `dmidecode | grep -i uuid` or from `/sys/class/dmi/id/product_uuid`. Each Veeam Agent that consumes a license installed in Veeam Backup & Replication must have a unique BIOS UUID. If a valid UUID cannot be obtained, Veeam Agent will generate it automatically.
- The `libmysqlclient` package is required to process MySQL database system located on the Veeam Agent server. For details, see [Backup of MySQL Database](#). Package version varies according to the MySQL database system version that you use.

- The `libpq5` package is required to process PostgreSQL database system located on the Veeam Agent server. For details, see [Backup of PostgreSQL Database](#).
- The `python3` package or another RPM package providing a `/usr/bin/python3` binary is required for CentOS, RHEL 7.0 and later distributions if a pre-built binary Veeam kernel module package is to be installed.
- The `btrfs-progs` package version 3.16 or later is required.

Backup Source

Any file systems and devices that are accessible from the host OS. To learn about limitations, see [File System](#).

Backup Target

Backup can be performed to the following types of storage:

- On-premises or cloud-based object storage.
- Local (internal) storage of the protected computer (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share. Requires `cifs-utils` or `nfs-utils` packages to be installed on the Veeam Agent for Linux computer, depending on a network storage type.
- Veeam Backup & Replication 12.3 backup repository (including deduplication appliances).

NOTE

To benefit from all recent enhancements and improvements, make sure to install the latest available build of Veeam Backup & Replication.

- Veeam Cloud Connect 12.0 or later backup repository.

IMPORTANT

Consider the following about backup repositories:

- [For local storage] A backup repository should be created on a separate volume from the volume that contains data you plan to back up.
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported.

Network

Consider the following:

- If you back up to a repository managed by a Veeam backup server, Veeam Agent for Linux must be able to establish a direct IP connection to the Veeam Backup & Replication server. Veeam Agent for Linux cannot work with Veeam Backup & Replication that is located behind the NAT gateway.
- Domain names of the Veeam Agent computer, Veeam Backup & Replication server and other servers in the Veeam backup infrastructure must be resolvable into IPv4 or IPv6 addresses.

System Requirements for Nosnap Veeam Agent for Linux

You can install Veeam Agent for Linux using a `veeam-nosnap` package. This package allows Veeam Agent to operate without Veeam kernel module.

The `veeam-nosnap` package can be useful in the following cases:

- You do not want to install kernel sources and compilers on your computer.
- You want to use third-party tools to create data snapshots.
- You want to perform bare metal restore, but Veeam Recovery Media does not work with your computer. In this case you can install the `veeam-nosnap` package on LiveCD of your choice and access the Veeam recovery UI.
- You want to back up machines that are used as cluster nodes.

Before you install Veeam Agent using the `veeam-nosnap` package, consider the following limitations:

- The RAM-based changed block tracking (CBT) mechanism is not supported. As a result, if you plan to back up a significant amount of data, the backup will require greater time.
- Veeam Agent can create a snapshot of LVM logical volumes and BTRFS subvolumes. To back up data that resides on other file systems and volumes, you can use only file-level backup in the snapshot-less mode. For details, see [Snapshot-Less File-Level Backup](#).
- For a successful backup, Veeam Agent requires unallocated extents on volume groups.
- For a successful bare metal restore, all disks of the Veeam Agent computer you want to restore must be available in the backup.

System Requirements for nosnap Veeam Agent for Linux

If you plan to use the `veeam-nosnap` package to install Veeam Agent, the protected Linux computer meet the requirements listed in the table below. To learn about system requirements for Veeam Agent installed using a Veeam kernel module package, see [System Requirements](#).

Specification	Requirement
Hardware	CPU: x86 or x64. Memory: 1 GB RAM or more. Memory consumption varies depending on the backup type and the total amount of backed-up data. For more information, see RAM Requirements for Backup of Large Number of Files. Disk Space: 100 MB free disk space for product installation. Network: 10 Mbps or faster network connection to a backup target. System firmware: BIOS or UEFI. Disk layout: MBR or GPT. For virtual machines: Only full virtualization type is supported. Containers and paravirtualized instances are not supported. Oracle VM virtual machines are supported with limitations.

Specification	Requirement
OS	Important! Check considerations and limitations that apply to the list of supported OSes. Veeam Agent supports the 64-bit versions of the following distributions: • Debian 10.13 – 12.11 • Ubuntu 16.04, 18.04, 20.04, 22.04, 22.10, 23.04, 23.10, 24.04, 24.10 and 25.04 • RHEL 6.4 – 9.6 and 10.0 • CentOS 7 • Oracle Linux 6 – 9.6 and 10.0 (RHCK) • Oracle Linux 6 (starting from UEK R2) – Oracle Linux 9 (up to 6.12.0-1.23.3.2.el9uek.x86_64) • SLES 12 SP4 and SP5, 15 SP1 – 15 SP6 • SLES for SAP 12 SP4 and SP5, 15 SP1 – 15 SP6 • openSUSE Leap 15.3 – 15.6 • openSUSE Tumbleweed has an experimental support status. For details about experimental support, see this Veeam KB article. • Rocky Linux 8.10, 9.3 – 9.6 and 10.0 • AlmaLinux 8.10, 9.3 – 9.6 and 10.0 Veeam Agent supports the 32-bit versions of the following distributions: • RHEL 6 • Oracle Linux 6 32-bit versions of other distributions are not supported.
File System	Important! Check considerations and limitations that apply to the list of the supported file systems. Veeam Agent for Linux supports consistent snapshot-based data backup for the following file systems: • All supported file systems that are built on top of LVM logical volumes. • BTRFS (for OSes that run Linux kernel 3.16 or later). BTRFS is supported only if it resides directly on a physical device with no additional abstraction layers (such as LVM, software RAID, dm-crypt and so on) below or above it. Supported file systems that are not located on logical volumes, other file systems and network file systems like NFS or SMB shares can be backed up using the snapshot-less mode only. For details, see Snapshot-Less File-Level Backup.

Specification	Requirement
Software	Important! Check considerations and limitations that apply to the list of supported components. Protected computer must have the following components installed: • libacl • libattr • lvm2 • libfuse • dmidecode • efibootmgr (for UEFI-based systems) • isolinux (for Debian-based systems) • syslinux (for RedHat-based systems) • btrfs-progs (for backup of BTRFS file system) • mksquashfs (for custom Veeam Recovery Media) • unsquashfs (for custom Veeam Recovery Media) • wget (for custom Veeam Recovery Media) • xorriso (for custom Veeam Recovery Media with EFI support) • tar (for file system indexing, log export and rotation) • gzip (for file system indexing, log export and rotation)

Considerations and Limitations

OS

- Only GA versions of the [supported distributions](#) that have been released before the current version of Veeam Agent for Linux are supported.

If a new version of a supported Linux distribution is released after the release of the current version of Veeam Agent, Veeam Agent may require a patch to support this new OS version. Customers with a valid contract can request a patch from Veeam Support; for other customers, the support of the new Linux distribution will be provided with the next release of Veeam Agent.

- The Linux OS must be set up to receive software updates from the default repositories enabled in the OS after installation.
- AlmaLinux OS Kitten 10 is not supported.
- [For RHEL 6 and 7] Veeam Agent supports the following versions of RHEL with Extended Life-cycle Support Add-On: 6.10 and 7.9.
- [For RHEL 8 and 9] Veeam Agent supports the following versions of RHEL with Extended Update Support Add-On: 8.4, 8.6, 8.8, 8.10, 9.0, 9.2, 9.4 and 9.6.

File System

- Veeam Agent for Linux does not back up volumes that reside on USB devices and SD cards.
- LVM volumes encrypted with dm-crypt software are not supported.
- Total size of all file systems must not exceed 216 TiB. This limitation applies to all file systems where files you plan to back up are located.

- Size of a file included in a file-level backup must not exceed 16 TiB.
- Name of a file must not be larger than 254 bytes.
Keep in mind that characters that you can use in the file name may be encoded in 2 bytes or more.
- The amount of space required for LVM snapshots largely depends on the IO intensity. Generally, from 10% to 20% of the system's occupied space should be enough for storing an LVM snapshot.
- Veeam Agent supports backup of extended attributes with the following limitations:
 - Veeam Agent backs up extended attributes only with the following public namespaces: `system`, `security`, `trusted`, and `user`.
 - All extended attribute names and values of a file must not exceed 4096 bytes (size of a default ext4 file system block). Veeam Agent does not back up attributes exceeding the limit.

For the kernel version 4.13 or later, if a value of extended attribute exceeds the limit, Veeam Agent uses the `ea_inodes` feature. Backups created using the `ea_inodes` feature cannot be mounted on kernel versions up to 4.12.
- Backup of file and directory attributes (for example, `a` – append only, `c` – compressed, and so on) is not supported.
- Each volume included in a backup must have a unique UUID.
- Consider the following about the backup of computers used as cluster nodes:
 - To back up data on local LVM volumes or BTRFS subvolumes, you can use [file-level](#) or [volume-level](#) backup.

NOTE

Consider the following about backup of LVM volumes:

- During volume-level backup, data from shared disks, clustered file systems or clustered LVM will not be backed up.
- To perform volume-level backup, Veeam Agent for Linux will create an LVM snapshot, which may cause instability of the cluster or cluster software. This can happen due to the failover conditions configured for the cluster. However, if the cluster instability is caused by creation of an LVM snapshot only during backup, please contact Veeam support for assistance.
- Backup of clustered file systems using a native file system snapshot is not supported. This includes snapshots created with the help of custom pre-job or post-job scripts.
- The following objects can be backed up only by [snapshot-less file-level backup](#):
 - Files on shared disks, clustered file systems or clustered LVM.
 - Files on local file systems that are not hosted by LVM or BTRFS.
- Certain limitations for Dell PowerPath configuration apply. To learn more, see [this Veeam KB article](#).
- Sparse files are not supported. Veeam Agent backs up and restores sparse files as regular files.
- Backup of pseudo file systems, such as `/proc`, `/sys`, `tmpfs`, `devfs` and others, is not supported.
- Backup of BTRFS volumes and subvolumes with enabled file-system compression is not supported.

Software

IMPORTANT

Linux user account used to work with Veeam Agent for Linux must have the `/bin/bash` shell set as the default shell.

- The `dmidecode` package is required for Veeam Agent management — a valid BIOS UUID must be obtainable either from `dmidecode | grep -i uuid` or from `/sys/class/dmi/id/product_uuid`. Each Veeam Agent that consumes a license installed in Veeam Backup & Replication must have a unique BIOS UUID. If a valid UUID cannot be obtained, Veeam Agent will generate it automatically.
- The `btrfs-progs` package version 3.16 or later is required.

Backup Source

Any file systems and devices that are accessible from the host OS. To learn about limitations, see [File System](#).

Backup Target

Backup can be performed to the following types of storage:

- On-premises or cloud-based object storage.
- Local (internal) storage of the protected computer (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share. Requires `cifs-utils` or `nfs-utils` packages to be installed on the Veeam Agent for Linux computer, depending on a network storage type.
- Veeam Backup & Replication 12.3 repository (including deduplication appliances).

NOTE

To benefit from all recent enhancements and improvements, make sure to install the latest available build of Veeam Backup & Replication.

- Veeam Cloud Connect 12.0 or later cloud repository.

IMPORTANT

Consider the following about backup repositories:

- [For local storage] A backup repository should be created on a separate volume from the volume that contains data you plan to back up.
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported.

Network

Consider the following:

- Veeam Agent for Linux should be able to establish a direct IP connection to the Veeam Backup & Replication server. Thus, Veeam Agent for Linux cannot work with Veeam Backup & Replication that is located behind the NAT gateway.
- Domain names of the Veeam Agent computer, Veeam Backup & Replication server and other servers in the Veeam backup infrastructure must be resolvable into IPv4 or IPv6 addresses.

System Requirements for Nosnap Veeam Agent for Linux on Power

If you want to back up Linux computers running on IBM Power Systems, you can install Veeam Agent for Linux using the `veeam-nosnap` package for Linux on Power. This package allows Veeam Agent to operate without Veeam kernel module using the native file system snapshots instead.

Before you install Veeam Agent using the `veeam-nosnap` package for Linux on Power, consider the following limitations:

- The RAM-based changed block tracking (CBT) mechanism is not supported. As a result, if you plan to back up a significant amount of data, the backup will require greater time.
- Veeam Agent can create a snapshot of LVM logical volumes and BTRFS subvolumes. To back up data that resides on other file systems and volumes, you can use only file-level backup in the snapshot-less mode. For details, see [Snapshot-Less File-Level Backup](#).
- For a successful backup, Veeam Agent requires unallocated extents on volume groups.
- For a successful bare metal restore, all disks of the Veeam Agent computer you want to restore must be available in the backup.

System Requirements for Veeam Agent for Linux on Power

If you plan to install Veeam Agent for Linux on Power, make sure the protected Linux computer meets the requirements listed in the table below. To learn about system requirements for Veeam Agent for Linux installed using Veeam kernel module package, see [System Requirements](#).

Specification	Requirement
Hardware	System: IBM Power System CPU: IBM POWER9 or POWER10 Memory: 1 GB RAM or more. Memory consumption varies depending on the backup type and the total amount of backed-up data. For more information, see RAM Requirements for Backup of Large Number of Files. Disk Space: 100 MB free disk space for product installation Network: 10 Mbps or faster network connection to a backup target Disk layout: MBR or GPT

Specification	Requirement
OS	Important! Check considerations and limitations that apply to the list of supported OSes. Veeam Agent supports little endian versions of the following Linux distributions for IBM Power: • SLES 15 SP3 – SP6 • SLES for SAP 12 SP5, 15 SP3 – SP6 • RHEL 8.4, 8.6, 8.8, 8.10, 9.0, 9.2, 9.4, 9.6 (including Extended Update Support Add-On) and 10.0 • RHEL for SAP 8.4, 8.6, 8.8, 8.10, 9.0, 9.2, 9.4, 9.6 and 10.0
File System	Important! Check considerations and limitations that apply to the list of supported file systems. Veeam Agent for Linux on Power supports consistent snapshot-based data backup for the following file systems: • All supported file systems that are built on top of LVM logical volumes. • BTRFS (for OSes that run Linux kernel 3.16 or later) If BTRFS has additional abstraction layers (such as LVM, software RAID, dm-crypt and so on) above it, only file-level restore operations are supported. Instant Recovery, restore verification (SureBackup), bare metal recovery and volume-level restore are not supported. Supported file systems that are not located on logical volumes, other file systems and network file systems like NFS or SMB shares can be backed up using the snapshot-less mode only. For details, see Snapshot-Less File-Level Backup.
Software	Important! Linux user account used to work with Veeam Agent for Linux on Power must have the <code>/bin/bash</code> shell set as the default shell. Protected computer must have the following components installed: • libacl • libattr • lvm2 • libfuse2 (FUSE libraries for SLES-based systems) • fuse-libs (FUSE libraries for RedHat-based systems) • syslinux (for RedHat-based systems) • btrfs-progs (version 3.16 or later, for backup of BTRFS file system) • tar (for file system indexing, log export and rotation) • gzip (for file system indexing, log export and rotation)

Considerations and Limitations

OS

- Only GA versions of the [supported distributions](#) that have been released before the current version of Veeam Agent for Linux for Power are supported.

If a new version of a supported Linux distribution is released after the release of the current version of Veeam Agent, Veeam Agent may require a patch to support this new OS version. Customers with a valid contract can request a patch from Veeam Support; for other customers, the support of the new Linux distribution will be provided with the next release of Veeam Agent.

- The Linux OS must be set up to receive software updates from the default repositories enabled in the OS after installation.

File System

- Veeam Agent does not back up volumes that reside on USB devices and SD cards.
- LVM volumes encrypted with dm-crypt software are not supported.
- Total size of all file systems must not exceed 216 TiB. This limitation applies to all file systems where files you plan to back up are located.
- Size of a file included in a file-level backup must not exceed 16 TiB.
- Name of a file must not be larger than 254 bytes.

Keep in mind that characters that you can use in the file name may be encoded in 2 bytes or more.

- The amount of space required for LVM snapshots largely depends on the IO intensity. Generally, from 10% to 20% of the system's occupied space should be enough for storing an LVM snapshot.
- Veeam Agent supports backup of extended attributes with the following limitations:
 - Veeam Agent backs up extended attributes only with the following public namespaces: `system`, `security`, `trusted`, and `user`.
 - All extended attribute names and values of a file must not exceed 4096 bytes (size of a default ext4 file system block). Veeam Agent does not back up attributes exceeding the limit.

For the kernel version 4.13 or later, if a value of extended attribute exceeds the limit, Veeam Agent uses the `ea_inodes` feature. Backups created using the `ea_inodes` feature cannot be mounted on kernel versions up to 4.12.

- Backup of file and directory attributes (for example, `a` – append only, `c` – compressed, and so on) is not supported.
- Each volume included in a backup must have a unique UUID.
- Consider the following about the backup of machines used as cluster nodes:
 - To back up data on local LVM volumes, you can use [file-level backup](#) or [volume-level backup](#).

NOTE

Consider the following:

- During volume-level backup, data from shared disks, clustered file systems or clustered LVM will not be backed up.
 - To perform volume-level backup, Veeam Agent for Linux will create an LVM snapshot, which can cause instability of the cluster or cluster software. This can happen due to the failover conditions configured for the cluster. However, if the cluster instability is caused by creation of an LVM snapshot only during backup, please contact Veeam support for assistance.
- Backup of clustered file systems using a native file system snapshot is not supported. This includes snapshots created with the help of custom pre-job or post-job scripts.
 - The following objects can be backed up only by [snapshot-less file-level backup](#):
 - Files on shared disks, clustered file systems or clustered LVM
 - Files on local file systems that are not hosted by LVM
 - Certain limitations for Dell PowerPath configuration apply. To learn more, see [this Veeam KB article](#).
 - Sparse files are not supported. Veeam Agent backs up and restores sparse files as regular files.
 - Backup of pseudo file systems, such as `/proc`, `/sys`, `tmpfs`, `devfs` and others, is not supported.
 - Backup of BTRFS volumes and subvolumes with enabled file-system compression is not supported.

Backup Source

Any file systems and devices that are accessible from the host OS. To learn about limitations, see [File System](#).

Backup Target

Backup can be performed to the following types of storage:

- On-premises or cloud-based object storage.
- Local (internal) storage of the protected computer (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share. Requires `cifs-utils` or `nfs-utils` packages to be installed on the Veeam Agent for Linux computer, depending on a network storage type.
- Veeam Backup & Replication 12.3 backup repository (including deduplication appliances).

NOTE

To benefit from all recent enhancements and improvements, make sure to install the latest available build of Veeam Backup & Replication.

IMPORTANT

Consider the following about backup repositories:

- [For local storage] A backup repository should be created on a separate volume from the volume that contains data you plan to back up.
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported.

Network

Consider the following:

- Veeam Agent for Linux should be able to establish a direct IP connection to the Veeam Backup & Replication server. Thus, Veeam Agent for Linux cannot work with Veeam Backup & Replication that is located behind the NAT gateway.
- Domain names of the Veeam Agent computer, Veeam Backup & Replication server and other servers in the Veeam backup infrastructure must be resolvable into IPv4 or IPv6 addresses.

RAM Requirements for Backup of Large Number of Files

Amount of RAM used by Veeam Agent to process backed-up files depends on the number of files included in the backup, the length of file names and depth of the directory structure. For large environments with great number of backed-up files, consider the following RAM sizing recommendations:

For large number of backed-up files with short names (7 characters) under a single-level root directory

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
1,000,000 Full	1.6
1,000,000 Incremental + 100,000	2.1
2,000,000 Full	2.0
2,000,000 Incremental + 100,000	3.5
5,000,000 Full	4.4
5,000,000 Incremental + 100,000	7.2
10,000,000 Full	7.3
10,000,000 Incremental + 100,000	13.6

For large number of backed-up files with long names (254 characters) under a single-level root directory

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
1,000,000 Full	2.8
1,000,000 Incremental + 100,000	4.6
2,000,000 Full	4.6
2,000,000 Incremental + 100,000	8.2
5,000,000 Full	10.0
5,000,000 Incremental + 100,000	19.2

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
10,000,000 Full	19.0
10,000,000 Incremental + 100,000	36.7

For large number of backed-up files with short names (7 characters) in a 7-level directory structure

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
1,000,000 Full	2.3
1,000,000 Incremental + 100,000	2.7
2,000,000 Full	2.7
2,000,000 Incremental + 100,000	3.4
5,000,000 Full	3.9
5,000,000 Incremental + 100,000	5.3
10,000,000 Full	5.8
10,000,000 Incremental + 100,000	8.5

For large number of backed-up files with long names (254 characters) in a 7-level directory structure

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
1,000,000 Full	2.8
1,000,000 Incremental + 100,000	3.5
2,000,000 Full	3.7
2,000,000 Incremental + 100,000	4.9
5,000,000 Full	6.4
5,000,000 Incremental + 100,000	9.3

Number of Backed-Up Files, Full / Incremental Backup	RAM (GB)
10,000,000 Full	11.1
10,000,000 Incremental + 100,000	16.2

Permissions

Depending on the scenario, the user accounts must have the permissions listed in the following subsections:

- [Permissions for Backup to Object Storage](#)
- [Permissions for Guest Processing](#)

Permissions for Backup to Object Storage

If you plan to back up data to object storage, make sure that the user account that you use to connect to the object storage has the required permissions. The list of required permissions differs depending on the selected object storage:

- [Amazon S3 or S3 compatible](#)
- [Google Cloud Storage](#)

Amazon S3 or S3 compatible

If you plan to back up data to the Amazon S3 or S3 compatible storage, make sure the user account that you plan to use has the following permissions:

Identity-based permission:

```
{
  "s3:ListAllMyBuckets"
}
```

Resource-based permissions:

```
{
  "s3:DeleteObject",
  "s3:GetBucketLocation",
  "s3:GetBucketObjectLockConfiguration",
  "s3:GetBucketVersioning",
  "s3:GetObject",
  "s3:ListBucket",
  "s3:PutObject"
}
```

TIP

For information about required permissions for Amazon S3 storage with immutability enabled, see the [Using Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.

Google Cloud Storage

If you plan to back up data to the Google Cloud storage, make sure the user account that you plan to use has the following permissions:

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

Permissions for Guest Processing

To use guest processing, make sure to configure user accounts according to the requirements listed in this section.

Consider the following general requirements when choosing a user account:

- The user account must have root privileges.
- The user account must have the home directory created.

Depending on the application you need to back up, the user account must have the permissions listed in the table below:

Application	Required Permission
MySQL	To process the MySQL database system, the MySQL user account must have the following privileges: ○ SELECT for all tables. This privilege is required to allow Veeam Agent to access table metadata. To learn more, see MySQL documentation. ○ LOCK TABLES. This privilege is required to allow Veeam Agent to process tables based on the MyISAM storage engine. ○ RELOAD. This privilege is required to allow the MySQL user account to perform FLUSH operations.
Oracle	To back up Oracle data, the user account must be granted <i>SYSDBA</i> privileges. You can use either the same account that was specified at the Guest Processing step if such an account is a member of the <i>OSDBA</i> and <i>OINSTALL</i> groups, or you can use any other account that has <i>SYSDBA</i> privileges. To perform guest processing for Oracle databases on Linux servers, make sure that the <code>/tmp</code> directory is mounted with the <code>exec</code> option. Otherwise, you will get a "<i>Permission denied</i>" error.
PostgreSQL	To back up PostgreSQL instances, the user account must have the superuser privileges for the PostgreSQL instance. For more information, see PostgreSQL documentation.

Ports

The following tables describe network ports that must be opened to ensure proper communication of Veeam Agent operating in the standalone mode with other infrastructure components.

To learn about ports required to enable proper work of Veeam Agent for Linux managed by Veeam Backup & Replication, see the [Ports](#) section in the Veeam Agent Management Guide.

IMPORTANT

The list of ports required for computers booted from the Veeam Recovery Media is the same as the list of ports required for Veeam Agent computers.

Communication Between Veeam Agent Components

The following table describes network ports that must be opened to enable proper communication between Veeam Agent for Linux components.

From	To	Protocol	Port	Notes
Veeam Agent computer	Veeam backup server	TCP	10002, 10006	Default ports used for communication with the Veeam backup server. Data between the Veeam Agent for Linux computer and backup repositories is transferred directly, bypassing Veeam backup servers.
	Veeam Agent computer	TCP	2500 to 3300	Default range of ports used for communication between Veeam Agent for Linux components during data transmission. For every TCP connection that a backup job uses, one port from this range is assigned. Ports must be open for incoming and outgoing traffic. Established connections must be allowed.
		TCP	10808	Port used locally on the Veeam Agent computer for communication via REST API between Veeam Agent components (such as control panel and command line interface) and Veeam Agent for Linux Service.

Communication with Veeam Backup & Replication Repositories

The following table describes network ports that must be opened to ensure proper communication with Veeam backup repositories.

From	To	Protocol	Port	Notes
Veeam Agent computer	Linux or Microsoft Windows server acting as a backup repository	TCP	2500 to 3300	Default range of ports used as data transmission channels. For every TCP connection that a backup job uses, one port from this range is assigned.
	Shared folder SMB (CIFS) share	TCP UDP	137 to 139, 445	Ports used as a data transmission channel from the Veeam Agent for Linux computer to the target SMB (CIFS) share. Ports 137 to 139 are used by backup infrastructure components to communicate using NetBIOS.
	Shared folder NFS share	TCP UDP	111, 2049	Standard NFS ports used as a data transmission channel from the Veeam Agent for Linux computer to the target NFS share.

Communication with Veeam Cloud Connect Repositories

The following table describes network ports that must be opened to ensure proper communication with Veeam Cloud Connect repositories.

From	To	Protocol	Port	Notes
Veeam Agent computer	Cloud gateway	TCP	6180	Port on the cloud gateway used to transport Veeam Agent data to the Veeam Cloud Connect repository.
	Certificate Revocation Lists	TCP	80 or 443 (most popular)	Veeam Agent computer needs access to CRLs (Certificate Revocation Lists) of the CA (Certification Authority) who issued a certificate to the Veeam Cloud Connect service provider. Generally, information about CRL locations can be found on the CA website.

Communication with Object Storage

The following table describes network ports that must be opened to ensure proper communication with object storage if you back up data to object storage directly or to object storage added as a Veeam backup repository with the direct connection mode. For more information about object storage connection modes, see [Types of Connection to Object Storage in Veeam Backup & Replication](#).

From	To	Protocol	Port	Notes
Veeam Agent computer	Amazon S3 object storage	TCP	443	Used to communicate with the Amazon S3 object storage through the following endpoints: • *.amazonaws.com (for both <i>Global</i> and <i>Government</i> regions) • *.amazonaws.com.cn (for <i>China</i> region) All AWS service endpoints are specified in the AWS documentation .
			80	Used to verify the certificate status through the following endpoints: • *.amazontrust.com • *.cloudfront.net Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.
	Microsoft Azure object storage	TCP	443	Used to communicate with the Microsoft Azure object storage through the following endpoints: • xxx.blob.core.windows.net (for <i>Global</i> region) • xxx.blob.core.chinacloudapi.cn (for <i>China</i> region) • xxx.blob.core.usgovcloudapi.net (for <i>Government</i> region) Consider that the <xxx> 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: • <code>ocsp.digicert.com</code> • <code>ocsp.msocsp.com</code> Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself. For more details, see also Microsoft documentation.
	Google Cloud storage	TCP	443	Used to communicate with Google Cloud storage through the following endpoints: • <code>storage.googleapis.com</code> All cloud endpoints are specified in this Google article.
			80	Used to verify the certificate status through the following endpoints: • <code>ocsp.pki.goog</code> • <code>pki.goog</code> • <code>crl.pki.goog</code> Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.
	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.

Live Patching Support

The live patching technology allows you to patch a running Linux kernel without the need to reboot or change kernel-related files on the disk.

If you plan to use live patching, consider the following limitations:

- Only `kGraft`, `kpatch` and `Ksplice` kernel extensions are supported.
- Live patching is supported only for the following Linux distributions:
 - SLES 12 SP4, 12 SP5, 15 SP1 – 15 SP6
 - SLES for SAP 12 SP4, 12 SP5, 15 SP1 – 15 SP6
 - RHEL 6.4 – 9.6 and 10.0
 - Oracle Linux 6.4 – 9.6 (RHCK)
 - Oracle Linux 6 (starting from UEK R2) – Oracle Linux 9 (up to 6.12.0-1.23.3.2.el9uek.x86_64)

Before live patching, check the following prerequisites:

- Before you start live patching on a production environment, make sure that a kernel patch does not harm your system in any way using a spare Veeam Agent computer.
- Back up an entire Veeam Agent computer before live patching.
- Make sure that there are no backup jobs running on the Veeam Agent computer during live patching.

Installation and Configuration

You can install Veeam Agent for Linux on any Linux-based endpoint whose data you plan to protect – virtual machine or physical device (server, desktop or laptop).

Before You Begin

Before you start the installation process, review the following information and prerequisites.

Types of Veeam Agent for Linux Installation Packages

You can install Veeam Agent using one of the following sets of installation packages:

- Veeam Agent for Linux — this set of packages depends on the Veeam kernel module for creating system snapshots. It works with the widest range of Linux distributions and file systems. For more information on system requirements and limitations, see [System Requirements](#).
- Nosnap Veeam Agent for Linux — this set of packages does not depend on the Veeam kernel module for creating system snapshots. Nosnap Veeam Agent for Linux leverages native file system snapshot capabilities on select Linux distributions. For more information on system requirements and limitations, see [System Requirements for Nosnap Veeam Agent for Linux](#).
- Nosnap Veeam Agent for Linux on Power — this set of nosnap packages is specifically designed for IBM Power Systems. For more information on system requirements and limitations, see [System Requirements for Nosnap Veeam Agent for Linux on Power](#).

For information on installation, see [Installing Veeam Agent for Linux](#) and [Installing Veeam Agent for Linux in Offline Mode](#).

General Prerequisites

Before you start the installation process, consider the following:

- The computer on which you plan to install Veeam Agent must satisfy system requirements specified in this document. To learn more, see [System Requirements](#).
- To install Veeam Agent software packages, you must use the `root` account or any user account that has super user (root) privileges on the computer where you plan to install the product.
- If you install Veeam Agent in a UEFI system with Secure Boot, you must configure UEFI Secure Boot to enable your system to work with Veeam Agent. You can configure UEFI Secure Boot before or after the installation of Veeam Agent, but before you run a backup job. For more information, see [Configuring UEFI Secure Boot](#).
- You must not install Veeam Agent on the server that is used as a [hardened repository](#) in the Veeam Backup & Replication infrastructure.
- If you have used the Beta version of Veeam Agent, you must remove Veeam Agent software packages prior to installing the release version of the product. To learn more, see [Uninstalling Veeam Agent for Linux](#).

Installing Veeam Agent for Linux

To install Veeam Agent for Linux on a computer with a connection to the internet, you must perform the following steps:

1. [Connect to the Veeam software repository](#).
2. Install Veeam Agent for Linux packages from the Veeam software repository.

Installation instructions depend on the type of the [packages](#) you want to use for Veeam Agent installation:

- [Installing Veeam Agent for Linux \(with Kernel Module\)](#)
- [Installing Nosnap Veeam Agent for Linux](#)
- [Installing Nosnap Veeam Agent for Linux on Power](#).

TIP

If the computer where you want to install Veeam Agent for Linux is not connected to the internet, you can download and install Veeam Agent for Linux packages manually. To learn more, see [Installing Veeam Agent for Linux in Offline Mode](#).

Connecting to Veeam Software Repository

To install Veeam Agent for Linux on a Linux computer, you must first connect the computer to the Veeam software repository. The Veeam software repository contains the Veeam Agent installation packages specific to the Linux distribution, version and architecture of the computer where you plan to install the product.

To connect to the Veeam software repository, do the following:

1. Download the Veeam software repository installation package (`veeam-release`) from the [this Veeam webpage](#), and save the downloaded package on your computer.
2. Navigate to the directory where you have saved the `veeam-release` package and install the package.

NOTE

To ensure authenticity and integrity of the `veeam-release` package, Veeam signs this package with a GPG (GNU Privacy Guard) key. When you install the `veeam-release` package for the first time, your package manager may prompt you to accept or import the Veeam GPG key. When prompted, you must accept or import the GPG key.

To install the `veeam-release` package, use the following commands:

For CentOS 7 / RHEL / Oracle Linux / Fedora / Rocky Linux / AlmaLinux

```
rpm -ivh ./veeam-release* && yum check-update
```

For openSUSE / SLES

```
zypper in ./veeam-release* && zypper refresh
```

For Debian / Ubuntu

```
dpkg -i ./veeam-release* && apt-get update
```

TIP

If the user account you use for Veeam Agent installation does not have root privileges, you can temporarily elevate this user account to root by using the `sudo` prefix in the install commands. If you run multiple commands in one command line, you must use the `sudo` prefix before each command – for example, `sudo rpm -ivh ./veeam-release* && sudo yum check-update`. Make sure the `sudo` user has sufficient privileges to run these commands.

Installing Veeam Agent for Linux with Kernel Module

To install Veeam Agent, you can use a package manager of your choice that works with software packages in your Linux distribution.

NOTE

Some dependency packages of the prerequisite software may require special handling. For details, see [Managing Package Dependencies](#).

To install Veeam Agent for Linux, use the following commands:

For CentOS 7 / RHEL / Fedora / Rocky Linux / AlmaLinux

```
yum install veeam
```

NOTE

[For CentOS 7 / RHEL / Rocky Linux / AlmaLinux] If the `dkms` package is already installed in the OS, you can install Veeam Agent with one of the following commands:

- `yum install veeam`
With this command, the Veeam kernel module will be installed from the source RPM package using `dkms`.
- [For CentOS 7 / RHEL 6 - 8 / Rocky Linux 8.10 / AlmaLinux 8.10] `yum install kmod-veeamsnap veeam` / [For RHEL 9 and 10 / Rocky Linux 9 and 10 / AlmaLinux 9 and 10] `yum install kmod-blksnap veeam`
With this command, the non-DKMS version of the Veeam kernel module will be installed from the pre-built `kmod` binary package.

For Oracle Linux 6 - 8

```
yum install veeamsnap  
yum install veeam
```

NOTE

If your system runs on Oracle Linux 8.x with UEK R7 kernel, you may need to rebuild the Veeam kernel module prior to its installation. For more information, see [this Veeam KB article](#).

For Oracle Linux 9

```
yum install blksnap  
yum install veeam
```

For openSUSE Tumbleweed

```
zypper in veeam
```

For openSUSE Leap 15.3 with default kernel, Leap 15.4 - 15.6

```
zypper in blksnap-kmp-default  
zypper in veeam
```

For openSUSE Leap 15.3 with preemptive kernel

```
zypper in blksnap-kmp-preempt  
zypper in veeam
```

For SLES 12 SP4 - SP5, 15 SP1 - SP2 with default kernel

```
zypper in veeamsnap-kmp-default  
zypper in veeam
```

For SLES 12 SP4 - SP5, 15 SP1 - SP2 with preemptive kernel

```
zypper in veeamsnap-kmp-preempt  
zypper in veeam
```

For SLES 15 SP3 with default kernel, 15 SP4 - SP6

```
zypper in blksnap-kmp-default  
zypper in veeam
```

For SLES 15 SP3 with preemptive kernel

```
zypper in blksnap-kmp-preempt  
zypper in veeam
```

For Debian 10 / Ubuntu 16.04, 18.04 and 20.04 (kernel 5.4)

```
apt-get install veeam
```

For Debian 11 and 12 / Ubuntu 22.04 and later

```
apt-get install blksnap veeam
```

Managing Package Dependencies

The following dependency packages may require special handling in case you see installation errors:

- The `dkms` package is not present in default repositories for some Linux distributions. You should obtain it from third-party repositories:
 - EPEL repository (for CentOS / RHEL / Oracle Linux / Fedora / Rocky Linux / AlmaLinux)
 - Packman repository (for openSUSE). To learn more, see [Installing dkms in openSUSE](#).

For SLES, the `dkms` package is not available in the Packman repository. You must use the package intended for openSUSE. To learn more, see [this Veeam KB article](#).

- Extended kernels, such as `kernel-pae`, `kernel-uek` and other, require appropriate `kernel-devel` packages to be installed, for example, `kernel-pae-devel`, `kernel-uek-devel`, and so on.

Version of the `kernel-devel` package must match your current kernel version. To check your current kernel version, run the `uname -r` command.

[For RHEL and derivatives] If the `yum` package manager installs packages that do not match your current kernel version, you should either update your system or fetch older versions of the required packages from the [CentOS Vault repository](#).

Installing dkms in openSUSE

In openSUSE systems, while installing the `dkms` package, you may see an error similar to the following:

```
Problem: nothing provides kernel-devel needed by dkms-2.2.0.3-14.1.noarch
Solution 1: do not install dkms-2.2.0.3-14.1.noarch
Solution 2: break dkms-2.2.0.3-14.1.noarch by ignoring some of its dependencies
```

To install the `dkms` package, do the following:

1. Make sure that you have an appropriate `kernel-devel` package installed and its version matches your kernel version. For example:

```
root@localhost:~> rpm -qa | grep kernel-default
kernel-default-devel-3.0.101-91.1
kernel-default-3.0.101-91.1
```

2. Install the `dkms` package ignoring dependencies:

```
zypper -n install --force dkms
```

3. Make sure that you have allowed unsupported modules. To learn more, see [SUSE documentation](#).

Installing Nosnap Veeam Agent for Linux

To install Veeam Agent, you can use a package manager of your choice that works with software packages in your Linux distribution. For example, use the following commands:

For CentOS 7 / RHEL / Oracle Linux / Rocky Linux / AlmaLinux

```
yum install veeam-nosnap
```

For openSUSE Tumbleweed / OpenSUSE Leap / SLES

```
zypper in veeam-nosnap
```

For Debian / Ubuntu

```
apt-get install veeam-nosnap
```

Installing Nosnap Veeam Agent for Linux on Power

To install Veeam Agent, you can use a package manager of your choice that works with software packages in your Linux distribution. For example, use the following commands:

For RHEL

```
yum install veeam-nosnap
```

For SLES

```
zypper in veeam-nosnap
```

Installing Veeam Agent for Linux in Offline Mode

If the computer where you want to install Veeam Agent for Linux has no connection to the internet, for example, for security reasons, you can install Veeam Agent in the offline mode. In this scenario, you do not need to download and install the Veeam software repository installation package (`veeam-release`). Instead, you need to download all Veeam Agent packages from the Veeam software repository and install them on the target computer.

Installation instructions depend on the type of the [packages](#) you want to use for Veeam Agent installation:

- [Install Veeam Agent for Linux \(with Kernel Module\) in Offline Mode](#)
- [Install Nosnap Veeam Agent for Linux in Offline Mode](#)
- [Install Nosnap Veeam Agent for Linux on Power in Offline Mode](#)

Installing Veeam Agent for Linux with Kernel Module in Offline Mode

To install Veeam Agent for Linux, do the following:

1. On a computer that is connected to the internet, download Veeam Agent packages intended for your Linux distribution from the [Veeam software repository](#).
 - Veeam Agent for Linux packages in the Debian format reside in the following folders of the Veeam software repository:
 - [/backup/linux/agent/dpkg/debian/public/pool/veeam/b/blksnap-dkms/](#)
 - [/backup/linux/agent/dpkg/debian/public/pool/veeam/v/veeamsnap/](#)
 - [/backup/linux/agent/dpkg/debian/public/pool/veeam/v/veeam-libs/](#)
 - [/backup/linux/agent/dpkg/debian/public/pool/veeam/v/veeam/](#)
 - For RPM packages, the Veeam Agent directory has the following structure: *Package format > Distribution > Version > Architecture*.

For example, Veeam Agent packages for 64-bit RHEL 9 reside in the [/rpm/el/9/x86_64/](#) folder of the Veeam software repository, and packages for 64-bit SLES 15 SP5 reside in the [/rpm/sles/SLE_15_SP5/x86_64/](#) folder.
2. Save Veeam Agent packages to a directory that can be accessed from the computer where you want to install the product, for example, a directory on a local drive or USB drive, or a network shared folder.
3. On the computer where you want to install Veeam Agent, navigate to the directory where you have saved the packages and install Veeam Agent:
 - [Installing Veeam Agent for Linux in CentOS 7 / RHEL / Rocky Linux / AlmaLinux](#)
 - [Installing Veeam Agent for Linux in Oracle Linux](#)
 - [Installing Veeam Agent for Linux in Fedora](#)
 - [Installing Veeam Agent for Linux in SLES](#)
 - [Installing Veeam Agent for Linux in openSUSE](#)
 - [Installing Veeam Agent for Linux in Debian / Ubuntu](#)

TIP

You can also set up a local mirror of the Veeam software repository in your internal network and add this repository to the list of software sources on a computer where you want to install the product. These operations may differ depending on the Linux distribution and package manager that you use. To learn more, refer to the documentation of your Linux distribution.

After you add a local repository to the list of software sources on a computer, you will be able to install and upgrade Veeam Agent in a regular way. To learn more, see [Installing Veeam Agent for Linux](#) and [Upgrading Veeam Agent for Linux](#).

Installing Veeam Agent for Linux in CentOS 7 / RHEL / Rocky Linux / AlmaLinux

To install Veeam Agent for Linux, use the following commands:

For 32-bit RHEL 6

```
rpm -i <...>/kmod-veeamsnap-6.3.2.1207-2.6.32_131.0.15.el6.i386.rpm
rpm -i <...>/veeam-libs-6.3.2.1207-1.i386.rpm
rpm -i <...>/veeam-6.3.2.1207-1.el6.i386.rpm
```

For 64-bit RHEL 6

```
rpm -i <...>/kmod-veeamsnap-6.3.2.1207-2.6.32_131.0.15.el6.x86_64.rpm
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
rpm -i <...>/veeam-6.3.2.1207-1.el6.x86_64.rpm
```

For CentOS 7 / RHEL 7

```
rpm -i <...>/kmod-veeamsnap-6.3.2.1207-1.el7.x86_64.rpm
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
rpm -i <...>/veeam-6.3.2.1207-1.el7.x86_64.rpm
```

RHEL 8 / Rocky Linux 8.10 / AlmaLinux 8.10

```
rpm -i <...>/kmod-veeamsnap-6.3.2.1207-1.el8.x86_64.rpm
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
rpm -i <...>/veeam-6.3.2.1207-1.el8.x86_64.rpm
```

RHEL 9 / Rocky Linux 9.3 - 9.5 / AlmaLinux 9.3 - 9.5

```
rpm -i <...>/kmod-blksnap-6.3.2.1207-1.el9.x86_64.rpm
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
rpm -i <...>/veeam-6.3.2.1207-1.el9.x86_64.rpm
```

RHEL 10 / Rocky Linux 10.0 / AlmaLinux 10

```
rpm -i <...>/kmod-blksnap-6.3.2.1207-1.el10.x86_64.rpm
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
rpm -i <...>/veeam-6.3.2.1207-1.el10.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

NOTE

The pre-built `veeamsnap` binaries require kernel 2.6.32-131.0.15 or later for RHEL 6 (excluding 2.6.32-279.el6.i686) and kernel 3.10.0-123 or later for CentOS / RHEL 7.0 – 7.7 to operate.

Installing Veeam Agent for Linux in Oracle Linux

To install Veeam Agent for Linux, use the following commands:

For 32-bit Oracle Linux 6

```
rpm -i <...>/veeamsnap-6.3.2.1207-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.3.2.1207-1.i386.rpm  
rpm -i <...>/veeam-6.3.2.1207-1.el6.i386.rpm
```

For 64-bit Oracle Linux 6

```
rpm -i <...>/veeamsnap-6.3.2.1207-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-6.3.2.1207-1.el6.x86_64.rpm
```

For Oracle Linux 7

```
rpm -i <...>/veeamsnap-6.3.2.1207-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-6.3.2.1207-1.el7.x86_64.rpm
```

For Oracle Linux 8

```
rpm -i <...>/veeamsnap-6.3.2.1207-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-6.3.2.1207-1.el8.x86_64.rpm
```

For Oracle Linux 9

```
rpm -i <...>/blksnap-6.3.2.1207-1.noarch.rpm  
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-6.3.2.1207-1.el9.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Veeam Agent for Linux in Fedora

To install Veeam Agent for Linux, use the following commands:

```
rpm -i <...>/blksnap-6.3.2.1207-1.noarch.rpm
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
rpm -i <...>/veeam-6.3.2.1207-1.fc34.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Veeam Agent for Linux in SLES

To install Veeam Agent for Linux, use the following commands:

For SLES 12 SP4

```
zypper in <...>/veeamsnap-kmp-default-6.3.2.1207_k4.12.14_94.41-sles12.4.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle12.x86_64.rpm
```

For SLES 12 SP5

```
zypper in <...>/veeamsnap-kmp-default-6.3.2.1207_k4.12.14_120-sles12.5.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle12.x86_64.rpm
```

For SLES 15 SP1

```
zypper in <...>/veeamsnap-kmp-default-6.3.2.1207_k4.12.14_195-sles15.1.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For SLES 15 SP2 with default kernel

```
zypper in <...>/veeamsnap-kmp-default-6.3.2.1207_k5.3.18_22-sles15.2.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For SLES 15 SP2 with preemptive kernel

```
zypper in <...>/veeamsnap-kmp-preempt-6.3.2.1207_k5.3.18_22-sles15.2.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For SLES 15 SP3 with default kernel

```
zypper in <...>/blksnap-kmp-default-6.3.2.1207_k5.3.18_57-sles15.3.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For SLES 15 SP3 with preemptive kernel

```
zypper in <...>/blksnap-kmp-preempt-6.3.2.1207_k5.3.18_57-sles15.3.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For SLES 15 SP4

```
zypper in <...>/blksnap-kmp-default-6.3.2.1207_k5.14.21_150400.22-sles15.4.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For SLES 15 SP5

```
zypper in <...>/blksnap-kmp-default-6.3.2.1207_k5.14.21_150500.53-sles15.5.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For SLES 15 SP6

```
zypper in <...>/blksnap-kmp-default-6.3.2.1207_k6.4.0_150600.21-sles15.6.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Veeam Agent for Linux in openSUSE

To install Veeam Agent for Linux, use the following commands:

For openSUSE Tumbleweed

```
zypper in <...>/blksnap-6.3.2.1207-1.sle.noarch.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.suse.x86_64.rpm
```

For openSUSE Leap 15.3 with default kernel

```
zypper in <...>/blksnap-kmp-default-6.3.2.1207_k5.3.18_59.10-opensuse_leap15.3.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For openSUSE Leap 15.3 with preemptive kernel

```
zypper in <...>/blksnap-kmp-preempt-6.3.2.1207_k5.3.18_59.10-opensuse_leap15.3.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For openSUSE Leap 15.4

```
zypper in <...>/blksnap-kmp-default-6.3.2.1207_k5.14.21_150400.22-opensuse_leap15.4.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For openSUSE Leap 15.5

```
zypper in <...>/blksnap-kmp-default-6.3.2.1207_k5.14.21_150500.53-opensuse_leap15.5.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

For openSUSE Leap 15.6

```
zypper in <...>/blksnap-kmp-default-6.3.2.1207_k6.4.0_150600.21-opensuse_leap15.6.x86_64.rpm
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm
zypper in <...>/veeam-6.3.2.1207-1.sle15.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Veeam Agent for Linux in Debian / Ubuntu

To install Veeam Agent for Linux, use the following commands:

For Debian 10 / Ubuntu 16.04, 18.04, 20.04 (kernel 5.4)

```
apt-get install <...>/veeamsnap_6.3.2.1207_all.deb
apt-get install <...>/veeam-libs_6.3.2.1207_amd64.deb
apt-get install <...>/veeam_6.3.2.1207_amd64.deb
```

For Debian 11 and 12 / Ubuntu 22.04 and later

```
apt-get install <...>/blksnap_6.3.2.1207_all.deb  
apt-get install <...>/veeam-libs_6.3.2.1207_amd64.deb  
apt-get install <...>/veeam_6.3.2.1207_amd64.deb
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Nosnap Veeam Agent for Linux in Offline Mode

To install nosnap Veeam Agent for Linux, do the following:

1. On a computer that is connected to the internet, download Veeam Agent packages intended for your Linux distribution from the [Veeam software repository](#).
 - For RPM packages of nosnap Veeam Agent for Linux, the Veeam Agent directory has the following structure: *Package format > Distribution > Version > Architecture*.

For example, Veeam Agent packages for 64-bit RHEL 9 reside in the [/rpm/el/9/x86_64/](#) folder of the Veeam software repository, and packages for 64-bit SLES 15 SP5 reside in the [/rpm/sles/SLE_15_SP5/x86_64/](#) folder.
 - Nosnap Veeam Agent for Linux packages in the Debian format reside in the following folders of the Veeam software repository:
 - [/backup/linux/agent/dpkg/debian/public//pool/veeam/v/veeam-nosnap/](#)
 - [/backup/linux/agent/dpkg/debian/public/pool/veeam/v/veeam-lib/](#)
2. Save the `veeam-nosnap` and `veeam-lib` packages to a directory that can be accessed from the computer where you want to install the product, for example, a directory on a local drive or USB drive, or a network shared folder.
3. On the computer where you want to install Veeam Agent, navigate to the directory where you have saved the packages and install Veeam Agent:
 - [Installing nosnap Veeam Agent for Linux in CentOS 7 / RHEL / Rocky Linux / AlmaLinux](#)
 - [Installing nosnap Veeam Agent for Linux in Oracle Linux](#)
 - [Installing nosnap Veeam Agent for Linux in SLES](#)
 - [Installing nosnap Veeam Agent for Linux in openSUSE](#)
 - [Installing nosnap Veeam Agent for Linux in Debian / Ubuntu](#)

TIP

You can also set up a local mirror of the Veeam software repository in your internal network and add this repository to the list of software sources on a computer where you want to install the product. These operations may differ depending on the Linux distribution and package manager that you use. To learn more, refer to the documentation of your Linux distribution.

After you add a local repository to the list of software sources on a computer, you will be able to install and upgrade Veeam Agent in a regular way. To learn more, see [Installing Veeam Agent for Linux](#) and [Upgrading Veeam Agent for Linux](#).

Installing Nosnap Veeam Agent for Linux in CentOS 7 / RHEL / Rocky Linux / AlmaLinux

To install nosnap Veeam Agent for Linux, use the following commands:

For 32-bit RHEL 6

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.i386.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el6.i386.rpm
```

For 64-bit RHEL 6

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el6.x86_64.rpm
```

For CentOS 7 / RHEL 7

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el7.x86_64.rpm
```

RHEL 8 / Rocky Linux 8.10 / AlmaLinux 8.10

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el8.x86_64.rpm
```

RHEL 9 / Rocky Linux 9.3 - 9.5 / AlmaLinux 9.3 - 9.5

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el9.x86_64.rpm
```

RHEL 10 / Rocky Linux 10 / AlmaLinux 10

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el10.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Nosnap Veeam Agent for Linux in Oracle Linux

To install nosnap Veeam Agent for Linux, use the following commands:

For 32-bit Oracle Linux 6

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.i386.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el6.i386.rpm
```

For 64-bit Oracle Linux 6

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el6.x86_64.rpm
```

For Oracle Linux 7

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el7.x86_64.rpm
```

For Oracle Linux 8

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el8.x86_64.rpm
```

For Oracle Linux 9

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el9.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Nosnap Veeam Agent for Linux in SLES

To install nosnap Veeam Agent for Linux, use the following commands:

For SLES 12 SP4 – SP5

```
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
zypper in <...>/veeam-nosnap-6.3.2.1207-1.sle12.x86_64.rpm
```

For SLES 15 SP1 – SP6

```
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
zypper in <...>/veeam-nosnap-6.3.2.1207-1.sle15.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Nosnap Veeam Agent for Linux in openSUSE

To install nosnap Veeam Agent for Linux, use the following commands:

For openSUSE Tumbleweed

```
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
zypper in <...>/veeam-nosnap-6.3.2.1207-1.suse.x86_64.rpm
```

For openSUSE Leap 15.3 - 15.6

```
zypper in <...>/veeam-libs-6.3.2.1207-1.x86_64.rpm  
zypper in <...>/veeam-nosnap-6.3.2.1207-1.sle15.x86_64.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Nosnap Veeam Agent for Linux in Debian / Ubuntu

To install nosnap Veeam Agent for Linux, use the following commands:

```
apt-get install <...>/veeam-libs_6.3.2.1207_amd64.deb  
apt-get install <...>/veeam-nosnap_6.3.2.1207_amd64.deb
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Nosnap Veeam Agent for Linux on Power in Offline Mode

To install nosnap Veeam Agent for Linux on Power, do the following:

1. On a computer that is connected to the internet, download Veeam Agent packages intended for your Linux distribution from the [Veeam software repository](#).

For RPM packages of nosnap Veeam Agent for Linux on Power, the Veeam Agent directory has the following structure: *Package format* > *Distribution* > *Version* > *Architecture*.

For example, Veeam Agent packages for RHEL 8 reside in the [/rpm/el/8/ppc64le/](#) folder of the Veeam software repository, and packages for SLES 15 SP4 reside in the [/rpm/sles/SLE_15_SP4/ppc64le/](#) folder.

2. Save the `veeam-nosnap` and `veeam-libs` packages to a directory that can be accessed from the computer where you want to install the product, for example, a directory on a local drive or USB drive, or a network shared folder.
3. On the computer where you want to install Veeam Agent, navigate to the directory where you have saved the packages and install Veeam Agent:
 - [Installing nosnap Veeam Agent for Linux on Power in RHEL](#)
 - [Installing nosnap Veeam Agent for Linux on Power in SLES](#)

TIP

You can also set up a local mirror of the Veeam software repository in your internal network and add this repository to the list of software sources on a computer where you want to install the product. These operations may differ depending on the Linux distribution and package manager that you use. To learn more, refer to the documentation of your Linux distribution.

After you add a local repository to the list of software sources on a computer, you will be able to install and upgrade Veeam Agent in a regular way. To learn more, see [Installing Veeam Agent for Linux](#) and [Upgrading Veeam Agent for Linux](#).

Installing Nosnap Veeam Agent for Linux on Power in RHEL

To install nosnap Veeam Agent for Linux on Power, use the following commands:

For RHEL 8.4, 8.6, 8.8 and 8.10

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.ppc64le.rpm
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el8.ppc64le.rpm
```

For RHEL 9.0, 9.2, 9.4 and 9.6

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.ppc64le.rpm
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el9.ppc64le.rpm
```

For RHEL 10

```
rpm -i <...>/veeam-libs-6.3.2.1207-1.ppc64le.rpm  
rpm -i <...>/veeam-nosnap-6.3.2.1207-1.el10.ppc64le.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Installing Nosnap Veeam Agent for Linux on Power in SLES

To install nosnap Veeam Agent for Linux on Power, use the following commands:

For SLES for SAP 12 SP5

```
zypper in <...>/veeam-libs-6.3.2.1207-1.ppc64le.rpm  
zypper in <...>/veeam-nosnap-6.3.2.1207-1.sle12.ppc64le.rpm
```

For SLES / SLES for SAP 15 SP3 – SP6

```
zypper in <...>/veeam-libs-6.3.2.1207-1.ppc64le.rpm  
zypper in <...>/veeam-nosnap-6.3.2.1207-1.sle15.ppc64le.rpm
```

where:

<...> – path to a directory where you have saved Veeam Agent packages.

Configuring UEFI Secure Boot

When you install Veeam Agent on a UEFI system with Secure Boot enabled, you must configure the UEFI Secure Boot to allow your system to run Veeam Agent and perform backups. You do this by enrolling a Machine Owner Key (MOK) for the Veeam kernel module in your system's firmware. To enroll MOK, perform the following steps:

1. Request enrollment of the key. For [pre-built](#) kernel modules, the key is provided by Veeam; for [DKMS](#) kernel modules, the key is generated by DKMS:
 - [Pre-built kernel module] To make UEFI system with Secure Boot work with pre-built Veeam kernel module, Veeam Agent requires Veeam public key to be enrolled to the system's MOK list. For more information on requesting enrollment of the Veeam kernel module key to your system, see [Importing MOK for Pre-Built Kernel Module](#).
 - [DKMS kernel module] If you install Veeam Agent in Ubuntu 22.04 and later, Debian 12.0 and later or Oracle Linux, DKMS generates a Machine Owner Key that allows third-party modules to be run on the system's firmware. Such key must also be enrolled to the system's MOK list. For more information on requesting enrollment of the key for the Veeam DKMS module, see [Importing MOK for Veeam DKMS Module](#).

NOTE

If UEFI system with Secure Boot enabled does not support automatic generation of the key for DKMS modules, you must either sign the Veeam kernel module yourself and enroll the Machine Owner Key to your system or disable Secure Boot.

2. Enroll the key using MOK management. For more information, see [Enrolling MOK](#).

Importing MOK for Pre-Built Kernel Module

The Veeam kernel module key is provided within the `ueficert` package that resides in the Veeam software repository.

NOTE

Depending on the Linux distribution and version, the `ueficert` package has one of the following two names:

- [For RHEL 9 and 10, AlmaLinux 9 and 10, Rocky Linux 9 and 10, SLES 15 SP3 – SP6 and openSUSE] `blksnap-ueficert`
- [For RHEL 6 – 8, SLES 12, 15 SP1 and SP2, CentOS] `veeamsnap-ueficert`

When you install the `ueficert` package, the Veeam kernel module key is imported into your system. Depending on how you install Veeam Agent – in the [online](#) or [offline](#) mode, installation of the Veeam `ueficert` package may vary.

Online Installation

If you install Veeam Agent for Linux by mounting the Veeam software repository directly on your Linux computer, install the `ueficert` package by using one of the following commands:

For RHEL 9 and 10 / Rocky Linux / AlmaLinux

```
yum install blksnap-ueficert
```

For RHEL 6 - 8 / CentOS 7

```
yum install veeamsnap-ueficert
```

For openSUSE / SLES 15 SP3 - SP6

```
zypper in blksnap-ueficert
```

For SLES 12, 15 SP1 and SP2

```
zypper in veeamsnap-ueficert
```

Offline Installation

If your Linux computer is not connected to the internet and you install Veeam Agent for Linux in the offline mode, to install the `ueficert` package and import MOK, do the following:

1. Download the Veeam `ueficert` package for your Linux distribution version from the [Veeam software repository](#).
 - [For RHEL, AlmaLinux, Rocky Linux and CentOS] The `ueficert` package is located in the `/backup/linux/agent/rpm/<Linux distr>/<ver>/x86_64/` directory, where `<Linux distr>` is the name of your Linux distribution and `<ver>` is your Linux OS version. For example, to download the `ueficert` package for RHEL 10, navigate to the [/backup/linux/agent/rpm/el/10/x86_64/](#) directory.
 - [For SLES and openSUSE] The `ueficert` package is located in the `/backup/linux/agent/rpm/<Linux distr>/<ver>/noarch/` directory, where `<Linux distr>` is the name of your Linux distribution and `<ver>` is your Linux OS version. For example, to download the `ueficert` package for SLES 15 SP6, navigate to the [/backup/linux/agent/rpm/sles/SLE_15_SP6/noarch/](#) directory.
2. Save the package to your computer.
3. Install the `ueficert` package by using the following command:

```
rpm -i <...>/blksnap-ueficert-6.3.2.1207-1.noarch.rpm
```

or

```
rpm -i <...>/veeamsnap-ueficert-6.3.2.1207-1.noarch.rpm
```

where `<...>` is the full path to the directory that contains the `ueficert` package.

Depending on the Linux distribution and version, after you install the `ueficert` package, the following will happen:

- [For all supported Linux distributions except RHEL 10, AlmaLinux 10 and Rocky Linux 10] The key is automatically imported into the enrollment request.

TIP

After the package is installed, you can verify that the key enrollment is planned for the next reboot using the following command: `mokutil -N`. If the command output shows that the key enrollment is not planned, import the public key [manually](#).

- [For RHEL 10, AlmaLinux 10 and Rocky Linux 10] According to enhanced security standards of the operating system, the Veeam kernel module key is not automatically imported during installation of the `ueficert` package. After you install the `ueficert` package, you must import the key manually using one of the following commands depending on the package you installed:

```
mokutil --import /etc/uefi/certs/blksnap-ueficert
```

or

```
mokutil --import /etc/uefi/certs/veeamsnap-ueficert
```

During manual key import, you will be prompted to provide a one-time MOK password that you will use to enroll the key.

After you install the `ueficert` package and import the key, you can [confirm the key enrollment](#).

Importing MOK for DKMS Kernel Module

Veeam does not provide a `ueficert` package for the DKMS module because it is not possible to sign such module automatically. Depending on the Linux distribution and version, you may have several options to make your system load the Veeam DKMS module properly – for more information, see [Linux documentation](#).

If your system runs on Ubuntu 22.04 and later or Debian 12.0 and later, after you install Veeam kernel module using DKMS, a new Machine Owner Key is generated. Depending on the Linux distribution, perform the following steps to request enrollment of the key to your system's firmware:

- [Debian 12.0 and later] By default, the key is stored in the `/var/lib/dkms/` directory. To import the key, run the following command:

```
mokutil --import /var/lib/dkms/mok.pub
```

- [Ubuntu 22.04 and later] After you install the Veeam kernel module, the key is generated and imported into your system automatically. By default, the key is stored in the `/var/lib/shim-signed/mok` directory.

When the key is imported into the enrollment request, you will be prompted to enter a password that you will use to confirm the enrollment of the key during MOK management. After you set the password, you can [confirm the key enrollment](#).


Enrolling MOK

To enroll the Veeam or DKMS-generated key to the MOK list, do the following:


1. Reboot the computer.
2. During reboot, when prompted, press any key to perform MOK management.

IMPORTANT

The prompt will time out in 10 seconds. If you don't press any key, the system will continue booting without enrolling the key. If you don't enroll the key at reboot, you will have to reconfigure the key by reinstalling the `ueficer` package and reboot again.




3. At the first step of the wizard, select **Enroll MOK** and press [Enter].




4. At the **Enroll MOK** step, select **Continue** and press [Enter].




5. At the **Enroll the key(s)** step, select **Yes** and press [Enter].



6. Depending on whether the key was imported automatically or manually, do the following:
- If the key was imported automatically, type in the password for the root account and press [Enter].
 - If the key was imported manually, type in the password you set when you imported the key and press [Enter].



- At the final step, select **Reboot** and press [Enter].



- After the system reboots, verify that the key is successfully enrolled with the following command:
`mokutil -l`. The system will list the enrolled keys.

Upgrading Veeam Agent for Linux

For Veeam Agent for Linux, upgrade to newer versions is supported. You can start the upgrade process when the new version becomes available.

During the upgrade process, configuration and backup files that were created with the previous version of Veeam Agent are not impacted in any way.

IMPORTANT

Before starting the upgrade process, make sure that there are no jobs running on the Veeam Agent computer.

Depending on the type of the [packages](#) you used for Veeam Agent installation, you can use the following upgrade procedures:

- [Upgrading Veeam Agent for Linux with Kernel Module](#)
- [Upgrading Nosnap Veeam Agent for Linux](#)
- [Upgrading Nosnap Veeam Agent for Linux on Power.](#)

TIP

If the computer where you want to upgrade Veeam Agent for Linux is not connected to the internet and does not have access to a local mirror of the [Veeam software repository](#), you can download and re-install Veeam Agent for Linux packages manually. To learn more, see [Installing Veeam Agent for Linux in Offline Mode](#).

Upgrading Veeam Agent for Linux with Kernel Module

The commands for the upgrade of Veeam Agent for Linux differ depending on the Linux distribution:

- [Upgrading Veeam Agent for Linux in CentOS 7 / RHEL 6 – 8 / Rocky Linux 8 / AlmaLinux 8](#)
- [Upgrading Veeam Agent for Linux in RHEL 9 and 10 / Rocky Linux 9.3 – 9.6 and 10 / AlmaLinux 9.3 – 9.6 and 10](#)
- [Upgrading Veeam Agent for Linux in Oracle Linux 6 – 8](#)
- [Upgrading Veeam Agent for Linux in Fedora / Oracle Linux 9](#)
- [Upgrading Veeam Agent for Linux in openSUSE](#)
- [Upgrading Veeam Agent for Linux in SLES 12 SP4 – SP5, 15 SP1 – SP2](#)
- [Upgrading Veeam Agent for Linux in SLES 15 SP3 – SP6](#)
- [Upgrading Veeam Agent for Linux in Debian 10 / Ubuntu 16.04, 18.04, 20.04 \(kernel 5.4\)](#)
- [Upgrading Veeam Agent for Linux in Debian 11 and 12 / Ubuntu 22.04 – 25.04](#)

Upgrading Veeam Agent for Linux in CentOS 7 / RHEL 6 – 8 / Rocky Linux 8 / AlmaLinux 8

To upgrade Veeam Agent for Linux, use the following command:

```
yum update veeam
```

With these commands, a pre-built binary package with Veeam kernel module will be installed in your system. To stay on the DKMS version of the Veeam kernel module, use the following command for upgrade:

```
yum update veeamsnap && yum update veeam
```

Upgrading Veeam Agent for Linux in RHEL 9 and 10 / Rocky Linux 9 and 10 / AlmaLinux 9 and 10

To upgrade Veeam Agent for Linux, use the following command:

```
yum install kmod-blksnap veeam --allowerase
```

With this command, a pre-built binary package with the Veeam kernel module will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the Veeam kernel module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
yum install blksnap veeam --allowerase
```

Upgrading Veeam Agent for Linux in Oracle Linux 6 – 8

To upgrade Veeam Agent for Linux, use the following command:

```
yum update veeam
```

Upgrading Veeam Agent for Linux in Fedora / Oracle Linux 9

To upgrade Veeam Agent for Linux, use the following command:

```
yum update veeam --allowerase
```

Upgrading Veeam Agent for Linux in openSUSE

To upgrade Veeam Agent for Linux, use the following commands:

For openSUSE Tumbleweed

```
zypper update veeam
```

For openSUSE Leap 15.3 with default kernel

```
zypper in -- replacefiles blksnap-kmp-default veeam
```

For openSUSE Leap 15.3 with preemptive kernel

```
zypper in -- replacefiles blksnap-kmp-preempt veeam
```

For openSUSE Leap 15.4 – 15.6

```
zypper in -- replacefiles blksnap-kmp-default veeam
```

With these commands, a pre-built binary package with the Veeam kernel module will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the Veeam kernel module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
zypper update veeam
```

Upgrading Veeam Agent for Linux in SLES 12 SP4 – SP5, 15 SP1 – SP2

To upgrade Veeam Agent for Linux, use the following commands:

For default kernel

```
zypper in veeamsnap-kmp-default veeam
```

For preemptive kernel

```
zypper in veeamsnap-kmp-preempt veeam
```

With these commands, a pre-built binary package with Veeam kernel module will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the Veeam kernel module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
zypper update veeam
```

Upgrading Veeam Agent for Linux in SLES 15 SP3 – SP6

To upgrade Veeam Agent for Linux, use the following commands:

For SLES 15 SP3 with default kernel

```
zypper in --replacefiles blksnap-kmp-default veeam
```

For SLES 15 SP3 with preemptive kernel

```
zypper in --replacefiles blksnap-kmp-preempt veeam
```

For SLES 15 SP4 – SP6

```
zypper in --replacefiles blksnap-kmp-default veeam
```


With these commands, a pre-built binary package with the Veeam kernel module will be installed in your system. The `--force` key is required to properly replace the missing link to `.ko` in case of update from the DKMS version of the Veeam kernel module to a pre-built binary. To stay on the DKMS version, use the following command for upgrade:

```
zypper in --replacefiles blksnap veeam
```

Upgrading Veeam Agent for Linux in Debian 10 / Ubuntu 16.04, 18.04, 20.04 (kernel 5.4)

To upgrade Veeam Agent for Linux, use the following commands:

```
apt-get update  
apt-get install veeam
```

Upgrading Veeam Agent for Linux in Debian 11 and 12 / Ubuntu 22.04 – 25.04

To upgrade Veeam Agent for Linux, use the following commands:

```
apt-get update  
apt-get install blksnap veeam
```

Upgrading Nosnap Veeam Agent for Linux

The commands for the upgrade of nosnap Veeam Agent for Linux differ depending on the Linux distribution:

For CentOS 7 / RHEL / Oracle Linux / Rocky Linux / AlmaLinux

```
yum update veeam-nosnap
```

For openSUSE Tumbleweed / OpenSUSE Leap / SLES

```
zypper in veeam-nosnap
```

For Debian / Ubuntu

```
apt-get update  
apt-get install veeam-nosnap
```

Upgrading Nosnap Veeam Agent for Linux on Power

The commands for the upgrade of nosnap Veeam Agent for Linux on Power differ depending on the Linux distribution:

For RHEL

```
yum update veeam-nosnap
```

For SLES

```
zypper in veeam-nosnap
```

Granting Permissions to Users

When you install Veeam Agent, the product program files are placed to the folders on the system volume. For full access to Veeam Agent files, super user (root) privileges are required. Rights to execute product files and run commands are also granted to users that belong to the `veeam` group.

The `veeam` group is automatically created by Veeam Agent at the process of the product installation. To let regular users work with Veeam Agent without the need to gain root privileges, you can add the necessary users to this group. Users in the `veeam` group will be able to execute Veeam Agent commands and perform backup and restore tasks under regular user account.

To add a user to the `veeam` group, in most of Linux distributions you can use the following command:

```
usermod -a -G veeam <username>
```

where:

`<username>` — name of the account to which you want to grant access to Veeam Agent.

For example:

```
root@srv01:~# usermod -a -G veeam user01
```

IMPORTANT

Consider the following:

- To add a user to the `veeam` group, you must have super user (root) privileges.
- After the user is added to the `veeam` group, they must re-log in to the system to apply the new permissions.
- Add only trusted users to the `veeam` group. Veeam Agent daemon runs and executes commands and scripts with the super user privileges. Thus, users who belong to this group can potentially escalate their privileges through the creative use of pre-freeze/post-thaw or pre-job/post-job scripts.

To check whether the user who is currently logged in to the Linux OS is added to the `veeam` group, you can use the following command:

```
groups
```

For example:

```
user@srv01:~$ groups
user adm cdrom dip plugdev lpadmin sambashare veeam
```

Performing Initial Setup

After you install Veeam Agent for Linux, you can use the [Veeam Agent for Linux control panel](#) to perform initial product setup. When you launch the control panel for the first time, Veeam Agent displays the initial setup wizard. The wizard offers you to accept license agreements, install a license and create a custom Veeam Recovery Media that will include drivers of your Veeam Agent computer.

To perform initial setup, launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command. Then use the initial setup wizard to complete the following steps:


1. [Accept Veeam and third-party license agreements.](#)
2. [Create a custom Veeam Recovery Media.](#)
3. [Install a license.](#)

Step 1. Accept License Agreements

At the **Agreements** step of the initial setup wizard, accept the terms of the product license agreement and license agreements for third-party components of the product. You must accept the license agreements to start using the product. Until you accept the license agreements, you will not be able to perform backup and data recovery tasks with the Veeam Agent for Linux control panel and command line interface.

To accept the license agreements:

1. Make sure that the **I accept the terms of the Veeam license agreement** option is selected and press [Space].
2. Select the **I accept the terms of the 3rd party components license agreements** option with the [Down] and [Up] key and press [Space].
3. Press [Enter].



Step 2. Create Custom Veeam Recovery Media

At the **Recovery ISO** step of the initial setup wizard, specify settings for the custom Veeam Recovery Media.

In addition to the generic Veeam Recovery Media that is available for download at the Veeam website, you can create a custom Veeam Recovery Media. This option may be helpful if your computer uses hardware that requires drivers not included in the generic Veeam Recovery Media. When you create the custom Veeam Recovery Media, Veeam Agent for Linux copies the Linux kernel running on your computer with its currently loaded modules and includes them into the custom recovery media.

IMPORTANT

Recovery Media patching is not supported by Veeam Agent for Linux on Power. Switch to the **Next** button with the [Tab] key and press [Enter]. You will proceed immediately to the [License](#) step of the initial setup wizard.

Before you create a custom Veeam Recovery Media, check the following prerequisites:

- The Linux system must have the `genisoimage` package installed. For openSUSE and SLES 15 SP1 – 15 SP5 distributions, the Linux system must have the `mkisofs` package installed instead.
- The Linux system must have the `mksquashfs` and `unsquashfs` utilities installed.
- For custom Veeam Recovery Media with EFI support, the Linux system must have the following packages installed:
 - `xorriso`
 - `isolinux` (or `syslinux`, if the software package repository of your Linux distribution lacks the `isolinux` package).
- For the scenario where you create a custom Veeam Recovery Media using the **Download and patch ISO** option, the Linux system must have the `wget` utility installed.

TIP

If you do not want to create a custom Veeam Recovery Media at the process of initial product setup, switch to the **Next** button with the [Tab] key and press [Enter]. You will proceed immediately to the [License](#) step of the initial setup wizard.

You can create the custom Veeam Recovery Media later, at any time you need, using the Veeam Agent for Linux command line interface. To learn more, see [Creating Custom Veeam Recovery Media](#).

To specify settings for the custom Veeam Recovery Media:

1. Make sure that the **Patch Veeam Recovery Media ISO** option is selected and press [Space].
2. If you want the Veeam Recovery Media to be able to boot on EFI-based systems, select the **EFI system** option with the [Tab] key and press [Space].

If you do not enable this option, the custom Veeam Recovery Media will be able to boot on BIOS-based systems only.

4. Press [Tab] and select how you want to create a custom Veeam Recovery Media depending on the location of the generic recovery media ISO file:

- If you have not downloaded the generic Veeam Recovery Media, make sure that the **Download and patch ISO** option is selected and press [Tab]. If you select this option, Veeam Agent will download the ISO file of the generic Veeam Recovery Media from the Veeam software repository to the directory of your choice on the Veeam Agent computer and use this image to create the custom Veeam Recovery Media.

Veeam Agent downloads the Veeam Recovery Media ISO file depending on the Veeam Agent computer architecture. For details, see [Veeam Recovery Media Versions](#).

- If you want only to download the generic Veeam Recovery Media, select the **Only download ISO** option with the [Down] key and press [Tab]. If you select this option, Veeam Agent will download the ISO file of the generic Veeam Recovery Media from the Veeam software repository to the directory of your choice on the Veeam Agent computer. You can use the downloaded ISO file later to boot your Veeam Agent computer or to create a custom Veeam Recovery Media.

Veeam Agent downloads the Veeam Recovery Media ISO file depending on the Veeam Agent computer architecture. For details, see [Veeam Recovery Media Versions](#).

- If you have already downloaded the generic Veeam Recovery Media to a local directory on the Veeam Agent computer or to a network shared folder, select the **Patch local ISO** option with the [Down] key and press [Tab]. If you select this option, Veeam Agent will use the generic Veeam Recovery Media ISO file to create the custom Veeam Recovery Media.

The name of the generic Veeam Recovery Media ISO file depends on the recovery image version, Veeam Agent computer architecture and the source from which you downloaded the ISO file: from the product download page or Veeam software repository. To learn more, see [Veeam Recovery Media Versions](#).

5. If you selected the **Download and patch ISO** or **Patch local ISO** option, the **EFI system** option is available. If you want to boot the Veeam Recovery Media on EFI-based systems, select the **EFI system** option with the [Tab] key and press [Space].

If you do not enable this option, the custom Veeam Recovery Media will be able to boot on BIOS-based systems only.

6. If you selected the **Patch local ISO** option, in the **Path to local ISO** field, specify a path to the ISO file of the generic Veeam Recovery Media:


- a. Select the **Browse** option with the [Tab] key and press [Enter].
- b. In the **Path to ISO** window, select the necessary directory and press [Enter].
- c. Repeat the step 'b' until a path to the directory in which the recovery media ISO file resides appears in the **Current directory** field.
- d. In the directory where the recovery media ISO file resides, select the ISO file and press [Enter].

7. Specify a path to the resulting ISO file of the Veeam Recovery Media.

If you selected the **Download and patch ISO** or **Patch local ISO** option, in the **Save patched ISO to** field, you can specify a path to the resulting ISO file of the custom Veeam Recovery Media; if you selected the **Only download ISO** option, in the **Save ISO to** field, specify a path to the resulting ISO file of the generic Veeam Recovery Media:

- a. Select the **Browse** option with the [Tab] key and press [Enter].
- b. In the **Save patched ISO to** window, select the necessary directory and press [Enter].
- c. Repeat the step 'b' until a path to the directory where you want to save the resulting custom recovery media ISO file appears in the **Current directory** field.

- d. Select the **OK** button with the [Tab] key and press [Enter].
8. To start the custom recovery media creation process, select the **Next** button with the [Tab] key and press [Enter].



Step 3. Install Product License

At the **License** step of the initial setup wizard, install the license. You can choose to install the license immediately or postpone this operation.

- If you choose to install the license, you can immediately browse for the license key on your computer and complete the license installation process.
- If you choose to postpone the license installation process, you will be able to install a license later at any time you need.

Until you install a license, Veeam Agent for Linux will operate in the Free edition. To learn more, see [Product Editions](#).

NOTE

If you choose not to install a license and use Veeam Agent in the Free edition, Veeam Agent will display a notification offering to install a license every time you open the control panel. The notification will appear in the control panel until Veeam Agent completes the first backup job session.


To install a license:

1. In the **File location** field, specify a path to the license key:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].
 - b. In the **Choose license file location** window, select the necessary directory and press [Enter].
 - c. Repeat the step 'a' until a path to the directory in which the license key resides appears in the **Current directory** field.
 - d. In the directory where the license key resides, select the license key and press [Enter].
2. In the **Choose agent edition to use on this computer** section, select the product edition in which Veeam Agent will operate and press [Enter] to install the license and finish working with the initial setup wizard.

TIP

Consider the following:

- If you do not want to install a license, to finish working with the initial setup wizard, switch to the **Finish** button with the [Tab] key and press [Enter].
- You can view information about the installed license (expiration date, status of the license, current edition of the product and so on) in the Veeam Agent control panel or using the Veeam Agent command line interface. To learn more, see [Viewing License](#).



Configuring Advanced Settings

Veeam Agent for Linux allows you to configure the following settings:

- [HTTP proxy settings for Veeam Cloud Connect repository](#)
- [Connection settings for Veeam backup server](#)

HTTP Proxy Settings for Veeam Cloud Connect Repository

If you want to use Veeam Agent for Linux to back up your data to a Veeam Cloud Connect repository, it might be required that you specify HTTP proxy settings for Veeam Agent.

Veeam Agent computer needs access to CRLs (Certificate Revocation Lists) of the CA (Certification Authority) who issued a certificate to the Veeam Cloud Connect service provider. In case it is not possible to establish a direct connection to CRLs, you must configure an HTTP proxy and specify settings to connect to the proxy in Veeam Agent.

To specify settings for an HTTP proxy, uncomment and edit the following lines in the `[cloudconnect]` section of the `/etc/veeam/veeam.ini` configuration file:

```
[cloudconnect]
...
# httpproxylogin= <username>
...
# httpproxypasswd= <password>
...
# httpproxyurl= <URL>
```

where:

- `<username>` – name of the account used to connect to the HTTP proxy.
- `<password>` – password of the account used to connect to the HTTP proxy.
- `<URL>` – URL of a proxy used for CRL checks.

NOTE

If the proxy does not require authentication, you do not need to specify the account name and password. Keep in mind that only the basic authentication method is supported for connection to a proxy.

For example:

```
[cloudconnect]
...
# HTTP proxy login
httpproxylogin= user01
# HTTP proxy password
httpproxypasswd= P@ssw0rd
# HTTP proxy URL for CRL checks
httpproxyurl= http://proxy.company.lan:3128
```

Connection Settings for Veeam Backup Server

If you want to connect Veeam Agent computer to Veeam backup server as a member of the protection group for pre-installed Veeam Agents, you must apply connection settings from the configuration file. The configuration file is one of the Veeam Agent for Linux setup files that you must obtain from your System Administrator. To learn more about protection group for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

To connect Veeam Agent for Linux to Veeam backup server:

1. Get the configuration file from your System Administrator and upload this file on the Veeam Agent computer.
2. Navigate to the directory where you have saved the configuration file and run the following command:

```
veeamconfig mode setvbrsettings --cfg <file_name>.xml
```

where <file_name> is a configuration file name.

Alternatively, you can specify the full path to the configuration file with the `--cfg` option.

For example:

```
user@srv01:~# veeamconfig mode setvbrsettings --cfg /home/Linux\ Servers\ Distrib  
ibs/Linux/LinuxServers.xml
```

Mind that the connection between Veeam backup server and Veeam Agent computer added as a member of the protection group for pre-installed Veeam Agents is not persistent. Veeam Agent synchronizes with Veeam Backup & Replication every 6 hours. To synchronize Veeam Agent immediately, run the following command from the Veeam Agent computer:

```
veeamconfig mode syncnow
```

Managing Veeam Agent Operation Mode

Veeam Agent for Linux can operate in different modes. Depending on the selected mode, Veeam Agent offers different features and limitations. To learn more, see [Standalone and Managed Operation Modes](#).

Veeam Agent allows you to perform the following actions to manage the operation mode:

- [View operation mode details](#)
- [Reset to the standalone operation mode](#)
- [Connect to Veeam backup server](#)
- [Synchronize with Veeam backup server](#)
- [Export logs to Veeam backup server](#)

Viewing Operation Mode

To view the current Veeam Agent operation mode, use the following command:

```
veeamconfig mode info
```

Veeam Agent displays the operation mode details:

Parameter	Description
Owner	Name of the Veeam backup server that manages Veeam Agent. If Veeam Agent operates in the standalone mode, Veeam Agent will display the <i>Not Set</i> value.
Mode	Current Veeam Agent operating mode. Possible values: • <i>Not Set</i> – Veeam Agent operates in the standalone mode. • <i>Job</i> – Veeam Agent operates in the managed mode. Veeam Agent computer is protected by a backup job managed by backup server. • <i>Policy</i> – Veeam Agent operates in the managed mode. Veeam Agent computer is protected by a backup job managed by Veeam Agent for Linux. Veeam Agent computer is connected to the Veeam backup server as a member of any protection group excluding protection group for pre-installed Veeam Agents. • <i>Pre-installed</i> – Veeam Agent operates in the managed mode. Veeam Agent computer is protected by a backup job managed by Veeam Agent for Linux. Veeam Agent computer is connected to the Veeam backup server as a member of a protection group for pre-installed Veeam Agents. Keep in mind that features and limitations of Veeam Agent operating in the managed mode are different from those in the standalone mode. To learn more about managed mode, see the Veeam Agent Management Guide .

For example:

```
user@srv01:~$ veeamconfig mode info
Owner: Backup server (backupserver001.tech.local)
Mode: Job
```

If Veeam Agent operates in the managed mode, you can reset it to the standalone mode at any time. To learn more, see [Resetting to Standalone Operation Mode](#).

Resetting to Standalone Operation Mode

If Veeam Agent operates in the managed mode, you can manually reset it to the standalone mode from the Veeam Agent side. To learn more about operation modes, see the [Standalone and Managed Operation Modes](#).

Before you reset Veeam Agent to the standalone mode, consider the following:

- All backup jobs configured on the Veeam Agent computer will be deleted. If you plan to protect this computer with a standalone Veeam Agent, you will need to create new backup jobs.
- Veeam backup server settings including protection group configuration settings will be deleted.
- Previously created backup files will remain in the target backup repository. If the target repository is managed by the Veeam backup server, in the Veeam Backup & Replication console, they will be marked as *Orphaned*.
- If you want to reset Veeam Agent that operates in the *Job* or *Policy* mode, we recommend that you do the following:
 - Remove Veeam Agent computer from the protection group using the Veeam Backup & Replication console. To learn more about removing computers from a protection group in the Veeam Backup & Replication console, see the [Removing Computer from Protection Group](#) section in the Veeam Agent Management Guide. Veeam Agent on the protected computer will automatically switch to the standalone operation mode.
 - If Veeam Agent does not automatically switch to the standalone mode after you remove the Veeam Agent computer from the Veeam Backup & Replication configuration database, reset the operating mode on the Veeam Agent computer.
 - If Veeam Agent operates in the *Pre-installed* mode, Veeam Agent computer will be automatically removed from the protection group for pre-installed Veeam Agents in Veeam Backup & Replication.

To reset Veeam Agent to the standalone operating mode, run the following command:

```
veeamconfig mode reset
```

You can use the `--force` option to override additional input prompts and error messages:

```
veeamconfig mode reset --force
```


Connecting to Veeam Backup & Replication

If you want to connect a Veeam Agent computer as a member of the protection group for pre-installed Veeam Agents to a Veeam backup server, you must apply connection settings from the protection group configuration file to Veeam Agent. The configuration file is one of the Veeam Agent setup files that you must obtain from your System Administrator. To learn more about deployment using external tools, see the [Deploying Veeam Agent for Linux](#) section in the Veeam Agent Management Guide.

To connect Veeam Agent to Veeam backup server:

1. Get the configuration file from your System Administrator and upload this file to the Veeam Agent computer.
2. Navigate to the directory where you have saved the configuration file and run the following command:

```
veeamconfig mode setvbrsettings --cfg <file_name>.xml --force
```

where:

- o `<file_name>` – configuration file name. Alternatively, you can specify the full path to the configuration file with the `--cfg` option.
- o `--force` – with this option enabled, Veeam Agent will override additional input prompts and error messages. This parameter is optional.

For example:

```
user@srv01:~$ veeamconfig mode setvbrsettings --cfg /home/Linux\ Servers\ Distribs/Linux/LinuxServers.xml
```

Synchronizing with Veeam Backup Server

When Veeam Agent is managed by Veeam backup server, the connection between Veeam backup server and Veeam Agent computer added to a protection group is not persistent. Veeam Agent synchronizes with Veeam Backup & Replication every 6 hours. During the synchronization, Veeam Agent gets updated backup policies and configuration settings from the Veeam backup server, the Veeam backup server gets certificate details and session logs from Veeam Agent. To synchronize Veeam Agent immediately, run the following command:

```
veeamconfig mode syncnow
```

Exporting Logs to Veeam Backup Server

If Veeam Agent is connected to the Veeam backup server as a member of the protection group for pre-installed Veeam Agents, Veeam Agent can collect the required logs, export them to an archive file and send to the Veeam backup server. This operation may be required if you want to report an issue and need to attach log files to the support case.

To export logs, use the following command:

```
veeamconfig mode exportdebuglogs
```

Veeam Agent will collect logs, export them to an archive file with the name `veeam_logs_<agent>_<date>_<time>.tar.gz` and save the archive to the following folder on the Veeam backup server:

```
C:\ProgramData\Veeam\Backup\Endpoint\Other\AgentLogs\<computer_name>
```

where `<computer_name>` – name of the computer with Veeam Agent installed.

TIP

If Veeam Agent operates in the standalone mode, you can export product logs only to a local directory on the Veeam Agent computer. To learn more, see [Exporting Product Logs](#).

Uninstalling Veeam Agent for Linux

Depending on the type of the [packages](#) you used for Veeam Agent installation, you can use the following procedures to uninstall Veeam Agent:

- [Uninstalling Veeam Agent for Linux with Kernel Module](#)
- [Uninstalling Nosnap Veeam Agent for Linux / Nosnap Veeam Agent for Linux on Power](#)

Uninstalling Veeam Agent for Linux with Kernel Module

To uninstall Veeam Agent for Linux, you need to remove the following packages from the computer: `veeam`, `veeam-libs` and `veeamsnap` or `blksnap`. To do this, run one of the following commands depending on the Veeam kernel module you used during installation and your Linux distribution:

For CentOS 7 / RHEL / Oracle Linux / Fedora / Rocky Linux / AlmaLinux

```
yum remove veeam veeam-libs veeamsnap
```

or

```
yum remove veeam veeam-libs blksnap
```

For openSUSE / SLES

```
zypper rm veeam veeam-libs veeamsnap
```

or

```
zypper rm veeam veeam-libs blksnap
```

For Debian / Ubuntu

```
apt-get remove veeam veeam-libs veeamsnap
```

or

```
apt-get remove veeam veeam-libs blksnap
```

Uninstalling Nosnap Veeam Agent for Linux / Nosnap Veeam Agent for Linux on Power

To uninstall nosnap Veeam Agent for Linux or nosnap Veeam Agent for Linux on Power, you need to remove the `veeam-nosnap` and `veeam-libs` packages from the computer. To do this, run one of the following commands depending on your Linux distribution:

For CentOS 7 / RHEL / Oracle Linux / Rocky Linux / AlmaLinux

```
yum remove veeam-nosnap veeam-libs
```

For openSUSE / SLES

```
zypper rm veeam-nosnap veeam-libs
```

For Debian / Ubuntu

```
apt-get remove veeam-nosnap veeam-libs
```

Getting Started

To protect your computer from a disaster of any kind, you must perform the following operations in Veeam Agent for Linux:

1. Define what data you want to back up and configure the backup job.

Before you configure the backup job, you should decide on the following backup details:

- Backup destination: where you want to store your backed-up data.
- Backup scope: entire computer image, individual computer volumes or specific computer folders and files.
- Backup schedule: how often you want to back up your data.

After that, you can configure one or several backup jobs. The backup job captures the data that you have added to the backup scope and creates a chain of restore points in the target location. If your data gets lost or corrupted, you can restore it from the necessary restore point.

In Veeam Agent, you can configure the backup job in one of the following ways:

- [With the Backup Job wizard](#)
- [With the command line interface](#)

2. Monitor backup task performance.

You can use the Veeam Agent Control Panel to check how backup tasks are being performed, what errors have occurred during backup job sessions and so on. You can also use Veeam Agent command line interface to get information on backup and restore sessions status and view session logs. To learn more, see [Reporting](#).

3. In case of a disaster, you can restore the entire computer image or specific data on the computer. With Veeam Agent, you can perform data recovery operations in several ways:
 - You can boot from the Veeam Recovery Media and perform volume-level restore or file-level restore.
 - You can perform volume-level restore with Veeam Agent command line interface.
 - You can perform file-level restore with the Veeam Agent File Level Restore wizard.
 - You can export backup to a VHD virtual disk and attach this disk to a virtual machine to recover your computer in virtual environment.

To learn more, see [Performing Restore](#).

Getting to Know User Interface

With Veeam Agent for Linux, you can perform backup, restore and configuration tasks in the following ways:

- [Using Veeam Agent control panel](#)

Veeam Agent control panel is a GUI-like user interface based on the `ncurses` programming library. With Veeam Agent control panel, you can perform all basic data protection tasks. You can configure a backup job, start and stop backup jobs, monitor backup job session performance and recover files and folders. When you perform restore tasks after booting from the Veeam Recovery Media, you can also perform volume-level restore with the Veeam Recovery Media wizard.

- [Using command line interface](#)

With Veeam Agent command line interface, in addition to operations that can be performed with the Veeam Agent control panel, you can perform a set of advanced tasks. For example, you can:

- Configure advanced settings for backup jobs: specify compression level and data block size.
- Perform operations with backup repositories.
- Perform volume-level restore without the need to boot from the Veeam Recovery Media.
- Export backups to VHD virtual disks.
- Monitor performance and status of any backup, restore and other data transfer session that was started in Veeam Agent.
- View detailed information on every backup that was created with Veeam Agent.
- Export/import Veeam Agent configuration database to/from a configuration file.

Veeam Agent for Linux Control Panel

Veeam Agent for Linux control panel is a GUI-like user interface that lets users perform main backup and restore tasks in an easy way. With Veeam Agent for Linux control panel, you do not need to work with Linux shell and remember numerous commands. However, some advanced Veeam Agent for Linux operations are not supported by the control panel and can be performed with the command line interface only.

IMPORTANT

You cannot use Veeam Agent for Linux control panel on terminals that do not support colors (for example, VT100).

To launch the Veeam Agent for Linux control panel, you can use the following commands:

```
veeamconfig ui
```

or

```
veeam
```

NOTE

Veeam Agent for Linux control panel is based on the `ncurses` programming library. To use the Veeam Agent for Linux control panel, you must have the `ncurses` library installed in your Linux OS. To learn more, see [System Requirements](#).

When you launch the Veeam Agent for Linux control panel for the first time, Veeam Agent for Linux offers you to perform initial product setup. To learn more, see [Performing Initial Setup](#).

After you perform initial product setup, before you configure the first backup job, you can use the Veeam Agent for Linux control panel to perform the following operations:

- [Configure a new backup job](#).
- [Restore files and folders from existing backup](#).
- Manage [license](#) and [product logs](#).
- [Create a custom Veeam Recovery Media](#)


After you configure one or several backup jobs, you can also use the control panel to [start a backup job](#) and [work with backup job sessions](#).

Navigating Veeam Agent for Linux Control Panel

In the Veeam Agent for Linux control panel, the use of a mouse is not supported. To start an operation, you need to use a specific key on your keyboard. For example, you can press the [C] key to start the backup job configuration, press the [S] key to start a backup job or press the [R] key to start the file-level restore process. Short help information on the currently available operations and keys is displayed at the bottom of the control panel.

To navigate the control panel, backup job configuration and file-level restore wizards, you can use the following keys:

- [Tab] – to switch between controls and buttons in the Backup Job wizard.
- [Up] and [Down] – to switch between items in a scrollable list.
- [Space] – to select the necessary item in a list. The selected item's mark may vary in different steps of the wizard.
- [Enter] – to proceed to the next step of a wizard or to view details of the backup job session selected in the list of sessions.
- [Backspace] – to return to the previous step of a wizard (you cannot use this button to change wizard steps when a text field is selected).
- [Esc] – to exit the currently used wizard or close the Veeam Agent for Linux control panel.



Command Line Interface

Veeam Agent command line interface is a powerful tool that lets users perform advanced operations that are not supported by the Veeam Agent control panel.

To work with Veeam Agent using command line interface, you can use a terminal console (TTY) or a terminal emulator of your choice. All tasks in Veeam Agent are performed with the `veeamconfig` command-line utility. To perform tasks with Veeam Agent, you should construct the necessary command and type it in the Linux shell prompt.

You can view short help information on every Veeam Agent command at any time you need. To learn more, see [Viewing Help](#).

You should construct a command in the following format:

```
veeamconfig <command_1> <command_2> --<parameter_1> --<parameter_2> --<parameter_n>
```

where:

- `<command_1>` – command that defines a type of an object with which you want to perform a task. Currently, the following commands are available in Veeam Agent:
 - aap
 - agreement
 - backup
 - cloud
 - config
 - downloadiso
 - gfs
 - grablogs
 - healthcheck
 - help
 - job
 - license
 - mode
 - objectstorage
 - patchiso
 - point
 - repository
 - schedule
 - session

- ui
- version
- vbrserver
- `<command_2>` – command that defines a task that you want to perform with an object of the specified type. For example, you can perform the following commands with backup repositories:
 - create
 - delete
 - disable
 - edit
 - enable
 - help
 - list
 - rescan
- `<parameter_1>`, `<parameter_2>`, `<parameter_n>` – parameters for the command that you want to execute. Commands may require one or several mandatory or optional parameters. Some commands, for example, `veeamconfig ui` and `veeamconfig [<command>] help` do not require parameters.

The following example shows the command that displays a list of backup repositories configured in Veeam Agent and the output of this command:

```
user@srv01:~$ veeamconfig repository list
Name      ID                               Location      Type
e Backup server
Repository_1 {818e3a0f-8155-4a51-9430-248a203a43d1} /home/backups local
l
Repository_2 {2155a2e7-a1e9-4347-9d8b-cf8f3a6f3fcb} 172.17.53.47/veeam cif
s
```

Viewing Help

You can view short help information on the specific Veeam Agent command. To view help, use the following command:

```
veeamconfig <command> help
```

where:

<command> — name of the command for which you want to view help information.

For example:

```
user@srv01:~$ veeamconfig help
```

or

```
user@srv01:~$ veeamconfig job help
```

or

```
user@srv01:~$ veeamconfig job create help
```

You can also view the manual page for the `veeamconfig` utility. Use the following command:

```
man veeamconfig
```

Licensing

To work with Veeam Agent, you must accept terms of the product license agreement and license agreements for third-party components operating as part of the product.

If you want to use a commercial version of Veeam Agent, you must obtain a license and install it on the protected computer. If you do not install a license, the product will operate in the Free edition.

You can use the Veeam Agent control panel or Veeam Agent command line interface to install a license, monitor status of the installed license or remove the license if necessary.

Product Editions

Veeam Agent for Linux offers three product editions that define product functionality and operation modes:

- *Server* – a commercial edition that provides access to all product functions and is intended for performing data protection tasks on computers that run Linux OS. Veeam Agent can operate in the Server edition if a commercial license that supports this edition is installed on the protected computer.

NOTE

Veeam Agent for Linux on Power requires a Server license.

- *Workstation* – a commercial edition that offers [limited](#) capabilities that are sufficient for performing data protection tasks on computers that run Linux OS. Veeam Agent can operate in the Workstation edition if a commercial license that supports this edition is installed on the protected computer.
- *Free* – a free edition that offers the capabilities of the Workstation edition with certain [limitations](#). Free edition does not come with a commercial support program. In contrast to the Workstation and Server editions, the Free edition does not require a license.

For more information about product editions, pricing and features available for them, see [this Veeam webpage](#).

TIP

To check in which edition Veeam Agent currently operates, you can use the Veeam Agent control panel or command line interface. To learn more, see [Viewing License Information](#).

When you install a license on the protected computer, you can select in which edition Veeam Agent will operate: Server edition or Workstation edition (if both editions are supported by the license). If you use Veeam Agent with Veeam Backup & Replication, you must manage product licenses and editions from the Veeam Backup & Replication console. To learn more, see [Managing License with Veeam Backup & Replication](#).

After the license expires, Veeam Agent automatically switches to the Free edition. To learn more, see [License Expiration](#).

Limitations for Free and Workstation Editions

Compared to the Server edition of Veeam Agent for Linux, Free and Workstation editions have the following limitations:

1. [Free edition] The number of backup jobs that you can configure in Veeam Agent for Linux is limited to one.
2. [Free edition] You cannot use a Veeam Cloud Connect repository as a target location for backup files.
3. [Free edition] You cannot perform direct backup to an object storage repository.
4. [Workstation edition] The number of backup jobs that you can configure in Veeam Agent for Linux is limited to one backup job targeted at a local drive, network shared folder, object storage repository or Veeam backup repository plus unlimited number of backup jobs targeted at a Veeam Cloud Connect repository.
5. [Free and Workstation editions] You cannot specify pre-freeze and post-thaw scripts in the backup job settings.
6. [Free and Workstation editions] You cannot specify database system processing settings.
7. Free and Workstation editions are not supported for Veeam Agent for Linux on Power.

License Agreement

After you install Veeam Agent for Linux, you must accept terms of the product license agreement and license agreements for third-party components operating as part of the product. Until you accept the license agreements, you will not be able to perform backup and data recovery tasks with the Veeam Agent control panel and command line interface.

License agreements are located in the `/usr/share/doc/veeam` directory of the machine where you installed the product.

The process of accepting license agreements differs depending on the way you work with Veeam Agent – using the control panel or command line interface.

- When you launch the Veeam Agent for Linux control panel for the first time, Veeam Agent prompts you to accept the license agreements at the **Agreements** step of the initial setup wizard. To learn more, see [Accept License Agreements](#).
- When you run a Veeam Agent for Linux command, for example, `veeamconfig repository create`, Veeam Agent prompts you to accept license agreements. To accept the license agreement, type `y` or `yes` in the command prompt and press [Enter].

Alternatively, you can accept license agreements using the dedicated commands. To learn more, see [Accepting License Agreements](#).

Installing License

When you launch the Veeam Agent for Linux control panel for the first time, Veeam Agent for Linux offers you to install a license at the [License](#) step of the initial setup wizard. You can choose to install the license immediately or postpone this operation.


If you choose to postpone the license installation process, you can install a license later at any time you need. Until you install a license, Veeam Agent for Linux will operate in the Free edition. To learn more, see [Product Editions](#).

NOTE

If you choose not to install a license and use Veeam Agent for Linux in the Free edition, Veeam Agent for Linux will display a notification offering to install a license every time you open the control panel. The notification will appear in the control panel until Veeam Agent for Linux completes the first backup job session.

To install a license:

1. Launch the Veeam Agent for Linux control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent for Linux control panel, press the [M] key to open the **Miscellaneous** menu.
3. In the menu, make sure that the **Manage License** option is selected and press [Enter].



4. In the **Manage license** window, make sure that the **Install** button is selected and press [Enter].

5. In the **Choose license** window, in the **File location** field, specify a path to the license key:
 - a. Select the **Browse** option with the [Tab] key and press [Space] or [Enter].
 - b. In the **Choose license file location** window, select the necessary directory and press [Enter].
 - c. Repeat the step 'b' until a path to the directory in which the license key resides appears in the **Current directory** field.


TIP

If you chose to install the license immediately from the Veeam Agent for Linux welcome screen notification, you will pass to the **Choose license** step right from the notification window.

6. In the **Choose agent edition to use on this computer** section, select the product edition in which Veeam Agent for Linux will operate and press [Enter]. To learn more about editions, see [Product Editions](#).
7. Veeam Agent for Linux will install the license and display a window notifying that the license is successfully installed. Press [Enter] to finish the license installation process.

TIP

After you install a license, you can view information about the license (expiration date, status of the license, current edition of the product and so on) in the **Manage license** window. You can also check information about the license using the Veeam Agent for Linux command line interface. To learn more, see [Viewing License](#).




Viewing License Information

To view information about the installed license, do the following:

1. Launch the Veeam Agent for Linux control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent for Linux control panel, press the [M] key to open the **Miscellaneous** menu.
3. In the menu, make sure that the **Manage License** option is selected and press [Enter].

Veeam Agent for Linux will display information about the license.



Removing License


You can remove the license if necessary. To remove a license:

1. Launch the Veeam Agent for Linux control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent for Linux control panel, press the [M] key to open the **Miscellaneous** menu.
3. In the menu, make sure that the **Manage License** option is selected and press [Enter].
4. In the **Manage license** window, press [Tab] to select the **Remove** button, then press [Enter].
5. Veeam Agent for Linux will remove the license and display a window notifying that the license is successfully removed. Press [Enter] to finish the license removal process.

NOTE

After you remove the license, Veeam Agent for Linux will continue to operate in the Free edition. Consider the following:

- If Veeam Agent for Linux operated in the Server edition and multiple backup jobs were configured, after switching to the Free edition, all backup jobs will fail.
- If pre-freeze or post-thaw scripts were specified for a backup job, after switching to the Free edition, this backup job will fail.
- If database system processing was set for a backup job, after switching to the Free edition, this backup job will fail.




License Expiration

30 days before the license expiration date, Veeam Agent will display a warning at the top of the control panel. After the license expires, Veeam Agent will switch to the Free edition.

Consider the following:

- If Veeam Agent operated in the Server edition and multiple backup jobs were configured, after switching to the Free edition, all backup jobs will be failing. You will have to configure a new backup job to protect the computer.
- If pre-freeze or post-thaw scripts were specified for a backup job, after switching to the Free edition, this backup job will be failing.
- If database system processing was set for a backup job, after switching to the Free edition, this backup job will be failing.

You can switch to the Free edition manually at any time if necessary. To learn more, see [Removing License](#).



Managing License with Command Line Interface

You can use the Veeam Agent for Linux command line interface to perform the following operations with the license:

- [Accept license agreements for the product itself and its third-party components.](#)
- [Install a license on the protected computer.](#)
- [View information about the license.](#)
- [Remove the license.](#)

Accepting License Agreements

To work with Veeam Agent for Linux, you must accept terms of the product license agreement and license agreements for third-party components operating as part of the product. Until you accept license agreements, you can use the `veeamconfig` utility to run the following commands only:

- `veeamconfig agreement show`
- `veeamconfig help` (or `veeamconfig -h` or `veeamconfig --help`)
- `veeamconfig mode info`
- `veeamconfig mode reset`
- `veeamconfig version` (or `veeamconfig -v` or `veeamconfig --version`)
- `veeamconfig ui`

To accept license agreements, use the following command:

```
veeamconfig agreement accepteula && veeamconfig agreement acceptthirdpartylicenses
```

TIP

To check whether license agreements are accepted, use the following command: `veeamconfig agreement show`.

Installing License

To install a license, use the following command:

```
veeamconfig license install --path <path> --workstation
```

or

```
veeamconfig license install --path <path> --server
```

where:

- `<path>` – path to the license key file in the local file system of your computer.
- `workstation` or `server` – edition in which Veeam Agent will operate. To learn more about editions, see [Product Editions](#).

Veeam Agent for Linux will install the license and display information about the license. You can also view this information later at any time. To learn more, see [Viewing License Information](#).

For example:

```
user@srv01:~$ veeamconfig license install --path /home/user/veeam/license/veeam
.lic --server
License was installed successfully.
License information:
  License source: Local license
  Mode: Server
  Support expiration: 2019/09/20 (649 days left)
  Status: License is valid.
  Issued to: TechCompany
  E-mail: administrators@tech.com
```

TIP

You can also install a license using the Veeam Agent control panel. To learn more, see [Installing License](#).

Viewing License Information

You can view information about the installed license. Use the following command:

```
veeamconfig license show
```

Veeam Agent for Linux will display information about the license. For example:

```
user@srv01:~$ veeamconfig license show
License information:
  License source: Local license
  Mode: Server
  Support expiration: 2019/09/20 (649 days left)
  Status: License is valid.
  Issued to: TechCompany
  E-mail: administrators@tech.com
```


Removing License

You can remove a license with the following command:

```
veeamconfig license remove
```

After you remove the license, Veeam Agent for Linux will continue to operate in the Free edition. Consider the following:

- If Veeam Agent operated in the Server edition and multiple backup jobs were configured, after switching to the Free edition, all backup jobs will fail.
- If pre-freeze or post-thaw scripts were specified for a backup job, after switching to the Free edition, this backup job will fail.
- If database system processing was set for a backup job, after switching to the Free edition, this backup job will fail.

Performing Backup

You can back up your data to protect the entire computer image, individual volumes or folders and files on your computer. To back up your data, you must configure a backup job. Depending on the product edition, Veeam Agent lets you configure one or several backup jobs targeted at the same or different backup repositories.

You can configure a backup job that will automatically back up your data by the defined schedule. You can also start a backup job manually at any time.

Creating Custom Veeam Recovery Media

In addition to the generic Veeam Recovery Media that is available for download at the Veeam website, you can create a custom Veeam Recovery Media. This option may be helpful if your computer uses hardware that requires drivers not included in the generic Veeam Recovery Media. When you create a custom Veeam Recovery Media, Veeam Agent updates the generic Veeam Recovery Media: copies the Linux kernel running on your computer with its currently loaded modules and includes them into the custom recovery image.

Prerequisite Software Requirements

Before you create custom Veeam Recovery Media, check the following prerequisites:

- The Linux system must have the `genisoimage` package installed. For openSUSE and SLES 15 SP1 – 15 SP6 distributions, the Linux system must have the `mkisofs` package installed.
- The Linux system must have the `mksquashfs` and `unsquashfs` utilities installed.
- For custom Veeam Recovery Media with EFI support, the Linux system must have the following packages installed:
 - `xorriso`
 - `isolinux` (or `syslinux`, if the software package repository of your Linux distribution lacks the `isolinux` package)
- For the scenario where you create a custom Veeam Recovery Media using the **Download and patch ISO** option, the Linux system must have the `wget` utility installed.

Considerations and Limitations

Before you create custom Veeam Recovery Media, consider the following:

- The custom recovery image contains an unsigned Linux kernel. As a result, you cannot use it for UEFI systems with Secure Boot enabled.
- If you plan to use live patching to create a custom Veeam Recovery Media, consider the following [limitations](#).
- You cannot create a custom Veeam Recovery Media for Veeam Agent computers that run Linux kernel version earlier than 3.10. For a workaround, see [this Veeam KB article](#).
- You cannot create a custom Veeam Recovery Media for Veeam Agent computers running on IBM Power Systems.

Creating Custom Veeam Recovery Media

You can create a custom Veeam Recovery Media in one of the following ways:

- With the Veeam Agent control panel. You can perform this operation in the following conditions:
 - During the process of initial product setup, at the [Recovery ISO](#) step of the initial setup wizard.
 - Any time you need, in the **Miscellaneous** menu. For details, see [Creating Custom Veeam Recovery Media with Control Panel](#).

- With the Veeam Agent command line interface. For details, see [Creating Custom Veeam Recovery Media with Command Line Interface](#).

If you create a custom Veeam Recovery Media using the command line interface, you can also specify a directory that contains additional drivers that you want to include in the recovery media.


Creating Custom Veeam Recovery Media with Control Panel

To create custom Veeam Recovery Media with the Veeam control panel, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, press the [M] key to open the **Miscellaneous** menu.
3. In the menu, select the **Patch Recovery Media** option and press [Enter].

IMPORTANT

Recovery Media patching is not supported by Veeam Agent for Linux on Power.



4. Press [Tab] and select how you want to create a custom Veeam Recovery Media depending on the location of the generic recovery media ISO file:
 - If you have not downloaded the generic Veeam Recovery Media, make sure that the **Download and patch ISO** option is selected and press [Tab]. If you select this option, Veeam Agent will download the ISO file of the generic Veeam Recovery Media from the Veeam software repository to the directory of your choice on the Veeam Agent computer and use this image to create the custom Veeam Recovery Media.

Veeam Agent downloads the Veeam Recovery Media ISO file depending on the Veeam Agent computer architecture. For details, see [Veeam Recovery Media Versions](#).

- If you want only to download the generic Veeam Recovery Media, select the **Only download ISO** option with the [Down] key and press [Tab]. If you select this option, Veeam Agent will download the ISO file of the generic Veeam Recovery Media from the Veeam software repository to the directory of your choice on the Veeam Agent computer. You can use the downloaded ISO file later to boot your Veeam Agent computer or to create a custom Veeam Recovery Media.

Veeam Agent downloads the Veeam Recovery Media ISO file depending on the Veeam Agent computer architecture. For details, see [Veeam Recovery Media Versions](#).

- If you have already downloaded the generic Veeam Recovery Media to a local directory on the Veeam Agent computer or to a network shared folder, select the **Patch local ISO** option with the [Down] key and press [Tab]. If you select this option, Veeam Agent will use the generic Veeam Recovery Media ISO file to create the custom Veeam Recovery Media.

The name of the generic Veeam Recovery Media ISO file depends on the recovery image version, Veeam Agent computer architecture and the source from which you downloaded the ISO file: from the product download page or Veeam software repository. To learn more, see [Veeam Recovery Media Versions](#).

5. If you selected the **Download and patch ISO** or **Patch local ISO** option, the **EFI system** option is available. If you want to boot the Veeam Recovery Media on EFI-based systems, select the **EFI system** option with the [Tab] key and press [Space].

If you do not enable this option, the custom Veeam Recovery Media will be able to boot on BIOS-based systems only.


6. If you selected the **Patch local ISO** option, in the **Path to local ISO** field, specify a path to the ISO file of the generic Veeam Recovery Media:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].
 - b. In the **Path to ISO** window, select the necessary directory and press [Enter].
 - c. Repeat the step 'b' until a path to the directory in which the recovery media ISO file resides appears in the **Current directory** field.
 - d. In the directory where the recovery media ISO file resides, select the ISO file and press [Enter].

7. Specify a path to the resulting ISO file of the Veeam Recovery Media.

If you selected the **Download and patch ISO** or **Patch local ISO** option, in the **Save patched ISO to** field, you can specify a path to the resulting ISO file of the custom Veeam Recovery Media; if you selected the **Only download ISO** option, in the **Save ISO to** field, specify a path to the resulting ISO file of the generic Veeam Recovery Media:

- a. Select the **Browse** option with the [Tab] key and press [Enter].
- b. In the **Save patched ISO to** window, select the necessary directory and press [Enter].
- c. Repeat step 'b' until a path to the directory where you want to save the resulting custom recovery media ISO file appears in the **Current directory** field.
- d. Select the **OK** button with the [Tab] key and press [Enter].

8. To start the custom recovery media creation process, select the **Next** button with the [Tab] key and press [Enter].



Creating Custom Veeam Recovery Media with Command Line Interface

To create a custom Veeam Recovery Media, you need to perform the following operations:

- [Download the ISO file of the generic Veeam Recovery Media](#). You can download this image from the [Veeam software repository](#) or from [this Veeam webpage](#)..
- [Using the downloaded ISO file, create the Custom Veeam Recovery Media](#).

IMPORTANT

Recovery Media patching is not supported by Veeam Agent for Linux on Power.

Downloading Generic Recovery Media

To download the generic Veeam Recovery Media with the command line interface, use the following command:

```
veeamconfig downloadiso --output <output_path>
```

where:

<output_path> – path to the downloaded ISO file of the generic Veeam Recovery Media.

Veeam Agent downloads the ISO file of the generic Veeam Recovery Media depending on the Veeam Agent computer architecture. For details, see [Veeam Recovery Media Versions](#).

For example:

```
$ veeamconfig downloadiso --output /mnt/veeam/iso
```

Creating Custom Recovery Media

To create a custom Veeam Recovery Media, use the following command:

```
veeamconfig patchiso --input <input_path> --output <output_path> [--copy <additional_path>] [--tempPath <temp_path>][--efi]
```

where:

- <input_path> – path to the ISO file of the generic Veeam Recovery Media.
- <output_path> – path to the resulting ISO file of the custom Veeam Recovery Media.
- <additional_path> – path to a directory with additional drivers that you want to include in the Veeam Recovery Media.

When you boot from the custom Veeam Recovery Media, the content of the directory specified with the <additional_path> parameter will be available in the root folder of the recovery environment.

- `<temp_path>` – path to a directory for storing temporary files that are created during generation of custom Veeam Recovery Media. By default, Veeam Agent stores temporary files in the `/tmp/veeam` directory. Use this option to change the default location – for example, if the `tmp` directory does not have enough space.
- `--efi` – option that defines whether custom Veeam Recovery Media should be able to boot on EFI-based systems. Without this option, the custom Veeam Recovery Media will be able to boot on BIOS-based systems only.

For example:

```
$ veeamconfig patchiso --input /mnt/veeam/iso/veeam-recovery-amd64-6.0.0.iso --  
output /mnt/veeam/iso/veeam-recovery-media-srv01.iso --copy /tmp/template --tem  
ppath /external/veeamtmp --efi
```

Creating Backup Jobs

You can choose one of the following backup modes:

- Backup of an entire computer image
- Backup of specific computer volumes, for example, a system volume or secondary volume
- Backup of individual files and folders

[For Server Edition] You can configure one or several backup jobs to back up your data. Configuring several backup jobs may be useful in the following situations:

- You can configure separate backup jobs for volume-level backup and file-level backup.
- You can configure backup jobs targeted at different backup repositories to keep several copies of your backed-up data at different locations.
- You can configure several backup jobs and define individual schedule for every job to back up necessary data at the desired time.

With Veeam Agent, you can configure the backup job in one of the following ways:

- [With the Backup Job wizard](#)
- [In the command line interface](#)

Before You Begin

Prerequisites

Before you configure the backup job, check the following prerequisites:

- The target location where you plan to store backup files must have enough free space.
- [For Veeam Backup & Replication repository targets] You can store created backups in a backup repository only if the backup server runs Veeam Backup & Replication 12.3.

NOTE

To benefit from all recent enhancements and improvements, make sure to install the latest available build of Veeam Backup & Replication.

- [For Veeam Backup & Replication repository targets] If you plan to use a Veeam Backup & Replication repository as a target for backups, you must pre-configure user access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

Considerations and Limitations

Before you configure the backup job, consider the following:

- Depending on the selected backup mode – entire machine or volume-level backup – during backup job, Veeam Agent behaves differently:
 - During entire machine backup, Veeam Agent backs up all the volumes on all the supported devices that are active in the system at the moment of backup. Inactive or unsupported devices are skipped, the backup job does not fail.
 - During volume-level backup, if the specified volume's device is not supported or is not active in the system at the moment of backup, the backup job will fail.
- When you configure the backup job with the Backup Job wizard, Veeam Agent creates the job with default values for the following advanced settings: compression level and data block size. You can specify custom values for these settings in the command line interface when you create a new backup job or edit an existing backup job. For more information, see [Creating Backup Job in Command Line Interface](#) and [Editing Backup Job Settings](#).
- You cannot save the backup of entire computer on the local computer disk. Use an external hard drive or USB drive, network shared folder or backup repository as a target location.
- [For nosnap Veeam Agent for Linux and Veeam Agent for Linux on Power] When you back up the entire computer or create a volume-level backup, Veeam Agent backs up the boot partition in the snapshot-less mode. To ensure consistency of the boot partition data during backup, Veeam Agent uses the `fsfreeze` or `umount` commands.
- Veeam Agent automatically stops running the backup job session if it has not finished after 21 days (504 hours).
- Veeam Agent does not back up data to which symbolic links are targeted. It only backs up the path information that the symbolic links contain. After restore, identical symbolic links are created in the restore destination.
- Veeam Agent does not support backup of bind mount points. In the scope of the backup job, you must specify the path to the original mount point instead.

Navigating Backup Job Wizard

The Backup Job wizard window comprises the following areas:

- The navigation pane, located on the left of the window, displays the list of wizard steps and currently selected step of the wizard.
- The working area displays controls relating to a specific step of the wizard.
- The buttons area, located at the bottom of the window, displays buttons that you can use to switch between steps of the wizard (**Previous** and **Next**) and close the wizard (**Cancel** and **Finish**).

In the Backup Job wizard, the use of a mouse is not supported. To navigate the Backup Job wizard and associated dialog windows, you can use the following keys:

- [Tab] – to switch between displayed controls in the working area and buttons in the buttons area. The currently selected control or button is highlighted with a green color.
- [Up] and [Down] – to switch between items in a scrollable list.
- [Space] – to select the necessary item in a list. The selected item's mark may vary in different steps of the wizard.
- [Enter] – to proceed to the next step of the wizard or to open a directory.
- [Backspace] – to return to the previous step of a wizard.
- [Esc] – to cancel the backup job configuration and exit the wizard.

TIP

You can switch between steps of the Backup Job wizard in two ways. The easier and more comfortable way is to use the [Enter] key to proceed to the next step and [Backspace] key to return to the previous step of the wizard. You can also use the [Tab] key to select the **Next** or **Previous** button in the buttons area and then press [Enter] to switch to the next or previous step of the wizard respectively.

Creating Backup Job with Backup Job Wizard

You can configure volume-level and file-level backup jobs with the Backup Job wizard.

Step 1. Launch Backup Job Wizard

To launch the **Backup Job** wizard, do the following:


1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. If you have not configured any backup jobs yet, Veeam Agent will display a welcome screen. Press the [C] key to proceed to the Backup Job wizard and configure the backup job.
3. If you have already configured and performed a backup job, Veeam Agent will display the list of backup job sessions. When you press the [C] key to launch the Backup Job wizard, Veeam Agent will display a list of configured backup jobs. To configure a new backup job, select the **Configure new job** option and press [Enter].

NOTE

The **Configure new job** option is not available if Veeam Agent for Linux operates in the Free edition and you have already configured one backup job.

To edit settings of a backup job that you have already configured, select the job in the list and press [Enter]. To learn more, see [Editing Backup Job Settings](#).

If you have decided not to create a backup job, press [Esc] to close the list of backup jobs and return to the welcome screen. After that, you can press [Esc] once again to return to the command line interface.




Step 2. Specify Backup Job Name

At the **Name** step of the wizard, in the **Job name** field, type the name for the backup job and press [Enter].

TIP

To proceed to the next step of the wizard, you can also select the **Next** button with the [Tab] key and then press [Enter].



Step 3. Select Backup Mode

At the **Backup mode** step of the wizard, select the mode in which you want to create a backup:

1. Select the necessary backup mode. You can select one of the following options:
 - **Entire machine** – select this option if you want to create a backup of the entire computer image. When you restore data from such backup, you will be able to recover the entire computer image as well as data on specific computer volumes: files, folders, application data and so on. With this option selected, you will pass to the [Destination](#) step of the wizard.
 - **Volume level backup** – select this option if you want to create a backup of specific computer volumes, for example, the system volume. When you restore data from such backup, you will be able to recover data on these volumes only: files, folders, application data and so on. With this option selected, you will pass to the [Volumes](#) step of the wizard.
 - **File level backup** – select this option if you want to create a backup of individual directories on your computer. With this option selected, you will pass to the [Files](#) step of the wizard.
2. [For file-level backup] If you want to perform backup in the snapshot-less mode, select **Disable snapshot**. With this option selected, Veeam Agent will not create a snapshot of the backed-up volumes during backup. This allows Veeam Agent to back up data residing in file systems that are not supported for snapshot-based backup with Veeam Agent. To learn more, see [Snapshot-Less File-Level Backup](#).


IMPORTANT

Consider the following:

- [For entire machine backup] Certain limitations for Dell PowerPath configuration apply. To learn more, see [this Veeam KB article](#).
- [For volume-level backup] Volume-level backup job relies on a device name in the `/dev` directory. Device names in the `/dev` directory (for example, `/dev/md-127`, `/dev/dm-1`) must stay persistent for backed-up volumes. Otherwise, the job will back up the wrong volume.
- [For file-level backup] If the backed-up file system has a complex folder structure with many hierarchy levels, during incremental backup, the inbound network traffic on the Veeam Agent computer may exceed by far the outbound traffic. Significant amount of data can be transferred to the Veeam Agent computer from the target backup location even if few files are changed since the previous job session.

TIP

File-level backup is typically slower than volume-level backup. If you plan to back up all folders with files on a specific volume, it is recommended that you configure volume-level backup instead of file-level backup.



Step 4. Specify Backup Scope Settings

Specify backup scope for the backup job:

- [Select volumes to back up](#) — if you have selected the **Volume level backup** option at the [Backup Mode](#) step of the wizard.
- [Select folders to back up](#) — if you have selected the **File level backup** option at the [Backup Mode](#) step of the wizard.

Selecting Volumes to Back Up

The **Volumes** step of the wizard is available if you have chosen to create a volume-level backup.

At this step of the wizard, you must specify the backup scope — define what volumes you want to include in the backup. Veeam Agent lets you include the following types of objects in the volume-level backup:

- [Block devices \(entire disks and individual volumes\)](#)
- [Mount points](#)
- [LVM logical volumes and volume groups](#)
- [BTRFS storage pools and subvolumes](#)

Selecting Devices

To add a block device to the backup scope, do the following:


1. At the **Volumes** step of the wizard, make sure that the **Device** option is selected and press [Enter].
2. In the **Add devices to scope** window, select individual volumes or entire computer disks that you want to include in the backup and press [Enter].
 - To include individual volumes of your computer in the backup, select block devices that represent volumes that you want to back up, for example: *sda1* or *sda6*.
 - To include all volumes on a computer disk in the backup, select block devices that represent disks whose volumes you want to back up, for example: *sda* or *sdb*. All volumes on the selected disk will be automatically selected, too.

To navigate the list of volumes and select the necessary items, use the [Up], [Down] and [Space] keys. To learn more, see [Navigating Backup Job Wizard](#).

If you have created several system partitions, for example, a separate partition for the `/boot` directory, you should remember to include all of these partitions in the backup. Otherwise, Veeam Agent does not guarantee that the OS will boot properly when you attempt to recover from such backup.

NOTE

If you include a block device in the backup, and this block device is a physical volume assigned to an LVM volume group, Veeam Agent will include the whole LVM volume group in the backup.



Selecting Mount Points

IMPORTANT


Veeam Agent does not support backup of bind mount points. You must specify the path to the original mount point instead.

To add a mount point to the backup scope, do the following:

1. At the **Volume** step of the wizard, select the **Mountpoint** option and press [Enter].

2. In the **Add mountpoints to scope** window, select mount points that you want to include in the backup and press [Enter].

To navigate the list of mount points and select the necessary mount points, use [Up], [Down] and [Space] keys. To learn more, see [Navigating Backup Job Wizard](#).



Selecting LVM Volumes

To add an LVM logical volume or volume group to the backup scope, do the following:


1. At the **Volume** step of the wizard, select the **LVM** option and press [Enter].
2. In the **Add LVM to scope** window, select LVM logical volumes or volume groups that you want to include in the backup and press [Enter].

To navigate the list of LVM volumes and select the necessary items, use [Up], [Down] and [Space] keys. To learn more, see [Navigating Backup Job Wizard](#).

If you include an LVM volume group in the backup, all LVM logical volumes in the selected volume group will be automatically selected, too.

NOTE

Veeam Agent does not back up LVM snapshots.



Selecting BTRFS Volumes

To add a BTRFS storage pool or subvolume to the backup scope, do the following:


1. At the **Volume** step of the wizard, select the **BTRFS** option and press [Enter].
2. In the **Add BTRFS to scope** window, select BTRFS storage pools or subvolumes that you want to include in the backup and press [Enter].

To navigate the list of BTRFS pools and subvolumes and select the necessary items, use [Up], [Down] and [Space] keys. To learn more, see [Navigating Backup Job Wizard](#).

Veeam Agent identifies BTRFS storage pools by UUIDs. If you include a BTRFS pool in the backup, all BTRFS subvolumes in the selected pool will be automatically selected, too.

NOTE

You cannot add read-only BTRFS snapshots to the backup scope.



Selecting Files and Directories to Back Up

The **Files** step of the wizard is available if you have chosen to create a file-level backup.

At this step of the wizard, you must specify the backup scope – define what directories with files you want to include in the backup.

In the file-level backup mode, you must include in the backup at least one directory. If you do not want to back up some subdirectories of the specified directory, you can exclude these directories from the backup.

You can also include or exclude files of a specific type in/from the backup. You can specify file names explicitly or use UNIX wildcard characters to define file name masks. Veeam Agent will apply the specified file name masks to files in directories that are included in the backup.

To specify the backup scope:


1. At the **Files** step of the wizard, make sure that the **Add directories** option is selected and press [Enter].
2. In the **Choose directories** window, select one or several directories that you want to include in the file-level backup.

IMPORTANT

Veeam Agent does not support backup of bind mount points. You must specify the path to the original mount point instead.


- To navigate the list of directories, use the [Up] and [Down] keys.
- To browse for subdirectories, navigate to the necessary directory and press [Enter].

- To include a directory in the backup, navigate to the necessary directory and press [Space]. The included directory will be marked with the '+' character. All subdirectories of the selected directory will be included in the backup too.



3. Specify directories that you want to exclude from the file-level backup. To exclude a directory:
 - a. Browse for subdirectories of a directory that you have included in the backup.

- b. Navigate to the directory that you want to exclude from the backup and press [Space]. The excluded directory will not be marked with the '+' character.




4. Switch to the **OK** button and press [Enter]. Veeam Agent will display a list of paths to the selected directories and the number of excluded subdirectories for each directory in the list.
5. Specify file name masks for files that you want to include or exclude in/from the backup:
 - a. Select the **File Masks** option with the [Tab] key and press [Enter].
 - b. In the **File masks** window, make sure that the **Create Mask** button is selected and press [Enter].
 - c. In the **Mask** field, enter the file name mask, for example, `report.pdf`, `*filename*` or `*.odt`.
 - d. In the **Type** field, select one of the following options:
 - **Exclude**— if you do not want to back up files whose names match the specified mask. Veeam Agent will back up all files in the directories selected for backup except for such files.
 - **Include**— if you want to back up files whose names match the specified mask. Veeam Agent will create a backup only for such files in the directories selected for backup.

You can use a combination of include and exclude masks. Keep in mind that exclude masks have a higher priority than include masks. For example, you can specify masks in the following way:

- Include mask: `report*.*`

- Exclude mask: *.odt


Veeam Agent will include in the backup all files whose name begins with `report` except for the files of the ODT format.



- Press [Enter]. Veeam Agent will display in the **File masks** window the specified file mask and its type: *Include* or *Exclude*.
- Repeat Steps **b** to **e** for each mask that you want to specify.
- After you specify all file masks, switch to the **OK** button and press [Enter].

TIP

To remove a file name mask, in the **File masks** window, select the necessary mask and press [Delete].




Step 5. Select Backup Destination

At the **Destination** step of the wizard, select a target location for the created backup.

You can select one of the following options:

- **Local** – select this option if you want to save the backup in a removable storage device attached to the computer or on a local computer drive. With this option selected, you will pass to the [Location](#) step of the wizard.
- **Object storage** – select this option if you want to create the backup in an object storage exposed to you by cloud service provider. With this option selected, you will pass to the [Storage](#) step of the wizard.
- **Shared Folder** – select this option if you want to save the backup in a network shared folder. With this option selected, you will pass to the [Network](#) step of the wizard.
- **Veeam Backup & Replication** – select this option if you want to save the backup in a backup repository managed by the Veeam backup server. With this option selected, you will pass to the [Veeam](#) step of the wizard.
- **Veeam Cloud Connect repository** – select this option if you want to create the backup in a cloud repository exposed to you by the Veeam Cloud Connect service provider. With this option selected, you will pass to the [Service Provider](#) step of the wizard.

It is recommended that you store backups in the external location like USB storage device or network shared folder. You can also keep your backup files on the separate non-system local drive.



Step 6. Specify Backup Storage Settings

Specify backup storage settings for the backup job:

- [Local storage settings](#) – if you have selected the **Local storage** option at the [Destination](#) step of the wizard.
- [Object storage settings](#) – if you have selected the **Object storage** option at the [Destination](#) step of the wizard.
- [Shared folder settings](#) – if you have selected the **Shared folder** option at the [Destination](#) step of the wizard.
- [Veeam backup repository settings](#) – if you have selected the **Veeam backup repository** option at the [Destination](#) step of the wizard.
- [Veeam Cloud Connect repository settings](#) – if you have selected the **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.

NOTE

The **Veeam Cloud Connect repository** option is available if Veeam Agent operates in the Workstation or Server edition.

Local Storage Settings

The **Location** step of the wizard is available if you have selected the **Local** option at the [Destination](#) step of the wizard. Specify location for the backup file and retention policy for the backup job:

1. To specify location for the backup file, browse to the directory where backup files must be saved:
 - a. Select the **Browse** option with the [Tab] key and press [Space] or [Enter].
 - b. In the **Choose backup location** window, select the necessary directory and press [Enter].
 - c. Repeat the step 'b' until a path to the directory in which you want to save backup files appears in the **Current directory** field.
 - d. To create a new directory, switch to the **Create Dir** button, press [Enter], then type a name for the new directory and press [Enter].
 - e. Switch to the **OK** button and press [Enter]. Veeam Agent will display the path to the specified directory in the **Location** field.

Alternatively, you can type a path to the directory in which you want to save backup files in the **Location** field.

After you specify location for the backup, Veeam Agent will display the following information on the volume where the directory selected for backup storage resides:

- **Space** – total size of the volume on which the selected directory resides.
 - **Free** – free space on the volume where the selected directory resides.
 - **Type** – file system type of the volume on which the selected directory resides.
2. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).

3. In the **Restore points** field, specify the number of backup files that you want to keep in the target location. By default, Veeam Agent keeps 7 latest backup files. When the number of restore points is exceeded, Veeam Agent for Linux will remove the earliest restore point from the backup chain.


To learn more, see the [Short-Term Retention Policy](#).

4. Select **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

IMPORTANT

Consider the following:

- The backup location must reside on a separate volume from a volume whose data you plan to back up.
- USB storage devices formatted as FAT32 do not allow storing files larger than 4 GB in size. For this reason, it is recommended that you do not use such USB storage devices as a backup target.



Object Storage Settings

The **Cloud Type** step of the wizard is available if you have selected the **Object storage** option at the [Destination](#) step of the wizard.

At the **Storage** step of the wizard, select the object storage. You can select one of the following options:


- **S3 compatible** – select this option if you want to create a backup in the S3 compatible storage. With this option selected, you will pass to the [Account](#) step of the wizard.

TIP

If you want to store backups in IBM Cloud or Wasabi Cloud, use the **S3 Compatible** option.

- **Amazon S3** – select this option if you want to create a backup in the Amazon S3 storage. With this option selected, you will pass to the [Account](#) step of the wizard.

- **Google Cloud storage** – select this option if you want to create a backup in the Google Cloud storage. With this option selected, you will pass to the [Account](#) step of the wizard.
- **Microsoft Azure Blob storage** – select this option if you want to create a backup in the Microsoft Azure storage. With this option selected, you will pass to the [Account](#) step of the wizard.



S3 Compatible Settings

If you have selected to store backup files in the S3 compatible storage, specify the following settings:

1. [Account settings](#).
2. [Bucket settings](#).

Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in the S3 compatible storage.

NOTE

You can store backups only in the S3 compatible storage repositories that are accessible over the HTTPs protocol.

To connect to the S3 compatible storage, specify the following:

1. In the **Service point** field, specify the address of your S3 compatible storage.

If you want to connect to the repository using an IPv6 address and port number, you must use the following format: `IPv6:port`, where:

- `IPv6` is the IPv6 address of the object storage.
- `port` is the number of the port that Veeam Agent will use to connect to the object storage.

For example: `https://2001-db8-1--bb8-c0b8-112.ipv6-literal.net:9000`


If you want to connect to the repository using AWS PrivateLink for Amazon S3, use the following format: `bucket.<VPC_endpoint_ID>.s3.<region>.vpce.amazonaws.com`, where:

- o `VPC_endpoint_ID` is the unique VPC endpoint ID that Veeam Agent will use to access the cloud storage.
- o `region` is the AWS Region where the VPC endpoint and the S3 bucket are located.

For example: `bucket.vpce-0abc123def456ghi7-01234567.s3.ap-northeast-1.vpce.amazonaws.com`

To learn more about AWS PrivateLink for Amazon S3, see [AWS Documentation](#).

2. In the **Region** field, specify the storage region based on your regulatory and compliance requirements.
3. In the **Access key** field, enter the access key ID.
4. In the **Secret key** field, enter the secret access key.



Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to save backup files in the S3 compatible storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

5. In the **Bucket** field, specify a bucket in the storage:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].
 - b. In the **Specify Bucket** window, select the necessary bucket and press [Enter].
6. In the **Folder** field, specify a folder in the bucket:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].
 - b. In the **Specify Folder** window, select the necessary folder and press [Enter].

TIP

You can also create a new folder. To do this, type a name for the new folder in the **Folder** field.

7. To prohibit modification and deletion of blocks of data in the object storage repository, select the **Make recent backups immutable for** check box and specify the immutability period in days. For more information, see [Backup Immutability](#).
8. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).


NOTE

If you use the GFS retention scheme and enable immutability for the backup, the restore points with GFS flags will become immutable for the whole GFS retention period. You will not be able to delete such restore points until the GFS retention period is over.

5. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

6. Select **Advanced** to specify additional backup job settings. For details, see [Specify Advanced Backup Settings](#).



After that, Veeam Agent will create a new repository in the object storage where you can store backups.

Amazon S3 Settings

If you have selected to store backup files in the Amazon S3 storage, specify the following settings:

1. [Account settings](#).


2. [Bucket settings](#).

Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in the Amazon S3 storage.

To connect to the Amazon S3 storage, specify the following:

1. In the **Access key** field, enter the access key ID.
2. In the **Secret key** field, enter the secret access key.
3. In the **AWS region** window, select the AWS region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Global** region. Switch to the **Ok** button and press [Enter].



Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to save backup files in the Amazon S3 storage and specified account settings to connect to the storage.

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. For example, it is not recommended that you use dots (.) in the bucket name. For more information on bucket naming rules, see [this AWS documentation article](#).

Specify settings for the bucket in the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups. Switch to the **Ok** button and press [Enter].
2. In the **Bucket** field, specify a bucket in the storage:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].

- b. In the **Specify Bucket** window, select the necessary bucket and press [Enter].
3. In the **Folder** field, specify the folder in the bucket:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].
 - b. In the **Specify Folder** window, select the necessary folder and press [Enter].

TIP

You can also create a new folder. To do this, type a name for the new folder in the **Folder** field.

4. To prohibit modification and deletion of blocks of data in the object storage repository, select the **Make recent backups immutable for** check box and specify the immutability period in days. For more information, see [Backup Immutability](#).
5. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).


NOTE

If you use the GFS retention scheme and enable immutability for the backup, the restore points with GFS flags will become immutable for the whole GFS retention period. You will not be able to delete such restore points until the GFS retention period is over.

6. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

7. Select **Advanced** to specify additional backup job settings. For details, see [Specify Advanced Backup Settings](#).



After that, Veeam Agent will create a new repository in the object storage where you can store backups.

Google Cloud Storage Settings

If you have selected to store backup files in a Google Cloud storage repository, specify the following settings:


1. [Account settings](#).
2. [Bucket settings](#).

Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in the Google Cloud storage.

To connect to the Google Cloud storage, in the **Access Key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) key associated with the Google Cloud account. Veeam Agent will use the HMAC key to authenticate requests to the Google Cloud storage. For more information about the Google Cloud accounts, see the [Google Cloud documentation](#).

If you have not created the HMAC key beforehand, you can create the key in the Google Cloud console, as described in [this Google Cloud documentation article](#).



Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to save backup files in the Google Cloud storage and specified account settings to connect to the storage.

Specify settings for the bucket in the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups. Switch to the **Ok** button and press [Enter].
2. In the **Bucket** field, specify the bucket in the storage:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].
 - b. In the **Specify Bucket** window, select the necessary bucket and press [Enter].


3. In the **Folder** field, specify the folder in the bucket:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].
 - b. In the **Specify Folder** window, select the necessary folder and press [Enter].

TIP

You can also create a new folder. To do this, type a name for the new folder in the **Folder** field.

4. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).
5. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).
6. Select **Advanced** to specify additional backup job settings. For details, see [Specify Advanced Backup Settings](#).



After that, Veeam Agent will create a new repository in the object storage where you can store backups.

Microsoft Azure Storage Settings

If you have selected to store backup files in the Microsoft Azure storage, specify settings to connect to the storage and container in this storage:

1. [Account settings](#).
2. [Container settings](#).


Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to save backup files in the Microsoft Azure storage.

NOTE

The **Allow storage account key access** option for Shared Key authorization must be enabled in the storage account. To learn how to find this option, see [this Microsoft Docs article](#).

1. In the **Account** field, enter the storage account name.
2. In the **Shared key** field, enter the storage account shared key.
3. In the **Azure Region** window, select the Microsoft Azure region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Azure Global (Standard)** region. Switch to the **Ok** button and press [Enter].



Specifying Container Settings

The **Container** step of the wizard is available if you have chosen to save backup files in the Microsoft Azure storage and specified account settings to connect to the storage.

Specify settings for the container in the storage:

1. In the **Container** field, specify the container in the storage:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].
 - b. In the **Specify Azure Container** window, select the necessary container and press [Enter].
2. In the **Folder** field, specify the folder in the container:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].

- b. In the **Specify Folder** window, select the necessary folder and press [Enter].

TIP

You can also create a new folder. To do this, type a name for the new folder in the **Folder** field.

3. To prohibit modification and deletion of blocks of data in the object storage repository, select the **Make recent backups immutable for** check box and specify the immutability period in days. For more information, see [Backup Immutability](#).
4. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).


NOTE

If you use the GFS retention scheme and enable immutability for the backup, the restore points with GFS flags will become immutable for the whole GFS retention period. You will not be able to delete such restore points until the GFS retention period is over.

5. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

6. Select **Advanced** to specify additional backup job settings. For details, see [Specify Advanced Backup Settings](#).



After that, Veeam Agent will create a new repository in the object storage where you can store backups.

Shared Folder Settings

The **Network** step of the wizard is available if you have selected the **Shared Folder** option at the [Destination](#) step of the wizard.

To save backup files in a remote network location, Veeam Agent mounts to the local file system of your computer the network shared folder that you specify as a location for the backup. When you specify the network shared folder settings, Veeam Agent saves information about the network shared folder and its mount point in the database.

You do not need to mount the network shared folder in advance before every backup job run. Veeam Agent will do it automatically when the backup job is started manually or upon schedule.

After the backup job completes, Veeam Agent will automatically unmount the network shared folder.

Specify shared folder settings:

1. Select the type of a network shared folder:
 - **NFS** – to connect to a network shared folder using the NFS protocol.
 - **SMB** – to connect to a network shared folder using the SMB (CIFS) protocol.
2. In the **Server** field, type the IP address or domain name of the server.
3. In the **Folder** field, type the name of the network shared folder in which you want to store backup files.
Every time the backup job starts, Veeam Agent will automatically mount the specified network shared folder to the `/tmp/veeam` directory in the computer file system. After the backup job completes, Veeam Agent will unmount the network shared folder.
4. [For SMB network shared folder] In the **Domain** field, type a name of the domain in which the account that has access permissions on the shared folder is registered, for example: *DOMAIN*.
5. [For SMB network shared folder] In the **Username** field, type a name of the account that has access permissions on the shared folder.
6. [For SMB network shared folder] In the **Password** field, type a password of the account that has access permissions on the shared folder.
7. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. To learn more, see [Specify GFS Retention Settings](#).
8. In the **Restore points** field, specify the number of backup files that you want to keep in the target location. By default, Veeam Agent keeps 7 latest backup files. When the number of restore points is exceeded, Veeam Agent will remove the earliest restore point from the backup chain.

To learn more, see the [Short-Term Retention Policy](#).

9. Select **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).

The screenshot shows the 'Specify a network location to backup to' dialog box in the Veeam Agent for Linux configuration utility. The window title is 'Veeam Agent for Linux [srv01]'. On the left is a sidebar with a tree view containing 'Name', 'Backup mode', 'Files', 'Destination', 'Network' (selected), 'Schedule', and 'Summary'. The main area contains the following fields and options:

- Name:** () NFS, (X) SMB
- Backup mode:** () NFS, (X) SMB
- Files:**
- Destination:** Server: 172.25.164.25
- Network:** Folder: VeeamBackups
- Schedule:** Domain: tech Username: richard.olson
- Summary:** Password: [masked]
- ☒ [X] Keep full backups for archival purposes [Configure]
Make sure that periodic fulls are enabled in Advanced
- Restore points: 7 [Advanced]

At the bottom right are buttons: [Prev], [Next] (highlighted in green), and [Cancel]. At the bottom left are 'Esc' and 'Cancel' (highlighted in green). At the bottom right are 'Enter' and 'Ok' (highlighted in green).

Veeam Backup Repository Settings

If you have selected to store backup files on a Veeam Backup & Replication repository, specify settings to connect to the backup repository:

1. [Specify backup server settings](#).
2. [Select the Veeam backup repository](#).

Specifying Backup Server Settings

The **Veeam** step of the wizard is available if you have chosen to store backup files on a Veeam Backup & Replication repository.

Specify settings for the Veeam backup server that manages the target backup repository:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent must communicate with the backup repository. By default, Veeam Agent uses port 10006.
3. In the **Login** field, type a name of the account that has access to the Veeam backup repository.
4. In the **Domain** field, type a name of the domain in which the account that has access to the Veeam backup repository is registered, for example: *DOMAIN*.

5. In the **Password** field, type a password of the account that has access to the Veeam backup repository.

Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

The screenshot shows a terminal window titled "Veeam Agent for Linux [srv01]". Inside, a configuration window titled "Specify a Veeam Backup & Replication server to backup to" is displayed. The window has a sidebar on the left with a tree view containing: Name, Backup mode, Destination, > Veeam (selected), Repository, Schedule, and Summary. The main area contains the following fields: Address: 172.24.31.136, Port: 10006, Login: Administrator, Domain: (empty), and Password: (masked with asterisks). At the bottom right of the window are buttons for [Prev], [Next] (highlighted in green), and [Cancel]. Below the window, at the bottom of the terminal, are keyboard shortcuts: Enter Next, Backspace Back, and Esc Cancel.

Selecting Backup Repository


The **Repository** step of the wizard is available if you have chosen to save backup files on a Veeam Backup & Replication repository.

Specify settings for the target backup repository:

1. From the list of available backup repositories, select a backup repository where you want to store backups. The list of backup repositories displays only those repositories on which you have permissions to store data. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
2. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).
3. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

4. Select **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).



Veeam Cloud Connect Repository Settings

If you have selected to store backup files on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. [Specify service provider settings](#).
2. [Verify the TLS certificate and specify user account settings](#).
3. [Select the cloud repository](#).

NOTE

The **Veeam Cloud Connect repository** option is available if Veeam Agent operates in the Workstation or Server edition.


Specifying Service Provider Settings

The **Provider** step of the wizard is available if you have chosen to save backup files on a Veeam Cloud Connect repository.

Specify settings for the cloud gateway that the Veeam Cloud Connect service provider (SP) or your backup administrator has provided to you:

1. In the **Address** field, enter a full DNS name or IP address of the cloud gateway.

2. In the **Port** field, specify the port over which Veeam Agent will communicate with the cloud gateway. By default, port 6180 is used.



Specifying User Account Settings

The **Credentials** step of the wizard is available if you have chosen to save backup files in a cloud repository and specified settings for the cloud gateway.

Verify TLS certificate settings and specify settings for the tenant account or subtenant account that you want to use to connect to the cloud repository.

1. In the **Certificate details** window, review information about the TLS certificate obtained from the SP side and verify the TLS certificate:
 - [Optional] To verify the TLS certificate with a thumbprint, do the following:
 - i. Select the **Verify thumbprint** button with the [Tab] key and press [Enter].
 - ii. Copy the thumbprint you obtained from the SP to the Clipboard and enter it to the **Thumbprint verification** field.
 - iii. Switch to the **Verify** button and press [Enter]. Veeam Agent will check if the thumbprint you entered matches the thumbprint of the obtained TLS certificate.

TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.

 - To accept the TLS certificate, select the **Accept** button with the [Tab] key and press [Enter].
2. In the **Username** field, enter the name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT|SUBTENANT* format.

3. In the **Password** field, provide a password for the tenant or subtenant account.

The screenshot shows a terminal window titled "Veeam Agent for Linux [srv01]". Inside, a configuration wizard is running. The left sidebar lists the steps: Name, Backup mode, Files, Destination, Provider, > Credentials (highlighted), Repository, Schedule, and Summary. The main area is titled "Specify credentials, and validate the certificate." and contains a "Username:" field with the value "ABC_Company\User_02" and a "Password:" field with masked characters "*****". Below these fields is a "[View certificate]" link. At the bottom of the wizard are three buttons: "[Prev]", "[Next]" (highlighted in green), and "[Cancel]". At the very bottom of the terminal window, there are keyboard shortcuts: "Enter [Next]", "Backspace [Back]", and "Esc [Cancel]".

Selecting Cloud Repository


The **Repository** step of the wizard is available if you have chosen to save backup files on a cloud repository and specified settings to connect to the SP.

Specify settings for the cloud repository:

1. From the **Repository** list, select a cloud repository where you want to store created backups. The **Repository** list displays only those cloud repositories that can be accessed by the tenant or subtenant account that you use to connect to the service provider.
2. To use the GFS (Grandfather-Father-Son) retention scheme, select the **Keep full backups for archival purposes** check box and select **Configure**. In the **Configure GFS** window, specify how weekly, monthly and yearly full backups must be retained. For details, see [Specify GFS Retention Settings](#).
3. In the **Restore points** field, specify the number of restore points that you want to store in the target location. By default, Veeam Agent keeps 7 latest restore points. After this number is exceeded, Veeam Agent will remove the earliest restore points from the backup chain.

To learn more, see [Short-Term Retention Policy](#).

4. Select **Advanced** to specify advanced settings for the backup job. To learn more, see [Specify Advanced Backup Settings](#).



Step 7. Specify GFS Retention Settings

This step of the wizard is available if you have chosen to use a long-term, or Grandfather-Father-Son (GFS), retention policy.

To configure GFS retention policy settings for the backup job:

1. Select the **Keep full backups for archival purposes** option and click **Configure** at one of the following steps of the wizard:
 - **Location** — if you have selected the **Local storage** option at the **Destination** step of the wizard.
 - **Network** — if you have selected the **Shared folder** option at the **Destination** step of the wizard.
 - **Repository** — if you have selected the **Veeam backup repository** option at the **Destination** step of the wizard.
 - **Repository** — if you have selected the **Veeam Cloud Connect repository** option at the **Destination** step of the wizard.
 - **Bucket** — if you have selected the **Object storage** option at the **Destination** step of the wizard, then selected the **S3 compatible** option at the **Storage** step of the wizard.
 - **Bucket** — if you have selected the **Object storage** option at the **Destination** step of the wizard, then selected the **Amazon S3** option at the **Storage** step of the wizard.
 - **Bucket** — if you have selected the **Object storage** option at the **Destination** step of the wizard, then selected the **Google Cloud storage** option at the **Storage** step of the wizard.
 - **Container** — if you have selected the **Object storage** option at the **Destination** step of the wizard, then selected the **Microsoft Azure Blob storage** option at the **Storage** step of the wizard.
2. In the **Configure GFS** window, do the following:
 - a. If you want to create weekly restore points for archival purposes, 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.


In the **If multiple full backups exist, use the one from** list, select a week day when Veeam Agent must assign the weekly GFS flag to a full restore point.
 - b. If you want to create monthly restore points for archival purposes, 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.

In the **Use weekly full backup for the following week of a month** list, select a week when Veeam Agent 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.
 - c. If you want to create yearly restore points for archival purposes, 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.

In the **Use monthly full backup for the following month** list, select a month when Veeam Agent must assign the yearly GFS flag to a full restore point.

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 [Assignment of GFS Flags](#) section in the Veeam Backup & Replication User Guide.
- To use a GFS retention policy, you must set Veeam Agent to create full backups. To learn more, see [Active Full Backup Settings](#).



Step 8. Specify Advanced Backup Settings

To configure advanced settings for the backup job, select **Advanced** at one of the following steps of the wizard:

- [Location](#) — if you have selected the Local storage option at the [Destination](#) step of the wizard.
- [Network](#) — if you have selected the Shared folder option at the [Destination](#) step of the wizard.
- [Repository](#) — if you have selected the Veeam backup repository option at the [Destination](#) step of the wizard.
- [Repository](#) — if you have selected the Veeam Cloud Connect repository option at the [Destination](#) step of the wizard.
- [Bucket](#) — if you have selected the Object storage option at the [Destination](#) step of the wizard, then selected the S3 compatible option at the Storage step of the wizard.
- [Bucket](#) — if you have selected the Object storage option at the [Destination](#) step of the wizard, then selected the Amazon S3 option at the Storage step of the wizard.
- [Bucket](#) — if you have selected the Object storage option at the [Destination](#) step of the wizard, then selected the Google Cloud storage option at the Storage step of the wizard.
- [Container](#) — if you have selected the Object storage option at the [Destination](#) step of the wizard, then selected the Microsoft Azure Blob storage option at the Storage step of the wizard.


In the **Specify advanced job settings** window, specify advanced settings for the backup job:

- [Data encryption settings](#)
- [File indexing settings](#)
- [Oracle database system processing settings](#)
- [MySQL database system processing settings](#)
- [PostgreSQL database system processing settings](#)
- [Active full backup settings](#)
- [Backup maintenance settings](#)
- [Script settings](#)
- [Health check settings](#)

NOTE

Consider the following:

- You cannot specify encryption settings for the backup job if you have chosen to save backup files on a Veeam backup repository. Encryption options for Veeam Agent backup jobs targeted at the backup repository are managed by a backup administrator working with Veeam Backup & Replication. To learn more about data encryption capabilities available in Veeam Backup & Replication, see the [Data Encryption](#) section in the Veeam Backup & Replication User Guide.
- You can specify file indexing settings only if Veeam Agent operates in the Workstation or Server edition.
- Nosnap Veeam Agent for Linux and nosnap Veeam Agent for Linux on Power do not support application-aware processing and cannot be used to back up database systems.
- You can specify settings for Oracle, MySQL or PostgreSQL database system processing only if Veeam Agent operates in the Server edition. The settings are available for a volume-level backup job only.
- You can specify backup maintenance settings only if you have selected the **Veeam backup repository** or **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.
- You can specify backup health check settings only if you have selected the **Object storage repository** option at the [Destination](#) step of the wizard.
- You cannot specify data compression settings when you configure a backup job with the Backup Job wizard. If you want to specify these settings, consider creating the backup job with the Veeam Agent command line interface. To learn more, see [Advanced Backup Job Settings](#).



Data Encryption Settings


If you want to encrypt the content of backup files, specify data encryption settings for the backup job:

1. In the **Specify advanced job settings** window, select the **Backup encryption** option with the [Tab] key and press [Enter].

NOTE

The **Backup encryption** option is unavailable if you have chosen to save backup files on a Veeam backup repository. Encryption options for Veeam Agent backup jobs targeted at the backup repository are managed by a backup administrator working with Veeam Backup & Replication. To learn more about data encryption capabilities available in Veeam Backup & Replication, see the [Data Encryption](#) section in the Veeam Backup & Replication User Guide.

2. In the **Encryption** window, make sure that the **Enable backup file encryption** option is selected and press [Space].
3. In the **Password** field, type a password that you want to use for encryption.
4. In the **Hint** field, type a hint for the password. In case you lose the password, the specified hint will help you to remember the lost password.
5. Switch to the **Ok** button and press [Enter].



File Indexing Settings

NOTE

The **File system indexing** option is available if Veeam Agent for Linux operates in the Workstation or Server edition and is not available in Veeam Agent for Linux on Power.


To specify file indexing settings for the backup job, do the following:

1. In the **Specify advanced job settings** window, select the **File system indexing** option with the [Tab] and [Down] keys and press [Enter].

2. In the **Indexing** window, specify the indexing scope:

- Select **Index everything** if you want to index all files within the backup scope that you have specified at the **Backup mode** step of the wizard. Veeam Agent for Linux will index all files that reside:
 - On your computer OS (for entire machine backup)
 - On the volumes that you have selected for backup (for volume-level backup)
 - In the directories that you have selected for backup (for file-level backup)
- [For entire machine and volume-level backups] Select **Index only following directories** to define directories that you want to index. Enter paths to the necessary directories. To separate several paths, use the ',' (comma) character.
- [For entire machine and volume-level backups] Select **Index everything except** if you want to index all files within the specified backup scope except those files that reside in specific directories. Enter paths to directories whose files you do not want to index. To separate several paths, use the ',' (comma) character.

3. Switch to the **Ok** button and press [Enter].



Oracle Database Processing Settings

To specify processing settings for the Oracle database system, do the following

1. In the **Specify advanced job settings** window, select the **Oracle processing** option with the [Tab] and [Down] keys and press [Enter].
2. In the **Oracle processing** section, select one of the following options:
 - **Require successful processing.** With this option selected, Veeam Agent will stop the backup process if an error occurs while processing the Oracle database system.
 - **Try application processing, ignore failures.** With this option selected, Veeam Agent will continue the backup process even if errors occur when processing the Oracle database system.

3. In the **Archived logs processing** section, specify how Veeam Agent will process archived logs on the Oracle database:

- Select **Do not delete archived logs** if you want Veeam Agent to keep archived logs. When the backup job completes, Veeam Agent will not delete archived logs.

It is recommended that you select this option when you do not have databases running in the ARCHIVELOG mode. If the database is running in the ARCHIVELOG mode, archived logs may grow large and consume all disk space. In this case, the database administrator must take care of archived logs themselves.


- Select **Delete logs older than <N> hours** or **Delete oldest logs exceeding <N> GB** if you want Veeam Agent to delete archived logs that are older than <N> hours or larger than <N> GB. Veeam Agent will wait for the backup job to complete successfully and then trigger archived logs truncation via Oracle Call Interface (OCI). If the backup job fails, the logs will remain untouched until the next successful backup job session.

TIP

If you configure backup job to back up archived logs, Veeam Agent for Linux will not trigger archived logs deletion after each log backup job session. To prevent Oracle database logs from overgrowing, run the backup job for the Veeam Agent computer more often.

4. In the **Specify Oracle account with SYSDBA privileges** section, specify which account type Veeam Agent will use to connect to the database system.

- Select **System account** if you want Veeam Agent to use an account of the Veeam Agent machine OS. The account must be a member of the group that owns configuration files for the Oracle database (for example, the install group).
- Select **Oracle account** if you want Veeam Agent to use an Oracle account. The account must have SYSDBA rights.



MySQL Database Processing Settings

IMPORTANT

MySQL tables that use the MyISAM storage engine must be locked to keep them in consistent state while Veeam Agent is creating the system snapshot. To correctly process such tables, MySQL account must have the following instance-wide privileges:


- **SELECT.** This privilege enables Veeam Agent to access tables' metadata and select for a lock the tables that use the MyISAM storage engine. Without this privilege, the processing of the MySQL database system will run successfully but MyISAM tables will not be locked, which may result in an inconsistent state of the backed up data.
- **LOCK TABLES.** This privilege is required for locking the selected MyISAM tables. If some MyISAM tables are selected but the MySQL account does not have the **LOCK TABLES** privilege, the processing of the MySQL database system will fail.
- **RELOAD or FLUSH_TABLES.** If some MyISAM tables are selected but the MySQL account does not have either **RELOAD** or **FLUSH_TABLES** privilege, the processing of the MySQL database system will fail.

To obtain information about the privileges that are assigned to an account, use MySQL functionality, for example, the **SHOW GRANTS** statement. To learn more, see [MySQL documentation](#).


To specify processing settings for the MySQL database system, do the following:

1. In the **Specify advanced job settings** window, select the **MySQL processing** option with the [Tab] and [Down] keys and press [Enter].
2. In the **MySQL processing** section, select one of the following options:
 - **Require successful processing.** With this option selected, Veeam Agent will stop the backup process if an error occurs when processing the MySQL database system.
 - **Try application processing, ignore failures.** With this option selected, Veeam Agent will continue the backup process even if errors occur when processing the MySQL database system.

3. In the **Authentication method** section, specify how Veeam Agent will connect to the MySQL database:
- Select **Password** if you want Veeam Agent to connect with the MySQL account name and password. With this option selected, you must specify account name and password in the backup job settings.



- Select **Password file** if you want Veeam Agent to connect with the MySQL account name and password that are stored in the `.my.cnf` password file. With this option selected, you must specify a path to the password file, but do not need to specify account credentials in the backup job settings. To learn more about password file configuration, see [Preparing Password File for MySQL Processing](#).



Preparing Password File for MySQL Processing

You can use MySQL account credentials that are stored in the password file to connect Veeam Agent for Linux to the MySQL database system.

NOTE

Consider the following:

- If you specify a custom path to the password file, specify a full path. Specifying relative paths is not supported.
- The password file can also contain user-specific connection settings that Veeam Agent will apply to connect to the MySQL database system. For example, if you want to connect to the MySQL database system using the custom socket, specify the socket path in the password file. To learn more, see [MySQL documentation](#).

If you want to use a password file for authentication, create a file. By default, Veeam Agent expects the password file to have the `.my.cnf` name and to be in the home directory of the `root` user. If the password file has a custom name or is stored in another directory, you can specify a custom path.

The password file must have the following contents:

```
[client]
user=<username>
password=<password>
```

where:

- `<username>` – name of the account that Veeam Agent will use to connect to the MySQL database system.
- `<password>` – password of the account that Veeam Agent will use to connect to the MySQL database system.

For example:

```
[client]
user=root
password=P@ssw0rd
```


PostgreSQL Database Processing Settings

To specify processing settings for the PostgreSQL database system, do the following:

1. In the **Specify advanced job settings** window, select the **PostgreSQL processing** option with the [Tab] and [Down] keys and press [Enter].
2. In the **PostgreSQL processing** section, select one of the following options:
 - **Require successful processing.** With this option selected, Veeam Agent will stop the backup process if an error occurs when processing the PostgreSQL database system.
 - **Try application processing, ignore failures.** With this option selected, Veeam Agent will continue the backup process even if errors occur when processing the PostgreSQL database system.

3. In the **Authentication method** section, specify how Veeam Agent will connect to the PostgreSQL database:

- Select **Database user with password** if you want Veeam Agent to connect with the PostgreSQL account name and password. With this method selected, you must specify account name and password in the backup job settings.
- Select **Database user with password file** if you want Veeam Agent to connect with the PostgreSQL account password that is stored in the `.pgpass` password file. With this method selected, you must specify account name only in the backup job settings. To learn more about password file configuration, see [Password File for PostgreSQL](#).
- Select **System user without password** if you want Veeam Agent to connect using a peer authentication method. In the peer authentication method, Veeam Agent uses the OS account as the PostgreSQL database user name. With this option selected, you must specify OS account in the backup job settings. To learn more about peer authentication, see [PostgreSQL documentation](#).



Preparing Password File for PostgreSQL Processing

You can use PostgreSQL account credentials that are stored in the password file to connect Veeam Agent to the PostgreSQL database system.

If you want to use a password file for authentication, create the `.pgpass` file in the home directory of the `root` user.

The password file must have the following contents:

```
<hostname>:<port>:<database>:<username>:<password>
```

where:

- `<hostname>` – name of the host where the PostgreSQL database system is located.

- `<port>` – number of the free port that Veeam Agent will use to connect to the PostgreSQL database system.
- `<database>` – name of the PostgreSQL database.
- `<username>` – name of the account that Veeam Agent will use to connect to the PostgreSQL database system.
- `<password>` – password of the account that Veeam Agent will use to connect to the PostgreSQL database system.

For example:

```
srv01:5432:mydb:postgres:P@ssw0rd
```

For more information about the password file, see [PostgreSQL documentation](#).

Active Full Backup Settings

To specify active full backup settings for the backup job, do the following:


1. In the **Specify advanced job settings** window, select the **Active full** option with the [Tab] and [Down] keys and press [Enter].
2. In the **Active full** window, make sure that the **Create active full backups periodically** option is selected and press [Space].

NOTE

If you plan to use a GFS retention policy, you must select the **Create active full backups periodically** option. Otherwise, Veeam Agent will not have full backups to mark with GFS flags. To learn more, see [Long-Term Retention Policy](#).


3. Specify schedule for periodic active full backups:
 - If you want active full backups to run monthly, do the following:
 - i. Select the **Monthly on this day** option and specify a day of a month when Veeam Agent will perform active full backup.
 - ii. To specify the months on which Veeam Agent will perform active full backups, select **Months** with the [Tab] key and press [Enter].
 - iii. In the **Months** window, specify the months on which Veeam Agent will perform active full backup. By default, Veeam Agent performs active full backup every month. To select months, use the [Up], [Down], [Right], [Left] and [Space] keys.

iv. Switch to the **Ok** button with the [Tab] key and press [Enter].



- If you want active full backups to run weekly, do the following:
 - i. Select the **Weekly on selected days** option, then select **Days** with the [Tab] key and press [Enter].
 - ii. In the **Days** window, specify the days on which Veeam Agent will perform active full backup. By default, Veeam Agent performs active full backup every Saturday. To select days, use the [Up], [Down], [Right], [Left] and [Space] keys.

- iii. Switch to the **Ok** button with the [Tab] key and press [Enter].



Maintenance Settings

You can specify the number of days for which you want to keep the backup created with the backup job in the target location. To do this:

1. In the **Specify advanced job settings** window, select the **Maintenance** option with the [Tab] and [Down] keys and press [Enter].


NOTE

The **Maintenance** option is available if you have selected the **Veeam backup repository** or **Veeam Cloud Connect repository** option at the [Destination](#) step of the wizard.

2. In the **Maintenance** window, make sure that the **Allow backup server to prune old backups** option is selected and press [Space].

3. In the **If this agent has been unavailable <N> days** field, specify the number of days for which you want to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the period that you have specified. When this period is over, the backup will be removed from the target location.

By default, the retention period for old backups is 30 days. Do not set this retention period to 1 day or a similar short interval. In the opposite case, the backup job may work not as expected and remove data that you still require.



Script Settings

To specify script settings for the backup job, do the following:


1. In the **Specify advanced job settings** window, select the **Scripts** option with the [Tab] and [Down] keys and press [Enter].
2. In the **Scripts** window, in the **Job scripts** section, specify custom scripts that you want to execute before and after the backup job:
 - In the **Pre-job** field, specify a path to the script that should be executed before the backup job starts.
 - In the **Post-job** field, specify a path to the script that should be executed after the backup job completes.
3. In the **Scripts** window, in the **Snapshot scripts** section, specify custom scripts that you want to execute before Veeam Agent creates a snapshot of the backed-up volume and after the snapshot is created:
 - In the **Pre-freeze** field, specify a path to the script that should be executed before Veeam Agent creates a volume snapshot.

- In the **Post-thaw** field, specify a path to the script that should be executed after Veeam Agent creates a volume snapshot.

4. Switch to the **Ok** button and press [Enter].

IMPORTANT

You can specify snapshot script settings only if Veeam Agent for Linux operates in the Server edition. To learn more about editions, see [Product Editions](#).



Specifying Path to Script

You can specify a path to the executable file of the job or snapshot script in one of the following ways:

1. Type a path to the executable file.
2. Browse to the executable file:
 - a. Select the **Browse** option with the [Tab] key and press [Enter].
 - b. In the **Choose script location** window, select the directory being a part of the path to the script and press [Enter].
 - c. Repeat the step 'b' until a path to the directory in which the executable file resides appears in the **Current directory** field.
 - d. Select the necessary executable file and press [Enter].

Alternatively, you can switch to the **Ok** button and press [Enter].

TIP

If you do not want to execute a script, you can leave the corresponding field blank and proceed to the next step of the wizard.

Health Check Settings

When you store backup files in an object storage repository, an automatic health check can help you avoid a situation when a restore point gets corrupted, making all dependent restore points corrupted, too. For more information, see [Health Check for Object Storage](#).

NOTE

When you schedule a health check, consider the following:

- Health check runs automatically during incremental backup job session on the days specified in the health check schedule. If the backup job runs several times on a specified day, health check is performed only with the first run of the backup job on that day.
Health check is not performed during the first full backup or subsequent active full backup jobs.
- If Veeam Agent does not run any backup jobs on the day specified in the health check schedule, health check will be performed during the first backup job session following that day.
For example, you may have scheduled to run health check every last day of a month, while the backup job is scheduled to run every day and create an active full backup on Sundays. If the last day of a month falls on a Sunday, health check will be performed on the following Monday with the first incremental backup job session on that day.

To specify backup health check settings, do the following:


1. In the **Specify advanced job settings** window, select the **Health check** option with the [Tab] and [Down] keys and press [Enter].

NOTE

The **Health check** option is available if you have selected the **Object storage repository** option at the [Destination](#) step of the wizard.

2. In the **Health check** window, make sure that the **Perform backup files health check** option is highlighted and select it by pressing [Space].

3. Use the **Monthly on** or **Weekly selected days** settings to define the schedule for the health check of the backup in the repository.



Step 9. Specify Backup Schedule

At the **Schedule** step of the wizard, specify the schedule according to which you want to perform backup.

Depending on the product edition, Veeam Agent provides the following scheduling options:

- [For Free and Workstation editions] You can set the backup job to run automatically on specific days of the week.
- [For Server edition] You can schedule the backup job to run on specific days of the week or month, as well as periodically.

To specify the schedule, do the following:

1. Make sure that the **Run the job automatically** check box is selected.


If you want to configure the backup job without schedule, you can clear the **Run the job automatically** check box. In this case you will be able to start the configured backup job manually at any time you need.

2. Define the schedule for the backup job:
 - To run the job at specific time daily or on specific weekdays, select the **Daily at** option. Use the fields of this option to configure the necessary schedule.
 - To run the job once a month on a specific day, select the **Monthly at this time** option. Use the fields of this option to configure the necessary schedule.
 - To run the job repeatedly throughout a day with a specific time interval, select the **Periodically at** option. Use the fields of this option to specify the time interval in hours or minutes.

NOTE

Veeam Agent 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.

Veeam Agent for Linux will save the scheduling settings for the backup job in its database Veeam Agent can start a backup job automatically regardless of the currently running user session. You can change schedule settings at any time in Veeam Agent.




Step 10. Review Backup Job Settings

At the **Summary** step of the wizard, complete the backup job configuration process.

1. Review settings of the configured backup job.
2. To start the job after you close the wizard, make sure that the **Start job now** check box is selected.

If you want to start the backup job later, you can clear the **Start job now** check box. You will be able to start the backup job manually at any time you need. To learn more, see [Starting Backup Job](#).

3. Press [Enter] to exit the wizard.



What You Do Next

After you configure the backup job, you can start the backup job at any time you need. To learn more, see [Starting Backup Job](#).

If some of your data gets lost or corrupted, you can do the following:

- [Recover all computer volumes or specific volumes from the backup.](#)
- [Recover individual files and folders from the backup.](#)

Creating Backup Job with Command Line Interface

You can configure the backup job with the command line interface. Using Veeam Agent for Linux commands, you can create volume-level and file-level backup jobs, specify advanced settings for the created backup job, define backup schedule and enable backup encryption.

Creating Volume-Level Backup Job

IMPORTANT

Volume-level backup job relies on a device name in the `/dev` directory. Device names in the `/dev` directory (for example, `/dev/md-127`, `/dev/dm-1`) must stay persistent for backed-up volumes. Otherwise, the job will back up the wrong volume.

You can create a volume-level backup of the entire computer image or specific volumes.

To back up the entire computer image, use the following command:

```
veeamconfig job create volumelevel --name <job_name> --reponame <repository_name> --backupallsystem  
<advanced_options> <schedule_options> <active_full_backup_options> <indexing_options>
```

To back up specific volumes, use the following command:

```
veeamconfig job create volumelevel --name <job_name> --reponame <repository_name> --objects <volume_to_backup> <advanced_options> <schedule_options> <active_full_backup_options> <indexing_options>
```

where:

- `--backupallsystem` – option that instructs Veeam Agent to backup up the entire computer image. Veeam Agent will back up all the supported devices detected on the protected computer. If you use this option, you do not need to add the `--objects` option to the command.
- `<job_name>` – name for the created backup job.
- `<repository_name>` – name of the backup repository that should be used as a target location for the backup job. The backup repository must be created in advance.

If you want to create Veeam Agent backups in local directory or network shared folder, you need to create a repository. To learn more, see [Creating Backup Repository](#).

If you want to create Veeam Agent backups in a Veeam backup repository of cloud repository, you need connect to the Veeam backup server or Veeam Cloud Connect service provider in advance, before configuring the backup job. To learn more, see [Connecting to Veeam Backup Server](#) and [Connecting to Service Provider](#).

If you want to create Veeam Agent backups in the object storage, you need to connect to an object storage and create a repository on this storage. To learn more, see [Creating Repository in Object Storage](#).

- `<volume_to_backup>` – object that should be included in backup:
 - For simple volumes – name of a block device that represents a volume or an entire disk that should be included in backup. You can specify entire disk to create backup of the entire computer image or individual computer volumes to create backup of specific volumes. If you want to back-up several disks or volumes, specify them one after another using the ',' (comma) character as a separator.

IMPORTANT

Veeam Agent does not support backup of bind mount points. You must specify the path to the original mount point instead.

NOTE

If you include a block device in the backup, and this block device is a physical volume assigned to an LVM volume group, Veeam Agent will include the whole LVM volume group in the backup.

- For LVM volumes – name of an LVM logical volume that should be included in backup. If you want to back-up several LVM logical volumes, specify them one after another using the ',' (comma) character as a separator.
- `<advanced_options>` – advanced options for the backup job. To learn more, see [Advanced Backup Job Settings](#).
- `<schedule_options>` – schedule options for the backup job. To learn more, see [Schedule Settings](#).
- `<active_full_backup_options>` – active full backup schedule options for the backup job. To learn more, see [Active Full Backup Schedule Settings](#).
- `<indexing_options>` – file system indexing options for the backup job. To learn more, see [File System Indexing Settings](#).

For example:

```
$ veeamconfig job create volumelevel --name SystemBackup --reponame Repository_01 --objects /dev/sda1 --weekdays Mon,Sun --weekdays-full Thu
```

TIP

After you create the backup job, you can additionally configure the following backup job settings:

- Backup schedule. For details, see [Configuring Backup Schedule](#).
- Active full backup schedule. For details, see [Configuring Active Full Backup Schedule](#).
- Long-term retention policy. For details, see [Configuring Long-Term Retention Policy](#).
- Database processing settings for a volume-level backup job. For details, see [Configuring Database Processing Settings](#).
- [For job targeted at an object storage repository] Schedule for backup health check. For details, see [Configuring Health Check Schedule](#).

Advanced Backup Job Settings

You can specify the following advanced options for the backup job:

Option	Description and values
--compressionlevel	Data compression level. Possible values are: • <i>0</i> – No compression • <i>1</i> – Rle • <i>2</i> – Lz4 • <i>3</i> – Zstd 3 • <i>4</i> – Zstd 9 The default value is <i>2</i>.
--blocksize	Data block size in kilobytes. The default value is <i>1024</i>. Possible values are 256, 512, 1024 or 4096. To learn more about available sizes of data blocks, see Data Compression.
--maxpoints	The number of restore points that you want to store in the backup location. By default, Veeam Agent for Linux keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent for Linux will remove the earliest restore point from the backup chain.
--immutabledays	The time period in days during which the backup stored in an object storage repository will be immutable to modification or deletion. For more information, see Backup Immutability.
--prefreeze	Path to the script that should be executed before the snapshot creation. This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see Product Editions.
--postthaw	Path to the script that should be executed after the snapshot creation. This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see Product Editions.
--prejob	Path to the script that should be executed at the start of the backup job.
--postjob	Path to the script that should be executed after the backup job completes.
--setencryption	Defines that data encryption option is enabled for the job. When you use the <code>veeamconfig job create</code> command with the <code>--setencryption</code> option, Veeam Agent for Linux will prompt you to specify a password for data encryption and hint for the password.

Option	Description and values
--deleteold	The number of days to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the specified number of days. When this period is over, the backup will be removed from the target location. Possible values are: 1-999. If you do not specify the <code>--deleteold</code> option, Veeam Agent will not apply this setting. The backup will be stored in the target location until you delete it manually.

Schedule Settings

You can specify a daily, monthly or periodic schedule for the backup job.

Option	Description and values
--weekdays	[For weekly schedules] Specifies the weekdays when the backup job must run. If you want to run the backup job more than once during the week, the list of weekdays must be separated by a comma (','). Possible values are: <i>Mon</i> – Monday <i>Tue</i> – Tuesday <i>Wed</i> – Wednesday <i>Thu</i> – Thursday <i>Fri</i> – Friday <i>Sat</i> – Saturday <i>Sun</i> – Sunday
--daily	[For weekly schedules] Defines that the backup job must start daily at specific time.
--thisday	[For monthly schedules] Specifies the day of the month when the backup job must run. Possible values: from 1 to 31 or <i>Last</i> .
--weeknumber	[For monthly schedules] Specifies the week of the month when the backup job must run. Possible values: <i>First</i> , <i>Second</i> , <i>Third</i> , <i>Fourth</i> or <i>Last</i> . This option must be used in combination with the <code>--monthlyweekday</code> option.

Option	Description and values
--monthlyweekday	[For monthly schedules] Specifies the day of the week when the backup job must run. You can select only one weekday. Possible values are: • <i>Mon</i> – Monday • <i>Tue</i> – Tuesday • <i>Wed</i> – Wednesday • <i>Thu</i> – Thursday • <i>Fri</i> – Friday • <i>Sat</i> – Saturday • <i>Sun</i> – Sunday
--months	[For monthly schedules] Specifies the months when the backup job must run. If you specify more than one month, the list must be separated by a comma (,) – for example: <i>Jan, Apr, Jul, Oct</i>. If you do not specify this option, the backup job will run every month.
--every	[For periodic schedules] Specifies the period of time in minutes or hours between the runs of the backup job. The period must be specified in the <i>HH:MM</i> format – for example, <i>06:00</i>.
--at	[For weekly and monthly schedules] Specifies the time of day in the <i>HH:MM</i> format when the backup job must start – for example: <i>20:00</i>.

After the backup job is created, Veeam Agent for Linux automatically enables backup schedule. To learn about how to configure backup schedule for an existing backup job, see [Configuring Backup Schedule](#).

Active Full Backup Schedule Settings

You can specify schedule options for the backup job to create active full backups on specific weekdays or days of the month.

You can specify schedule options for the backup job to create active full backups on specific days of the week or month.

Option	Description and values
--weekdays-full	[For weekly schedules] Specifies the weekdays when the backup job must create an active full backup. If you want to create an active full backup more than once during the week, the list of weekdays must be separated by a comma (','). Possible values are: • <i>Mon</i> – Monday • <i>Tue</i> – Tuesday • <i>Wed</i> – Wednesday • <i>Thu</i> – Thursday • <i>Fri</i> – Friday • <i>Sat</i> – Saturday • <i>Sun</i> – Sunday
--thisday-full	[For monthly schedules] Specifies the day of the month when the backup job must create an active full backup. Possible values: from 1 to 31 or <i>Last</i>.
--weeknumber-full	[For monthly schedules] Specifies the week of the month when the backup job must create an active full backup.. Possible values: <i>First</i>, <i>Second</i>, <i>Third</i>, <i>Fourth</i> or <i>Last</i>. This option must be used in combination with the <code>--monthlyweekday</code> option.
--monthlyweekday-full	[For monthly schedules] Specifies the day of the week when the backup job must create an active full backup. You can select only one weekday. Possible values are: • <i>Mon</i> – Monday • <i>Tue</i> – Tuesday • <i>Wed</i> – Wednesday • <i>Thu</i> – Thursday • <i>Fri</i> – Friday • <i>Sat</i> – Saturday • <i>Sun</i> – Sunday
--months-full	[For monthly schedules] Specifies the months when the backup job must create an active full backup. If you specify more than one month, the list must be separated by a comma (,) – for example: <i>Jan</i>, <i>Apr</i>, <i>Jul</i>, <i>Oct</i>. If you do not specify this option, the backup job will create an active full backup. every month.

After the backup job is created, Veeam Agent automatically enables active full backup schedule. To learn about how to configure active full backup schedule for an existing backup job, see [Configuring Active Full Backup Schedule](#).

File System Indexing Settings

You can specify one the following file system indexing options for the backup job:

Option	Description and values
--indexall	Defines that Veeam Agent for Linux must index all files on the volumes included in backup.
--indexonly	Path to a directory that contains files that you want to index. Enter paths to the necessary directories. To separate several paths, use the ',' (comma) character.
--indexexcept	Path to a directory that contains files that you do not want to index. You can specify one or more paths. To separate several paths, use the ',' (comma) character.

To learn more about file indexing, see [File System Indexing](#).

Creating File-Level Backup Job

To create a file-level backup job, use the following command:

```
veeamconfig job create filelevel --name <job_name> --reponame <repository_name>
--includedirs <objects> <include/exclude options> <advanced_options> <schedule_
options> <active_full_backup_options> <indexing_options> [--nosnap]
```

where:

- **<job_name>** – name for the created backup job.
- **<repository_name>** – name of the backup repository that should be used as a target location for the backup job. The backup repository must be created in advance. To learn more, see [Creating Backup Repository](#).

If you want to create Veeam Agent backups in the Veeam backup repository, you should connect to the Veeam backup server in advance, before configuring the backup job. To learn more, see [Connecting to Veeam Backup Server](#).
- **<objects>** – directories that will be included in the backup scope. File-level backup job must have at least one directory added to the backup scope.
- **<include/exclude options>** – files and directories inclusion/exclusion options. To learn more, see [File Inclusion Options](#).
- **<advanced_options>** – advanced options for the backup job. To learn more, see [Advanced Backup Job Settings](#).
- **<schedule_options>** – schedule options for the backup job. To learn more, see [Schedule Settings](#).
- **<active_full_backup_options>** – active full backup schedule options for the backup job. To learn more, see [Active Full Backup Schedule Settings](#).

- `<indexing_options>` – file system indexing options for the backup job. To learn more, see [File System Indexing Settings](#).
- `--nosnap` – option that instructs Veeam Agent for Linux to perform backup in the snapshot-less mode. With this option enabled, Veeam Agent for Linux will not create a snapshot of the backed-up volumes during backup. This allows Veeam Agent to back up data residing in file systems that are not supported for snapshot-based backup with Veeam Agent for Linux. Keep in mind that the snapshot-less file-level backup does not guarantee that data in the backup is consistent. To learn more, see [Snapshot-Less File-Level Backup](#).

For example:

```
$ veeamconfig job create filelevel --name HomeFolderBackup --reponame NetworkRepository --includedirs /home/user --excludedirs /home/user/temp --excludemasks "*.pdf"
```

TIP

After you create the backup job, you can additionally configure the following backup job settings:

- Backup schedule. For details, see [Configuring Backup Schedule](#).
- Active full backup schedule. For details, see [Configuring Active Full Backup Schedule](#).
- Long-term retention policy. For details, see [Configuring Long-Term Retention Policy](#).
- [For job targeted at an object storage repository] Schedule for backup health check. For details, see [Configuring Health Check Schedule](#).

File Inclusion Options

When you create a file-level backup job, you must specify at least one directory that should be included in backup. If you do not want to back up some files and directories in the specified directory, you can exclude specific files and directories from backup.

IMPORTANT

Veeam Agent does not support backup of bind mount points. You must specify the path to the original mount point instead.

To define the backup scope for the file-level backup job, you can use the following command-line options:

Option	Description and values
--includedirs	Full path to a directory that should be included in backup, for example: <code>/home/user</code>. File-level backup job must have at least one directory added to the backup scope. You can specify one or several paths to directories in the computer file system. To separate several paths, use the <code>,</code> (comma) character, for example: <code>/home/user/Documents,/home/user/reports</code>. Tip: If you want to back up the root directory and specify the <code>/</code> (slash) character, Veeam Agent will not automatically include the mount points in the backup scope. To include the mount points, you can do either of the following: • Enable automatic inclusion of local mount points when the root directory is added into the backup scope. To do this, in the Veeam Agent configuration file, enable the <code>rootRecursion</code> option and set it to <code>true</code>: <code>rootRecursion = true</code>. Note that even if you enable this configuration option, network file systems will not be included into backup automatically; you will need to specify paths to such mount points manually. • Specify paths to the mount points manually. For example, you have a network file system mounted to the <code>/home/media</code> directory. If you add <code>/</code> as an object to the backup scope, Veeam Agent will not back up the mounted network file system. To back up the root directory and the mounted network file system, add the following objects to the backup scope: <code>/,/home/media</code>.
--excludedirs	Full path to a directory that should be excluded from backup. The directory specified with this option must be a subdirectory of the directory specified with the <code>--includedirs</code> option. To separate several paths, use the <code>,</code> (comma) character, for example, <code>/home/user/Documents,/home/user/reports</code>.

Option	Description and values
--includemasks	A name mask for the files that should be included in the backup. You can use the following UNIX wildcard characters for file name masks: • '*' — a substitution for one or more characters in the file name. Can be used for any sequence of characters (including no characters). For example, *.pdf. • '?' — a substitution of one character in the file name. For example, repor?.pdf. • '[' — a substitution of one character in the file name with any of the characters enclosed in square brackets (or a range of characters defined with the '-' character). For example: report_201[3456].pdf or report_201[3-6].pdf. Keep in mind that you must specify each name mask in double quotation marks (""). For example: <code>--includemasks "*.bak"</code>. If you want to use several file name masks, you must specify them in double quotation marks ("") and separate them with a comma (,). For example: <code>--includemasks "*.bak, *.pdf"</code>. File inclusion option is applied to all directories that are specified with the <code>--includedirs</code> option. For example, if you include in backup the <code>/home/user/Documents</code> directory and files that match the <code>repor?.pdf</code> file name mask, Veeam Agent for Linux will back up the <code>/home/user/Documents/report.pdf</code> file and will not back up the <code>/home/user/reports/report.pdf</code> file.

Option	Description and values
--excludemasks	A name mask for the files that should be excluded from the backup. You can use the following UNIX wildcard characters for file name masks: • '*' – a substitution for one or more characters in the file name. Can be used for any sequence of characters (including no characters). For example, *.pdf. • '?' – a substitution of one character in the file name. For example, repor?.pdf. • '[' – a substitution of one character in the file name with any of the characters enclosed in square brackets (or a range of characters defined with the '-' character). For example: report_201[3456].pdf or report_201[3-6].pdf. Keep in mind that you must specify each name mask in double quotation marks (""). For example: --excludemasks "*.bak". If you want to use several file name masks, you must specify them in double quotation marks ("") and separate them with a comma (,). For example: --excludemasks *.bak, *.pdf". File exclusion option is applied to all directories that are specified with the --includedirs option and files that match file name masks specified with the --includemasks option. For example, you may want to specify the following backup scope for the backup job: • Include in backup the /home/user/Documents directory • Include files that match the report.* file name mask • Exclude files that match the *.odt file name mask. In this case, Veeam Agent for Linux will back up the /home/user/Documents/report.pdf file and will not back up /home/user/Documents/report.odt and /home/user/reports/report.pdf files.

Advanced Backup Job Settings

You can specify the following advanced options for the backup job:

Option	Description and values
--compressionlevel	Data compression level. Possible values are: • 0 – No compression • 1 – Rle • 2 – Lz4 • 3 – Zstd 3 • 4 – Zstd 9 To learn more about available data compression levels, see Data Compression.
--blocksize	Data block size in kilobytes. The default value is 1024. Possible values are 256, 512, 1024 or 4096. To learn more about available sizes of data blocks, see Data Compression.
--maxpoints	The number of restore points that you want to store in the backup location. By default, Veeam Agent for Linux keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent for Linux will remove the earliest restore point from the backup chain.
--immutabledays	The time period in days during which the backup stored in an object storage repository will be immutable to modification or deletion. For more information, see Backup Immutability.
--prefreeze	Path to the pre-freeze script that should be executed before the snapshot creation. This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see Product Editions.
--postthaw	Path to the post-thaw script that should be executed after the snapshot creation. This option is available only if Veeam Agent for Linux operates in the Server edition. To learn about editions, see Product Editions.
--prejob	Path to the script that should be executed at the start of the backup job.
--postjob	Path to the script that should be executed after the backup job completes.
--setencryption	Defines that data encryption option is enabled for the job. When you use the <code>veeamconfig job create</code> command with the <code>--setencryption</code> option, Veeam Agent for Linux will prompt you to specify a password for data encryption and hint for the password.

Option	Description and values
--deleteold	The number of days to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the specified number of days. When this period is over, the backup will be removed from the target location. Possible values are: 1-999. If you do not specify the <code>--deleteold</code> option, Veeam Agent will not apply this setting. The backup will be stored in the target location until you delete it manually.

Schedule Settings

You can specify a daily, monthly or periodic schedule for the backup job.

Option	Description and values
--weekdays	[For weekly schedules] Specifies the weekdays when the backup job must run. If you want to run the backup job more than once during the week, the list of weekdays must be separated by a comma (','). Possible values are: <i>Mon</i> – Monday <i>Tue</i> – Tuesday <i>Wed</i> – Wednesday <i>Thu</i> – Thursday <i>Fri</i> – Friday <i>Sat</i> – Saturday <i>Sun</i> – Sunday
--daily	[For weekly schedules] Defines that the backup job must start daily at specific time.
--thisday	[For monthly schedules] Specifies the day of the month when the backup job must run. Possible values: from 1 to 31 or <i>Last</i> .
--weeknumber	[For monthly schedules] Specifies the week of the month when the backup job must run. Possible values: <i>First</i> , <i>Second</i> , <i>Third</i> , <i>Fourth</i> or <i>Last</i> . This option must be used in combination with the <code>--monthlyweekday</code> option.

Option	Description and values
--monthlyweekday	[For monthly schedules] Specifies the day of the week when the backup job must run. You can select only one weekday. Possible values are: • <i>Mon</i> – Monday • <i>Tue</i> – Tuesday • <i>Wed</i> – Wednesday • <i>Thu</i> – Thursday • <i>Fri</i> – Friday • <i>Sat</i> – Saturday • <i>Sun</i> – Sunday
--months	[For monthly schedules] Specifies the months when the backup job must run. If you specify more than one month, the list must be separated by a comma (,) – for example: <i>Jan, Apr, Jul, Oct</i>. If you do not specify this option, the backup job will run every month.
--every	[For periodic schedules] Specifies the period of time in minutes or hours between the runs of the backup job. The period must be specified in the <i>HH:MM</i> format – for example, <i>06:00</i>.
--at	[For weekly and monthly schedules] Specifies the time of day in the <i>HH:MM</i> format when the backup job must start – for example: <i>20:00</i>.

After the backup job is created, Veeam Agent for Linux automatically enables backup schedule. To learn about how to configure backup schedule for an existing backup job, see [Configuring Backup Schedule](#).

Active Full Backup Schedule Settings

You can specify schedule options for the backup job to create active full backups on specific weekdays or days of the month.

You can specify schedule options for the backup job to create active full backups on specific days of the week or month.

Option	Description and values
--weekdays-full	[For weekly schedules] Specifies the weekdays when the backup job must create an active full backup. If you want to create an active full backup more than once during the week, the list of weekdays must be separated by a comma (','). Possible values are: • <i>Mon</i> – Monday • <i>Tue</i> – Tuesday • <i>Wed</i> – Wednesday • <i>Thu</i> – Thursday • <i>Fri</i> – Friday • <i>Sat</i> – Saturday • <i>Sun</i> – Sunday
--thisday-full	[For monthly schedules] Specifies the day of the month when the backup job must create an active full backup. Possible values: from 1 to 31 or <i>Last</i>.
--weeknumber-full	[For monthly schedules] Specifies the week of the month when the backup job must create an active full backup.. Possible values: <i>First</i>, <i>Second</i>, <i>Third</i>, <i>Fourth</i> or <i>Last</i>. This option must be used in combination with the <code>--monthlyweekday</code> option.
--monthlyweekday-full	[For monthly schedules] Specifies the day of the week when the backup job must create an active full backup. You can select only one weekday. Possible values are: • <i>Mon</i> – Monday • <i>Tue</i> – Tuesday • <i>Wed</i> – Wednesday • <i>Thu</i> – Thursday • <i>Fri</i> – Friday • <i>Sat</i> – Saturday • <i>Sun</i> – Sunday
--months-full	[For monthly schedules] Specifies the months when the backup job must create an active full backup. If you specify more than one month, the list must be separated by a comma (,) – for example: <i>Jan</i>, <i>Apr</i>, <i>Jul</i>, <i>Oct</i>. If you do not specify this option, the backup job will create an active full backup. every month.

After the backup job is created, Veeam Agent automatically enables active full backup schedule. To learn about how to configure active full backup schedule for an existing backup job, see [Configuring Active Full Backup Schedule](#).

File System Indexing Settings

You can specify the following file system indexing option for the backup job:

Option	Description and values
--indexall	Defines that Veeam Agent for Linux must index all files in the directories included in backup.

To learn more about file indexing, see [File System Indexing](#).

Configuring Backup Schedule

To run a backup job periodically without the user intervention, you can schedule it to start automatically. You can specify schedule settings individually for every job created in Veeam Agent. You can perform the following actions with the backup job schedule via command line interface:

- [Specify schedule settings for the job.](#)
- [Enable schedule for the job.](#)
- [View the schedule defined for the job.](#)
- [Disable schedule for the job.](#)

TIP

You can also specify backup schedule for the backup job when you create the job. For details, see [Creating Volume-Level Backup Job](#) and [Creating File-Level Backup Job](#).

Specifying Backup Schedule

Depending on the product edition, Veeam Agent allows you to set [daily](#), [monthly](#) or [periodic](#) schedule for a backup job. Daily schedules are available for the Free and Workstation editions of Veeam Agent. In the Server edition of Veeam Agent, you can additionally set monthly and periodic schedules for backup jobs. For details on Veeam Agent editions, see [Product Editions](#).

After you define the schedule, Veeam Agent automatically enables this schedule for the specified backup job.

Specifying Daily Schedules

You can set the backup job to run automatically on specific weekdays or every day.

- To run the backup job on specific days of the week, use the following command:

```
veeamconfig schedule set --jobid <job_id> --weekdays <days> --at <time>
```

or

```
veeamconfig schedule set --jobname <job_name> --weekdays <days> --at <time>
```

where:

- <job_id> – ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job_name> – name of the backup job for which you want to configure the schedule.
- <days> – days when the backup job must start separated by a comma (',') – for example: Monday, Tuesday, Wednesday, Thursday, Friday or Mon, Tue, Wed, Thu, Fri.
- <time> – time of day when the backup job must start specified in the HH:MM format – for example, 20:00.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --weekdays Monday, Tuesday, Wednesday, Thursday, Friday --at 20:00
```

- To run the backup job every day, use the following command:

```
veeamconfig schedule set --jobid <job_id> --daily --at <time>
```

or

```
veeamconfig schedule set --jobname <job_name> <daily options> --at <time>
```

where:

- <job_id> – ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job_name> – name of the backup job for which you want to configure the schedule.
- <time> – time of day when the backup job must start specified in the HH:MM format – for example, 20:00.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f --daily --at 20:00
```

Specifying Monthly Schedules

You can set the backup job to run automatically on specific months or every month.

- To run the backup job monthly on a specific day of the specific week, use the following command:

```
veeamconfig schedule set --jobid <job_id> --monthlyweekday <day> --weeknumber <week> [--months <months>] --at <time>
```

or

```
veeamconfig schedule set --jobname <job_name> --monthlyweekday <day> --weeknumber <week> [--months <months>] --at <time>
```

where:

- <job_id> – ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job_name> – name of the backup job for which you want to configure the schedule.
- <day> – day of the week when the backup job must start – for example, Tuesday or Tue.
- <week> – week of the month when the backup job must run. Possible values: First, Second, Third, Fourth or Last.
- <months> – months when the backup job must run separated by a comma (',') – for example: Jan, Apr, Jul, Oct. If you do not specify this option, the backup job will run every month.
- <time> – time of day when the backup job must start specified in the HH:MM format, – for example, 20:00.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --monthlyweekday Mon --weeknumber Second --months Jan,Jul --at 20:00
```

- To run the backup job monthly on a specific day of the month, use the following command:

```
veeamconfig schedule set --jobid <job_id> --thisday <day> [--months <months>] --at <time>
```

or

```
veeamconfig schedule set --jobname <job_name> --thisday <day> [--months <months>] --at <time>
```

where:

- <job_id> – ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).

- `<job_name>` – name of the backup job for which you want to configure the schedule.
- `<day>` – day of the month when the backup job must start. Possible values range from 1 to 31 or Last.
- `<months>` – months when the backup job must run separated by a comma (',') – for example: Jan, Apr, Jul, Oct. If you do not specify this option, the backup job will run every month.
- `<time>` – time of day when the backup job must start specified in the HH:MM format, – for example, 20:00.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --thisday 21
--months Jan,Jul --at 20:00
```

Specifying Periodic Schedules

To run the job periodically, run the following command:

```
veeamconfig schedule set --jobid <job_id> --every <time>
```

or

```
veeamconfig schedule set --jobname <job_name> --every <time>
```

where:

- `<job_id>` – ID of the backup job for which you want to configure the schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to configure the schedule.
- `<time>` – period of time when the backup job must start specified in the HH:MM format, – for example, 06:00.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --every 12:00
```

Viewing Backup Schedule

To view the schedule defined for the backup job, use the following command:

```
veeamconfig schedule show --jobid <job_id>
```

or

```
veeamconfig schedule show --jobname <job_name>
```

where:

- <job_id> – ID of the backup job for which you want to view the schedule.
- <job_name> – name of the backup job for which you want to view the schedule.

Veeam Agent will display the details and the status (enabled or disabled) of the job schedule – for example:

```
user@srv01:~$ veeamconfig schedule show --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
Days: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday
At: 20:00
Run automatically: enabled
```

or

```
user@srv01:~$ veeamconfig schedule show --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
Every 12 hours
Run automatically: enabled
```

Disabling Backup Schedule

To disable the schedule for the backup job, use the following command:

```
veeamconfig schedule disable --jobid <job_id>
```

or

```
veeamconfig schedule disable --jobname <job_name>
```

where:

- <job_id> – ID of the backup job for which you want to disable the schedule.
- <job_name> – name of the backup job for which you want to disable the schedule.

For example:

```
user@srv01:~$ veeamconfig schedule disable --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
```

Enabling Backup Schedule

After you define the schedule, Veeam Agent automatically enables this schedule for the specified backup job. If you disable the schedule for the backup job, you can enable it again by using the following command:

```
veeamconfig schedule enable --jobid <job_id>
```

or

```
veeamconfig schedule enable --jobname <job_name>
```

where:

- `<job_id>` – ID of the backup job for which you want to enable the schedule. You should look up the job ID in advance, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to enable the schedule.

For example:

```
user@srv01:~$ veeamconfig schedule enable --jobid 4849a3ae-1935-4969-98a3-d8acd  
2f6c73f
```

You can disable the schedule for the job at any time. To learn more, see [Disabling Backup Schedule](#).

Configuring Long-Term Retention Policy

You can configure the backup job to store backup files for long periods of time – for weeks, months and even years, you can set the long-term or Grandfather-Father-Son (GFS) retention policy. This policy uses backup files created while backup job is enabled and marks these backups with specific GFS flags. You can perform the following actions with the long-term retention policy in command line interface:

- [Specify long-term retention policy for the job.](#)
- [Enable long-term retention policy for the job.](#)
- [View the long-term retention policy defined for the job.](#)
- [Disable long-term retention policy for the job.](#)

Specifying Long-Term Retention Policy

You can configure long-term retention policy to keep [weekly](#), [monthly](#) or [yearly](#) full backups. After you configure a long-term retention policy, Veeam Agent will automatically enable long-term retention for the backup job.

Configuring Long-Term Retention Policy to Keep Weekly Full Backups

To configure long-term retention policy to keep weekly full backups, use the following command:

```
veeamconfig gfs set --jobid <job_id> weekly --on <weekday> --keep <weeks>
```

or

```
veeamconfig gfs set --jobname <job_name> weekly --on <weekday> --keep <weeks>
```

where:

- **<job_id>** – ID of the backup job for which you want to configure the long-term retention policy. You should look up the job ID in advance, before configuring the schedule, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job_name>** – name of the backup job for which you want to configure the long-term retention policy.
- **<weekday>** – week day when Veeam Agent must assign a weekly GFS flag to a full restore point – for example, `Tue` or `Tuesday`.
- **<weeks>** – number of weeks to keep the weekly GFS flag on the full restore point.

For example:

```
user@srv01:~$ veeamconfig gfs set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f  
weekly --on Saturday --keep 1
```

Configuring Long-Term Retention Policy to Keep Monthly Full Backups

To configure long-term retention policy to keep monthly full backups, use the following command:

```
veeamconfig gfs set --jobid <job_id> monthly --on <week_number> --keep <months>
```

or

```
veeamconfig gfs set --jobname <job_name> monthly --on <week_number> --keep <months>
```

where:

- **<job_id>** – ID of the backup job for which you want to configure the long-term retention policy. You should look up the job ID in advance, before configuring the schedule, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).

- `<job_name>` – name of the backup job for which you want to configure the long-term retention policy.
- `<week_number>` – number of the week when Veeam Agent must assign a monthly GFS flag to a full restore point. Possible values: `First`, `Second`, `Third`, `Fourth` or `Last`.
- `<months>` – number of months to keep the monthly GFS flag on the full restore point.

For example:

```
user@srv01:~$ veeamconfig gfs set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
monthly --on Second --keep 6
```

Configuring Long-Term Retention Policy to Keep Yearly Full Backups

To configure long-term retention policy to keep yearly full backups, use the following command:

```
veeamconfig gfs set --jobid <job_id> yearly --on <month> --keep <years>
```

or

```
veeamconfig gfs set --jobname <job_name> yearly --on <month> --keep <years>
```

where:

- `<job_id>` – ID of the backup job for which you want to configure the long-term retention policy. You should look up the job ID in advance, before configuring the schedule, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to configure the long-term retention policy.
- `<month>` – month when Veeam Agent must assign a yearly GFS flag to a full restore point – for example, `Jan` or `January`.
- `<years>` – number of years to keep the yearly GFS flag on the full restore point.

For example:

```
user@srv01:~$ veeamconfig gfs set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
yearly --on January --keep 3
```

Enabling Long-Term Retention Policy

To start marking backups with specific GFS flags, you must enable the long-term retention policy for the job. Use the following command:

```
veeamconfig gfs enable --jobid <job_id> [--type <period>]
```

or

```
veeamconfig gfs enable --jobname <job_name> [--type <period>]
```

where:

- **<job_id>** – ID of the backup job for which you want to enable the long-term retention policy. You should look up the job ID in advance, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job_name>** – name of the backup job for which you want to enable the long-term retention policy.
- **<period>** – type of the long-term retention policy. Possible values: `weekly`, `monthly` or `yearly`. This parameter is optional. You can use it to enable a specific type of long-term retention policy. To enable several types of retention at once, specify all necessary retention types separated by a comma (',') – for example: `weekly,monthly`.

For example:

```
user@srv01:~$ veeamconfig gfs enable --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f --type monthly
```

You can disable the long-term retention policy for the job at any time. To learn more, see [Disabling Long-Term Retention Policy](#).

Viewing Long-Term Retention Policy

To view the long-term retention policy defined for the backup job, use the following command:

```
veeamconfig gfs show --jobid <job_id>
```

or

```
veeamconfig gfs show --jobname <job_name>
```

where:

- **<job_id>** – ID of the backup job for which you want to view the long-term retention policy.
- **<job_name>** – name of the backup job for which you want to view the long-term retention policy.

Veeam Agent displays the following information about the backup job long-term retention policy:

Parameter	Description
GFS state	State of long-term retention policy. Possible values: • GFS is enabled • GFS is disabled • GFS is not set
Enabled	Possible values: <code>true</code> or <code>false</code> .
Desired time	Weekday, week number or month when Veeam Agent will set the GFS flag on the full restore point.
Keep for	Period of time for retaining the GFS flag on the full restore point.

The information listed in the table above is displayed for weekly, monthly and yearly retention policies.

For example:

```
user@srv01:~$ veeamconfig gfs show --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
GFS is enabled
Weekly:
  Enabled: true
  Desired time: Friday
  Keep for: 1 weeks
Monthly:
  Enabled: false
  Desired time: First
  Keep for: 1 months
Yearly:
  Enabled: false
  Desired time: January
  Keep for: 1 years
```

Disabling Long-Term Retention Policy

You can disable all or specific types of the long-term retention policy: weekly, monthly or yearly.

Disabling All Types of Long-Term Retention

To disable the long-term retention policy for the backup job, use the following command:

```
veeamconfig gfs disable --jobid <job_id>
```

or

```
veeamconfig gfs disable --jobname <job_name>
```

where:

- <job_id> – ID of the backup job for which you want to disable the long-term retention policy.
- <job_name> – name of the backup job for which you want to disable the long-term retention policy.

For example:

```
user@srv01:~$ veeamconfig gfs disable --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f
```

Disabling Specific Types of Long-Term Retention

To disable a specific type of the long-term retention policy for the backup job, use the following command:

```
veeamconfig gfs set --jobid <job_id> <period> --disable
```

or

```
veeamconfig gfs set --jobname <job_name> <period> --disable
```

where:

- <job_id> – ID of the backup job for which you want to disable the long-term retention policy.
- <job_name> – name of the backup job for which you want to disable the long-term retention policy.
- <period> – single type of the long-term retention policy you want to disable. Possible values: `weekly`, `monthly` or `yearly`.

For example:

```
user@srv01:~$ veeamconfig gfs set --jobid 4849a3ae-1935-4969-98a3-d8acd2f6c73f weekly --disable
```

Configuring Active Full Backup Schedule

You can schedule a backup job to create active full backups periodically. You can specify active full schedule settings individually for every job created in Veeam Agent. You can perform the following actions with the active full backup schedule via the command-line interface:

- [Specify active full backup schedule.](#)
- [Enable active full backup schedule.](#)
- [View active full backup schedule.](#)

- [Disable active full backup schedule.](#)

TIP

You can also specify active full backup schedule for the backup job when you create the job. For details, see [Creating Volume-Level Backup Job](#) and [Creating File-Level Backup Job](#).

Specifying Active Full Backup Schedule

You can configure the backup job to create active full backups on a [weekly](#) or [monthly](#) schedule.

After you define the active full backup schedule, Veeam Agent automatically enables this schedule for the specified backup job.

Specifying Weekly Schedules

To instruct Veeam Agent to create an active full backup on specific week days, use the following command:

```
veeamconfig schedule set --jobid <job_id> --weekdays <days>
```

or

```
veeamconfig schedule set --jobname <job_name> --weekdays <days>
```

where:

- **<job_id>** – ID of the backup job for which you want to configure the active full backup schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job_name>** – name of the backup job for which you want to configure the active full backup schedule.
- **<days>** – days when the backup job must create an active full backup separated by a comma (',') – for example: `Monday, Friday` or `Mon, Fri`.

For example:

```
user@srv01:~$ veeamconfig schedule activefull set --jobname DailyBackup --weekdays Monday, Friday
```

Specifying Monthly Schedules

You can configure the backup job to create active full backups on specific months or every month.

- To create an active full backup monthly on a specific day of a specific week, use the following command:

```
veeamconfig schedule activefull set --jobid <job_id> --monthlyweekday <day> --weeknumber <week> [--months <months>]
```

or

```
veeamconfig schedule set --jobname <job_name> --monthlyweekday <day> --weeknumber <week> [--months <months>]
```

where:

- <job_id> – ID of the backup job for which you want to configure the active full backup schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job_name> – name of the backup job for which you want to configure the active full backup schedule.
- <day> – days when the backup job must create an active full backup separated by a comma (','). For example: `Monday, Friday`. The backup job will create an active full backup on the specified days at the time specified in the backup job schedule settings.
- <week> – week of the month when the backup job must create an active full backup. Possible values: `First, Second, Third, Fourth` or `Last`.
- <months> – months when the backup job must create an active full backup separated by a comma (',') – for example: `Jan, Apr, Jul, Oct`. If you do not specify this option, the backup job will create an active full backup every month.

For example:

```
user@srv01:~$ veeamconfig schedule activefull set --jobname DailyBackup --monthlyweekday Mon --weeknumber Second --months Jan,Jul
```

- To configure the backup job to create an active full backup monthly on a specific day of the month, use the following command:

```
veeamconfig schedule set --jobid <job_id> --thisday <day> [--months <months>]
```

or

```
veeamconfig schedule set --jobname <job_name> --thisday <day> [--months <months>]
```

where:

- <job_id> – ID of the backup job for which you want to configure the active full backup schedule. You should look up the job ID in advance, before configuring the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job_name> – name of the backup job for which you want to configure the active full backup schedule.
- <day> – day of the month when Veeam Agent must create an active full backup. Possible values range from 1 to 31 or `Last`.

- **<months>** – months when the backup job must create an active full backup separated by a comma (,) – for example: Jan, Apr, Jul, Oct. If you do not specify this option, the backup job will create an active full backup every month.

For example:

```
user@srv01:~$ veeamconfig schedule set --jobname DailyBackup --thisday 21
--months Jan,Jul
```

Viewing Active Full Backup Schedule

To view the active full backup schedule defined for the backup job, use the following command:

```
veeamconfig schedule activefull show --jobid <job_id>
```

or

```
veeamconfig schedule activefull show --jobname <job_name>
```

where:

- **<job_id>** – ID of the backup job for which you want to view the active full backup schedule.
- **<job_name>** – name of the backup job for which you want to view the active full backup schedule.

Veeam Agent for Linux displays the following information about the active full backup schedule:

Parameter	Description
Every <value>	Days on which the backup job creates active full backups. For example: <i>Every Sat</i> or <i>Every 1 day of every month</i> .
Run automatically	State of the active full backup schedule. Possible values: • Enabled • Disabled

For example:

```
user@srv01:~$ veeamconfig schedule activefull show --jobname DailyBackup
Every second Monday of every month
Run automatically: enabled
```

Disabling Active Full Backup Schedule

To disable the active full backup schedule for the backup job, use the following command:

```
veeamconfig schedule activefull disable --jobid <job_id>
```

or

```
veeamconfig schedule activefull disable --jobname <job_name>
```

where:

- <job_id> – ID of the backup job for which you want to disable the active full backup schedule.
- <job_name> – name of the backup job for which you want to disable the active full backup schedule.

For example:

```
user@srv01:~$ veeamconfig schedule activefull disable --jobname DailyBackup
```

Enabling Active Full Backup Schedule

After you specify active full backup schedule settings for the backup job, Veeam Agent automatically enables active full backup schedule for the job. You can also enable active full backup schedule manually, for example, if you previously disabled it for the backup job. To enable active full backup schedule, use the following command:

```
veeamconfig schedule activefull enable --jobid <job_id>
```

or

```
veeamconfig schedule activefull enable --jobname <job_name>
```

where:

- <job_id> – ID of the backup job for which you want to enable the active full backup schedule. You should look up the job ID in advance, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- <job_name> – name of the backup job for which you want to enable the active full backup schedule.

For example:

```
user@srv01:~$ veeamconfig schedule activefull enable --jobname DailyBackup
```

You can disable the schedule for the job at any time. To learn more, see [Disabling Backup Schedule](#).

Configuring Health Check Schedule

You can schedule a periodic [health check](#) of a backup that resides in an object storage repository. You can specify backup health check schedule settings individually for every backup job created in Veeam Agent for Linux or backup policy created in Veeam Backup & Replication. For more information on configuring backup health check in a backup policy, see the [Maintenance Settings](#) topic of the Veeam Agent Management Guide.

You can perform the following actions with backup health check schedule in command line interface:

- [Specify health check schedule.](#)
- [Enable health check schedule.](#)
- [View health check schedule.](#)
- [Disable health check schedule.](#)

Specifying Health Check Schedule

You can schedule a backup health check to run on a specific week day of a specific month or on specific days of the week.

Specifying Monthly Health Check Schedule

To instruct Veeam Agent to perform backup health check on a specific week day of a month, use the following command:

```
veeamconfig healthcheck set --light --jobid <job_id> --monthlyweekday <day> --weeknumber <week> [--months <months>]
```

or

```
veeamconfig healthcheck set --light --jobname <job_name> --monthlyweekday <day> --weeknumber <week> [--months <months>]
```

where:

- **<job_id>** – ID of the backup job for which you want to configure health check schedule. You should look up the backup job ID before you configure the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job_name>** – name of the backup job for which you want to configure the health check schedule.

TIP

If the name of the job consists of several words and contains spaces, use quote marks around the name – for example, `--jobName "Files Backup"`.

- **<day>** – day of the week when the backup job must perform health check – for example, `Tuesday` or `Tue`.
- **<week>** – week of the month when the backup job must perform health check. Possible values: `First`, `Second`, `Third`, `Fourth` or `Last`.

- `<months>` – months when the backup job must perform health check, separated by a comma (',') – for example: Jan, Apr, Jul, Oct. If you do not specify this option, the health check will run every month.

For example:

```
user@srv01:~$ veeamconfig healthcheck set --light --jobname SystemBackup --monthlyweekday Fri --weeknumber Last --months Mar,Jun,Sep,Dec
```

Specifying Weekly Health Check Schedule

To instruct Veeam Agent to perform backup health check on specific week days, use the following command:

```
veeamconfig healthcheck set --light --jobid <job_id> --weekdays <days>
```

or

```
veeamconfig healthcheck set --light --jobname <job_name> --weekdays <days>
```

where:

- `<job_id>` – ID of the backup job for which you want to configure health check schedule. You should look up the backup job ID before you configure the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to configure the health check schedule.

TIP

If the name of the job consists of several words and contains spaces, use quote marks around the name – for example, `--jobName "Files Backup"`.

- `<days>` – comma-separated list of days when the backup job must run backup health check. For example: Mon, Fri. The backup job will run the health check on the specified days at the time specified in the backup job schedule settings.

For example:

```
user@srv01:~$ veeamconfig healthcheck set --light --jobname "System Backup" --weekdays mon,fri
```

Enabling Health Check Schedule

After you set health check schedule for a backup job, Veeam Agent automatically enables this backup health check schedule for the job. You can also enable health check schedule manually – for example, if you previously disabled it. To enable health check schedule, use the following command:

```
veeamconfig healthcheck enable --light --jobid <job_id>
```

or

```
veeamconfig healthcheck enable --light --jobname <job_name>
```

where:

- `<job_id>` – ID of the backup job for which you want to configure health check schedule. You should look up the backup job ID before you run the command to enable the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to configure the health check schedule.

TIP

If the name of the job consists of several words and contains spaces, use quote marks around the name – for example, `--jobName "Files Backup"`.

For example:

```
user@srv01:~$ veeamconfig healthcheck enable --light --jobname SystemBackup
```

You can disable health check schedule for a job at any time. To learn more, see [Disabling Health Check Schedule](#).

Viewing Health Check Schedule

To view the health check schedule defined for a backup job, use the following command:

```
veeamconfig healthcheck show --jobid <job_id>
```

or

```
veeamconfig healthcheck show --jobname <job_name>
```

where:

- `<job_id>` – ID of the backup job for which you want to configure health check schedule. You should look up the backup job ID before you run the command to view the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to configure the health check schedule.

TIP

If the name of the job consists of several words and contains spaces, use quote marks around the name – for example, `--jobName "Files Backup"`.

Veeam Agent for Linux displays the following information about the health check schedule:

Parameter	Description
Every <value>	Days on which the backup job runs the health check – for example, <i>Every last Fri of every month</i> .
Run health-check automatically	State of the backup health check schedule. Possible values: • Enabled • Disabled

For example:

```
user@srv01:~$ veeamconfig healthcheck show --jobname SystemBackup
Every last Fri of Mar, Jun, Sep, Dec
Run health check automatically: enabled (light)
```

Disabling Health Check Schedule

To disable the health check schedule for a backup job, use the following command:

```
veeamconfig healthcheck disable --jobid <job_id>
```

or

```
veeamconfig healthcheck disable --jobname <job_name>
```

where:

- **<job_id>** – ID of the backup job for which you want to configure health check schedule. You should look up the backup job ID before you run the command to disable the schedule – for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job_name>** – name of the backup job for which you want to configure the health check schedule.

TIP

If the name of the job consists of several words and contains spaces, use quote marks around the name – for example, `--jobName "Files Backup"`.

For example:

```
user@srv01:~$ veeamconfig healthcheck disable --jobname SystemBackup
```

Configuring Database Processing Settings

You can enable database processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux. With database processing settings enabled, Veeam Agent will create transactionally consistent backups of Veeam Agent machines that run database systems. For information on the limitations of database processing, see [Backup of Database Systems](#).

You can perform the following actions with database processing settings via the command-line interface:

- [Specify Oracle database processing settings](#)
- [Specify MySQL database processing settings](#)
- [Specify PostgreSQL database processing settings](#)
- [View database processing settings](#)
- [Disable database processing settings](#)

Specifying Oracle Processing Settings

You can enable Oracle processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux.

To enable Oracle processing settings for the backup job, use the following command:

```
veeamconfig aap set oracle --jobid <job_id> <oracle_options>
```

or

```
veeamconfig aap set oracle --jobname <job_name> <oracle_options>
```

where:

- `<job_id>` – ID of the backup job for which you want to enable Oracle processing settings. You should look up the job ID in advance, before configuring Oracle processing settings, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to enable Oracle processing settings.
- `<oracle_options>` – Oracle processing settings for the backup job. To learn more, see [Oracle Processing Settings](#).

TIP

To view IDs or names of all existent backup jobs, you can press the [Tab] key right after you type the `--jobid` or `--jobname` option.

Oracle Processing Settings

You can specify the following Oracle processing settings for the backup job:

Option	Description and values
--tryprocess	Defines that Veeam Agent must continue the backup process if errors occur when processing the Oracle database system. If you do not specify this option, Veeam Agent will stop the backup process if an error occurs when processing the Oracle database system.
--prunelogs	The number of hours to keep Oracle archived logs or the size of archived logs to keep. • If you want Veeam Agent to delete archived logs that are older than <N> hours, specify the necessary value in the <N>H format. For example, 10H. • If you want Veeam Agent to delete archived logs that are larger than <N> GB, specify the necessary value in the <N>G format. For example: 10G. Veeam Agent will wait for the backup job to complete successfully and then trigger archived logs truncation via Oracle Call Interface (OCI). If the backup job fails, the logs will remain untouched until the next successful backup job session.
--usroracleos	Name of the Veeam Agent machine OS account. To connect to the Oracle database system, the account must be a member of the group that owns configuration files for the Oracle database (for example, the oinstall group). You do not need this option if you want to use the Oracle account to connect to the database. Instead, specify the necessary account with the <code>--usroracledb</code> option.
--usroracledb	Name of the Oracle account. To connect to the Oracle database system, the account must have SYSDBA rights on the databases to be processed. You do not need this option if you want to use the OS account to connect to the database. Instead, specify the necessary account with the <code>--usroracleos</code> option.

For example:

```
user@srv01:~$ veeamconfig aap set oracle --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75 --tryprocess --prunelogs 10G --usroracledb system
```

Specifying MySQL Processing Settings

You can enable MySQL processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux.

IMPORTANT

MySQL tables that use the MyISAM storage engine must be locked to keep them in consistent state while Veeam Agent is creating the system snapshot. To correctly process such tables, MySQL account must have the following instance-wide privileges:

- `SELECT`. This privilege enables Veeam Agent to access tables' metadata and select for a lock the tables that use the MyISAM storage engine. Without this privilege, the processing of the MySQL database system will run successfully but MyISAM tables will not be locked, which may result in an inconsistent state of the backed up data.
- `LOCK TABLES`. This privilege is required for locking the selected MyISAM tables. If some MyISAM tables are selected but the MySQL account does not have the `LOCK TABLES` privilege, the processing of the MySQL database system will fail.
- `RELOAD` or `FLUSH TABLES`. If some MyISAM tables are selected but the MySQL account does not have either `RELOAD` or `FLUSH TABLES` privilege, the processing of the MySQL database system will fail.

To obtain information about the privileges that are assigned to an account, use MySQL functionality, for example, the `SHOW GRANTS` statement. To learn more, see [MySQL documentation](#).

To enable MySQL processing settings for the backup job, use the following command:

```
veeamconfig aap set mysql --jobid <job_id> <mysql_options>
```

or

```
veeamconfig aap set mysql --jobname <job_name> <mysql_options>
```

where:

- `<job_id>` – ID of the backup job for which you want to enable MySQL processing settings. You should look up the job ID in advance, before configuring MySQL processing settings, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- `<job_name>` – name of the backup job for which you want to enable MySQL processing settings.
- `<mysql_options>` – MySQL processing settings for the backup job. To learn more, see [MySQL Processing Settings](#).

TIP

To view IDs or names of all existent backup jobs, you can press the [Tab] key right after you type the `--jobid` or `--jobname` option.

MySQL Processing Settings

You can specify the following MySQL processing settings for the backup job:

Option	Description and values
--tryprocess	Defines that Veeam Agent must continue the backup process if errors occur when processing the MySQL database system. If you do not specify this option, Veeam Agent will stop the backup process if an error occurs when processing the MySQL database system.
--usrmysqldb	Name of the MySQL account. Veeam Agent can connect to the MySQL database system in one of the following ways: • If you specify account name (--usrmysqldb option) only, Veeam Agent will prompt you to specify a password to access the MySQL database system. • If you specify account name and password (--usrmysqldb and --password options), Veeam Agent will access the MySQL database system. • If you do not specify account credentials (--usrmysqldb and --password options), Veeam Agent will use a password file to connect to the MySQL database system. To learn more about password file configuration, see Preparing Password File for MySQL Processing.
--password	Password of the MySQL account. If you do not specify the --password value, Veeam Agent will prompt you to specify a password to access the MySQL database. Keep in mind, if you specify the password using the --password option, password is stored in terminal in plain text.
--defaults-file	Path to a password file. You must specify a full path to a password file if you want Veeam Agent to use a password file located in specific directory. Specifying relative paths is not supported. With this method selected, you do not need to specify account credentials in the backup job settings. You do not need this option in the following cases: • Veeam Agent uses account name and password that are specified in the backup job settings to connect to the MySQL database. • Veeam Agent uses account credentials that are stored in the password file in <code>/root/.my.cnf</code>.

Examples

Authentication with password:

```
user@srv01:~$ veeamconfig aap set mysql --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75 --tryprocess --usrmysqldb root --password P@ssw0rd
```


Authentication with password file:

```
user@srv01:~$ veeamconfig aap set mysql --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8
ea75 --tryprocess
--defaults-file /data/root/.my.cnf --password P@ssw0rd
```

Preparing Password File for MySQL Processing

You can use MySQL account credentials that are stored in the password file to connect Veeam Agent for Linux to the MySQL database system.

NOTE

Consider the following:

- If you specify a custom path to the password file, specify a full path. Specifying relative paths is not supported.
- The password file can also contain user-specific connection settings that Veeam Agent will apply to connect to the MySQL database system. For example, if you want to connect to the MySQL database system using the custom socket, specify the socket path in the password file. To learn more, see [MySQL documentation](#).

If you want to use a password file for authentication, create a file. By default, Veeam Agent expects the password file to have the `.my.cnf` name and to be in the home directory of the `root` user. If the password file has a custom name or is stored in another directory, you can specify a custom path.

The password file must have the following contents:

```
[client]
user=<username>
password=<password>
```

where:

- `<username>` – name of the account that Veeam Agent will use to connect to the MySQL database system.
- `<password>` – password of the account that Veeam Agent will use to connect to the MySQL database system.

For example:

```
[client]
user=root
password=P@ssw0rd
```

Specifying PostgreSQL Processing Settings

You can enable PostgreSQL processing settings in the properties of a volume-level backup job configured in Veeam Agent for Linux.

To enable PostgreSQL processing settings for the backup job, use the following command:

```
veeamconfig aap set postgres --jobid <job_id> <postgres_options>
```

or

```
veeamconfig aap set postgres --jobname <job_name> <postgres_options>
```

where:

- **<job_id>** – ID of the backup job for which you want to enable PostgreSQL processing settings. You should look up the job ID in advance, before configuring PostgreSQL processing settings, for example, with the `veeamconfig job list` command. To learn more, see [Viewing List of Backup Jobs](#).
- **<job_name>** – name of the backup job for which you want to enable PostgreSQL processing settings.
- **<postgres_options>** – PostgreSQL processing settings for the backup job. To learn more, see [PostgreSQL Processing Settings](#).

TIP

To view IDs or names of all existent backup jobs, you can press the [Tab] key right after you type the `--jobid` or `--jobname` option.

PostgreSQL Processing Settings

You can specify the following PostgreSQL processing settings for the backup job:

Option	Description and values
--tryprocess	Defines that Veeam Agent must continue the backup process if errors occur when processing the PostgreSQL database system. If you do not specify this option, Veeam Agent will stop the backup process if an error occurs when processing the PostgreSQL database system.
--usrpgdb	Name of the PostgreSQL account. If you use a password file to connect to the PostgreSQL database system, the <code>--usrpgdb</code> option allows to select the user from the password file. You do not need this option if you want to use a Veeam Agent machine OS account to connect to the PostgreSQL database system. Instead, specify the OS account with the <code>--usrpgos</code> option.

Option	Description and values
--password	Password of the PostgreSQL account. If you do not specify this option, Veeam Agent will prompt to enter the password. If you do not specify the password in prompt, Veeam Agent uses a password file to connect to the PostgreSQL database system. To learn more about password file configuration, see Preparing Password File for PostgreSQL Processing. Keep in mind, if you specify the password using the <code>--password</code> option, password is stored in terminal in plain text.
--usrpgos	Name of the OS account. Veeam Agent will use the name to connect to the PostgreSQL database system using the peer authentication method. In the peer authentication method, Veeam Agent uses the OS account as the PostgreSQL database user name. With this option selected, you must specify OS account only. To learn more about peer authentication, see PostgreSQL documentation. You do not need this option if you want to use a PostgreSQL account to connect to the database system. Instead, specify the PostgreSQL account with the <code>--usrpgdb</code> option.

For example:

```
user@srv01:~$ veeamconfig aap set postgres --jobid 29bc2e1a-e35c-4efb-8d37-b717
7b8ea75 --tryprocess --usrpgdb postgres --password P@ssw0rd
```

Preparing Password File for PostgreSQL Processing

You can use PostgreSQL account credentials that are stored in the password file to connect Veeam Agent to the PostgreSQL database system.

If you want to use a password file for authentication, create the `.pgpass` file in the home directory of the `root` user.

The password file must have the following contents:

```
<hostname>:<port>:<database>:<username>:<password>
```

where:

- `<hostname>` – name of the host where the PostgreSQL database system is located.
- `<port>` – number of the free port that Veeam Agent will use to connect to the PostgreSQL database system.
- `<database>` – name of the PostgreSQL database.
- `<username>` – name of the account that Veeam Agent will use to connect to the PostgreSQL database system.
- `<password>` – password of the account that Veeam Agent will use to connect to the PostgreSQL database system.

For example:

```
srv01:5432:mydb:postgres:P@ssw0rd
```

For more information about the password file, see [PostgreSQL documentation](#).

Viewing Database Processing Settings

To view database processing settings defined for the backup job, use the following command:

```
veeamconfig aap show --jobid <job_id>
```

or

```
veeamconfig aap show --jobname <job_name>
```

where:

- `<job_id>` – ID of the backup job for which you want to view database processing settings.
- `<job_name>` – name of the backup job for which you want to view database processing settings.

Veeam Agent for Linux displays the following information about database processing settings:

- [Oracle processing settings](#)
- [MySQL processing settings](#)
- [PostgreSQL processing settings](#)

Oracle Processing Settings

Parameter	Description
Oracle processing	Oracle processing settings status. Possible values: • <i>Required</i> – Oracle processing settings are enabled for the job. If an error occurs when processing the Oracle database system, Veeam Agent will stop the backup process. • <i>Try</i> – Oracle processing settings are enabled for the job. If an error occurs when processing the Oracle database system, Veeam Agent will continue the backup process. • <i>Disabled</i> – Oracle processing settings are disabled for the job using command line interface.

Parameter	Description
Account used for processing	Account used to connect to the Oracle database. Possible values: <i>System account (username: <username>)</i> – if Veeam Agent connects to the Oracle database system with the account of the Veeam Agent machine OS. <i>Oracle account (username: <username>)</i> – if Veeam Agent connects to the Oracle database system with the Oracle account. where <username> is a name of the user account that Veeam Agent will use to connect to the Oracle database.
Delete logs over <N> Gb	Veeam Agent displays this information if Veeam Agent is set to delete archived logs that are larger than <N> GB.
Delete logs older <N> Hr	Veeam Agent displays this information if Veeam Agent is set to delete archived logs that are older than <N> hours.

For example:

```
user@srv01:~$ veeamconfig aap show --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75
Oracle processing: required
  Account used for processing: Oracle account (username: sys)
  Delete logs over 10 Gb
```

MySQL Processing Settings

Parameter	Description
MySQL processing	MySQL processing settings status. Possible values: <i>Required</i> – MySQL processing settings are enabled for the job. If an error occurs when processing a MySQL database, Veeam Agent will stop the backup process. <i>Try</i> – MySQL processing settings are enabled for the job. If an error occurs when processing a MySQL database, Veeam Agent will continue the backup process. <i>Disabled</i> – MySQL processing settings are disabled for the job using command line interface.
Account used for processing	Veeam Agent displays this information if Veeam Agent is set to connect to the MySQL database system with the account name and password.
Path to a password file	Veeam Agent displays this information if Veeam Agent is set to connect to the MySQL database system with the account credentials that are stored in the password file.

For example:

```
user@srv01:~$ veeamconfig aap show --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75
MySQL processing: required
Account used for processing: username: root
```

PostgreSQL Processing Settings

Parameter	Description
PostgreSQL processing	PostgreSQL processing settings status. Possible values: • <i>Required</i> – PostgreSQL processing settings are enabled for the job. If an error occurs when processing a PostgreSQL database, Veeam Agent will stop the backup process. • <i>Try</i> – PostgreSQL processing settings are enabled for the job. If an error occurs when processing a PostgreSQL database, Veeam Agent will continue the backup process. • <i>Disabled</i> – PostgreSQL processing settings are disabled for the job using command line interface.
Account used for processing	Account used to connect to the PostgreSQL database. Possible values: • <i><username> (password)</i> – Veeam Agent displays this information if Veeam Agent is set to connect to the PostgreSQL database with the PostgreSQL account. • <i><username> (file)</i> – Veeam Agent displays this information if Veeam Agent is set to connect to the PostgreSQL database with the password file. • <i><username> (peer)</i> – Veeam Agent displays this information if Veeam Agent is set to connect to the PostgreSQL database with the Veeam Agent machine OS account. where <i><username></i> is a name of the account that Veeam Agent will use to connect to the PostgreSQL database.

For example:

```
user@srv01:~$ veeamconfig aap show --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea75
PostgreSQL processing: required
Account used for processing: postgres (password)
```

Disabling Database Processing Settings

To disable database processing settings defined for the backup job, use the following command:

```
veeamconfig aap disable <db_sys> --jobid <job_id>
```

or

```
veeamconfig aap disable <db_sys> --jobname <job_name>
```

where:

- `<db_sys>` – name of the database system that you want to disable. Possible values:
 - *oracle* – Oracle database processing to be disabled.
 - *mysql* – MySQL database processing to be disabled.
 - *postgres* – PostgreSQL database processing to be disabled.
- `<job_id>` – ID of the backup job for which you want to disable database processing settings.
- `<job_name>` – name of the backup job for which you want to disable database processing settings.

For example:

```
user@srv01:~$ veeamconfig aap disable oracle --jobid 29bc2e1a-e35c-4efb-8d37-b7177b8ea759
Oracle processing was disabled.
```

Starting and Stopping Backup Jobs

You can start a backup job manually at any time you need, for example, if you want to create an additional restore point for Veeam Agent backup and do not want to change the job schedule. You can also stop the running backup job before the job session completes, if necessary.


You can start and stop backup jobs in one of the following ways:

- [With the Veeam Agent control panel.](#)
- [With the Veeam Agent command line interface.](#)

Starting Backup Job from Control Panel

To start a backup job with the Veeam Agent control panel, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. Press the [S] key to open the **Select job to start** dialog window.
3. Select the necessary backup job in the list and start the job in on of the following ways:
 - To start an incremental backup job session, press [Enter].
 - To create an active full backup, press [F].




4. Veeam Agent will immediately start the backup job and display a notification window informing that the job has been started. Press [Enter] to close the window and proceed to the list of backup job sessions.

You can monitor the backup job performance in the Veeam Agent control panel. To learn more, see [Viewing Real-Time Job Session Statistics](#).

If you start the backup job while another backup job is running, Veeam Agent will perform the backup job immediately after the current job is completed. For details, see [Job Queue](#).

Job Queue

If another backup job is running when you start the backup job, Veeam Agent will submit this backup job to job queue. Veeam Agent will perform the job in the queue as soon as the previous job is completed.



The queued backup job creates a new session with the *Pending* status. You can view all jobs in the queue in the **Latest backup sessions** list in the Veeam Agent control panel.

Veeam Agent for Linux [srv01]

Latest backup sessions:

Job name	State	Started at	Finished at
DailyBackup	Pending (0%)	---	---
SystemBackup	Running (19%)	2023-08-07 18:58:55	---

Enter Show C Configure S Start Job R Recover Files M Misc Esc Quit

- NOTE**
- Consider the following about job queue:
- Job queue can contain up to 3 backup jobs besides the job that is already running.
 - You cannot submit the same backup job to the queue if it is already running.

Starting Backup Job from Command Line Interface

You can start a backup job with the command line interface. When you start a backup job, Veeam Agent initiates a new backup job session and provides you with a Session ID. You can monitor the progress of the backup job session or view the session status. To start a backup job, use the following command:

```
veeamconfig job start --name <job_name>
```

or

```
veeamconfig job start --id <job_id>
```

where:

- `<job_name>` – name of the backup job that you want to start.
- `<job_id>` – ID of the backup job that you want to start.

TIP

Consider the following:

- You can use the `veeamconfig job start` command with the `--nosnap` option to start a file-level backup job. In this case, Veeam Agent will not create a snapshot of the backed-up volume during the backup job session. Keep in mind that the snapshot-less file-level backup does not guarantee that data in the backup is consistent. To learn more, see [Snapshot-Less File-Level Backup](#).
- You can use the `veeamconfig job start` command with the `--activefull` option to create active full backups. To learn more, see [Creating Active Full Backups](#).

For example:

```
$ veeamconfig job start --name SystemBackup
Backup job has been started.
Session ID: [{381532f7-426a-4e89-b9fc-43d98942c71a}].
Logs stored in: [/var/log/veeam/Backup/SystemBackup/Session_20250307_162608_{381532f7-426a-4e89-b9fc-43d98942c71a}].
```

Veeam Agent will start the backup job. You can [check the backup job session status](#) or [view the backup job session log](#) using the Veeam Agent command line interface.

You can also monitor the backup job performance in the Veeam Agent control panel. To learn more, see [Viewing Real-Time Job Session Statistics](#).

If you start the backup job while another backup job is running, Veeam Agent will perform the backup job immediately after the current job is completed. For details, see [Job Queue](#).

Job Queue

If another backup job is running when you start the backup job, Veeam Agent will submit this backup job to job queue. Veeam Agent will perform the job in the queue as soon as the previous job is completed.

```
$ veeamconfig job start --name DailyBackup
The backup job has been added to the queue.
Session ID: [{10e8c599-b2aa-4008-89d9-af9b6e04aeba}].
Logs stored in: [/var/log/veeam/Backup/DailyBackup/Session_20250314_153342_{10e8c599-b2aa-4008-89d9-af9b6e04aeba}].
```

The queued backup job creates a new session with the *Pending* status. You can view all jobs in the queue by running the `veeamconfig session list` command.

```
$ veeamconfig session list
Job name      Type      ID                               State      Started
at           Finished at
SystemBackup  Backup    {37427202-b139-4b36-9982-e0c33894d0cc} Running    2025-03
-14 15:33
DailyBackup   Backup    {10e8c599-b2aa-4008-89d9-af9b6e04aeba} Pending
```

NOTE

Consider the following about job queue:

- Job queue can contain up to 3 backup jobs besides the job that is already running.
- You cannot submit the same backup job to the queue if it is already running.

Creating Active Full Backups in Command Line Interface

You can create an ad-hoc full backup – active full backup, and add it to the backup chain on the target storage. 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 on the target storage until it is removed from the backup chain according to the retention policy.

Before you create an active full backup, check the following prerequisites:

- The backup job must be configured.
- You cannot create an active full backup if a backup task of any type is currently running.

To perform active full backup, use the following command:

```
veeamconfig job start --name <job_name> --activefull
```

or

```
veeamconfig job start --id <job_id> --activefull
```

where:

- `<job_name>` – name of the backup job that you want to start to create an active full backup.
- `<job_id>` – ID of the backup job that you want to start to create an active full backup.

For example:

```
$ veeamconfig job start --name SystemBackup --activefull
Backup job has been started.
Session ID: [{ce864e24-8211-4df7-973a-741adce96fe7}].
Logs stored in: [/var/log/veeam/Backup/SystemBackup/Session_20250311_150046_{ce
864e24-8211-4df7-973a-741adce96fe7}].
```

Veeam Agent will start the backup job. You can view the progress for the active full backup session in the same way as for any other backup job session: you can [check the backup job session status](#) or [view the backup job session log](#) using the Veeam Agent command line interface.

You can also monitor the backup job performance in the Veeam Agent control panel. To learn more, see [Viewing Real-Time Job Session Statistics](#).

Stopping Backup Job

You can stop the running backup job before the job session completes, for example, if the backup process is about to take long, and you do not want the job to increase server load in the production environment during business hours.

When you stop a backup job, the job session will finish immediately. Veeam Agent will not produce a new restore point during the session, and the session will finish with the *Failed* status.


You can stop a job in one of the following ways:

- [Using the Veeam Agent control panel](#)
- [In the command line interface](#)

Stopping Job from Control Panel

To stop a backup job:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, in the list of backup job sessions, select the currently running session with the [Up] and [Down] keys and press [Enter].
3. In the session statistics window, press [S].
4. In the displayed window, make sure that the **Yes** button is selected and press [Enter].



Stopping Job in Command Line Interface

To stop a backup job, use the following command:

```
veeamconfig session stop --id <session_id>
```

or

```
veeamconfig session stop --force --id <session_id>
```

where:

- `<session_id>` – ID of the currently running backup job session that you want to stop.
- `--force` – with this option enabled, Veeam Agent will immediately stop the backup session even if it is unable to stop the *veeamjobman* process for some reason.

For example:

```
$ veeamconfig session stop --id 381532f7-426a-4e89-b9fc-43d98942c71b  
Session has stopped.
```

TIP

You can use the `veeamconfig session stop` command to remove a backup job with the *Pending* status from the job queue.

Managing Backup Jobs

You can perform the following actions with backup jobs configured in Veeam Agent for Linux:

- [View the list of configured backup jobs.](#)
- [View information about the backup job settings.](#)
- [Edit the backup job settings.](#)
- [Delete a backup job.](#)

Viewing List of Backup Jobs

To view the list of backup jobs configured in Veeam Agent, use the following command:

```
veeamconfig job list
```

In the list of backup jobs, Veeam Agent displays the following information:

Parameter	Description
Name	Name of the backup job.
ID	ID of the backup job.
Type	Backup type. Possible values: • All System – the backup job is set to back up the entire system (the root mount point). • VolumeLevel – the backup job is set to back up specific system volumes. • FileLevel – the backup job is set to back up specific files and directories.
Repository	Name of the backup repository that is specified as backup storage for the backup job.

For example:

```
user@srv01:~$ veeamconfig job list
Name                               ID                               Type      Repos
itory
SystemBackup                      {2495911e-58db-4452-b4d1-f53dcfbc600e} All System Repos
itory_1
DocumentsBackup                   {bcf821e6-b35f-4d57-b1c3-d3a477605cb9} FileLevel  Repos
itory_1
HomePartitionBackup              {2aaa8c71-2434-4f12-a168-3d8e225fa416} VolumeLevel Repos
itory_2
```

Viewing Backup Job Settings

To view detailed information about the backup jobs settings, use the following command:

```
veeamconfig job info --name <job_name>
```

or

```
veeamconfig job info --id <job_id>
```

where:

- **<job_name>** – name of the backup job for which you want to view settings.
- **<job_id>** – ID of the backup job for which you want to view settings.

Veeam Agent for Linux displays the following information about the backup job:

Parameter	Description
ID	ID of the backup job.
Name	Name of the backup job.
Repository ID	ID of the backup repository that is specified as a backup storage for the backup job.
Repository name	Name of the backup repository that is specified as a backup storage for the backup job.
Creation time	Date and time of the backup job creation.
Compression	Data compression level. Possible values are: • <i>0</i> – No compression • <i>1</i> – Rle • <i>2</i> – Lz4 • <i>3</i> – Zstd 3 • <i>4</i> – Zstd 9
Max Points	Number of restore points to keep on disk. By default, Veeam Agent for Linux keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent for Linux will remove the earliest restore point from the backup chain.
Index	File system indexing options defined for the backup job.

Parameter	Description
Objects for backup	Backup scope specified for the backup job.

For example:

```

user@srv01:~$ veeamconfig job info --name SystemBackup
Backup job
  ID: {2495911e-58db-4452-b4d1-f53dcfbc600e}
  Name: SystemBackup
  Repository ID: {4557ef7a-9c44-4f28-b8d0-44d78e5ddd5d}
  Repository name: Repository_1
  Creation time: 2017-04-06 13:29:03
  Options:
    Compression: Lz4
    Max Points: 7
    Index all mounted filesystems on the volumes selected for backup
  Objects for backup:
    Include Disk: sda1

```

Editing Backup Job Settings

If you want to change settings of the backup job, you can edit it at any time. For example, you may want to edit the backup job to add a new directory to the backup scope or change the target location.

To edit a backup job, use the following command:

For volume-level backup jobs

```
veeamconfig job edit volumelevel <option> for --name <job_name>
```

or

```
veeamconfig job edit volumelevel <option> for --id <job_id>
```

For file-level backup jobs

```
veeamconfig job edit filelevel <option> for --name <job_name>
```

or

```
veeamconfig job edit filelevel <option> for --id <job_id>
```

where:

- `<option>` – option that you want to edit for the job. You can specify one or several options at a time. To learn more about available options, see [Backup Job Options](#).
- `<job_name>` – name of the backup job that you want to edit.
- `<job_id>` – ID of the backup job that you want to edit.

For example:

```
user@srv01:~$ veeamconfig job edit volumelevel --name SystemVolumeBackup for --  
name SystemVolume
```

Backup Job Options

You can use the following options to edit parameters for the backup job:

Option	Description and values
--compressionlevel	Data compression level. Possible values are: • 0 – No compression • 1 – Rle • 2 – Lz4 • 3 – ZlibLow • 4 – ZlibHigh
--blocksize	Data block size in kilobytes. The default value is 1024. Possible values are 256, 512, 1024 or 4096. To learn more about available sizes of data blocks, see Data Compression .
--maxpoints	Number of restore points that you want to store in the backup location. By default, Veeam Agent keeps 7 latest restore points. When the new restore point that exceeds the specified number is created, Veeam Agent will remove the earliest restore point from the backup chain.
--prefreeze	Pre-freeze command that should be executed before the snapshot creation.
--postthaw	Post-thaw command that should be executed after the snapshot creation.
--objects	Object that should be included in backup: • For simple volumes – name of a block device that represents a volume or an entire disk that should be included in backup. You can specify entire disk to create backup of the entire computer image or individual computer volumes to create backup of specific volumes. If you want to back-up several disks or volumes, specify them one after another using a ',' (comma) character as a separator. • For LVM volumes – name of an LVM logical volume that should be included in backup. If you want to back-up several LVM logical volumes, specify them one after another using a ',' (comma) character as a separator. This option is available for volume-level backup jobs only.
--includedirs	Full path to a directory that should be included in backup, for example: <code>/home/user</code>. The option is available for file-level backup jobs only. You can specify one or several paths to directories in the computer file system. To separate several paths, use a ',' (comma) character, for example: <code>/home/user/Documents, /home/user/reports</code>.

Option	Description and values
--excludedirs	Full path to a directory that should be excluded from backup. The option is available for file-level backup jobs only. The directory specified with this option must be a subdirectory of the directory specified with the <code>--includedirs</code> option. To separate several paths, use a ',' (comma) character, for example, <code>/home/user/Documents,/home/user/reports</code>.
--includemasks	Mask for file name or path that should be included in backup. The option is available for file-level backup jobs only. You can use the following UNIX wildcard characters for file name masks: '*' — a substitution for one or more characters in the file name or path. Can be used for any sequence of characters (including no characters). For example, <code>*.pdf</code>. '?' — a substitution of one character in the file name or path. For example, <code>repor?.pdf</code>. '[' — a substitution of one character in the file name or path with any of the characters enclosed in square brackets (or a range of characters defined with the '-' character). For example: <code>report_201[3456].pdf</code> or <code>report_201[3-6].pdf</code>. To separate several masks, use a ',' (comma) character, for example, <code>report.*,reports.*</code>. File inclusion option is applied to all directories that are specified with the <code>--includedirs</code> option. For example, if you include in backup the <code>/home/user/Documents</code> directory and files that match the <code>repor?.pdf</code> file name mask, Veeam Agent will back up the <code>/home/user/Documents/report.pdf</code> file and will not back up the <code>/home/user/reports/report.pdf</code> file.

Option	Description and values
--excludemasks	Mask for file name or path that should be excluded from backup. The option is available for file-level backup jobs only. You can use the following UNIX wildcard characters for file name masks: '*' – a substitution for one or more characters in the file name or path. 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. '[' – a substitution of one character in the file name or path with any of the characters enclosed in square brackets (or a range of characters defined with the '-' character). For example: report_201[3456].pdf or report_201[3-6].pdf. To separate several masks, use a ',' (comma) character, for example, report.*,reports.*. File exclusion option is applied to all directories that are specified with the --includedirs option and files that match file name masks specified with the --includemasks option. For example, you may want to specify the following backup scope for the backup job: - Include in backup the /home/user/Documents directory - Include files that match the report.* file name mask - Exclude files that match the *.odt file name mask. In this case, Veeam Agent will backup the /home/user/Documents/report.pdf file and will not backup /home/user/Documents/report.odt and /home/user/reports/report.pdf files. If you want to use several name masks, you must specify them in double quotation marks, for example: veeamconfig job create filelevel --name BackupJob1 --reponame vault13 --includedirs /home --includemasks "*.bak,*.pdf".
--indexnothing	Defines that file system indexing options are disabled for the backup job.
--indexall	Defines that Veeam Agent must index all files on the volumes included in backup.
--indexonly	Path to a directory that contains files that you want to index. Enter paths to the necessary directories. To separate several paths, use the ',' (comma) character. The option is available for volume-level backup jobs only.
--indexexcept	Path to a directory that contains files that you do not want to index. You can specify one or more paths. To separate several paths, use the ',' (comma) character. The option is available for volume-level backup jobs only.

Option	Description and values
--setencryption	Defines that data encryption option is enabled for the job. You can use this option to enable encryption for the existing backup job or change a password used for encryption for the backup job. When you use the <code>veeamconfig job edit</code> command with the <code>--setencryption</code> option, Veeam Agent for Linux will prompt you to specify a password for data encryption and hint for the password.
--resetencryption	Defines that data encryption option is disabled for the job. You can use this option to disable encryption for the existing backup job.
--deleteold	The number of days to keep the backup created with the backup job in the target location. If Veeam Agent for Linux does not create new restore points for the backup, the backup will remain in the target location for the specified number of days. When this period is over, the backup will be removed from the target location. Possible values are: 1-999. If you do not specify the <code>--deleteold</code> option, Veeam Agent will not apply this setting. The backup will be stored in the target location until you delete it manually. If you specified the value earlier and want to disable this setting, specify the <i>false</i> value for this option: <code>--deleteold false</code>. After the next successful backup session, this setting will be disabled for the backup in the target location.
--nosnap	Defines whether Veeam Agent must perform backup in the snapshot-less mode. Possible values: <i>true</i> – if you use this option, Veeam Agent will create a snapshot of the backed-up volumes during file-level backup. <i>false</i> – if you use this option, Veeam Agent will not create a snapshot of the backed-up volumes during file-level backup. Keep in mind that the snapshot-less file level backup does not guarantee that data in the backup is consistent. To learn more, see Snapshot-Less File-Level Backup.

NOTE

Consider the following:

- If you change the target location for the backup job, during the next backup job session Veeam Agent for Linux will perform full data backup. All subsequent backup sessions will produce incremental backups – Veeam Agent for Linux will copy only changed data to the target location and add a new incremental backup file to the backup chain.
- If you change the backup scope for the backup job, during the next backup job session Veeam Agent for Linux will create a new incremental backup. The backup will contain all data blocks pertaining to new data added to the backup scope and changed data blocks pertaining to original data in the backup scope (data that was processed by the job at the time before you changed the backup scope).
- If you enable or disable encryption for the existing backup job that has already created one or more restore points, during the next job session, Veeam Agent for Linux will create active full backup.
- Full backup takes much more time than incremental backup. If you change the target location, you can copy an existing backup chain to the new location manually. In this case, the new backup job session will produce an incremental backup file and add it to the backup chain.

Deleting Backup Job

You can delete a backup job configured in Veeam Agent for Linux. When you delete a backup job, backup files created by this job remain intact in the backup repository.


You can delete backup jobs in one of the following ways:

- [In the Veeam Agent control panel](#)
- [In the Veeam Agent command line interface](#)

Deleting Backup Job in Control Panel

To delete a backup job using the Veeam Agent control panel, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. Press the [C] key to open the **Select job to configure** dialog window or the [S] key to open the **Select job to start** dialog window.
3. Select the necessary backup job in the list and press [Delete].
4. In the displayed notification window, make sure that the **Yes** button is selected and press [Enter].



Deleting Backup Job in Command Line Interface

To delete a backup job in the Veeam Agent command line interface, use the following command:

```
veeamconfig job delete --name <job_name>
```

or

```
veeamconfig job delete --id <job_id>
```

where:

- <job_name> – name of the backup job that you want to delete.
- <job_id> – ID of the backup job that you want to delete.

For example:

```
$ veeamconfig job delete --name SystemBackup
```

Managing Backup Repositories

A backup repository is a storage location where Veeam Agent for Linux keeps backup files. You can use the following types of storage as a target location for a backup job:

- Local (internal) storage of the protected machine (not recommended).
- Direct attached storage (DAS), such as USB, eSATA or Firewire external drives.
- Network Attached Storage (NAS) able to represent itself as SMB (CIFS) or NFS share.
- Object storage repository, such as S3 Compatible storage, Amazon S3, Google Cloud or Microsoft Azure Blob.
- 12.1 or later backup repository (including deduplication appliances).
- Veeam Backup & Replication 12.0 or later cloud repository.

IMPORTANT

Consider the following about backup repositories:

- [For local storage] A backup repository should be created on a separate volume from the volume that contains data you plan to back up.
- [For Veeam backup repository] Backup repositories with enabled KMS encryption are not supported.

Creating Backup Repository

Veeam Agent for Linux creates backup files in a backup repository. When you create a backup job with the Backup Job wizard, you must specify a target location for the backup. Veeam Agent will create a backup repository in the specified location and save information about this repository in the database.

IMPORTANT

A backup repository must be created on a separate volume from a volume whose data you plan to back up.

If you want to create backups in local directory, network shared folder or object storage, you must create a repository. To learn more, see the following sections:

- [Creating a repository in a local directory.](#)
- [Creating a repository in an NFS network shared folder.](#)
- [Creating a repository in an SMB network shared folder.](#)
- [Creating a repository in object storage.](#)

If you want to create Veeam Agent backups in a Veeam backup repository or cloud repository, you do not need to create repositories. Before configuring the backup job, you need to connect to the Veeam backup server or Veeam Cloud Connect service provider. To learn more, see the following sections:

- [Connecting to Veeam Backup Server](#)
- [Connecting to Service Provider](#)

Creating Repository in Local Directory

To create a repository in a local directory, use the following command:

```
veeamconfig repository create --name <repository_name> --location <path_to_repository>
```

where:

- `<repository_name>` – name of the repository.
- `<path_to_repository>` – path to the directory in which backup files will be stored.

For example:

```
$ veeamconfig repository create --name VeeamBackup --location /home/backups
```

Creating Repository in NFS Share

To create a repository in an NFS share, use the following command:

```
veeamconfig repository create --name <repository_name> --type nfs --location <path_to_repository>
--options <mounting_options>
```

where:

- `<repository_name>` – name of the backup repository.
- `<path_to_repository>` – path to the network shared folder where backup files will be stored in the *SERVER:/DIRECTORY* format.
- `<mounting_options>` – additional options that Veeam Agent will use to mount the network shared folder to the Veeam Agent machine file system. You can use the standard Linux `mount` command content as mounting options. This parameter is optional.

For example:

```
$ veeamconfig repository create --name VeeamBackup --type nfs --location srv01:/VeeamRepository --options vers=3,hard,retry=1
```

TIP

If you mount a network shared folder to a directory in the Veeam Agent machine file system in advance, you can create the backup repository in the same way as in a local directory. For details, see [Creating Repository in Local Directory](#).

Creating Repository in SMB Share

To create a repository in an SMB share, use the following command:

```
veeamconfig repository create --name <repository_name> --type smb --location <path_to_repository>
--username <user_name> --password <password> --domain <domain> --options <mounting_options>
```

where:

- `<repository_name>` – name for the backup repository.
- `<path_to_repository>` – path to the network shared folder where backup files will be stored in the *//SERVER/DIRECTORY* format.
- `<user_name>` – account name that Veeam Agent will use to access the SMB network shared folder.
- `<domain>` – domain in which the account that has access permissions on the shared folder is registered.

- `<mounting_options>` – options that Veeam Agent will use to mount the network shared folder to the Veeam Agent machine file system. You can use the standard Linux `mount` command content as mounting options. This parameter is optional.

You can specify account name and domain for the SMB network shared folder using the `--username` and `--domain` parameters. If a password is required to access the network shared folder, you must also specify the `--password` parameter. When you run the `veeamconfig repository create` command, Veeam Agent will prompt you to type a password of the specified account.

Alternatively, you can specify account name, password and domain for the network shared folder as values for the `--options` parameter. Mind that these values will override values of the `--username`, `--password` and `--domain` parameters.

Examples

Command with `--username`, `--password` and `--domain` parameters:

```
$ veeamconfig repository create --name VeeamBackup --type smb --location //srv02/VeeamRepository --username Administrator --password --domain srv02
```

Command with `--options` parameter:

```
$ veeamconfig repository create --name VeeamBackup --type smb --location //srv02/VeeamRepository --options username=Administrator,password=P@ssw0rd,domain=srv02,port=666
```

TIP

If you mount a network shared folder to a directory in the Veeam Agent machine file system in advance, you can create the backup repository in the same way as in a local directory. For details, see [Creating Repository in Local Directory](#).

Creating Repository in Object Storage

To create a repository in an object storage location, you must specify a storage provider name, a name for the backup repository and settings for the object storage account and bucket or container.

Before You Begin

Before you start creating an object storage repository, consider the following:

- [Microsoft Azure Blob storage] The soft delete feature for blobs and containers must be disabled in the storage account.
- [Microsoft Azure Blob storage] To use the [Veeam backup immutability feature](#), you must enable blob versioning and version-level immutability support in the storage account. For more information, see [this Microsoft Azure documentation](#).
- [S3 Compatible and Amazon S3 storage] To use the [Veeam backup immutability feature](#), you must enable versioning and the S3 Object Lock feature in the storage account. For more information, see [this Amazon S3 documentation](#).

- [Google Cloud storage] The Veeam backup immutability feature is not supported for repositories configured in Google Cloud storage.

Creating Object Storage Repository

To create an object storage repository, use the following command:

```
veeamconfig objectstorage createrepository <provider_type> <options>
```

where:

- `<provider_type>` – name of the object storage provider. Veeam Agent supports the following options:
 - `azureblob` – for creating a Microsoft Azure Blob repository.
 - `google` – for creating a Google Cloud repository.
 - `amazons3` – for creating an Amazon S3 repository.
 - `s3compatible` – for creating an S3 Compatible repository (including repositories in IBM Cloud and Wasabi Cloud).
- `<options>` – options necessary to connect to the target object storage. For more information, see the following subsections:
 - [Specifying options for S3 Compatible repository](#)
 - [Specifying options for Amazon S3 repository](#)
 - [Specifying Options for Google Cloud repository](#)
 - [Specifying options for Microsoft Azure Blob repository](#)

After Veeam Agent creates a new backup repository in the object storage location, you can specify object storage as a destination for the backup job.

Specifying Options for S3 Compatible Repository

To create a backup repository in an S3 compatible storage bucket, use the following command:

```
veeamconfig objectstorage createrepository s3compatible --name <repository_name> --servicepoint <address> --region <storage_region> --accesskeyid <id> [--fingerprint <ssl_thumbprint>] --bucketname <bucket_name> --folder <folder_name>
```

where:

- `<repository_name>` – name for the backup repository.
- `<address>` – address of the service point for the object storage.

NOTE

If you want to connect to the repository using the IPv6 address and the port number, you must use the following format: `IPv6:port`, where:

- `IPv6` is the IPv6 address of the object storage.
 - `port` is the number of the port that Veeam Agent will use to connect to the object storage.
- `<storage_region>` – region associated with the bucket.

NOTE

You can find the list of supported regions in the documentation of the selected storage provider.

- `<id>` – access key associated with the object storage account.
- `<ssl_thumbprint>` – fingerprint to verify the SSL certificate.
- `<bucket_name>` – name of the bucket.
- `<folder_name>` – name of the folder in the bucket.

If Veeam Agent does not find a folder with the name specified in the command, Veeam Agent will create a new folder in the bucket under `Veeam/Backup/` – for example, `Veeam/Backup/folder01`.

For example:

```
$ veeamconfig objectstorage createrepository s3compatible --name s3comp --servicepoint fd00:ca19:0:18b0:0:ac8a:abca:c942:9000 --accesskeyid S3ertlD9EIO9DjnZjuD4 --region us-east-1 --fingerprint <value> --bucketname backup01 --folder folder01
```

After you use the command, Veeam Agent will prompt you to specify a secret key associated with the object storage account. Enter the secret key to complete the creation of the repository.

Specifying Options for Amazon S3 Repository

To create a backup repository in an Amazon S3 bucket, use the following command:

```
veeamconfig objectstorage createrepository amazons3 --name <repository_name> --accesskeyid <id> --region <storage_region> --bucketname <bucket_name> --folder <folder_name>
```

where:

- `<repository_name>` – name for the backup repository.
- `<id>` – access key associated with the Amazon S3 storage account.
- `<storage_region>` – region associated with the bucket.

NOTE

You can find the full list of supported regions by storage provider in the `PublicCloudRegions.xml` file located in the `/Library/Application Support/Veeam` folder on Veeam Agent computer.

- `<bucket_name>` – name of the 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. For example, it is not recommended that you use dots (.) in the bucket name. For more information on bucket naming rules, see [this AWS documentation article](#).

- `<folder_name>` – name of the folder in the bucket.

If Veeam Agent does not find a folder with the name specified in the command, Veeam Agent will create a new folder in the bucket under `Veeam/Backup/` – for example, `Veeam/Backup/folder01`.

For example:

```
$ veeamconfig objectstorage createrepository amazons3 --name amazon --accesskey  
id AMAZONKIAWHDY4BDYCJC --region us-east-1 --bucketname bucket01 --folder folde  
r01
```

After you use the command, Veeam Agent will prompt you to specify a secret key associated with the Amazon S3 storage account. Enter the secret key to complete the creation of the repository.

Specifying Options for Google Cloud Repository

To create a backup repository in a Google Cloud storage bucket, use the following command:

```
veeamconfig objectstorage createrepository google --name <repository_name> --ac  
cesskeyid <id> --region <storage_region> --bucketname <bucket_name> --folder <f  
older_name>
```

where:

- `<repository_name>` – name for the backup repository.
- `<id>` – access key associated with the Google Cloud storage account.
- `<storage_region>` – region associated with the bucket.

NOTE

You can find the full list of supported regions by storage provider in the `PublicCloudRegions.xml` file located in the `/Library/Application Support/Veeam` folder on Veeam Agent computer.

- `<bucket_name>` – name of the bucket.

- `<folder_name>` – name of the folder in the bucket.

If Veeam Agent does not find a folder with the name specified in the command, Veeam Agent will create a new folder in the bucket under `Veeam/Backup/` – for example, `Veeam/Backup/folder01`.

For example:

```
$ veeamconfig objectstorage createrepository google --name google --accesskeyid
GOOGLE56L5ATTDKJCLWUQG3E --region europe-west3 --bucketname backup01 --folder f
older01
```

After you use the command, Veeam Agent will prompt you to specify a secret key associated with the Google Cloud storage account. Enter the secret key to complete the creation of the repository.

Specifying Options for Microsoft Azure Blob Repository

To create a backup repository in a Microsoft Azure Blob container, use the following command:

```
veeamconfig objectstorage createrepository azureblob --name <repository_name> -
-account <storage_account_name> --region <storage_region> --bucketname <bucket_
name> --folder <folder_name>
```

- `<repository_name>` – name of the backup repository for the Veeam Agent database.
- `<account>` – name of the Microsoft Azure Blob storage account.
- `<storage_region>` – region associated with the container.

NOTE

Veeam Agent supports specification of 3 generic Microsoft Azure Blob storage locations:

- **Azure Global (Standard)** – can be used for any data center region, except the regions in China and the regions intended for US governments. To specify this region in the command to create the repository, use the following value: `AzureCloud`.
- **Asia China** – can be used for any region in China. To specify this region in the command to create the repository, use the following value: `AzureChinaCloud`.
- **Azure Government** – can be used for Azure Government regions only. To specify this region in the command to create the repository, use the following value: `AzureGovernmentCloud`.

You can find the full list of supported regions by storage provider in the `PublicCloudRegions.xml` file located in the `/Library/Application Support/Veeam` folder on Veeam Agent computer.

- `<bucket_name>` – name of the container.
- `<folder_name>` – name of the folder in the container.

If Veeam Agent does not find a folder with the name specified in the command, Veeam Agent creates a new folder in the container under `Veeam/Backup/` – for example, `Veeam/Backup/folder01`.

For example:

```
$ veeamconfig objectstorage createrepository azureblob --name azure --account my-account --region azurecloud --bucketname backup01 --folder folder01
```

After you use the command, Veeam Agent will prompt you to specify the shared key associated with the object storage account. Enter the shared key to complete the creation of the repository.

Viewing List of Backup Repositories

To view backup repositories configured in Veeam Agent for Linux, use the following command:

```
veeamconfig repository list
```

Veeam Agent will display a list of backup repositories.

You can view the following information about backup repositories:

Parameter	Description
Name	Name of the backup repository.
ID	ID of the backup repository.
Location	Directory in the local file system specified as a target location for backup files.
Type	Type of the backup repository. Possible values: - Local - Backup server
Backup server	Backup server on which Veeam backup repository added to Veeam Agent is configured.

For example:

```
user@srv01:~$ veeamconfig repository list
Name      ID                                     Location      Type      Backup server
BackupVol01 {818e3a0f-8155-4a51-9430-248a203a43d1} /home/backups local
BackupVol02 {2155a2e7-a1e9-4347-9d8b-cf8f3a6f3fcb} /home/backups2 local
```

Editing Backup Repository Settings

In command line interface, you can edit settings for a backup repository created with Veeam Agent for Linux in a local or network shared folder.

You can edit properties of the following repository types only in the backup job settings in the Veeam Agent control panel:

- Veeam backup repository.
- Veeam Cloud Connect repository.
- Object storage repository.

You can edit the following parameters for the backup repository:

- [Name of the backup repository](#)
- [Location of the backup repository](#)

NOTE

Consider the following:

- If you change location for the backup repository that is already used by a backup job and contains backup files, during the next backup job run, Veeam Agent will create a new backup chain in the new repository location.
- You can temporarily change backup repository location if you want to create an ad hoc full backup in addition to the backup chain created by the backup job in the original repository location.

Changing Backup Repository Name

To change a name for the backup repository, use the following command:

```
veeamconfig repository edit --name <new_name> for --name <old_name>
```

or

```
veeamconfig repository edit --name <new_name> for --id <id>
```

where:

- `<old_name>` – current name of the backup repository.
- `<new_name>` – desired name for the backup repository.
- `<id>` – ID of the backup repository.

For example:

```
user@srv01:~$ veeamconfig repository edit --name LocalRepository for --name Repository_1
```

Changing Backup Repository Location

To change location for the backup repository, use the following command:

```
veeamconfig repository edit --location <path> for --name <name>
```

or

```
veeamconfig repository edit --location <path> for --id <id>
```

where:

- <path> – desired path for the backup repository.
- <name> – current name of the backup repository.
- <id> – ID of the backup repository.

For example:

```
user@srv01:~$ veeamconfig repository edit --location /home/veeam for --id 3458797-3ffe-45bc-870e-c5628643bbb3
```

Changing Backup Repository Name and Location

You can change a name and location for the backup repository at the same time, for example:

```
user@srv01:~$ veeamconfig repository edit --name LocalRepository --location /home/veeam for --name Repository_1
```


Rescanning Veeam Backup Repository

If Veeam Agent for Linux fails to display backups stored in the Veeam Backup & Replication backup repository for some reason, you can rescan the Veeam backup repository. Veeam Agent will try to reconnect to the Veeam backup server and refresh the list of backups in the backup repository.

To rescan a Veeam backup repository, use the following command:

```
veeamconfig repository rescan --id <repository_id>
```

or

```
veeamconfig repository rescan --name <repository_name>
```

where:

- `<repository_id>` – ID of the backup repository that you want to rescan.
- `<repository_name>` – name of the backup repository that you want to rescan.

For example:

```
user@srv01:~$ veeamconfig repository rescan --name [vbr01]BackupVol01
```

You can also rescan all Veeam backup repositories managed by the backup server to which Veeam Agent is connected with the following command:

```
veeamconfig repository rescan --all
```

NOTE

When you use the `veeamconfig repository rescan` command with the `--all` option, consider the following:

- Rescanning can take significant amount of time if there are multiple repositories configured in Veeam Agent.
- Rescanning multiple object storage repositories may result in greater storage costs due to additional volume of data transactions.

TIP

You can also use the `veeamconfig repository rescan` command to rescan local backup repositories. This may be useful, for example, after information about a backup stored in the local repository is deleted from the Veeam Agent configuration database, or after you copy a backup to the local repository.

Deleting Backup Repository

You can delete a backup repository configured with Veeam Agent for Linux. When you delete a backup repository, Veeam Agent removes record of the deleted repository from its database. Backup files created by a backup job targeted at the deleted backup repository remain intact on the backup storage.

To delete a backup repository, use the following command:

```
veeamconfig repository delete --id <repository_id>
```

or

```
veeamconfig repository delete --name <repository_name>
```

where:

- `<repository_id>` – ID of the backup repository that you want to delete.
- `<repository_name>` – name of the backup repository that you want to delete.

For example:

```
user@srv01:~$ veeamconfig repository delete --name Repository_1
```

NOTE

You cannot delete a backup repository that is specified as a backup storage location in the backup job settings.

Managing Veeam Backup & Replication Servers

You can store backup files created with Veeam Agent for Linux on backup repositories managed by Veeam Backup & Replication. To do this, you must [connect to a Veeam backup server](#). After that, you can specify a Veeam backup repository as a target location for backup files [in the properties of the backup job](#).

Connecting to Veeam Backup Server

To create Veeam Agent backups on a backup repository managed by Veeam Backup & Replication, you must connect to a Veeam backup server.

IMPORTANT

Currently, Veeam Agent for Linux can be connected to one Veeam Backup & Replication server only. If you want to create backups on the backup repository managed by another Veeam backup server, you need to delete currently used backup server and all jobs targeted at backup repositories managed by this backup server. To learn more, see [Deleting Connection to Veeam Backup Server](#).

If you add a connection to another backup server, backup jobs targeted at the original backup server will fail, and backups created on the Veeam backup repository will become unavailable in Veeam Agent. To continue using the original backup server, you need to delete the connection to the new backup server and re-create all backup jobs that use the original backup server.

If you change an account to connect to the Veeam backup server and then start a backup job targeted at the backup repository managed by this backup server, Veeam Agent will start a new backup chain on the backup repository.

To connect Veeam Agent for Linux to a Veeam backup server, use the following command:

```
veeamconfig vbrserver add --name <vbr_name> --address <vbr_address> --port <vbr_port> --login <username> --domain <domain> --password <password>
```

where:

- `<vbr_name>` – name of the Veeam backup server that manages the backup repository.
- `<vbr_address>` – DNS name or IP address of the Veeam backup server.
- `<vbr_port>` – port over which Veeam Agent must communicate with Veeam Backup & Replication. The default port used for communication with the Veeam backup server is 10006.
- `<username>` – a name of the account that has access to the Veeam backup repository.
- `<domain>` – a name of the domain in which the account that has access to the Veeam backup repository is registered.
- `<password>` – password of the account that has access to the Veeam backup repository.

Permissions on the backup repository managed by the target Veeam backup server must be granted beforehand. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

For example:

```
user@srv01:~$ veeamconfig vbrserver add --name vbr01 --address 172.17.53.1 --port 10006 --login veeam --domain tech --password P@ssw0rd
```

When Veeam Agent for Linux connects to a Veeam Backup & Replication server, Veeam Agent retrieves information about backup repositories managed by this Veeam backup server and displays them in the list of available backup repositories. You can then specify a Veeam backup repository as a target for a backup job.

TIP

To view the list of backup repositories, use the `veeamconfig repository list` command. To learn more, see [Viewing List of Backup Repositories](#).

Viewing List of Veeam Backup Servers

To view a list of Veeam backup servers to which Veeam Agent for Linux is connected, use the following command:

```
veeamconfig vbrserver list
```

Veeam Agent will display the list of Veeam backup servers.

For the Veeam backup server in the list, Veeam Agent for Linux displays the following information:

Parameter	Description
Name	Name of the Veeam backup server.
ID	ID of the Veeam backup server in the Veeam Agent database.
Endpoint	IP address of the Veeam backup server and port over which Veeam Agent for Linux communicates with Veeam Backup & Replication.
Username	Name of the account that has access to the Veeam backup repository.
Domain	Name of the domain in which the account that has access to the Veeam backup repository is registered.

For example:

```
user@srv01:~$ veeamconfig vbrserver list
Name      ID                               Endpoint      Usernam
e         Domain
vbr01     {0fc87c11-6a8d-48c1-8aeb-7f7655738796} 172.17.53.1:10006 veea
m         tech
```

Viewing Backup Server Details

You can view detailed information about the Veeam backup server to which Veeam Agent for Linux is connected. Use the following command:

```
veeamconfig vbrserver info --name <vbr_name>
```

or

```
veeamconfig vbrserver info --id <vbr_id>
```

where:

- <vbr_name> – name of the Veeam backup server.
- <vbr_id> – ID of the Veeam backup server in the Veeam Agent database.

Veeam Agent for Linux displays the following information about the Veeam backup server:

Parameter	Description
ID	ID of the Veeam backup server in the Veeam Agent database.
Name	Display name of the Veeam backup server.
Endpoint	IP address of the Veeam backup server and port over which Veeam Agent for Linux communicates with Veeam Backup & Replication.
Login	Name of the account that has access to the Veeam backup repository.
Domain	Name of the domain in which the account that has access to the Veeam backup repository is registered.

For example:

```
user@srv01:~$ veeamconfig vbrserver info --name vbr01
VBR server
  ID: {0fc87c11-6a8d-48c1-8aeb-7f7655738796}
  Name: vbr01
  Endpoint: 172.17.53.1:10006
  Login: veeam
  Domain: tech
```

Editing Connection to Veeam Backup Server

You can edit the following parameters for a connection to a Veeam backup server:

- [Display name of the Veeam backup server](#)
- [IP address and port used to connect to the Veeam backup server](#)
- [Account to connect to the Veeam backup server](#)

Changing Veeam Backup Server Name

To change a name for the Veeam backup server, use the following command:

```
veeamconfig vbrserver edit --name <new_vbr_name>
```

where:

<new_vbr_name> – desired name for the backup server.

For example:

```
user@srv01:~$ veeamconfig vbrserver edit --name vbr01
```

Changing IP Address and Port for Veeam Backup Server

To change the IP address and port used to connect to the Veeam backup server, use the following command:

```
veeamconfig vbrserver edit --address <vbr_address> --port <vbr_port>
```

where:

- <vbr_address> – DNS name or IP address of the Veeam backup server.
- <vbr_port> – port over which Veeam Agent for Linux must communicate with Veeam Backup & Replication.

For example:

```
user@srv01:~$ veeamconfig vbrserver edit --address 172.17.53.1 --port 10006
```


Changing Account to Connect to Veeam Backup Server

NOTE

If you change an account to connect to the Veeam backup server and then start a backup job targeted at the backup repository managed by this backup server, Veeam Agent will start a new backup chain on the backup repository.

To change an account whose credentials will be used to connect to the Veeam backup server, use the following command:

```
veeamconfig vbrserver edit --login <username> --domain <domain> --password
```

where:

- `<username>` – name of the account that has access to the Veeam backup repository.
- `<domain>` – name of the domain in which the account that has access to the Veeam backup repository is registered.

When you run the command, Veeam Agent will prompt you to enter the password of the specified account.

For example:

```
user@srv01:~$ veeamconfig vbrserver edit --login veeam --domain tech --password  
Enter password:
```

Changing Several Backup Server Parameters

You can change several parameters for the connection to the Veeam backup server simultaneously. For example, the following command changes the name and connection settings for the Veeam backup server:

```
user@srv01:~$ veeamconfig vbrserver edit --name vbr02 --address 172.17.53.2 --p  
ort 10006
```

Updating List of Veeam Backup Repositories

When you connect to a Veeam backup server, Veeam Agent for Linux retrieves information about backup repositories managed by this Veeam backup server and displays them in the list of available backup repositories. You can refresh information about available Veeam backup repositories manually at any time. This may be useful, for example, after a new backup repository was added on the Veeam backup server.

To update the list of backup repositories managed by the Veeam backup server, use the following command:

```
veeamconfig vbrserver resync
```

TIP

To view updated list of available Veeam backup repositories after resync, use the `veeamconfig repository list` command. To learn more, see [Viewing List of Backup Repositories](#).

Deleting Connection to Veeam Backup Server

You can delete a connection to the Veeam backup server to which Veeam Agent is currently connected. When you delete a connection to a Veeam backup server, Veeam Agent removes record on the deleted backup server from its database. Veeam backup repositories managed by the deleted backup server are removed from the list of available backup repositories. Backup files created by backup jobs targeted these repositories remain intact on the backup storage.

You cannot delete a connection to a Veeam backup server in the following situations:

- Veeam Agent operates in the managed mode. To delete connection to a Veeam backup server, reset Veeam Agent to the standalone mode. For details, see [Resetting to Standalone Operation Mode](#).
- Veeam Agent has a backup job that saves backup files to a repository managed by this backup server. To remove such connection to a Veeam backup server, you first need to delete reference to the Veeam backup repository in the job settings.

To delete a connection to the Veeam backup server, use the following command:

```
veeamconfig vbrserver delete --name <vbr_name>
```

or

```
veeamconfig vbrserver delete --id <vbr_id>
```

where:

- <vbr_name> – name of the Veeam backup server.
- <vbr_id> – ID of the Veeam backup server.

For example:

```
user@srv01:~$ veeamconfig vbrserver delete --name vbr01
```

Managing Service Providers

You can store backup files created with Veeam Agent for Linux on a cloud repository exposed to you by a Veeam Cloud Connect service provider. To do this, you must [connect to a service provider](#). After that, you can specify a cloud repository as a target location for backup files [in the properties of the backup job](#).

Connecting to Service Provider

To create Veeam Agent backups on a cloud repository, you must connect to a Veeam Cloud Connect service provider.

To connect Veeam Agent for Linux to a service provider, use the following command:

```
veeamconfig cloud add --name <sp_name> --address <sp_address> --port <sp_port>
--login <username> --password <password> --fingerprint <sp_thumbprint>
```

where:

- `<sp_name>` – name of the service provider to which you want to connect.
- `<sp_address>` – IP address or full DNS name of the cloud gateway that the SP or your backup administrator has provided to you.
- `<sp_port>` – port over which Veeam Agent must communicate with the cloud gateway. The default port used for communication with the cloud gateway is 6180.
- `<username>` – name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT/SUBTENANT* format.
- `<password>` – password of the tenant or subtenant account used to connect to the service provider.
- `<sp_thumbprint>` – thumbprint used to verify the TLS certificate that the SP has provided to you.

For example:

```
user@srv01:~$ veeamconfig cloud add --name SP --address 172.17.53.15 --port 6180
--login TechCompany/User01 --password P@ssw0rd --fingerprint 92FA988A3D9E80EE
095DDAB75BF06B05DF6F205B
```

NOTE

When you enter the `veeamconfig cloud add` command, Veeam Agent will display information about the TLS certificate obtained from the SP. To accept the certificate, type `yes` in the command prompt and press [Enter].

When Veeam Agent connects to the service provider, Veeam Agent retrieves information about cloud repositories available to the tenant or subtenant and displays them in the list of available backup repositories. You can then specify a cloud repository as a target for a backup job.

TIP

To view the list of available cloud repositories, use the `veeamconfig repository list` command. To learn more, see [Viewing List of Backup Repositories](#).

Viewing List of Service Providers

To view a list of service providers to which Veeam Agent is connected, use the following command:

```
veeamconfig cloud list
```

Veeam Agent will display the list of service providers.

For the service provider in the list, Veeam Agent for Linux displays the following information:

Parameter	Description
Name	Name of the service provider.
ID	ID of the service provider in the Veeam Agent database.
Address	IP address of the cloud gateway and port over which Veeam Agent communicates with the cloud gateway.
Gate servers	IP address of the cloud gateway and port over which Veeam Agent communicates with the cloud gateway.
Username	Name of the tenant or subtenant account used for connection to the service provider.

For example:

```
user@srv01:~$ veeamconfig cloud list
Name      ID                               Address      Gate
servers   Username
SP         {0840f770-354d-426a-b5ce-1aa80f56cc08}  172.17.53.15:618
0          TechCompany
```

Editing Connection to Service Provider

You can edit the following parameters for a connection to a Veeam Cloud Connect service provider:

- [Name of the Veeam Cloud Connect service provider](#)
- [IP address and port used to connect to the cloud gateway](#)
- [Account to connect to the service provider](#)
- [Thumbprint to connect to the service provider](#)

Changing SP Name

To change a name for the SP, use the following command:

```
veeamconfig cloud edit --name <new_sp_name> for --name <old_sp_name>
```

or

```
veeamconfig cloud edit --name <new_sp_name> for --id <sp_id>
```

where:

- `<old_sp_name>` – current name of the SP.
- `<new_sp_name>` – desired name for the SP.
- `<sp_id>` – ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --name SP for --id 7d3022de-4f4d-4c70-85eb-e8a946a555cd
```

Changing IP Address and Port for Cloud Gateway

To change the IP address and port of the cloud gateway provided by the SP, use the following command:

```
veeamconfig cloud edit --address <sp_address> --port <sp_port> for --name <sp_name>
```

or

```
veeamconfig cloud edit --address <sp_address> --port <sp_port> for --id <sp_id>
```

where:

- `<sp_address>` – IP address or full DNS name of the cloud gateway that the SP or your backup administrator has provided to you.
- `<sp_port>` – port over which Veeam Agent must communicate with the cloud gateway. The default port used for communication with the cloud gateway is 6180.
- `<sp_name>` – name of the SP.
- `<sp_id>` – ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --address 172.17.53.67 --port 6180 for --name SP
```

Changing Account to Connect to SP

To change an account whose credentials will be used to connect to the SP, use the following command:

```
veeamconfig cloud edit --login <username> --password <password> for --name <sp_name>
```

or

```
veeamconfig cloud edit --login <username> --password <password> for --id <sp_id>
```

where:

- `<username>` – name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT/SUBTENANT* format.
- `<password>` – password of the tenant or subtenant account used to connect to the service provider.
- `<sp_name>` – name of the SP.
- `<sp_id>` – ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --login ABC_Compan/User01 --password P@ssw0rd for --name SP
```


Changing Thumbprint to Connect to SP

To change a thumbprint that will be used to connect to the SP, use the following command:

```
veeamconfig cloud edit --fingerprint <sp_thumbprint> for --name <sp_name>
```

or

```
veeamconfig cloud edit --fingerprint <sp_thumbprint> for --id <sp_id>
```

where:

- <sp_thumbprint> – thumbprint used to verify the TLS certificate and connect to the service provider.
- <sp_name> – name of the SP.
- <sp_id> – ID of the SP.

For example:

```
user@srv01:~$ veeamconfig cloud edit --fingerprint 92FA988A3D9E80EE095DDAB75BF0  
6B05DF6F205B for --name SP
```

Updating List of Cloud Repositories

When you connect to the Veeam Cloud Connect service provider, Veeam Agent for Linux retrieves and saves to the database information about cloud repositories available to the tenant or subtenant whose account you use to connect to the SP. You can refresh information about available cloud repositories manually at any time. This may be useful, for example, after the SP changes backup resource settings for the tenant.

To update the list of cloud repositories, use the following command:

```
veeamconfig cloud resync
```

If the cloud repository currently used as a target location for Veeam Agent backups becomes unavailable, and Veeam Agent fails to reflect this change in its database for some reason, the `veeamconfig cloud resync` command may finish with errors. In this case, you can use the `--force` option to refresh information about available cloud repositories. For example:

```
veeamconfig cloud resync --force
```

With the `--force` option, Veeam Agent will retrieve the list of available cloud repositories from the service provider and save the new information about cloud repositories in the Veeam Agent database.

TIP

To view updated list of available cloud repositories after resync, use the `veeamconfig cloud list` command. To learn more, see [Viewing List of Service Providers](#).

Deleting Connection to Service Provider

You can delete a connection to the service provider to which Veeam Agent for Linux is currently connected. When you delete a connection to a service provider, Veeam Agent removes the record on the deleted service provider from the database. Cloud repositories managed by the deleted service provider are removed from the list of available backup repositories. Backup files created by backup jobs targeted at these repositories remain intact on the cloud repository.

You cannot delete a connection to the service provider if a cloud repository managed by this service provider is used by a backup job. To remove such connection to a service provider, you first need to delete a reference to the cloud repository in the job settings.

To delete a connection to the service provider, use the following command:

```
veeamconfig cloud delete --name <sp_name>
```

or

```
veeamconfig cloud delete --id <sp_id>
```

where:

- <sp_name> – name of the service provider.
- <sp_id> – ID of the service provider.

For example:

```
user@srv01:~$ veeamconfig cloud delete --name SP
```

Managing Backups

You can perform the following operations with backups created by backup jobs configured in Veeam Agent for Linux:

- [View backups](#)
- [View backup details](#)
- [View restore points in backup](#)
- [Export backup to a virtual disk](#)
- [Import backup to the Veeam Agent database](#)
- [Delete backup](#)

Viewing Backups

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list [--all]
```

Where `--all` is an option that instructs Veeam Agent to display information about all Veeam Agent for Linux backups in the backup repositories configured in the product. If you do not use this option, Veeam Agent will display information about the backups of the current Veeam Agent computer only.

If you work with Veeam Agent connected to a Veeam backup server as a member of a protection group for pre-installed Veeam Agents, for security reasons, the `veeamconfig backup list --all` command will display backups created only by the current Veeam Agent computer with the current connection settings. To learn more about protection groups for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

NOTE

If you cannot locate the backup from which you want to restore data, make sure Veeam Agent has access to the backup repository that contains this backup. To learn more about configuring backup repositories, see [Managing Backup Repositories](#), [Managing Veeam Backup & Replication Servers](#) and [Managing Service Providers](#).

You can also import a backup if it is stored on the Veeam Agent computer or in a network shared folder. For example, this can be a backup created with another instance of Veeam Agent. To learn more about backup import, see [Importing Backups](#).

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as <i>Imported</i> in the Repository column. For information about the import procedure, see Importing Backups .
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name      Backup ID      Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocumentsBackup {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
wrk01 SystemBackup {951ac571-dd29-45ac-8624-79b8ccb45863} Repository_
2    2023-11-13 15:26
wrk02 SystemBackup {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167} Repository_
3    2023-11-13 15:59
```

Viewing Backup Details

You can view detailed information about specific backup. To view backup details, use the following command:

```
veeamconfig backup show --id <backup_id>
```

where:

<backup_id> – ID of the backup for which you want to view detailed information.

For a volume-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Name	Name of the volume in the backup.
Device	Path to the block device file that represents the volume.
FS UUID	File system ID.
Offset	Position of the volume on the computer disk.
Size	Size of the volume in the backup.

For example:

```
user@srv01:~$ veeamconfig backup show --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
Machine name: srv01 SystemBackup
  Name:      [sda1]
  Device:    [/dev/sda1]
  FS UUID:   [6945f2eb-e8bb-48fe-a276-5ba67b9030a5]
  Offset:    [1048576] bytes (2048 sectors)
  Size:      [9999220736] bytes (19529728 sectors)
```

For a file-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Backed up	Backup scope for the file-level backup job.

For example:

```
user@srv01:~$ veeamconfig backup show --id ea64a7e5-038a-4c86-970a-6d59d4cf3968
Machine name: srv01 DocsBackup
File-level backup
Backed up:
/home/user/Documents
```


Viewing Restore Points in Backup

To view information about restore points in the backup, you can use one of the following commands:

```
veeamconfig backup info --id <backup_id>
```

or

```
veeamconfig point list --backupid <backup_id>
```

where:

<backup_id> – ID of the backup for which you want to view information on restore points.

For example:

```
user@srv01:~$ veeamconfig backup info --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
```

or

```
user@srv01:~$ veeamconfig point list --backupid 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
```

Veeam Agent for Linux displays the following information about restore points in the backup:

Parameter	Description
Job name	Name of the backup job by which the backup was created.
OIB ID	ID of the restore point in the backup.
Type	Type of the restore point. Possible values: • Full • Increment
Created at	Date and time of the restore point creation.
Is corrupt	Indicates whether restore point in the backup is corrupted. Possible values: • True • False
Retention	Displays information about enabled long-term retention per each type: weekly (W), monthly (M) and yearly (Y).

For example:

```
user@srv01:~$ veeamconfig backup info --id 535c93a1-2d4e-4475-afb1-a5d46114dca8
Job name      OIB ID      Type      Created
at           Is corrupt  Retention
srv01 SystemBackup {83b76b05-d0a6-4c2e-ae73-727c9aa2cee8} Full      2024-11-
29 00:30 false      WM
srv01 SystemBackup {5b4b7e2b-eab7-401f-8528-5510c4dcc88c} Increment 2024-11-
30 00:30 false
srv01 SystemBackup {054dfac4-4dda-4b16-b24c-acf1f474b3e6} Increment 2024-12-
01 00:30 false
srv01 SystemBackup {0d06ae01-e483-4878-9162-a1765a07df55} Increment 2024-12-
02 00:30 false
```

Importing Backups

You can import a backup created by Veeam Agent into the Veeam Agent database. For example, you may want to import a previously deleted backup or backup that was created in a network shared folder by Veeam Agent installed on another computer.

To import a backup:

1. Start the import process with the following command:

```
veeamconfig backup import --path <path>
```

where:

<path> – path to the VBM file of the backup that you want to import.

For example:

```
user@srv01:~$ veeamconfig backup import --path /home/share/BackupJob/BackupJob.vbm
Backup has been imported successfully.
Session ID: [{4031f058-766c-4f2c-a7ae-7257adb2929f}].
Logs stored in: [/var/log/veeam/Import/Session_{4031f058-766c-4f2c-a7ae-7257adb2929f}].
```

2. You can monitor the import process and result by viewing the import session log with the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session_id> – ID of the import session.

For example:

```
user@srv01:~$ veeamconfig session log --id 4031f058-766c-4f2c-a7ae-7257adb2929f
2023-11-19 13:21:33 UTC {765af178-a9cc-4596-8bf2-03850c5dalac} [info] Job started at 2023-11-19 16:21:33
2023-11-19 13:21:33 UTC {6ae2922d-454b-4a8d-a11b-2b5c7a85029d} [info] Importing backup
2023-11-19 13:21:33 UTC {783f40a7-ead7-4555-9c35-545d875990ee} [info] Backup has been imported.
```

3. Imported backup will be displayed in the list of backups. To view the list of backups, use the following command:

```
veeamconfig backup list
```

For example:

```
user@srv01:~$ veeamconfig backup list
Job name          Backup ID          Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocsBackup   {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomeBackup   {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
BackupJob          {64957b1d-d219-456c-a9cd-9598292c10cd} Importe
d    2023-11-19 19:12
```

Importing Encrypted Backups

You can import an encrypted backup created by Veeam Agent into the Veeam Agent database. This operation is required if you want to use the Veeam Agent command line interface to restore data from an encrypted backup created by Veeam Agent running on another computer.

To import an encrypted backup:

1. Start the import process with the following command:

```
veeamconfig backup import --path <path>
```

where:

<path> – path to the VBM file of the backup that you want to import.

For example:

```
user@srv01:~$ veeamconfig backup import --path /home/share/srv15\ Backup/B
ackup.vbm
```

2. Veeam Agent will prompt you to provide a password for the backup file. Type in the password and press [Enter] **key** to import the backup.

Veeam Agent displays a hint for the password that was used to encrypt the backup file. Use the hint to recall the password.

If you enter the correct password, Veeam Agent will decrypt the backup file and import it into the database.

```
user@srv01:~$ veeamconfig backup import --path /home/share/srv15\ Backup/B
ackup.vbm
[Info] Backup srv15 Backup encrypted
[Info] Press "Enter" to skip. Enter password to decrypt the backup:
[Info] Hint: Standard password
Password:
Backup imported successfully
```

3. Imported backup will be displayed in the list of backups. To view the list of backups, use the following command:

```
veeamconfig backup list
```

For example:

```
user@srv01:~$ veeamconfig backup list
Job name          Backup ID          Repository
y    Created at
srv15 Backup      {4b1f873c-857d-b984-4f22-6ce66bf62570} Importe
d      2018-06-12 20:20
srv01 ServerBackup {f212f641-54aa-40de-a0eb-8727be56760b} Importe
d      2018-06-12 20:04
```

Deleting Backups

Backup files created with Veeam Agent are removed automatically according to the retention policy settings. You can also remove backups from the target location and Veeam Agent configuration database manually if necessary.

Removing Backup from Configuration

To remove a backup from the Veeam Agent configuration database, use the following command:

```
veeamconfig backup delete --id <backup_id>
```

where `<backup_id>` is an ID of the backup that you want to delete.

The way Veeam Agent removes a backup from configuration depends on the backup location:

- If the backup resides in a local directory or network shared folder, Veeam Agent removes records about the deleted backup from the Veeam Agent database. Backup files themselves (VBK, VIB, VBM) remain in the backup repository.

You can import information about the removed backup later to Veeam Agent and perform restore operations with the imported backup. To import information about the removed backup, use the `veeamconfig repository rescan --all` command.

- If the backup resides in a Veeam Backup & Replication repository, Veeam Agent removes records about the deleted backup from the Veeam Agent database and Veeam Backup & Replication database. Backup files themselves (VBK, VIB, VBM) remain in the backup repository.

If you want to import information about the removed backup later to Veeam Agent and perform restore operations with this backup, you must contact backup administrator working with Veeam Backup & Replication. The administrator must rescan the backup repository that contained the backup in the Veeam Backup & Replication console. For details, see the [Rescanning Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

After rescan, the backup will be displayed in the list of backups on the Veeam Agent machine connected to the Veeam backup server.

Deleting Backup Files

To delete backup files from the target location and Veeam Agent database, use the following command:

```
veeamconfig backup delete --id <backup_id> --purge
```

where `<backup_id>` is an ID of the backup that you want to delete.

Veeam Agent for Linux will remove records about the deleted backup from the Veeam Agent database and, additionally, delete backup files themselves from the destination storage.

Performing Restore

If you experience a problem with your computer, your data gets lost or corrupted, you can use one of the following options to recover your data or bring the computer back to work:

- [Restore from the Veeam Recovery Media](#)
 - [Restore volumes](#)
 - [Restore files and folders](#)
- [Restore volumes with the command line interface](#)
- Restore files and folders:
 - [Restore files and folders with the File Level Restore wizard](#)
 - [Restore files and folders with the command line interface](#)
- [Export data as VHD disks](#)
- [Restore data from encrypted backups](#)

Restoring from Veeam Recovery Media

If the OS on your computer fails to start, you can use the Veeam Recovery Media to recover your computer. The Veeam Recovery Media will help you boot the computer in the limited mode. After booting, you can use a backup created with Veeam Agent for Linux to restore the whole system image of your computer, specific volumes on your computer or specific files and folders. You can also use standard Linux command line utilities to diagnose problems and fix errors.

IMPORTANT

If you plan to use the custom Veeam Recovery Media, Veeam Agent requires 3 GB RAM or more installed on the target computer or virtual machine. Memory consumption varies depending on size and number of modules included into the recovery media. To learn more, see [Creating Custom Veeam Recovery Media](#).

Restoring Volumes

You can restore a specific computer volume or all volumes from the volume-level backup.

Volumes can be restored to their original location or to a new location.

- If you restore a volume to its original location, Veeam Agent will overwrite the data on the original volume with the data restored from the backup.
- If you restore volume data to a new location, Veeam Agent will restore data from the backup and write it to the selected destination. If necessary, you can specify new disk mapping settings for the restored volume.

Before You Begin

Before you boot from the recovery image and restore your data, check the following prerequisites and limitations:

- You must have a recovery image on any type of media: CD/DVD/BD or removable storage device.
- To recover data on your computer, you must have both the Veeam Recovery Media and data backup. For volume-level restore, you can use a volume-level backup created with Veeam Agent for Linux. Make sure that the backup or system image is available on the computer drive (local or external), on a network shared folder or on the backup repository managed by a Veeam backup server.
- The media type on which you have created the recovery image must be set as a primary boot source on your computer.
- The volume-level backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a Veeam backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
- You cannot restore a volume to the volume where the backup file that you use for restore is located.
- If you restore to a virtual environment, note that the current version of Veeam Recovery Media supports only the VMware and Hyper-V virtualization solutions. To resolve possible issues during bare metal recovery of Oracle VM virtual machines, use instructions in the second section of [this Veeam KB article](#).

Step 1. Boot from Veeam Recovery Media

To boot from the Veeam Recovery Media:

1. [For CD/DVD/BD] Power on your computer. Insert the media with the recovery image to the drive and power off the computer.

[For removable storage device] Attach the removable storage device with the recovery image to your computer.
2. Start your computer.
3. If you boot from a regular recovery image, in the boot menu, select what Linux kernel version to use to boot your computer and specify boot options if necessary.

You can select a Linux kernel version if you use generic Veeam Recovery Media downloaded from [the Veeam website](#) or [Veeam software repository](#). If you created a custom Veeam Recovery Media, you will be prompted to boot using the Linux kernel of your Veeam Agent computer included in the recovery image.

If you boot from a [legacy recovery image](#), the boot menu is unavailable. After you start your computer, Veeam Agent will immediately start loading files from the Veeam Recovery Media.



TIP

On legacy BIOS-based systems, the boot menu has a different interface but provides the same options.


4. Wait for Veeam Agent to load files from the Veeam Recovery Media.

5. After the recovery image OS has loaded, choose whether you want to start the SSH server. The SSH server allows you to connect to the Veeam Recovery Media from a remote machine.

The Veeam Recovery Media starts the SSH server automatically after a time-out. The default value for the time-out is 60 seconds.

If you do not want to start the SSH server, make sure that the **Proceed without SSH** button is selected and press [Enter]. You will proceed immediately to the step 7.

To override the default time-out and start the SSH server immediately, select the **Start SSH now** button using the [Tab] key and press [Enter].



6. After the SSH server has started, review settings to connect to the Veeam Recovery Media and press [Enter].

The Veeam Recovery Media displays the following connection settings:

- IP address of the computer booted from the Veeam Recovery Media
- User name and password of the account used to connect to the Veeam Recovery Media
- Fingerprints of the computer booted from the Veeam Recovery Media

NOTE

The user name of the account used to work with the Veeam Recovery Media is *veeamuser*.

If you want to use command-line utilities built in the regular recovery image, use the `sudo` command to provide the *veeamuser* account with privileges of the *root* account.

```
SSH Connection Info

Credentials
login: veeamuser
passwd: kaAnL


NetConfigs
ens160
  IP: 172.24.28.72
  IPv6: fd00:ac18:0:1810:0:b5f9:ab46:e4c5

Fingerprints
ecdsa-sha2-nistp256
  SHA1:BaXFVwjaWKUf6Rvv2gAwR+g+knI
  MD5:6d:9c:56:1d:62:d3:f6:56:f0:0e:62:25:31:da:3c:a2
ssh-ed25519
  SHA1:618oSzFazLsSUaMDD/EQJCymqjc
  MD5:2b:2b:5d:78:14:66:55:da:cc:7e:6a:bb:29:a3:01:da
ssh-rsa
  SHA1:6PsfT1Vv+Gkn8dgdR7420HasGBQ
  MD5:ff:96:15:0b:e4:30:86:67:08:8e:7b:21:47:0c:b4:0a
```

[Continue]

7. Accept the terms of the product license agreement and license agreements for third-party components of the product:
 - a. Make sure that the **I accept Veeam End User Software License Agreement** option is selected and press [Space].
 - b. Select the **I accept the terms of the following 3rd party software components license agreements** option with the [Tab] key and press [Space].


c. Switch to the **Continue** button with the [Tab] key and press [Enter].



8. Make sure that network settings are specified correctly and configure the network adapter if necessary. To learn more, see [Configure Network Settings](#).
9. Choose the necessary recovery option. Veeam Agent offers the following tools:
 - **Restore volumes** – the Veeam Recovery wizard to recover data on the original computer or perform bare metal recovery.
 - **Restore files** – the File Level Restore wizard to restore files and folders to the original location or to a new location.
 - **Exit to shell** – Linux shell prompt with standard utilities to diagnose problems and fix errors.

TIP


To stop working with the Veeam Recovery Media and shut down or restart your computer, in the Veeam Recovery Media main menu, select the **Reboot** or **Shutdown** option and press [Enter].




Step 2. Configure Network Settings

If there is a DHCP server in your network, Veeam Agent will configure the network settings automatically. To verify or configure network settings manually, use **nmtui**, a text-based user interface network manager tool provided with Veeam Recovery Media. To learn more about working with nmtui, see [Linux documentation](#).

1. In the Veeam Recovery Media main menu, select **Configure network** and press [Enter].




2. To add new or modify existing connection, in NetworkManager, select **Edit a connection**.



3. After you add or edit a connection, in the main menu of the NetworkManager, select **Activate a connection**.
 - a. If the connection is new, choose it in the list of connections; then select **Activate**.
 - b. If the connection was modified, you must reactivate it. To do this, choose it in the list of connections and select **Deactivate**; then choose the connection again and select **Activate**.
4. After you finish working with Network Manager, press [Esc] to return to the Veeam Recovery Media main menu and launch the Volume Restore wizard.

Step 3. Launch Volume Restore Wizard

To launch the volume restore wizard, in the Veeam Recovery Media main menu, select **Restore volumes** and press [Enter].




Step 4. Select Backup Location

At the **Select backup location** step of the wizard, specify where the backup file that you want to use for data recovery is located.

To recover data from backup, you need to mount the backup storage on which the backup file resides to the recovery image OS file system. Veeam Agent for Linux automatically mounts external USB drives that are connected to the computer and displays them in the list of available backup locations. You can select the necessary device and press [Enter] to pass to the [Browse for backup files](#) step of the wizard.

If the backup file is located in a network shared folder, on a local drive or on a Veeam backup repository, select one of the following options:

- **Mount local disk** – select this option if the backup file resides on the local computer drive, external drive or removable storage device that is currently connected to your computer. With this option selected, you will pass to the [Select local disk](#) step of the wizard.
- **Add object storage repository** – select this option if the backup file resides in an object storage repository. With this option selected, you will pass to the [Select cloud storage type](#) step of the wizard.
- **Add shared folder** – select this option if the backup file is located in a network shared folder. With this option selected, you will pass to the [Mount shared folder](#) step of the wizard.
- **Add VBR server** – select this option if the backup file resides on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the [Specify backup server parameters](#) step of the wizard.
- **Add Cloud Connect provider** – select this option if the backup file resides in a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the [Specify Cloud provider parameters](#) step of the wizard.



Step 5. Specify Backup Location Settings

Specify settings for the target storage that contains a backup file from which you plan to restore data:

- [Specify shared folder settings](#) — if you have selected the **Add shared folder** option at the [Select backup location](#) step of the wizard.
- [Select local drive](#) — if you have selected the **Mount local disk** option at the [Select backup location](#) step of the wizard.
- [Specify Veeam backup repository settings](#) — if you have selected the **Add VBR server** option at the [Select backup location](#) step of the wizard.
- [Specify Veeam Cloud Connect repository settings](#) — if you have selected the **Add Cloud provider** option at the [Select backup location](#) step of the wizard.
- [Specify object storage repository settings](#) - if you have selected the **Add object storage repository** option at the [Select backup location](#) step of the wizard.

Shared Folder Settings


The **Mount shared folder** step of the wizard is available if you have selected to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. Select the type of a network shared folder:
 - **NFS** — to connect to a network shared folder using the NFS protocol.
 - **SMB** — to connect to a network shared folder using the SMB (CIFS) protocol.
2. In the **Path** field, specify the network shared folder name in the *SERVER/DIRECTORY* format: type an IP address or domain name of the server and the name of the network shared folder in which the backup file resides.
3. [For SMB network shared folder] In the **Domain** field, type a name of the domain in which the account that has access permissions on the shared folder is registered, for example: *DOMAIN*.
4. [For SMB network shared folder] In the **Username** field, type a name of the account that has access permissions on the shared folder.
5. [For SMB network shared folder] In the **Password** field, type a password of the account that has access permissions on the shared folder.
6. Press [Enter] to connect to the network shared folder. Veeam Agent will mount the specified network shared folder to the `/media` directory of the recovery image OS file system and display content of the network shared folder.

TIP

You can mount several network shared folders to work with backup files that are stored in different locations if needed. To do this, return to the [Select Backup Location](#) step of the wizard and select the **Add shared folder** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted network shared folders and browse for a backup file located on the necessary storage.




Local Backup Repository Settings

The **Select local disk** step of the wizard is available if you have selected to restore data from a backup file located on a computer drive.

In the list of devices, select the necessary disk or disk partition and press [Enter]. Veeam Agent will mount the selected device to the `/media` directory of the recovery image OS file system and display content of the directory.

TIP

You can mount several devices to work with backup files that are stored in different locations if needed. To do this, return to the [Select Backup Location](#) step of the wizard and select the **Mount local disk** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted devices and browse for a backup file located on the necessary storage.



Veeam Backup Repository Settings

The **Specify Backup Server parameters** step of the wizard is available if you have selected to restore data from a backup repository managed by the Veeam backup server.

Specify settings for the Veeam backup server that manages the backup repository where the backup file resides:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent must communicate with the backup repository. By default, Veeam Agent uses port 10006.
3. Select the type of **Authentication** to access the Veeam backup server:
 - **Login and password.** With this option selected, specify the following settings:
 - i. In the **Login** field, type a name of the account that has access to the Veeam backup repository.
 - ii. In the **Domain** field, type a name of the domain in which the account that has access to the Veeam backup repository is registered, for example: *DOMAIN*.

- iii. In the **Password** field, type a password of the account that has access to the Veeam backup repository.

NOTE

If you want to perform restore from a backup created by Veeam Agent operating in the managed mode, you must use an account that has the Veeam Backup Administrator or Veeam Restore Operator role on the Veeam backup server. For more information about user roles, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.

The screenshot shows the Veeam Recovery Media interface. At the top, it says "Veeam Recovery Media". The main area is a dark gray rectangle. In the center, there is a light gray dialog box titled "Specify Backup Server parameters:". Inside the dialog box, the following fields are visible: "Address: 172.24.31.136", "Port: 10006", "Authentication:" with two options: "(X) Login and password" and "() Recovery token", "Login: Administrator", "Domain:", and "Password: *****". At the bottom right of the dialog box, there are two buttons: "[Prev]" and "[Next]". Below the dialog box, at the bottom of the screen, there are three keyboard shortcuts: "Enter Connect", "Backspace Back", and "Esc Main menu".

- **Recovery token:** With this option selected, in the **Token** field, enter the value of the recovery token generated in the Veeam Backup & Replication console. For more information on generating recovery tokens, see [Creating Recovery Token](#) in the Veeam Agent Management guide.

Veeam Recovery Media

Specify Backup Server parameters:

Address: 172.24.31.136

Port: 10006

Authentication:
() Login and password
(X) Recovery token


Token: 5e6d-45aA-DBBd-d0ec

Recovery tokens can be created using the Veeam Backup & Replication console.

[Prev] [Next]

Enter Connect Backspace Back Esc Main menu

4. Press [Enter]. Veeam Agent will connect to the Veeam backup server. If prompted, accept the self-signed certificate of the Veeam backup server to continue.



After successful connection to the Veeam backup server, you will pass immediately to the [Backup](#) step of the wizard.

Veeam Cloud Connect Repository Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. [Specify service provider settings.](#)
2. [Verify the TLS certificate.](#)
3. [Specify user account settings.](#)


Specifying Service Provider Settings

The **Specify Cloud provider parameters** step of the wizard is available if you have selected to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent will communicate with the cloud gateway. By default, port 6180 is used.

3. Press [Enter]. Veeam Agent will connect to the service provider and display the [Certificate details](#) window.




Verifying TLS Certificate

In the **Certificate details** window, review information about the TLS certificate obtained from the SP side and verify the TLS certificate.

- To accept the TLS certificate, select the **Accept** button with the [Tab] key and press [Enter].
- [Optional] To verify the TLS certificate with a thumbprint, do the following:
 - a. Select the **Verify thumbprint** button with the [Tab] key and press [Enter].
 - b. In the **Thumbprint verification** field, enter the thumbprint that you obtained from the SP.

- c. Switch to the **Verify** button and press [Enter]. Veeam Agent will check if the thumbprint that you entered matches the thumbprint of the obtained TLS certificate.

TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.




Specifying User Account Settings

The **Specify Cloud provider credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

1. In the **Username** field, enter the name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT|SUBTENANT* format.
2. In the **Password** field, provide a password for the tenant or subtenant account.

3. Press [Enter]. Veeam Agent will connect to the cloud repository, and you will pass immediately to the [Backup](#) step of the wizard.



Object Storage Repository Settings

If you have selected to restore data from a backup file located in a object storage repository, specify settings to connect to the repository:

At the **Select cloud storage type** step of the wizard, select one of the following options:


- **S3 Compatible** – select this option if you want to import a backup from an S3 compatible storage repository.

TIP

If you plan to restore data from a backup that resides in IBM Cloud or Wasabi Cloud, use the **S3 Compatible** option.

- **Amazon S3** – select this option if you want to import a backup from an Amazon S3 storage repository.
- **Google Cloud Storage** – select this option if you want to import a backup from a Google Cloud storage repository.

- **Microsoft Azure Blob Storage** – select this option if you want to import a backup from a Microsoft Azure storage repository.



Specifying Settings for S3 Compatible Repository

If you have selected to import backup from an S3 Compatible storage repository, specify settings to connect to the storage:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from an S3 compatible storage repository.

To connect to the S3 compatible storage, specify the following:

1. In the **Service point** field, specify the address of your S3 compatible storage.

If you want to connect to the repository using an IPv6 address and port number, you must use the following format: `IPv6:port`, where:

- `IPv6` is the IPv6 address of the object storage.
- `port` is the number of the port that Veeam Agent will use to connect to the object storage.

For example: `https://2001-db8-1--bb8-c0b8-112.ipv6-literal.net:9000`

If you want to connect to the repository using AWS PrivateLink for Amazon S3, use the following format: `bucket.<VPC_endpoint_ID>.s3.<region>.vpce.amazonaws.com`, where:

- o `VPC_endpoint_ID` is the unique VPC endpoint ID that Veeam Agent will use to access the cloud storage.
- o `region` is the AWS Region where the VPC endpoint and the S3 bucket are located.

For example: `bucket.vpce-0abc123def456ghi7-01234567.s3.ap-northeast-1.vpce.amazonaws.com`

To learn more about AWS PrivateLink for Amazon S3, see [AWS Documentation](#).

2. In the **Region** field, specify a storage region based on your regulatory and compliance requirements.
3. In the **Access key** field, enter an access key ID.
4. In the **Secret key** field, enter a secret access key.

The screenshot shows the Veeam Recovery Media interface with a dark background. A central dialog box titled 'Specify S3 compatible storage:' contains the following fields and controls:

- Service point: `https://myservicepoint.com:9000`
- Region: `reg-1`
- S3 compatible account:
- Access key: `Access_Key`
- Secret key: `*****`

At the bottom of the dialog box are two buttons: '[Prev]' and '[Next]'. The '[Next]' button is highlighted in green. At the bottom of the main window, there are two keyboard shortcuts: 'Enter' followed by a green 'Connect' button, and 'Esc' followed by a green 'Main menu' button.


Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from an S3 compatible storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Bucket** field, specify a bucket on the storage:
 - a. Click **Browse**.
 - b. In the **Buckets** window, select the necessary bucket and click **OK**.

2. In the **Folder** field, specify a folder in the bucket:
 - a. Click **Browse**.
 - b. In the **Folders** window, select the necessary folder and click **OK**.



Specifying Settings for Amazon S3 Repository

If you have selected to store backup files on an Amazon S3 storage, specify settings to connect to the storage:

1. [Specify account settings](#).
2. [Specify bucket settings](#).


Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from an Amazon S3 storage repository.

To connect to the Amazon S3 storage, specify the following:

1. In the **Access key** field, enter an access key ID.
2. In the **Secret key** field, enter a secret access key.

3. In the **AWS region** window, select an AWS region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Global** region.




Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from an Amazon S3 storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups.
2. In the **Bucket** field, specify a bucket on the storage:
 - a. Click **Browse**.
 - b. In the **Buckets** window, select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
 - a. Click **Browse**.

b. In the **Folders** window, select the necessary folder and click **OK**.



Specifying Settings for Google Cloud Repository


If you have selected to import backup from a Google Cloud storage repository, specify settings to connect to the storage:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from a Google Cloud storage repository.

To connect to the Google Cloud storage, in the **Access key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) key associated with the Google Cloud account. Veeam Agent will use the HMAC key to authenticate requests to the Google Cloud storage. For more information on Google Cloud accounts, see the [Google Cloud documentation](#).




Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from a Google Cloud storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups.
2. In the **Bucket** field, specify a bucket on the storage:
 - a. Click **Browse**.
 - b. In the **Buckets** window, select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
 - a. Click **Browse**.

b. In the **Folders** window, select the necessary folder and click **OK**.



Specifying Settings for Microsoft Azure Repository

If you have selected to import backup from a Microsoft Azure storage repository, specify settings to connect to the storage:

1. [Specify account settings](#).
2. [Specify container settings](#).

Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from a Microsoft Azure storage repository.


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](#).

To connect to the Microsoft Azure storage, specify the following:

1. In the **Account** field, enter the storage account name.
2. In the **Shared key** field, enter the storage account shared key.

3. In the **Region** window, select a Microsoft Azure region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Azure Global (Standard)** region.




Specifying Container Settings

The **Container** step of the wizard is available if you have chosen to import backup from a Microsoft Azure storage repository and specified account settings to connect to the storage.

Specify settings for the container on the storage:

1. In the **Container** field, specify a container on the storage:
 - a. Click **Browse**.
 - b. In the **Containers** window, select the necessary container and click **OK**.
2. In the **Folder** field, specify a folder in the bucket:
 - a. Click **Browse**.


b. In the **Folders** window, select the necessary folder and click **OK**.



Step 6. Browse for Backup File

At the **Browse for backup files** step of the wizard, select the backup file that you plan to use for volume-level restore:

1. In the file system tree, select a directory in which the backup file you plan to use for restore resides:
 - Use [Up] and [Down] arrow keys to select a directory.
 - Use the [Enter] key to open the necessary directory.
2. In the directory where the backup file resides, select the backup file and press [Enter].



Step 7. Select Backup and Restore Point

At the **Backup** step of the wizard, select a backup and restore point from which you want to recover data.


The **Backup** step window comprises two panes:

- The **Imported backups** pane on the left displays information about backup: host name of the computer whose data is stored in the backup file, backup job name and number of restore points.
- The **Restore points** pane on the right displays a list of restore points in the backup.

To select backup and restore point:

1. In the **Imported backups** pane, ensure that the backup from which you want to recover data is selected and press [Enter].


If you want to select another backup, press the [i] key and browse for the necessary backup file. To learn more, see [Locate Backup File](#).



2. In the **Restore points** pane, select with [Up] and [Down] keys the restore point from which you want to recover data and press [Enter].

NOTE

If you selected an encrypted backup for data restore, Veeam Agent will prompt you to provide a password to unlock the encrypted file. To learn more, see [Restoring Data from Encrypted Backups](#).



Step 8. Map Restored Disks

At the **Disk Mapping** step of the wizard, select what volumes you want to restore and map volumes from the backup to volumes on your computer.

IMPORTANT

It is strongly recommended that you change disk mapping settings only if you have experience in working with Linux disks and partitions. If you make a mistake, your computer data may get corrupted.

You can map volumes in the backup (source volumes) and volumes on your computer (target volumes) in one of the following ways:

- [Map a source volume to a target volume](#)
- [Map a target volume to a source volume](#)

As well as individual volumes, you can also map entire disks:

- [Map a source disk to a target disk](#)
- [Map a target disk to a source disk](#)

If you choose to restore an entire disk, Veeam Agent will try to map all volumes that reside on this disk.


If you want to restore BTRFS subvolumes, you must map subvolumes in the backup to a BTRFS pool on the Veeam Agent computer. To learn more, see [Mapping BTRFS Subvolumes](#).

Mapping Source Volume to Target Volume


The **In backup** pane of the **Veeam Recovery Media** wizard contains a list of disks and volumes in the backup. You can select volumes in the backup that you want to restore to your computer and specify mapping rules for these volumes.

To map a source volume to a target volume:


1. In the **In backup** pane, select a volume in the backup whose data you want to recover and press [Enter].




2. Veeam Agent for Linux will display a window with information on the selected volume (partition type, file system type, mount point and volume size) and a list of available operations:
 - **Restore volume to** – select this option if you want to restore the selected volume to your computer.
 - **Close** – select this option if you want to close the window and select another volume.
3. Select the **Restore volume to** option and press [Enter].



- Veeam Agent for Linux will display a list of volumes on your computer. Select the volume that you want to restore and press [Enter].



- In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volume from backup will be restored to the target volume.



- Repeat steps 1-5 for all volumes that you want to restore.
- Press [S] to start the restore process.

Mapping Target Volume to Source Volume

The **Current system** pane of the **Veeam Recovery Media** wizard displays a partition table of your computer booted from the Veeam Recovery Media. In this pane, you can select volumes on your computer which you want to restore and specify mapping rules for these volumes. If necessary, you can edit the disk layout before restoring volumes.

To map a target volume to a source volume:


1. In the **Current system** pane, select a volume on your computer whose data you want to recover and press [Enter].

CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
sda1		18.63G	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lvm)		10.00G	sdb1 (lvm)	10.00G	(LVM2_mem...
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	


Enter **Select** Backspace **Back** Esc **Main menu**

2. Veeam Agent for Linux will display a window with information on the selected volume (partition type, file system type, mount point and volume size) and a list of available operations:
 - **Restore volume from** – select this option if you want to recover the selected volume from the backup.
 - **Delete partition** [for simple volumes] or **Delete volume** [for LVM volumes] – select this option if you want to change the disk layout before restoring a volume. After you delete a partition or volume, you will be able to create a new partition or volume of the desired size and map a volume in the backup to the volume on your computer.
 - [For simple volumes] **Create LVM physical volume** – select this option if you want to create an LVM physical volume on the selected disk partition. In the created physical volume, you will be able to create a volume group and restore to this volume group LVM logical volumes from the backup.
 - **Close** – select this option if you want to close the window and select another volume.

3. Select the **Restore volume from** option and press [Enter].



4. Veeam Agent for Linux will display a window with a list of volumes in the backup. Select the volume that you want to restore and press [Enter].



- In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volume from backup will be restored to the target volume.

Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
sda1	sda1 (/)	18.63G	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lum)		10.00G	sdb1 (lum)	10.00G	(LVM2_mem...
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	


- Repeat steps 1-5 for all volumes that you want to restore.
- Press [S] to start the restore process.

Mapping Source Disk to Target Disk

The **In backup** pane of the **Veeam Recovery Media** wizard contains a list of disks and volumes in the backup. As well as individual volumes, you can select for restore entire computer disks.

To map a source disk to a target disk:

1. In the **In backup** pane, select a disk in the backup volumes on which you want to recover and press [Enter].




The screenshot shows the Veeam Recovery Media interface. It features two side-by-side tables comparing disk information. The left table, titled 'CURRENT SYSTEM', lists disks with their 'Restore' status and 'Size'. The right table, titled 'IN BACKUP', lists disks with their 'Size' and 'Usage'. At the bottom, there are navigation instructions: 'Enter Select', 'Backspace Back', and 'Esc Main menu'.

CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
sda1		18.63G	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lvm)		10.00G	sdb1 (lvm)	10.00G	(LVM2_mem...)
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	


Enter Select Backspace Back Esc Main menu

2. Veeam Agent for Linux will display a window with information on the selected disk (partition table type, bootloader type and disk size) and a list of available operations:
 - **Restore whole disk to** – select this option if you want to restore all volumes on the selected disk in the backup to your computer.
 - **Restore bootloader to** – select this option if you want to restore a bootloader from the disk in the backup to your computer.
 - **Close** – select this option if you want to close the window and select another disk or volume.

3. To restore volumes that reside on the selected disk, select the **Restore whole disk to** option and press [Enter].



4. Veeam Agent for Linux will display a list of disks and volumes on your computer. Select the disk whose volumes you want to restore and press [Enter].



- In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volumes from the disk in the backup will be restored to the target disk.

Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)	loader (sda)	60.00G	sda (boot)	60.00G	
sda1	sda1 (/)	18.63G	sda1	18.63G	/ (ext4)
sda2	sda2 (/home)	38.31G	sda2	38.31G	/home (ext4)
sda3	sda3 (swap)	3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lum)		10.00G	sdb1 (lum)	10.00G	(LUM2_mem...
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	


- Repeat steps 1-5 for all computer disks whose volumes you want to restore.
- Press [S] to start the restore process.

Mapping Target Disk to Source Disk

The **Current system** pane of the **Veeam Recovery Media** wizard displays a partition table of your computer booted from the Veeam Recovery Media. As well as individual volumes, you can select for restore entire computer disks. If necessary, you can edit the disk layout before restoring volumes.


To map a target disk to a source disk:

1. In the **Current system** pane, select a disk on your computer to which you want to restore volumes and press [Enter].




2. Veeam Agent for Linux will display a window with information on the selected disk (partition table type, bootloader type and disk size) and a list of available operations:
 - **Restore whole disk from** – select this option if you want to restore to the selected disk all volumes from a disk in the backup.
 - **Restore bootloader from** – select this option if you want to restore to the selected disk a bootloader from a disk in the backup.
 - **Delete partition table** – select this option if you want to change the disk layout before restoring volumes. After you delete a partition table, you will be able to create a new partition table, create disk partitions and volumes of the desired size, and map volumes in the backup to volumes on your computer.
 - **Close** – select this option if you want to close the window and select another disk or volume.

3. To restore volumes to the selected disk, select the **Restore whole disk from** option and press [Enter].



4. Veeam Agent for Linux will display a list of disks and volumes in the backup. Select the disk whose volumes you want to restore and press [Enter].



5. In the **Current system** pane, in the **Restore** column, Veeam Agent will display which volumes from the disk in the backup will be restored to the target disk.

Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)	loader (sda)	60.00G	sda (boot)	60.00G	
sda1	sda1 (/)	18.63G	sda1	18.63G	/ (ext4)
sda2	sda2 (/home)	38.31G	sda2	38.31G	/home (ext4)
sda3	sda3 (swap)	3.06G	sda3	3.06G	(swap)
sdb (boot)		10.00G	sdb (boot)	10.00G	
sdb1 (lum)		10.00G	sdb1 (lum)	10.00G	(LUM2_mem...
vg		10.00G	vg	10.00G	
lv1		5.00G	lv1	5.00G	
lv2		5.00G	lv2	5.00G	

6. Repeat steps 1-5 for all disks whose volumes you want to restore.
7. Press [S] to start the restore process.

Mapping Btrfs Subvolumes

If the backup contains BTRFS file system data, in the **In backup** pane of the **Veeam Recovery Media** wizard, Veeam Agent displays the list of backed-up BTRFS subvolumes. Information about the original BTRFS pool that contained these subvolumes is not included in the backup.

You can restore from the backup all BTRFS subvolumes or selected subvolumes. To restore a subvolume, you must specify a target BTRFS pool – a BTRFS pool on the computer where you perform restore using the Veeam Recovery Media.

You can restore BTRFS subvolumes to the original BTRFS pool or new BTRFS pool. If the target BTRFS pool contains a subvolume with the same name as the name of the subvolume that you selected for restore, Veeam Agent will automatically map these subvolumes. During the restore process, Veeam Agent will overwrite data on the target subvolume with the data retrieved from the backup.

NOTE

Veeam Agent for Linux does not check whether the target BTRFS pool has enough disk space to restore the selected subvolumes. If the total size of the restored data is larger than the size of the target BTRFS pool, after the restore process completes, the restored data will be corrupted.

To map a source BTRFS subvolume to a target BTRFS pool:

1. In the **In backup** pane, select a subvolume in the backup whose data you want to restore and press [Enter].


You can also choose to restore all subvolumes from the backup at once. To do this, in the **In backup** pane, select **btrfs** and press [Enter].

Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
sda1		18.63G	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sdb1 (lvm)	10.00G	(LVM2_men...
sdb		10.00G	vg	10.00G	
sdb1 (lvm)		10.00G	lv1	5.00G	
sdc		30.00G	lv2	5.00G	
sdc1 (bt...		10.00G	btrfs	30.00G	
sdc2 (bt...		10.00G	/		/btrfs (b...
sdc3 (bt...		10.00G	/sub1		/btrfs (b...
vg		10.00G	/sub2		/btrfs (b...
lv1		5.00G	/sub2/sub3		/btrfs (b...
lv2		5.00G			
btrfspool		30.00G			
/					
/sub1					
/sub2					
/sub2/sub3					


Enter **Select** Backspace **Back** Esc **Main menu**

2. In the displayed window, select the necessary option for BTRFS restore and press [Enter]. The available options depend on what BTRFS subvolumes you selected for restore: all subvolumes or specific subvolume.
 - **Restore subvolume to** – this option is available if you chose to restore a specific BTRFS subvolume from the backup. Select this option to restore the selected subvolume to your computer.
 - **Restore btrfs to** – this option is available if you chose to restore all BTRFS subvolumes from the backup. Select this option to restore subvolumes to your computer.

- **Close** – select this option if you want to close the window and select another subvolume.



3. Veeam Agent for Linux will display a list of BTRFS pools on your computer. Select the BTRFS pool where you want to restore data from the backup and press [Enter].



- In the **Current system** pane, in the **Restore** column, Veeam Agent for Linux will display which subvolume from backup will be restored to the target BTRFS pool.


Veeam Recovery Media					
CURRENT SYSTEM			IN BACKUP		
Device	Restore	Size	Device	Size	Usage
sda (boot)		60.00G	sda (boot)	60.00G	
sda1		18.63G	sda1	18.63G	/ (ext4)
sda2		38.31G	sda2	38.31G	/home (ext4)
sda3		3.06G	sdb1 (lvm)	10.00G	(LVM2_men...
sdb		10.00G	vg	10.00G	
sdb1 (lvm)		10.00G	lv1	5.00G	
sdc		30.00G	lv2	5.00G	
sdc1 (bt...		10.00G	btrfs	30.00G	
sdc2 (bt...		10.00G	/		/btrfs (b...
sdc3 (bt...		10.00G	/sub1		/btrfs (b...
vg		10.00G	/sub2		/btrfs (b...
lv1		5.00G	/sub2/sub3		/btrfs (b...
lv2		5.00G			
btrfs pool		30.00G			
/					
/sub1					
/sub2	/sub2				

- If you want to restore more than one subvolume, repeat steps 1-4 for all subvolumes that you want to restore.
- Press [S] to start the restore process.


Step 9. Complete Restore Process

At the **Recovery summary** step of the wizard, complete the procedure of volume-level restore.

1. Review the specified recovery settings.




2. Press [Enter] to start the volume-level restore process. Veeam Agent for Linux will perform partition re-allocation operations if necessary, restore the necessary data from the backup and overwrite data on your computer with it.



Step 10. Finish Working with Veeam Recovery Media

When the restore operation completes, finish working with the Veeam Recovery Media and start your operating system.

1. Press [Esc] to return to the Veeam Recovery Media main menu.
2. Eject the media or removable storage device with the recovery image.
3. In the Veeam Recovery Media main menu, select the **Reboot** option and press [Enter].



4. Wait for your Linux operating system to start.

Restoring Files and Folders

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore with the Veeam Recovery Media, Veeam Agent publishes the backup content directly into the recovery image OS file system and displays it in the file browser. You can restore files and folders to their initial location or copy files and folders to a new location.

Before You Begin

Before you boot from the recovery image and recover your data, check the following prerequisites:

- You must have a recovery image on any type of media: CD/DVD/BD or removable storage device.
- To recover data on your computer, you must have both the Veeam Recovery Media and data backup. For data recovery, you can use a volume-level or file-level backup created with Veeam Agent for Linux. Make sure that the backup or system image is available on the computer drive (local or external), on a network shared folder or on the backup repository managed by a Veeam backup server.
- The media type on which you have created the recovery image must be set as a primary boot source on your computer.
- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a Veeam backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
- If you restore to a virtual environment, note that the current version of Veeam Recovery Media supports only the VMware and Hyper-V virtualization solutions. To resolve possible issues during bare metal recovery of Oracle VM virtual machines, use instructions in the second section of [this Veeam KB article](#).

Step 1. Boot from Veeam Recovery Media


To boot from the Veeam Recovery Media:

1. [For CD/DVD/BD] Power on your computer. Insert the media with the recovery image to the drive and power off the computer.

[For removable storage device] Attach the removable storage device with the recovery image to your computer.
2. Start your computer.
3. If you boot from a regular recovery image, in the boot menu, select what Linux kernel version to use to boot your computer and specify boot options if necessary.

You can select a Linux kernel version if you use generic Veeam Recovery Media downloaded from [the Veeam website](#) or [Veeam software repository](#). If you created a custom Veeam Recovery Media, you will be prompted to boot using the Linux kernel of your Veeam Agent computer included in the recovery image.

If you boot from a [legacy recovery image](#), the boot menu is unavailable. After you start your computer, Veeam Agent will immediately start loading files from the Veeam Recovery Media.



TIP

On legacy BIOS-based systems, the boot menu has a different interface but provides the same options.


4. Wait for Veeam Agent to load files from the Veeam Recovery Media.

5. After the recovery image OS has loaded, choose whether you want to start the SSH server. The SSH server allows you to connect to the Veeam Recovery Media from a remote machine.

The Veeam Recovery Media starts the SSH server automatically after a time-out. The default value for the time-out is 60 seconds.

To override the default time-out and start the SSH server immediately, select the **Start SSH now** button and press [Enter].

If you do not want to start the SSH server, make sure that the **Proceed without SSH** button is selected and press [Enter]. You will proceed immediately to the step 7.



6. After the SSH server has started, review settings to connect to the Veeam Recovery Media and press [Enter].

The Veeam Recovery Media displays the following connection settings:

- IP address of the computer booted from the Veeam Recovery Media
- User name and password of the account used to connect to the Veeam Recovery Media
- Fingerprints of the computer booted from the Veeam Recovery Media

NOTE

The user name of the account used to work with the Veeam Recovery Media is *veeamuser*.

If you want to use command-line utilities built in the regular recovery image, use the `sudo` command to provide the *veeamuser* account with privileges of the *root* account.

```
SSH Connection Info

Credentials
login: veeamuser
passwd: kaAnL


NetConfigs
ens160
IP: 172.24.28.72
IPv6: fd00:ac18:0:1810:0:b5f9:ab46:e4c5

Fingerprints
ecdsa-sha2-nistp256
SHA1:BaXFVwjaWKUf6Rvv2gAwR+g+knI
MD5:6d:9c:56:1d:62:d3:f6:56:f0:0e:62:25:31:da:3c:a2
ssh-ed25519
SHA1:618oSzFazLsSUaMDD/EQJCymqjc
MD5:2b:2b:5d:78:14:66:55:da:cc:7e:6a:bb:29:a3:01:da
ssh-rsa
SHA1:6PsfT1Vv+Gkn8dgdR7420HAsGBQ
MD5:ff:96:15:0b:e4:30:86:67:08:8e:7b:21:47:0c:b4:0a

[Continue]
```

7. Accept the terms of the product license agreement and license agreements for third-party components of the product:
 - a. Make sure that the **I accept Veeam End User Software License Agreement** option is selected and press [Space].
 - b. Select the **I accept the terms of the following 3rd party software components license agreements** option with the [Tab] key and press [Space].


c. Switch to the **Continue** button with the [Tab] key and press [Enter].



8. Make sure that network settings are specified correctly and configure the network adapter if necessary. To learn more, see [Configure Network Settings](#).
9. Choose the necessary recovery option. Veeam Agent offers the following tools:
 - **Restore volumes** – the Veeam Recovery wizard to recover data on the original computer or perform bare metal recovery.
 - **Restore files** – the File Level Restore wizard to restore files and folders to the original location or to a new location.
 - **Exit to shell** – Linux shell prompt with standard utilities to diagnose problems and fix errors.

TIP


To stop working with the Veeam Recovery Media and shut down or restart your computer, in the Veeam Recovery Media main menu, select the **Reboot** or **Shutdown** option and press [Enter].




Step 2. Configure Network Settings

If there is a DHCP server in your network, Veeam Agent will configure the network settings automatically. To verify or configure network settings manually, use **nmtui**, a text-based user interface network manager tool provided with Veeam. To learn more about working with nmtui, see [Linux documentation](#).

1. In the Veeam Recovery Media main menu, select **Configure network** and press [Enter].




2. To add new or modify existing connection, in NetworkManager, select **Edit a connection**.



3. After you add or edit a connection, in the main menu of the NetworkManager, select **Activate a connection**.
 - a. If the connection is new, choose it in the list of connections; then select **Activate**.
 - b. If the connection was modified, you must reactivate it. To do this, choose it in the list of connections and select **Deactivate**; then choose the connection again and select **Activate**.
4. After you finish working with Network Manager, press [Esc] to return to the Veeam Recovery Media main menu and launch the File Level Restore wizard.

Step 3. Launch File Level Restore Wizard

To launch the file-level restore wizard, in the Veeam Recovery Media main menu, select **Restore files** and press [Enter].




Step 4. Select Backup Location

At the **Select backup location** step of the wizard, specify where the backup file that you want to use for data recovery is located.

To recover data from backup, you need to mount the backup storage on which the backup file resides to the recovery image OS file system. Veeam Agent automatically mounts external USB drives that are connected to the computer and displays them in the list of available backup locations. You can select the necessary device and press [Enter] to pass to the [Browse for backup files](#) step of the wizard.

If the backup file is located in a network shared folder or on a local drive, select one of the following options:

- **Mount local disk** – select this option if the backup file resides on the local computer drive, external drive or removable storage device that is currently connected to your computer. With this option selected, you will pass to the [Select local disk](#) step of the wizard.
- **Add object storage repository** – select this option if the backup file resides in an object storage repository. With this option selected, you will pass to the [Select cloud storage type](#) step of the wizard.
- **Add shared folder** – select this option if the backup file is located in a network shared folder. With this option selected, you will pass to the [Mount shared folder](#) step of the wizard.
- **Add VBR server** – select this option if the backup file resides on a backup repository managed by the Veeam backup server. With this option selected, you will pass to the [Specify backup server parameters](#) step of the wizard.
- **Add Cloud provider** – select this option if the backup file resides on a cloud repository exposed to you by a Veeam Cloud Connect service provider. With this option selected, you will pass to the [Specify Cloud provider parameters](#) step of the wizard.



Step 5. Specify Backup Location Settings

Specify settings for the target storage that contains a backup file from which you plan to restore data:

- [Specify shared folder settings](#) — if you have selected the **Add shared folder** option at the [Select backup location](#) step of the wizard.
- [Select local drive](#) — if you have selected the **Mount local disk** option at the [Select backup location](#) step of the wizard.
- [Specify Veeam backup repository settings](#) — if you have selected the **Add VBR server** option at the [Select backup location](#) step of the wizard.
- [Specify Veeam Cloud Connect repository settings](#) — if you have selected the **Add Cloud Connect provider** option at the [Select backup location](#) step of the wizard.
- [Specify object storage repository settings](#) — if you have selected the **Add object storage repository** option at the [Select backup location](#) step of the wizard.

Shared Folder Settings


The **Mount shared folder** step of the wizard is available if you have selected to restore data from a backup file located in a network shared folder.

Specify settings for the network shared folder:

1. Select the type of a network shared folder:
 - **NFS** — to connect to a network shared folder using the NFS protocol.
 - **SMB** — to connect to a network shared folder using the SMB (CIFS) protocol.
2. In the **Path** field, specify the network shared folder name in the *SERVER/DIRECTORY* format: type an IP address or domain name of the server and the name of the network shared folder in which the backup file resides.
3. [For SMB network shared folder] In the **Domain** field, type a name of the domain in which the account that has access permissions on the shared folder is registered, for example: *DOMAIN*.
4. [For SMB network shared folder] In the **Username** field, type a name of the account that has access permissions on the shared folder.
5. [For SMB network shared folder] In the **Password** field, type a password of the account that has access permissions on the shared folder.

TIP

You can mount several network shared folders to work with backup files that are stored in different locations if needed. To do this, return to the [Select Backup Location](#) step of the wizard and select the **Add shared folder** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted network shared folders and browse for a backup file located on the necessary storage.




Local Backup Repository Settings

The **Select local disk** step of the wizard is available if you have selected to restore data from a backup file located on a computer drive.

In the list of devices, select the necessary disk or disk partition and press [Enter]. Veeam Agent will mount the selected device to the `/media` directory of the recovery image OS file system and display content of the directory.

TIP

You can mount several devices to work with backup files that are stored in different locations if needed. To do this, return to the [Select Backup Location](#) step of the wizard and select the **Mount local disk** option once again. For every mounted location, Veeam Agent displays its name, type and mount point. You can view the list of mounted devices and browse for a backup file located on the necessary storage.



Veeam Backup Repository Settings

The **Specify Backup Server parameters** step of the wizard is available if you have selected to restore data from a backup repository managed by the Veeam backup server.

Specify settings for the Veeam backup server that manages the backup repository where the backup file resides:

1. In the **Address** field, specify a DNS name or IP address of the Veeam backup server.
2. In the **Port** field, specify a number of the port over which Veeam Agent must communicate with the backup repository. By default, Veeam Agent uses port 10006.
3. Select the type of **Authentication** to access the Veeam backup server:
 - **Login and password.** With this option selected, specify the following settings:
 - i. In the **Login** field, type a name of the account that has access to the Veeam backup repository.
 - ii. In the **Domain** field, type a name of the domain in which the account that has access to the Veeam backup repository is registered, for example: *DOMAIN*.

- iii. In the **Password** field, type a password of the account that has access to the Veeam backup repository.

NOTE

If you want to perform restore from a backup created by Veeam Agent operating in the managed mode, you must use an account that has the Veeam Backup Administrator or Veeam Restore Operator role on the Veeam backup server. For more information about user roles, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.

The screenshot displays the Veeam Recovery Media interface. At the top, a header bar reads "Veeam Recovery Media". The main area is a dark gray background. In the center, a light gray dialog box titled "Specify Backup Server parameters:" contains the following fields and options:

- Address: 172.24.31.136
- Port: 10006
- Authentication:
 - ☒ Login and password
 - ☐ Recovery token
- Login: Administrator
- Domain:
- Password: *****

At the bottom right of the dialog box, there are two buttons: "[Prev]" and "[Next]". The "[Next]" button is highlighted in green.

At the bottom of the screen, there are three keyboard shortcuts displayed in a light blue box:

- Enter Connect
- Backspace Back
- Esc Main menu

- **Recovery token:** With this option selected, in the **Token** field, enter the value of the recovery token generated in the Veeam Backup & Replication console. For more information on generating recovery tokens, see [Creating Recovery Token](#) in the Veeam Agent Management guide.

Veeam Recovery Media

Specify Backup Server parameters:

Address: 172.24.31.136

Port: 10006

Authentication:
() Login and password
(X) Recovery token


Token: 5e6d-45aA-DBBd-d0ec

Recovery tokens can be created using the Veeam Backup & Replication console.

[Prev] [Next]

Enter Connect Backspace Back Esc Main menu

4. Press [Enter]. Veeam Agent will connect to the Veeam backup server. If prompted, accept the self-signed certificate of the Veeam backup server to continue.



After successful connection to the Veeam backup server, you will pass immediately to the [Backup](#) step of the wizard.

Veeam Cloud Connect Repository Settings

If you have selected to restore data from a backup file located on a Veeam Cloud Connect repository, specify settings to connect to the cloud repository:

1. [Specify service provider settings.](#)
2. [Verify the TLS certificate.](#)
3. [Specify user account settings.](#)


Specifying Service Provider Settings

The **Specify Cloud provider parameters** step of the wizard is available if you have selected to restore data from a cloud repository exposed to you by a Veeam Cloud Connect service provider.

Specify service provider settings that the SP or your backup administrator has provided to you:

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the cloud gateway.
2. In the **Port** field, specify the port over which Veeam Agent will communicate with the cloud gateway. By default, port 6180 is used.

3. Press [Enter]. Veeam Agent will connect to the service provider and display the [Certificate details](#) window.




Verifying TLS Certificate

In the **Certificate details** window, review information about the TLS certificate obtained from the SP side and verify the TLS certificate.

- To accept the TLS certificate, select the **Accept** button with the [Tab] key and press [Enter].
- [Optional] To verify the TLS certificate with a thumbprint, do the following:
 - a. Select the **Verify thumbprint** button with the [Tab] key and press [Enter].
 - b. In the **Thumbprint verification** field, enter the thumbprint that you obtained from the SP.

- c. Switch to the **Verify** button and press [Enter]. Veeam Agent for Linux will check if the thumbprint that you entered matches the thumbprint of the obtained TLS certificate.

TLS certificate verification is optional. You can use this option to verify self-signed TLS certificates. TLS certificates signed by the CA do not require additional verification.




Specifying User Account Settings

The **Specify Cloud provider credentials** step of the wizard is available if you have chosen to restore data from a cloud repository and specified settings for the cloud gateway.

1. In the **Username** field, type a name of the tenant or subtenant account that the SP or your backup administrator has provided to you. The name of the subtenant account must be specified in the *TENANT|SUBTENANT* format.
2. In the **Password** field, provide a password for the tenant or subtenant account.

3. Press [Enter]. Veeam Agent for Linux will connect to the cloud repository, and you will pass immediately to the [Backup](#) step of the wizard.



Object Storage Repository Settings

If you have selected to restore data from a backup file located in a object storage repository, specify settings to connect to the repository:

At the **Select cloud storage type** step of the wizard, select one of the following options:


- **S3 Compatible** – select this option if you want to import a backup from an S3 compatible storage repository.

TIP

If you plan to restore data from a backup that resides in IBM Cloud or Wasabi Cloud, use the **S3 Compatible** option.

- **Amazon S3** – select this option if you want to import a backup from an Amazon S3 storage repository.
- **Google Cloud Storage** – select this option if you want to import a backup from a Google Cloud storage repository.

- **Microsoft Azure Blob Storage** – select this option if you want to import a backup from a Microsoft Azure storage repository.



Specifying Settings for S3 Compatible Repository

If you have selected to import backup from an S3 Compatible storage repository, specify settings to connect to the storage:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from an S3 compatible storage repository.

To connect to the S3 compatible storage, specify the following:

1. In the **Service point** field, specify the address of your S3 compatible storage.

If you want to connect to the repository using an IPv6 address and port number, you must use the following format: `IPv6:port`, where:

- `IPv6` is the IPv6 address of the object storage.
- `port` is the number of the port that Veeam Agent will use to connect to the object storage.

For example: `https://2001-db8-1--bb8-c0b8-112.ipv6-literal.net:9000`

If you want to connect to the repository using AWS PrivateLink for Amazon S3, use the following format: `bucket.<VPC_endpoint_ID>.s3.<region>.vpce.amazonaws.com`, where:

- o `VPC_endpoint_ID` is the unique VPC endpoint ID that Veeam Agent will use to access the cloud storage.
- o `region` is the AWS Region where the VPC endpoint and the S3 bucket are located.

For example: `bucket.vpce-0abc123def456ghi7-01234567.s3.ap-northeast-1.vpce.amazonaws.com`

To learn more about AWS PrivateLink for Amazon S3, see [AWS Documentation](#).

2. In the **Region** field, specify a storage region based on your regulatory and compliance requirements.
3. In the **Access key** field, enter an access key ID.
4. In the **Secret key** field, enter a secret access key.

The screenshot shows a terminal window titled "Veeam Recovery Media". In the center, a dialog box titled "Specify S3 compatible storage:" is displayed. The dialog box contains the following fields and values:

- Service point: `https://myservicepoint.com:9000`
- Region: `reg-1`
- S3 compatible account:
- Access key: `Access_Key`
- Secret key: `*****`

At the bottom right of the dialog box, there are two buttons: "[Prev]" and "[Next]". The "[Next]" button is highlighted in green. At the bottom of the terminal window, there are two keyboard shortcuts: "Enter" followed by a green button labeled "Connect", and "Esc" followed by a green button labeled "Main menu".


Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from an S3 compatible storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Bucket** field, specify a bucket on the storage:
 - a. Click **Browse**.
 - b. In the **Buckets** window, select the necessary bucket and click **OK**.

2. In the **Folder** field, specify a folder in the bucket:
 - a. Click **Browse**.
 - b. In the **Folders** window, select the necessary folder and click **OK**.



Specifying Settings for Amazon S3 Repository

If you have selected to store backup files on an Amazon S3 storage, specify settings to connect to the storage:

1. [Specify account settings](#).
2. [Specify bucket settings](#).


Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from an Amazon S3 storage repository.

To connect to the Amazon S3 storage, specify the following:

1. In the **Access key** field, enter an access key ID.
2. In the **Secret key** field, enter a secret access key.

3. In the **AWS region** window, select an AWS region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Global** region.




Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from an Amazon S3 storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups.
2. In the **Bucket** field, specify a bucket on the storage:
 - a. Click **Browse**.
 - b. In the **Buckets** window, select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
 - a. Click **Browse**.

b. In the **Folders** window, select the necessary folder and click **OK**.



Specifying Settings for Google Cloud Repository


If you have selected to import backup from a Google Cloud storage repository, specify settings to connect to the storage:

1. [Specify account settings.](#)
2. [Specify bucket settings.](#)

Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from a Google Cloud storage repository.

To connect to the Google Cloud storage, in the **Access key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) key associated with the Google Cloud account. Veeam Agent will use the HMAC key to authenticate requests to the Google Cloud storage. For more information on Google Cloud accounts, see the [Google Cloud documentation](#).




Specifying Bucket Settings

The **Bucket** step of the wizard is available if you have chosen to import backup from a Google Cloud storage repository and specified account settings to connect to the storage.

Specify settings for the bucket on the storage:

1. In the **Data center** window, select the geographic region where Veeam Agent will store backups.
2. In the **Bucket** field, specify a bucket on the storage:
 - a. Click **Browse**.
 - b. In the **Buckets** window, select the necessary bucket and click **OK**.
3. In the **Folder** field, specify a folder in the bucket:
 - a. Click **Browse**.

b. In the **Folders** window, select the necessary folder and click **OK**.



Specifying Settings for Microsoft Azure Repository

If you have selected to import backup from a Microsoft Azure storage repository, specify settings to connect to the storage:

1. [Specify account settings](#).
2. [Specify container settings](#).

Specifying Account Settings

The **Account** step of the wizard is available if you have chosen to import backup from a Microsoft Azure storage repository.


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](#).

To connect to the Microsoft Azure storage, specify the following:

1. In the **Account** field, enter the storage account name.
2. In the **Shared key** field, enter the storage account shared key.

3. In the **Region** window, select a Microsoft Azure region based on your regulatory and compliance requirements. By default, Veeam Agent uses the **Azure Global (Standard)** region.




Specifying Container Settings

The **Container** step of the wizard is available if you have chosen to import backup from a Microsoft Azure storage repository and specified account settings to connect to the storage.

Specify settings for the container on the storage:

1. In the **Container** field, specify a container on the storage:
 - a. Click **Browse**.
 - b. In the **Containers** window, select the necessary container and click **OK**.
2. In the **Folder** field, specify a folder in the bucket:
 - a. Click **Browse**.


b. In the **Folders** window, select the necessary folder and click **OK**.



Step 6. Browse for Backup File

At the **Browse for backup files** step of the wizard, select the backup file that you plan to use for volume-level restore:

1. In the file system tree, select a directory in which the backup file you plan to use for restore resides:
 - Use the [Up] and [Down] keys to select a directory.
 - Press [Enter] to open the necessary directory.
2. In the directory where the backup file resides, select the backup file and press [Enter].



Step 7. Select Backup and Restore Point

At the **Backup** step of the wizard, select a backup and restore point from which you want to recover data.


The **Backup** step window comprises two panes:

- The **Imported backups** pane on the left displays information about backup: host name of the computer whose data is stored in the backup file, backup job name and number of restore points.
- The **Restore points** pane on the right displays a list of restore points in the backup.

To select backup and restore point:

1. In the **Imported backups** pane, ensure that the backup from which you want to recover data is selected and press [Enter].


If you want to select another backup, press the [i] key and browse for the necessary backup file. To learn more, see [Locate Backup File](#).



2. In the **Restore points** pane, select with the [Up] and [Down] keys the restore point from which you want to recover data and press [Enter].


NOTE

If you selected an encrypted backup for data restore, Veeam Agent will prompt you to provide a password to unlock the encrypted file. To learn more, see [Restoring Data from Encrypted Backups](#).



3. Veeam Agent will mount the content of the backup file to the `/mnt/backup` directory in the recovery image OS file system and display a notification window with the corresponding message. Press [Enter] to proceed to the File Level Restore wizard menu, open the file manager and save restored files.

When you perform file-level restore with the File Level Restore wizard, Veeam Agent always mounts the backup to the `/mnt/backup` directory. If you want to specify another directory for backup mount, you can perform file-level restore with the Veeam Agent command line interface. To learn more, see [Restoring Files and Folders with Command Line Interface](#).



Step 8. Save Restored Files

When the backup file content is mounted to the recovery image OS file system, Veeam Agent opens the File Level Restore wizard menu displaying a list of available operations.

NOTE

If a backed up file system was mounted to multiple mount points, during restore, depending on the file system type, Veeam Agent will behave as follows:


- [root file system] Veeam Agent will restore all mount points to the root directory.
- [non-root file system] Veeam Agent will restore all mount points to a single mount point randomly chosen from the mount points to which it was originally mounted.

You can perform the following operations with file-level backup:

- **Start file browser** – select this option if you want to start the file manager and work with restored files and folders. To learn more, see [Working with Midnight Commander](#).
- **Stop backup mount** – select this option if you want to stop the backup mount session and unmount the backup file content from the `/mnt/backup` directory of the recovery image OS file system. To learn more, see [Stopping Backup Mount Session](#).
- **Exit to shell** – select this option if you want to open the Linux shell prompt and use common Linux command-line tools.

TIP

To stop working with the Veeam Recovery Media and shut down or restart your computer, in the File Level Restore wizard menu, select the **Reboot** or **Shutdown** option and press [Enter].



Working with Midnight Commander

To work with restored files and folders, you can use Midnight Commander – a file manager that is included into the Veeam Recovery Media. With the Midnight Commander file manager, you can browse the mounted backup content and file system on your computer, and save restored files and folders to the original location or to a new location.

To launch the file manager, in the File Level Restore wizard menu, select **Start file browser** and press [Enter].

When you launch Midnight Commander, Veeam Agent displays in the file manager the directory with the backup content and your computer's file system:

- In the left pane, Veeam Agent displays a directory of your computer's file system mounted under the `/mnt/system` directory of the recovery image OS file system. By default, Veeam Agent mounts to the recovery image OS file system the following volumes of your computer:
 - If you use a volume-level backup for file-level restore, Veeam Agent detects the partition table in the backup, mounts to the `/mnt/system` directory block devices that represent volumes of your computer with the same names as volumes in the backup. For example, if your volume-level backup contains `/dev/sda1` and `/dev/sda6` volumes with `/` and `/home` mount points, Veeam Agent will mount to the `/mnt/system` directory both root (`/`) and `/home` partitions.

- If you use a file-level backup for file-level restore, Veeam Agent mounts to the `/mnt/system` directory only the system volume of your computer, for example, `/dev/sda1`. If you want to save restored files and folders to a directory on another computer volume or to a network shared folder, you need to mount this volume or folder manually. To mount a target storage for restored files:

- In Midnight Commander, press [F10] to close the file manager.
- In the **File Level Restore** wizard menu, select the **Exit to shell** option and press [Enter].
- Mount the target storage for the restored files and folders with the `mount` command.

- In the right pane, Veeam Agent displays a directory in which the backup content is mounted. Veeam Agent mounts the backup content under the `/mnt/backup` folder.

While the Midnight Commander file manager is open, you can perform the following operations with restored files and folders:

- [Save files to initial location](#)
- [Save files to a new location](#)

After you finish working with files and folders, [finish working with the Veeam Recovery Media](#).


Left	File	Command	Options	Right	File	Command	Options												
<-	/mnt/system		[.^]>	<-	/mnt/backup		[.^]>												
.n	Name	Size	Modify time	.n	Name	Size	Modify time												
/..		UP--DIR	Nov 23 20:09	/..		UP--DIR	Nov 23 20:09												
/bin		4096	Nov 22 14:40	/FileLevelBackup_0		4096	Nov 23 18:06												
/boot		4096	Nov 22 14:43																
/dev		4096	Mar 25 2016																
/etc		12288	Nov 23 17:48																
/home		4096	Nov 23 18:20																
/lib		4096	Mar 25 2016																
/lib64		4096	Nov 22 14:40																
/lost+found		16384	Mar 25 2016																
/media		4096	Nov 21 19:21																
/mnt		4096	Nov 2 20:23																
/opt		4096	Mar 25 2016																
/proc		4096	Aug 26 2015																
/root		4096	Nov 21 20:39																
/run		4096	Mar 25 2016																
/sbin		12288	Nov 22 14:41																
/sda5		4096	Apr 20 2016																
/srv		4096	Mar 25 2016																
/sys		4096	Apr 6 2015																
/tmp		4096	Nov 23 18:23																
/usr		4096	Mar 25 2016																
UP--DIR				UP--DIR															
4867M/9258M (52%)				658M/955M (68%)															
Hint: Completion works on all input lines in all dialogs. Just press M-Tab.																			
#																			
1	Help	2	Menu	3	View	4	Edit	5	Copy	6	RenMov	7	Mkdir	8	Delete	9	FullDn	10	Quit

Saving Files to Initial Location


To save restored files or folders to their initial location on your computer, do the following:

1. In the left pane of the file manager window, open the directory in your computer's file system in which the backed-up file or folder that you want to restore originally resided.
2. In the right pane of the file manager window, open the directory that contains the file or folder in the backup that you want to restore to its original location.
3. Select the file or folder that you want to restore and press [F5].

4. In the **Copy** dialog window, review the file or folder copy settings, select **Ok** and press [Enter].



5. If the file or folder you want to restore exists in its original location, Midnight Commander will display a warning. In the warning window, select the necessary operation with the target file or folder and press [Enter]. Midnight Commander will save the file or folder in its original location.




6. After you finish working with files and folders, press [F10] to close the file manager.

Saving Files to New Location

To save restored files or folders to a new location on your computer or to a network shared folder, do the following:

1. In the left pane of the file manager window, open the directory in your computer's file system in which you want to restore a file or folder.
2. In the right pane of the file manager window, open the directory that contains the file or folder in the backup that you want to restore.
3. Select the file or folder that you want to restore and press [F5].
4. In the **Copy** dialog window, review the file or folder copy settings, select **Ok** and press [Enter].



- Midnight Commander will save the file or folder to the specified location.


Left				Right			
<- ...ystem/home/user/Documents -.[^I>				<- ...er/Documents/Reports 2015 -.[^I>			
.n	Name	Size	Modify time	.n	Name	Size	Modify time
./..		UP--DIR	Nov 23 17:48	./..		UP--DIR	Apr 23 2016
/1		4096	Apr 23 2016	Report1.pdf		356864	Apr 21 2016
/Reports 2015		4096	Apr 14 2016	Report2.pdf		356864	Sep 25 2015
/Reports 2016		4096	Apr 14 2016				
Report1.pdf		356864	Apr 21 2016				
/Reports 2015				Report1.pdf			
26G/38G (67%)				42G/42G (99%)			
Hint: M-t changes quickly the listing mode.							
#							
1	Help	2	Menu	3	View	4	Edit
5	Copy	6	RenMov	7	Mkdir	8	Delete
9	PullDn	10	Quit				

- After you finish working with files and folders, press [F10] to close the file manager.

Stopping Backup Mount Session

When Veeam Agent mounts a backup for file-level restore, Veeam Agent starts a new backup mount session. To unmount a backup, you need to stop the backup mount session. This may be required, for example, if you want to stop working with files and folders in one backup and mount another backup for file-level restore.


To stop the backup mount session with the Veeam Recovery Media, in the File Level Restore wizard menu, select the **Stop backup mount** option and press [Enter]. Veeam Agent will stop the backup mount session, unmount the backup from the `/mnt/backup` directory of the recovery image OS file system, exit the File Level Restore wizard and display the Veeam Recovery Media main menu.



Step 9. Finish Working with Veeam Recovery Media

When the restore operation completes, finish working with the Veeam Recovery Media and start your operating system.

1. Eject the media or removable storage device with the recovery image.
2. In the File Level Recovery wizard menu or Veeam Recovery Media main menu, select the **Reboot** option and press [Enter].



3. Wait for your Linux operating system to start.

Restoring Volumes with Command Line Interface

You can restore a specific computer volume or all volumes from the volume-level backup.

NOTE

You cannot use the Veeam Agent for Linux command line interface to restore BTRFS subvolumes.

Volumes can be restored to their original location or to a new location.

- If you restore a volume to its original location, Veeam Agent will overwrite the data on the original volume with the data restored from the backup.
- If you restore volume data to a new location, Veeam Agent will restore data from the backup and write it to the selected destination. If necessary, you can specify new disk mapping settings for the restored volume.

You can use Veeam Agent commands to restore volumes from a backup or restore point:

- [Restore from backup](#)

When you restore a volume from the backup, Veeam Agent will automatically select the latest restore point in the backup. The volume will be restored to the state in which the volume was at the time when the latest restore point was created.

- [Restore from a restore point](#)

When you restore a volume from the restore point, you can select the necessary restore point in the backup to recover data to a specific point in time.

Before You Begin

Before you begin the volume-level restore process, check the following prerequisites:

- The volume-level backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders and on Veeam backup repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).

Volume-level restore has the following limitations:

- You cannot restore the system volume to its original location.
- You cannot restore a volume to the volume on which the Linux swap space is hosted.
- You cannot restore a volume to the volume where the backup file that you use for restore is located.

To overcome the first two limitations, you can boot from the recovery image and use the Veeam Recovery Media tools for volume-level restore. To learn more, see [Restoring from Veeam Recovery Media](#).

Restoring from Backup

With Veeam Agent command line interface, you can restore volumes from the backup. When you restore a volume from the backup, Veeam Agent automatically selects the latest restore point in the backup and restores the volume to the state in which the volume was at the time when the latest restore point was created.

Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list [--all]
```

Where `--all` is an option that instructs Veeam Agent to display information about all Veeam Agent for Linux backups in the backup repositories configured in the product. If you do not use this option, Veeam Agent will display information about the backups of the current Veeam Agent computer only.

If you work with Veeam Agent connected to a Veeam backup server as a member of a protection group for pre-installed Veeam Agents, for security reasons, the `veeamconfig backup list --all` command will display backups created only by the current Veeam Agent computer with the current connection settings. To learn more about protection groups for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

NOTE

If you cannot locate the backup from which you want to restore data, make sure Veeam Agent has access to the backup repository that contains this backup. To learn more about configuring backup repositories, see [Managing Backup Repositories](#), [Managing Veeam Backup & Replication Servers](#) and [Managing Service Providers](#).

You can also import a backup if it is stored on the Veeam Agent computer or in a network shared folder. For example, this can be a backup created with another instance of Veeam Agent. To learn more about backup import, see [Importing Backups](#).

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as <i>Imported</i> in the Repository column. For information about the import procedure, see Importing Backups .
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name      Backup ID      Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocumentsBackup {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
wrk01 SystemBackup {951ac571-dd29-45ac-8624-79b8ccb45863} Repository_
2    2023-11-13 15:26
wrk02 SystemBackup {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167} Repository_
3    2023-11-13 15:59
```

Step 2. Explore Backup Content

To view detailed information about specific backup, use the following command:

```
veeamconfig backup show --id <backup_id>
```

where:

<backup_id> – ID of the backup for which you want to view detailed information.

For a volume-level backup, Veeam Agent displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Name	Name of the volume in the backup.
Device	Path to the block device that represents the volume.
FS UUID	File system ID.
Offset	Position of the volume on the computer disk.
Size	Size of the volume in the backup.

For example:

```
user@srv01:~$ veeamconfig backup show --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
Machine name: srv01 HomePartitionBackup
  Name:      [sda6]
  Device:    [/dev/sda6]
  FS UUID:   [4967f2eb-e8bb-48fe-a694-5ba67b9030a5]
  Offset:    [11813257216] bytes (23072768 sectors)
  Size:      [41872785408] bytes (81782784 sectors)
```

Step 3. Start Restore Process

To start the process of volume-level restore from the backup, use the following command:

```
veeamconfig backup restore --id <backup_id> --targetdev <target_volume> --backupdev <volume_in_backup>
```

where:

- `<backup_id>` – ID of the backup.
- `<target_volume>` – path to a block device that represents a volume on your computer that you want to recover.
- `<volume_in_backup>` – path to a block device that represents a volume in the backup.

This parameter is optional. If you do not specify this parameter, Veeam Agent will restore from the backup a volume that has the same name as a `<target_volume>`.

For example:

```
user@srv01:~$ veeamconfig backup restore --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b --targetdev /dev/sdb --backupdev /dev/sda6
Restoring backup.
Backup: 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
Devices:
  Device in current system: [/dev/sdb]   In backup: [/dev/sda6];
You are sure? (y/n)
y
Volume restore from backup has been started.
Session ID: [{0b72ef45-4c88-4639-b940-ad3828b1cd4e}].
Logs stored in: [/var/log/veeam/Restore/Session_{0b72ef45-4c88-4639-b940-ad3828b1cd4e}].
```

IMPORTANT

You can restore a backed-up volume only to a target volume that is not used by your Linux OS (that does not have file system mount points). For example, you can add a new disk to your computer and restore a volume in the backup to this disk.

If you want to restore a volume to the location that is crucial for the OS running, you should boot from the Veeam Recovery Media and perform volume-level restore with the Volume Restore wizard. For example, this approach is helpful when you restore the root (/) partition.

Alternatively, if the volume is backed-up in the unmounted state, it can be restored without booting from the Veeam Recovery Media.

Step 4. Monitor Restore Process

You can monitor the restore process by viewing the restore session log in the command line interface.

To view Veeam Agent for Linux session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session_id> – ID of the restore session.

For example:

```
user@srv01:~$ veeamconfig session log --id 0b72ef45-4c88-4639-b940-ad3828b1cd4e
2023-11-27 11:04:04 UTC {b141f32a-3e77-45a6-b55a-c100a1464d67} [info] Job start
ed at 2023-11-27 14:04:04
2023-11-27 11:04:04 UTC {9b60ac03-2de0-4fe2-a00e-bec556d98ee8} [info] Starting
volume restore
2023-11-27 11:04:07 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [processing] sdb
2023-11-27 11:04:15 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 512.0
kB at 58.6kB/s (0%)
...
2023-11-27 11:14:35 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 6.5GB
at 10.6MB/s (97%)
2023-11-27 11:14:37 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb 6.5GB
at 10.6MB/s (100%)
2023-11-27 11:14:37 UTC {00add723-cbfa-4cc8-b299-d2349a051d6f} [warn] /dev/sdb
has a duplicate filesystem UUID
2023-11-27 11:14:37 UTC {ced9af4a-6af1-4756-8ffb-8ec1325e18ec} [info] sdb resto
red 6.5GB at 10.6MB/s
2023-11-27 11:14:37 UTC {8b8742a2-1c80-4e14-bbf1-45a3612bc3a7} [info] Volume re
store completed
```

TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

Restoring from Restore Point

With Veeam Agent command line interface, you can restore volumes from the specific restore point. When you restore a volume from the restore point, you can select the necessary restore point in the backup to recover data to a desired point in time.

Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list [--all]
```

Where `--all` is an option that instructs Veeam Agent to display information about all Veeam Agent for Linux backups in the backup repositories configured in the product. If you do not use this option, Veeam Agent will display information about the backups of the current Veeam Agent computer only.

If you work with Veeam Agent connected to a Veeam backup server as a member of a protection group for pre-installed Veeam Agents, for security reasons, the `veeamconfig backup list --all` command will display backups created only by the current Veeam Agent computer with the current connection settings. To learn more about protection groups for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

NOTE

If you cannot locate the backup from which you want to restore data, make sure Veeam Agent has access to the backup repository that contains this backup. To learn more about configuring backup repositories, see [Managing Backup Repositories](#), [Managing Veeam Backup & Replication Servers](#) and [Managing Service Providers](#).

You can also import a backup if it is stored on the Veeam Agent computer or in a network shared folder. For example, this can be a backup created with another instance of Veeam Agent. To learn more about backup import, see [Importing Backups](#).

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as <i>Imported</i> in the Repository column. For information about the import procedure, see Importing Backups .
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name      Backup ID      Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocumentsBackup {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
wrk01 SystemBackup {951ac571-dd29-45ac-8624-79b8ccb45863} Repository_
2    2023-11-13 15:26
wrk02 SystemBackup {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167} Repository_
3    2023-11-13 15:59
```

Step 2. Explore Restore Points

To view information about restore points in the backup, use the following command:

```
veeamconfig backup info --id <backup_id>
```

or

```
veeamconfig point list --backupid <backup_id>
```

where

<backup_id> – ID of the backup for which you want to view information on restore points.

You can view the following information about restore points in the backup:

Parameter	Description
Job name	Name of the backup job by which the backup was created.
OIB ID	ID of the restore point in the backup.
Type	Type of the restore point. Possible values: • Full • Increment.
Created at	Date and time of the restore point creation.
Is corrupt	Indicates whether restore point in the backup is corrupted. Possible values: • True • False
Retention	Displays information about enabled long-term retention per each type: weekly (W), monthly (M) and yearly (Y).

For example:

```
user@srv01:~$ veeamconfig backup info --id 4f75bb20-a6b6-4323-9287-1c6c8ceccb6b
Job name      OIB ID      Type      C
reated at    Is corrupt  Retention
srv01 HomePartitionBackup {23cb927d-5e2d-42fe-a4a4-e5f254a6413e} Full      2
023-11-15 11:28 false      WM
srv01 HomePartitionBackup {25e31075-4c30-4d67-86a6-293c0887f4eb} Increment 2
023-11-15 11:58 false
srv01 HomePartitionBackup {9375140d-720a-4d3e-a69b-ab9cf60d53fa} Increment 2
023-11-27 13:15 false
```

or

```
user@srv01:~$ veeamconfig point list --backupid 4f75bb20-a6b6-4323-9287-1c6c8ce
ccb6b
Job name          OIB ID          Type          C
reated at        Is corrupt      Retention
srv01 HomePartitionBackup {23cb927d-5e2d-42fe-a4a4-e5f254a6413e} Full          2
023-11-15 11:28 false      WM
srv01 HomePartitionBackup {25e31075-4c30-4d67-86a6-293c0887f4eb} Increment     2
023-11-15 11:58 false
srv01 HomePartitionBackup {9375140d-720a-4d3e-a69b-ab9cf60d53fa} Increment     2
023-11-27 13:15 false
```

Step 3. Start Restore Process

To start the process of volume-level restore from the specific restore point, use the following command:

```
veeamconfig point restore --id <point_id> --targetdev <target_volume> --backupdev <volume_in_backup>
```

where:

- `<point_id>` – ID of the restore point.
- `<target_volume>` – path to a block device that represents a volume on your computer that you want to recover.
- `<volume_in_backup>` – path to a block device that represents a volume in the backup.

This parameter is optional. If you do not specify this parameter, Veeam Agent will restore from the backup a volume that has the same name as a `<target_volume>`.

For example:

```
user@srv01:~$ veeamconfig point restore --id 9375140d-720a-4d3e-a69b-ab9cf60d53fa --backupdev /dev/sda6 --targetdev /dev/sdb
Restoring point.
Restore point: 9375140d-720a-4d3e-a69b-ab9cf60d53fa
Devices:
  Device in current system: [/dev/sdb]   In backup: [/dev/sda6];
You are sure? (y/n)
y
Volume restore by point has been started.
Session ID: [{697d9348-9001-4845-8764-3cc4fb3f296b}].
Logs stored in: [/var/log/veeam/Restore/Session_{697d9348-9001-4845-8764-3cc4fb3f296b}].
```

IMPORTANT

You can restore a backed-up volume only to a target volume that is not used by your Linux OS (that does not have file system mount points). For example, you can add a new disk to your computer and restore a volume in the backup to this disk.

If you want to restore a volume to the location that is crucial for the OS running, you should boot from the Veeam Recovery Media and perform volume-level restore with the Volume Restore wizard. For example, this approach is helpful when you restore the root (/) partition.

Alternatively, if the volume is backed-up in the unmounted state, it can be restored without booting from the Veeam Recovery Media.

Step 4. Monitor Restore Process

You can monitor the restore process by viewing the restore session log in the command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session_id> – ID of the restore session.

For example:

```
user@srv01:~$ veeamconfig session log --id 697d9348-9001-4845-8764-3cc4fb3f296b
2023-11-27 10:35:47 UTC {b9604775-d265-4537-b98e-848fd77c7375} [info] Job started at 2023-11-27 13:35:47
2023-11-27 10:35:47 UTC {ed66a1f6-5216-4596-a7b5-be10dd10c32f} [info] Starting volume restore
2023-11-27 10:35:50 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [processing] sdb
2023-11-27 10:35:59 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [info] sdb 512.0 kB at 59.1kB/s (0%)
...
2023-11-27 10:46:27 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [info] sdb 6.5GB at 10.5MB/s (100%)
2023-11-27 10:46:28 UTC {dae118c8-eb7c-4e14-9832-f0bfd089b329} [warn] /dev/sdb has a duplicate filesystem UUID
2023-11-27 10:46:28 UTC {2e37de47-c4e2-46f9-8b70-f24fbff3697d} [info] sdb restored 6.5GB at 10.5MB/s
2023-11-27 10:46:28 UTC {a21a89d9-d0ca-4f5c-8399-28ae599f2f1c} [info] Volume restore completed
```

TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

Restoring Files and Folders with Recovery Wizard

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore, Veeam Agent publishes the backup content directly into the computer file system. You can browse to files and folders in the backup, restore files and folders to their initial location, copy files and folders to a new location or simply target applications to restored files and work with them as usual.

Before You Begin

Before you begin the file-level restore process, check the following prerequisites:


- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
- [For backups of BTRFS file system] A machine on which you perform file-level restore must run the same or later Linux kernel version as the machine on which the backup was created.

For example, you created a backup of a machine that runs Linux kernel version 4.14. If you perform file-level restore from this backup on another machine that runs Linux kernel 2.6, the file-level restore process will fail.

Step 1. Launch File Level Restore Wizard

To launch the **File Level Restore** wizard, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, press the [R] key to proceed to the File Level Restore wizard.



Step 2. Select Backup and Restore Point

At the **Backup** step of the wizard, select a backup and restore point from which you want to recover data.

The **Backup** step window comprises two panes:

- The **Imported Backups** pane displays available backups and information about each backup: host name of the computer whose data is stored in the backup file, backup job name and number of restore points. By default, Veeam Agent displays the list of backups that were created by the backup jobs configured in Veeam Agent on your computer only.

If Veeam Agent is connected to a Veeam Backup & Replication server or Veeam Cloud Connect service provider, and the account you used to connect to the Veeam backup server has the Veeam Backup Administrator role, you can also list all other compatible Veeam Agent for Linux backups in that Veeam backup repository. To display such backups, switch to the **Show all** link at the bottom of the backup list and press [Enter].

TIP

If Veeam Agent fails to display backups stored in the Veeam backup repository for some reason, you can press the [R] key to rescan the backup repository. Veeam Agent will try to reconnect to the Veeam backup server and refresh the list of backups.

You can also restore data from a backup that is created with this or another instance of Veeam Agent and stored in a backup repository to which Veeam Agent is not currently connected. In this case, you can import such backup to your computer. When you import a backup, you connect Veeam Agent to a backup repository and select a backup file. This information is then saved in the Veeam Agent database and the specified backup file becomes available for restore.

To connect to another backup repository, press the [i] key. At the **Select Backup Location** step of the restore wizard, select backup location and specify connection settings to the backup repository. Connecting to a backup repository is very similar to the same procedure during bare metal recovery. For details, see [Specify Backup Location Settings](#).


After you connect to the repository, [select backup and restore point](#).

To view backups imported from other repositories, switch to the **Show all** link at the bottom of the backup list and press [Enter].

- The **Restore Points** pane on the right displays a list of restore points in the backup.

To select backup and restore point:

1. In the **Imported Backups** pane, select with [Up] and [Down] keys the backup from which you want to recover data and press [Enter].



2. In the **Restore Points** pane, select with [Up] and [Down] keys the restore point from which you want to recover data and press [Enter].

NOTE

If you selected an encrypted backup for data restore, Veeam Agent will prompt you to provide a password to unlock the encrypted file. To learn more, see [Restoring Data from Encrypted Backups](#).

Veeam Agent for Linux [srv02]			
IMPORTED BACKUPS			RESTORE POINTS
Job name	Hostname	Points	Created at
srv02 DailyBackup	srv02	4	06:10 14-12-2022
srv02 srv02CloudBackup	srv02	5	06:10 13-12-2022
srv02 srv02CloudBackup(encrypted)srv02		2	06:00 12-12-2022
Show all...			22:14 11-12-2022
I Import backup Enter Next Esc Main menu			

- Veeam Agent will mount the content of the backup file to the `/mnt/backup` directory in the computer's file system and display a notification window with the corresponding message. Press [Enter] to close the window and return to the Veeam Agent control panel.

IMPORTANT

When you finish working with restored data, you must unmount the backup from the `/mnt/backup` folder. To unmount the backup, in the Veeam Agent for Linux control panel, press [U]. For details, see [Stop Backup Mount Session](#).

Veeam Agent for Linux
[srv02]

Latest backup sessions:

Job name	State	Started at	Finished at
DailyBackup	Success	2022-12-14 06:10:04	2022-12-14 06:16:30
DailyBackup	Failed	2022-12-14 06:00:03	2022-12-14 06:00:04
srv02CloudBackup	Success	2022-12-14 06:00:03	2022-12-14 06:07:11
srv02CloudBackup	Success	2022-12-13 21:54:12	2022-12-13 22:00:22
srv02CloudBackup	Success	2022-12-13 20:44:16	2022-12-13 20:46:50
DailyBackup	Success	2022-12-13 06:10:01	2022-12-13 06:10:42
DailyBackup	Failed	2022-12-13 06:00:01	2022-12-13 06:00:01
srv02CloudBackup			2-12-13 06:00:44
srv02CloudBackup			2-12-12 21:20:21
srv02CloudBackup			2-12-12 19:32:21
srv02CloudBackup			2-12-12 15:43:40
DailyBackup			2-12-12 06:00:53
DailyBackup			2-12-11 22:19:24

Info

Backup content has been mounted successfully into /mnt/backup

OK

Enter
Show
C
Configure
S
Start Job
U
Unmount
M
Misc
Esc
Quit

Step 3. Save Restored Files

When the backup file content is mounted to the `/mnt/backup` directory in the computer's file system, you can use Linux command line utilities or preferred file browser to work with restored files and directories. You can browse for files and directories in the mounted backup and copy files and directories that you want to restore to their initial location or to a new location.

NOTE

If a backed up file system was mounted to multiple mount points, during restore, depending on the file system type, Veeam Agent will behave as follows:

- [root file system] Veeam Agent will restore all mount points to the root directory.
- [non-root file system] Veeam Agent will restore all mount points to a single mount point randomly chosen from the mount points to which it was originally mounted.

In the following example, the restored file `Report1.pdf` is copied from the mounted backup to the new location with Linux command line utilities:


```
user@srv01:~$ ls Documents/
Reports
user@srv01:~$ ls /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/
Report1.pdf Report2.pdf
user@srv01:~$ cp /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/Report1.pdf Documents/
user@srv01:~$ ls Documents/
Report1.pdf Reports
```


Step 4. Stop Backup Mount Session

When Veeam Agent mounts a backup for file-level restore, Veeam Agent starts a new backup mount session. To unmount a backup, you need to stop the backup mount session. This may be required, for example, if you want to stop working with files and folders in one backup and mount another backup for file-level restore. You can also stop the backup mount session to unmount a backup after you have finished working with restored files and folders.

To stop the backup mount session, do the following:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command:
2. In the Veeam Agent control panel, press the [U] key to unmount a backup.
3. Veeam Agent will stop the backup mount session and display a notification window. Press [Enter] to close the window and return to the Veeam Agent control panel.



Restoring Files and Folders with Command Line Interface

If some files and folders on your computer get lost or corrupted, you can restore them from backups. For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

When you perform file-level restore, Veeam Agent publishes the backup content directly into the computer file system. You can browse to files and folders in the backup, restore files and folders to their initial location, copy files and folders to a new location or simply target applications to restored files and work with them as usual.

With the Veeam Agent command line interface, you can restore files and folders in a more flexible way than with the use of the File Level Restore wizard. In particular, you can specify a directory in which Veeam Agent should mount the backup file content for file-level restore. You can also mount several backups to different directories to work with files and folders restored from different backups simultaneously.

You can use Veeam Agent commands to restore files and folders from backup or from specific restore point:

- [Restore from backup](#)

When you restore files and folders from the backup, Veeam Agent will automatically select the latest restore point in the backup. You can restore files and folders to the state in which they were at the time when the latest restore point was created.

- [Restore from a restore point](#)

When you restore files and folders from the restore point, you can select the necessary restore point in the backup to recover data to a specific point in time.

Before You Begin

Before you begin the file-level restore process, check the following prerequisites:

- The backup from which you plan to restore data must be successfully created at least once.
- [For backups stored in network shared folders, on Veeam backup repositories and Veeam Cloud Connect repositories] You must have access to the target location where the backup file resides.
- [For Veeam backup repository targets] If you plan to restore data from a backup stored on a backup repository, you must have access permissions on this backup repository. To learn more, see [Setting Up User Permissions on Backup Repositories](#).
- [For backups of BTRFS file system] A machine on which you perform file-level restore must run the same or later Linux kernel version as the machine on which the backup was created.

For example, you created a backup of a machine that runs Linux kernel version 4.14. If you perform file-level restore from this backup on another machine that runs Linux kernel 2.6, the file-level restore process will fail.

Restoring from Backup

With Veeam Agent command line interface, you can restore files and folders from the backup. When you perform file-level restore from the backup, Veeam Agent for Linux automatically selects the latest restore point in the backup. You can restore files and folders to the state in which they were at the time when the latest restore point was created.

Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list [--all]
```

Where `--all` is an option that instructs Veeam Agent to display information about all Veeam Agent for Linux backups in the backup repositories configured in the product. If you do not use this option, Veeam Agent will display information about the backups of the current Veeam Agent computer only.

If you work with Veeam Agent connected to a Veeam backup server as a member of a protection group for pre-installed Veeam Agents, for security reasons, the `veeamconfig backup list --all` command will display backups created only by the current Veeam Agent computer with the current connection settings. To learn more about protection groups for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

NOTE

If you cannot locate the backup from which you want to restore data, make sure Veeam Agent has access to the backup repository that contains this backup. To learn more about configuring backup repositories, see [Managing Backup Repositories](#), [Managing Veeam Backup & Replication Servers](#) and [Managing Service Providers](#).

You can also import a backup if it is stored on the Veeam Agent computer or in a network shared folder. For example, this can be a backup created with another instance of Veeam Agent. To learn more about backup import, see [Importing Backups](#).

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as <i>Imported</i> in the Repository column. For information about the import procedure, see Importing Backups .
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name      Backup ID      Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocumentsBackup {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
wrk01 SystemBackup {951ac571-dd29-45ac-8624-79b8ccb45863} Repository_
2    2023-11-13 15:26
wrk02 SystemBackup {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167} Repository_
3    2023-11-13 15:59
```

Step 2. Explore Backup Content

For file-level restore, you can use backups of any type:

- Volume-level backups (backups of the entire computer or specific volumes)
- File-level backups

To view detailed information about specific backup, use the following command:

```
veeamconfig backup show --id <backup_id>
```

where:

<backup_id> – ID of the backup for which you want to view detailed information.

For a volume-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Name	Name of the volume in the backup.
Device	Path to the block device that represents the volume.
FS UUID	File system ID.
Offset	Position of the volume on the computer disk.
Size	Size of the volume in the backup.

For a file-level backup, Veeam Agent for Linux displays the following information:

Parameter	Description
Machine name	Host name of the machine on which the backup job is configured and the name of the job.
Backed up	Backup scope for the file-level backup job.

For example:

```
user@srv01:~$ veeamconfig backup show --id ea64a7e5-038a-4c86-970a-6d59d4cf3968
Machine name: srv01 DocumentsBackup
File-level backup
Backed up:
/home/user/Documents
```


Step 3. Mount Backup

To mount a backup for file-level restore, use the following command:

```
veeamconfig backup mount --id <backup_id> --mountdir <path>
```

where:

- `<backup_id>` – ID of the backup that you want to mount to the computer file system for file-level restore.
- `<path>` – path to the directory to which you want to mount the backup file content.

For example:

```
user@srv01:~$ veeamconfig backup mount --id ea64a7e5-038a-4c86-970a-6d59d4cf3968 --mountdir /mnt/backup
Backup is mounted.
Session ID: [{2a313184-32d0-4d3a-a1b0-2eebac986047}].
Logs stored in: [/var/log/veeam/Mount/Session_{2a313184-32d0-4d3a-a1b0-2eebac986047}].
```

Step 4. Monitor Mount Process and Result

You can monitor the backup mount process by viewing the mount session log in the command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session_id> – ID of the backup mount session.

For example:

```
user@srv01:~$ veeamconfig session log --id 2a313184-32d0-4d3a-a1b0-2eebac986047
2023-11-22 17:30:34 UTC {30878c82-27d0-45dc-ab21-6f27d5082fd4} [info] Job started at 2023-11-22 20:30:34
2023-11-22 17:30:34 UTC {714b21d0-0d20-486e-b1e5-22d5fb5a8ee9} [info] Mounting restore point
2023-11-22 17:30:35 UTC {d331f038-5b7c-4549-85cf-5e1b54dbaf71} [info] Restore point has been mounted
```

To ensure that the backup is successfully mounted, you can browse to the directory that you specified in the `veeamconfig backup mount` command. For example:

```
user@srv01:~$ ls /mnt/backup/
FileLevelBackup_0
```

TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

Step 5. Save Restored Files

When the backup file content is mounted to the computer file system, you can use Linux command line utilities or preferred file browser to work with restored files and folders. You can browse for files and folders in the mounted backup and copy files and folders that you want to restore to their initial location or to a new location.

NOTE

If a backed up file system was mounted to multiple mount points, during restore, depending on the file system type, Veeam Agent will behave as follows:

- [root file system] Veeam Agent will restore all mount points to the root directory.
- [non-root file system] Veeam Agent will restore all mount points to a single mount point randomly chosen from the mount points to which it was originally mounted.

In the following example, the restored file `Report1.pdf` is copied from the mounted backup to a new location with the Linux command line utilities:

```
user@srv01:~$ ls Documents/
Reports
user@srv01:~$ ls /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/
Report1.pdf  Report2.pdf
user@srv01:~$ cp /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/Report1.pdf Documents/
user@srv01:~$ ls Documents/
Report1.pdf  Reports
```

Step 6. Stop Backup Mount Session

When Veeam Agent mounts a backup for file-level restore, Veeam Agent starts a new backup mount session. After you have finished working with restored files and folders, you can stop the backup mount session to unmount the backup.

To stop the backup mount session, use the following command:

```
veeamconfig session stop --id <session_id>
```

where:

<session_id> – ID of the backup mount session that you want to stop.

Veeam Agent will stop the mount session and unmount the backup from the computer file system. For example:

```
user@srv01:~$ veeamconfig session stop --id 2a313184-32d0-4d3a-a1b0-2eebac98604
7
Session has stopped.
user@srv01:~$ ls /mnt
user@srv01:~$
```

Restoring from Restore Point

With Veeam Agent command line interface, you can restore files and folders from the specific restore point. When you restore files and folders from the restore point, you can select the necessary restore point in the backup to recover data to a specific point in time.

Step 1. Locate Backup

To view a list of backups created by Veeam Agent, use the following command:

```
veeamconfig backup list [--all]
```

Where `--all` is an option that instructs Veeam Agent to display information about all Veeam Agent for Linux backups in the backup repositories configured in the product. If you do not use this option, Veeam Agent will display information about the backups of the current Veeam Agent computer only.

If you work with Veeam Agent connected to a Veeam backup server as a member of a protection group for pre-installed Veeam Agents, for security reasons, the `veeamconfig backup list --all` command will display backups created only by the current Veeam Agent computer with the current connection settings. To learn more about protection groups for pre-installed Veeam Agents, see the [Protection Group Types](#) section in the Veeam Agent Management Guide.

NOTE

If you cannot locate the backup from which you want to restore data, make sure Veeam Agent has access to the backup repository that contains this backup. To learn more about configuring backup repositories, see [Managing Backup Repositories](#), [Managing Veeam Backup & Replication Servers](#) and [Managing Service Providers](#).

You can also import a backup if it is stored on the Veeam Agent computer or in a network shared folder. For example, this can be a backup created with another instance of Veeam Agent. To learn more about backup import, see [Importing Backups](#).

For each backup, Veeam Agent displays the following information:

Parameter	Description
Job name	Host name of the computer on which the backup job was configured and name of the job by which the backup was created.
Backup ID	ID of the backup.
Repository	Name of the backup repository in which the backup was created. Imported backups are marked as <i>Imported</i> in the Repository column. For information about the import procedure, see Importing Backups .
Created at	Date and time of the backup creation.

For example:

```
user@srv01:~$ veeamconfig backup list --all
Job name      Backup ID      Repository
y    Created at
srv01 SystemBackup {45f074d2-d2d9-423d-84e9-8f1798b08d4c} Repository_
1    2023-11-11 17:37
srv01 DocumentsBackup {ea64a7e5-038a-4c86-970a-6d59d4cf3968} Repository_
1    2023-11-11 18:30
srv01 HomePartitionBackup {4f75bb20-a6b6-4323-9287-1c6c8ceccb6b} Repository_
2    2023-11-15 11:28
wrk01 SystemBackup {951ac571-dd29-45ac-8624-79b8ccb45863} Repository_
2    2023-11-13 15:26
wrk02 SystemBackup {8d6d4d39-51b2-48b1-ac7a-84f2d6dbc167} Repository_
3    2023-11-13 15:59
```

Step 2. Explore Restore Points

To view information about restore points in the backup, use the following command:

```
veeamconfig backup info --id <backup_id>
```

or

```
veeamconfig point list --backupid <backup_id>
```

where:

<backup_id> – ID of the backup for which you want to view information on restore points.

You can view the following information about restore points in the backup:

Parameter	Description
Job name	Name of the backup job by which the backup was created.
OIB ID	ID of the restore point in the backup.
Type	Type of the restore point. Possible values: • Full • Increment.
Created at	Date and time of the restore point creation.
Is corrupt	Indicates whether restore point in the backup is corrupted. Possible values: • True • False
Retention	Displays information about enabled long-term retention per each type: weekly (W), monthly (M) and yearly (Y).

For example:

```
user@srv01:~$ veeamconfig backup info --id ea64a7e5-038a-4c86-970a-6d59d4cf3968
Job name      OIB ID      Type      Creat
ed at      Is corrupt  Retention
srv01 DocumentsBackup {0f3c9f3e-3985-4dc9-8cd6-979dba810c2f} Full      2023-
11-11 18:31 false      M
srv01 DocumentsBackup {ff0c6969-8b9b-4865-b4f9-d686faf41d50} Increment 2023-
11-14 13:35 false
srv01 DocumentsBackup {a9e420df-d749-4b9a-b675-19d8e94c3bf1} Increment 2023-
11-15 13:43 false
```


or

```
user@srv01:~$ veeamconfig point list --backupid ea64a7e5-038a-4c86-970a-6d59d4c
f3968
Job name          OIB ID          Type          Creat
ed at          Is corrupt Retention
srv01 DocumentsBackup {0f3c9f3e-3985-4dc9-8cd6-979dba810c2f} Full          2023-
11-11 18:31 false      M
srv01 DocumentsBackup {ff0c6969-8b9b-4865-b4f9-d686faf41d50} Increment 2023-
11-14 13:35 false
srv01 DocumentsBackup {a9e420df-d749-4b9a-b675-19d8e94c3bf1} Increment 2023-
11-15 13:43 false
```

Step 3. Mount Restore Point

To mount a backup for file-level restore, use the following command:

```
veeamconfig point mount --id <point_id> --mountdir <path>
```

where:

- `<point_id>` – ID of the restore point that you want to mount to the computer file system for file-level restore.
- `<path>` – path to the directory to which you want to mount the backup file content.

For example:

```
user@srv01:~$ veeamconfig point mount --id b127e64e-1f1c-4e0b-bb36-b087761267b3
--mountdir /mnt/backup
Restore point is mounted.
Session ID: [{4d69dd85-ac60-4cff-883d-50f25f49a9c8}].
Logs stored in: [/var/log/veeam/Mount/Session_{4d69dd85-ac60-4cff-883d-50f25f49a9c8}].
```

Step 4. Monitor Mount Process and Result

You can monitor the restore point mount process by viewing the mount session log in the command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session_id> – ID of the restore point mount session.

For example:

```
user@srv01:~$ veeamconfig session log --id 4d69dd85-ac60-4cff-883d-50f25f49a9c8
2023-11-23 12:44:55 UTC {9c5c8ece-cb88-4742-bb90-1f8ff79b4bdc} [info] Job started at 2023-11-23 15:44:55
2023-11-23 12:44:55 UTC {4ac10045-a74b-4a41-9c5e-53521cba1045} [info] Mounting restore point
2023-11-23 12:44:56 UTC {540a61f7-5d5c-47d5-a2b8-51daa694d5ec} [info] Restore point has been mounted
```

To ensure that the restore point is successfully mounted, you can browse to the directory that you specified in the `veeamconfig point mount` command. For example:

```
user@srv01:~$ ls /mnt/backup/
FileLevelBackup_0
```

TIP

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

Step 5. Save Restored Files

When the restore point is mounted to the computer file system, you can use Linux command line utilities or preferred file browser to work with restored files and folders. You can browse for files and folders in the mounted backup and copy files and folders that you want to restore to their initial location or to a new location.

NOTE

If a backed up file system was mounted to multiple mount points, during restore, depending on the file system type, Veeam Agent will behave as follows:

- [root file system] Veeam Agent will restore all mount points to the root directory.
- [non-root file system] Veeam Agent will restore all mount points to a single mount point randomly chosen from the mount points to which it was originally mounted.

In the following example, the restored file `Report1.pdf` is copied from the mounted restore point to a new location with the Linux command line utilities:

```
user@srv01:~$ ls Documents/  
Reports  
user@srv01:~$ ls /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/  
Report1.pdf  Report2.pdf  
user@srv01:~$ cp /mnt/backup/FileLevelBackup_0/home/user/Documents/Reports/Report1.pdf Documents/  
user@srv01:~$ ls Documents/  
Report1.pdf  Reports
```

Step 6. Stop Backup Mount Session

When Veeam Agent mounts a restore point for file-level restore, Veeam Agent starts a new restore point mount session. After you have finished working with restored files and folders, you can stop the mount session to unmount the restore point.

To stop the restore point mount session, use the following command:

```
veeamconfig session stop --id <session_id>
```

where:

<session_id> – ID of the restore point mount session that you want to stop.

Veeam Agent will stop the mount session and unmount the restore point from the computer file system. For example:

```
user@srv01:~$ veeamconfig session stop --id 4d69dd85-ac60-4cff-883d-50f25f49a9c
8
Session has stopped.
user@srv01:~$ ls /mnt
user@srv01:~$
```

Exporting Backup to Virtual Disk

You can export a backup to a virtual disk in the VHD format. You can then attach the created VHD disk to a virtual machine to recover your computer in a virtual environment.

- [Exporting Backups](#)
- [Exporting Restore Points](#)

Exporting Backups

You can export the backup file to a virtual disk in the VHD format. When you export a backup, you export to a virtual disk data pertaining to the latest restore point in the backup. The created VHD disk will reflect the state in which backed-up volumes were at the time when the latest restore point was created.

To export backup to a VHD disk:

1. Start the export process with the following command:

```
veeamconfig backup export --id <backup_id> --outdir <path>
```

where:

- o <backup_id> – ID of the backup that you want to export to a virtual disk.
- o <path> – full path to a directory in which you want to save the created virtual disk. Specifying relative paths is not supported.

For example:

```
user@srv01:~$ veeamconfig backup export --id 45f074d2-d2d9-423d-84e9-8f179
8b08d4c --outdir /home/user/disk
Export has been started.
Session ID: [{5f001367-8937-46e0-a756-449bf9f1a182}].
Logs stored in: [/var/log/veeam/Export/Session_{5f001367-8937-46e0-a756-44
9bf9f1a182}].
```

2. You can monitor the export process and result by viewing the export session log with the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session_id> – ID of the export session.

For example:

```
user@srv01:~$ veeamconfig session log --id 5f001367-8937-46e0-a756-449bf9f
1a182
2023-11-27 11:20:56 UTC {b54af37c-35a6-4807-80d2-0f070f024e69} [info] Job
started at 2023-11-27 14:20:56
2023-11-27 11:20:56 UTC {48d699d2-86cf-4a32-b9c8-ab51b8325f3c} [info] Expo
rting virtual disks content
2023-11-27 11:20:57 UTC {0e2e7d97-f067-4823-8dde-084c401eb62b} [processing
] Restoring device: [30460cb5].
2023-11-27 11:22:59 UTC {0e2e7d97-f067-4823-8dde-084c401eb62b} [info] Devi
ce [30460cb5] has been exported
2023-11-27 11:23:00 UTC {36f0d0c5-2af7-48d8-abc2-c8ef9aaffc54} [info] Virt
ual disks content has been exported
```

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

3. Exported backup will be saved as a virtual disk file in the specified directory. You can check this with a file browser or with the following command:

```
ls <path>
```

where:

<path> – path to the directory in which the virtual disk with the backup is saved.

For example:

```
user@srv01:~$ ls disk/  
dev_30460cb5.vhd
```


Exporting Restore Points

You can export the specific restore point to a virtual disk in VHD format. When you export a restore point, you select the necessary restore point in the backup to recover data to a desired point in time. The created VHD disk will reflect the state in which backed-up volumes were at the time when the selected restore point was created.

To export restore point to a VHD disk:

1. Start the export process with the following command:

```
veeamconfig point export --id <point_id> --outdir <path>
```

where:

- o <point_id> – ID of the restore point that you want to export to a virtual disk.
- o <path> – full path to a directory in which you want to save the created virtual disk. Specifying relative paths is not supported.

For example:

```
user@srv01:~$ veeamconfig point export --id b319ealf-59a2-41ea-9ca3-b668e86ac941 --outdir /home/user/veeam/  
Export has been started.  
Session ID: [{aeb9c549-a660-4a0e-b89c-cb076b8bfa85}].  
Logs stored in: [/var/log/veeam/Export/Session_{aeb9c549-a660-4a0e-b89c-cb076b8bfa85}].
```

2. You can monitor the export process and result by viewing the export session log with the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session_id> – ID of the export session.

For example:

```
user@srv01:~$ veeamconfig session log --id aeb9c549-a660-4a0e-b89c-cb076b8bfa85  
2023-05-05 11:15:21 UTC {b950503d-55c9-435f-946e-1078184f5a86} [info] Job started at 2023-05-05 14:15:21  
2023-05-05 11:15:21 UTC {32d56391-9002-431e-ae6b-2285537a67e5} [info] Exporting virtual disks content  
2023-05-05 11:15:22 UTC {ba3dabe0-0556-430c-9671-9448a6dc4bcb} [processing] Restoring device: [30460cb5].  
2023-05-05 11:17:26 UTC {ba3dabe0-0556-430c-9671-9448a6dc4bcb} [info] Device [30460cb5] has been exported  
2023-05-05 11:17:26 UTC {9e945c29-900e-4a07-9e3b-ccf7f156807d} [info] Virtual disks content has been exported
```

You can also check the restore session status with the `veeamconfig session info` command. To learn more, see [Viewing Session Status](#).

3. Exported backup will be saved as a virtual disk file in the specified directory. You can check this with a file browser or with the following command:

```
ls <path>
```

where

<path> – path to the directory in which the virtual disk with the backup is saved.

For example:

```
user@srv01:~$ ls /home/user/veeam/  
dev_30460cb5.vhd
```

Restoring Data from Encrypted Backups

When you restore data from an encrypted backup, Veeam Agent performs data decryption automatically in the background or requires you to specify a password.

IMPORTANT

You cannot restore data from a backup encrypted in Veeam Backup & Replication using KMS keys.

Consider the following about unlocking encrypted backup files:

- If encryption keys required to unlock the backup file are available in the Veeam Agent database, that is, if you encrypt and decrypt the backup file on the same Veeam Agent computer, you do not need to specify the password. Veeam Agent uses keys from the database to unlock the backup file. Data decryption is performed in the background, and data restore from the encrypted backup does not differ from that from an unencrypted one.
- If encryption keys are not available in the Veeam Agent database, you need to provide a password to unlock the encrypted file. The password must be the same as the password that was used to encrypt the backup file. If the password has changed once or several times, you need to specify the latest password. In Veeam Agent, you can use the latest password to restore data from all restore points in the backup chain, including restore points that were encrypted with an old password and restore points that were created before you have enabled the encryption option for the job.

The process of unlocking an encrypted backup file differs depending on what Veeam Agent user interface you use for data restore.


- [Veeam Agent graphical user interface](#)
- [Veeam Agent command line interface](#)

Restoring Data from Encrypted Backups Using GUI

To restore data from an encrypted backup using the Veeam Agent graphical user interface:


1. Launch the necessary data restore wizard:
 - If you want to perform file-level restore from an encrypted backup that was created on another Veeam Agent computer, launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command. To learn more, see [Restoring Files and Folders](#).
 - If you want to perform volume-level restore or file-level restore recovery from an encrypted backup, boot from the Veeam Recovery Media and launch the necessary data restore wizard. To learn more, see [Restoring from Veeam Recovery Media](#).
2. Follow the steps of the wizard to specify where the encrypted backup file that you plan to use for restore resides. If the backup file resides in a remote location, select the backup location type and specify settings to connect to the backup location.

3. Select the encrypted backup and restore point from which you want to restore data.



4. Veeam Agent will display the **Encryption** window. Enter the password for the backup file.
In the **Hint** field of the **Encryption** window, Veeam Agent displays a hint for the password that was used to encrypt the backup file. Use the hint to recall the password.
If you changed the password one or several times while the backup chain was created, you need to specify the latest password. In Veeam Agent, you can use the latest password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

If you enter correct password, Veeam Agent will decrypt the backup metadata. You will be able to continue the restore operation in a regular manner.



Restoring Data from Encrypted Backups Using Command Line Interface

To restore data from an encrypted backup using the Veeam Agent command line interface, complete the following steps:

1. Import the encrypted backup file to the Veeam Agent database. To learn more, see [Importing Encrypted Backups](#).
2. Perform the necessary restore operation in a regular manner. To learn more, see [Restoring Volumes with Command Line Interface](#) and [Restoring Files and Folders with Command Line Interface](#).

Reporting

Veeam Agent for Linux provides several ways to get information about performed operations:

- With the Veeam Agent control panel
- With the Veeam Agent command line interface

For every data transfer operation, for example data backup and restore, backup import and export, Veeam Agent starts a new session. You can monitor performance of sessions started by Veeam Agent in the following ways:

- [Monitor backup job session progress with the control panel.](#)
- [View real-time backup job session statistics with the control panel.](#)
- [View backup job sessions results with the control panel.](#)
- [View the session status using the command line interface.](#)
- [View session logs.](#)

Viewing Job Session Progress

You can monitor the backup job session progress in the list of sessions in the Veeam Agent control panel. For the currently running backup job session, Veeam Agent shows session status and percentage of session completion in the **State** column of the list of sessions.

To view backup job session progress, do the following:

1. If you have started the backup job from the command line, launch the Veeam Agent control panel with the `veeam` command.
2. In the Veeam Agent control panel, in the list of backup job sessions, monitor progress of the currently running session.

If you have started the backup job from the Veeam Agent control panel, Veeam Agent will immediately display the list of backup job sessions with the currently running session.

TIP

You can stop the backup job session at any time. To stop the backup job session, press the [S] key.

```
Veeam Agent for Linux [ srv01 ]

Latest backup sessions:

Job name          State          Started at      Finished at
-----
DailyBackup       Running (93%)  2017-12-10 20:23:25  ---

Enter Show  C Configure  S Start Job  R Recover Files  M Misc  Esc Quit
```

Viewing Real-Time Job Session Statistics

You can view real-time statistics for a job session in the Veeam Agent control panel. Veeam Agent shows detailed data for every backup job session: job progress, duration, processing rate, performance bottlenecks, amount of processed data, read and transferred data and details of the session performance, for example, warnings and errors that have occurred in the process of operation.

To view detailed information on the currently running backup job session, do the following:

1. If you have started the backup job from the command line, launch the Veeam Agent control panel with the `veeam` command.
2. In the Veeam Agent control panel, in the list of backup job sessions, select the currently running session with the [Up] and [Down] keys and press [Enter].

If you have started the backup job from the Veeam Agent control panel, the current session will be already selected in the list of backup job sessions.

TIP

You can stop the backup job session at any time. To stop the backup job session, press the [S] key.

Statistics Counters

Veeam Agent for Linux displays jobs statistics for the following counters:

- The pane at the top of the control panel shows information on the job session type, percentage of the job completion and session status. If Veeam Agent operates in the Server edition and you have created more than one backup job, the job name also appears on the pane.
- The **Summary** box shows general information about the job:
 - **Duration** – time from the job start till the job end.
 - **Processing rate** – average speed of data processing. This counter is a ratio between the amount of processed data (**Processed** counter) and job duration (**Duration** counter).
 - **Bottleneck** – bottleneck in the data transmission process.
- The **Data** box shows information about processed data:
 - **Processed** – total size of all volumes processed by the job.
 - **Read** – amount of data read from the backed-up volume by Veeam Agent for Linux prior to applying compression. For incremental job runs, the value of this counter is typically lower than the value of the **Processed** counter. Veeam Agent reads only data blocks that have changed since the last job session, processes and copies these data blocks to the target location.
 - **Transferred** – amount of data transferred from the backed-up volume to the backup location after applying compression. This counter does not directly indicate the size of the resulting files. Depending on the backup infrastructure and job settings, Veeam Agent can perform additional activities with data, for example, decompress data prior to writing the file to disk. The activities can impact the size of the resulting file.
- The box in the center of the control panel shows a list of operations performed during the job session, their start time and duration time. To scroll the list of operations, use **Up** and **Down** arrow keys on the keyboard.

- The pane at the lower side of the control panel shows help information on how to navigate the control panel.

Veeam Agent for Linux [srv01]

Backup [SystemBackup]
56%
Status: Running

Summary

Duration: 00:01:11

Processing rate: 42.1 MB/s

Bottleneck: Agent

Data

Processed: 2.8 GB (56%)

Read: 2.8 GB

Transferred: 1.3 GB (2.1x)

Time	Action	Duration
14:27:14	Job SystemBackup started at 2016-12-07 14:27:14 MSK	
14:27:14	Preparing to backup	
14:27:15	Creating volume snapshot	00:00:00
14:27:15	Starting full backup to Repository_1	
14:27:15	Backing up BIOS bootloader on /dev/sda	00:00:01
14:27:17	Backing up sda 2.8 GB at 43.3 MB/s (56%)	00:01:06

S Stop
Esc Back

Viewing Job Session Result

You can view detailed statistics on every backup job session performed by Veeam Agent for Linux.

To view statistics for a specific job session:

1. Open the Veeam Agent control panel with one of the following commands:

```
veeam
```

or

```
veeamconfig ui
```


or

```
veeamconfig session ui
```

2. In the **Latest backup sessions** list, select the necessary backup job session with the [Up] and [Down] keys and press [Enter].

To return to the list of backup job sessions, press [Esc]. You can then select another backup job session or exit the Veeam Agent control panel in one of the following ways:

- with the [Esc] key – if you opened the control panel with the `veeam` or `veeamconfig ui` command.
- with the [Q] key – if you opened the control panel with the `veeamconfig session ui` command.



Viewing Session Status

You can view status of every session that was started by Veeam Agent for Linux. To view the session status, use the following command:

```
veeamconfig session info --id <session_id>
```

where:

<session_id> – ID of the session for which you want to check status.

Veeam Agent displays the following information about sessions:

Parameter	Description
ID	ID of the session.
Job name	Name of the backup job parent to the session. Veeam Agent displays value for this parameter only for backup job sessions.
Job ID	ID of the backup job parent to the session. Veeam Agent displays value for this parameter only for backup job sessions.
State	Current status of the session.
Start time	Date and time of the session start.
End time	Date and time of the session completion. Veeam Agent displays value for this parameter only for completed sessions.

The following example shows status information on the completed backup job session:

```
user@srv01:~$ veeamconfig session info --id 1592755d-3a2b-40a9-a036-5c81853b369e
Backup session
  ID: {1592755d-3a2b-40a9-a036-5c81853b369e}
  Job name: SystemBackup
  Job ID: {2495911e-58db-4452-b4d1-f53dcfbc600e}
  State: Success
  Start time: 2023-11-11 14:37:21 UTC
  End time: 2023-11-11 14:40:02 UTC
```

The following example shows status information on the running volume restore session:

```
user@srv01:~$ veeamconfig session info --id 697d9348-9001-4845-8764-3cc4fb3f296
b
Restore session
  ID: {697d9348-9001-4845-8764-3cc4fb3f296b}
  State: Running
  Start time: 2023-11-27 10:35:47 UTC
  End time:
```

Viewing Session Logs

You can monitor the backup and restore process by viewing the backup job session and restore session logs in the Veeam Agent command line interface.

To view Veeam Agent session log, use the following command:

```
veeamconfig session log --id <session_id>
```

where:

<session_id> – ID of the backup job or restore session.

For example:

```
user@srv01:~$ veeamconfig session log --id 0b72ef45-4c88-4639-b940-ad3828b1cd4e
2023-11-27 11:04:04 UTC [info] Job started at 2023-11-27 11:04:04
2023-11-27 11:04:04 UTC [info] Starting volume restore
2023-11-27 11:04:07 UTC [processing] sdb
2023-11-27 11:04:15 UTC [info] sdb 512.0kB at 58.6kB/s (0%)
2023-11-27 11:04:25 UTC [info] sdb 125.0MB at 6.7MB/s (0%)
2023-11-27 11:04:35 UTC [info] sdb 238.5MB at 8.3MB/s (1%)
...
2023-11-27 11:14:32 UTC [info] sdb 6.5GB at 10.7MB/s (92%)
2023-11-27 11:14:35 UTC [info] sdb 6.5GB at 10.6MB/s (97%)
2023-11-27 11:14:37 UTC [info] sdb 6.5GB at 10.6MB/s (100%)
2023-11-27 11:14:37 UTC [warn] /dev/sdb has a duplicate filesystem UUID
2023-11-27 11:14:37 UTC [info] sdb restored 6.5GB at 10.6MB/s
2023-11-27 11:14:37 UTC [info] Volume restore completed
```

Exporting Product Logs

Veeam Agent offers a simple and convenient way to collect product logs and export them to an archive file. This operation may be required if you want to report an issue and need to attach log files to the support case.

When you export logs, Veeam Agent collects its log files and configuration files, exports them to an archive file in the `tar.gz` format and saves this archive file to a directory on the Veeam Agent computer.

You can perform the export logs operation in one of the following ways:

- [With the Veeam Agent control panel](#) — in this case, you can specify a directory to which Veeam Agent should save the log archive.
- [With the command line interface](#) — in this case, Veeam Agent will save the log archive to the current working directory.

TIP

When you perform restore operations after booting from the Veeam Recovery Media, Veeam Agent also saves restore logs to the backup location. Restore logs are saved to an archive file with the name `veeam_logs_<date>_<time>.tar.gz`. The archive is placed to the folder that contains the backup file from which you restored data.


If you encounter problems after restoring from the Veeam Recovery Media, it is recommended that you attach restore logs, as well as product logs collected by Veeam Agent, to the support case.

Exporting Logs with Control Panel

You can use the Veeam Agent control panel to collect and export product logs. When you export logs with the control panel, you can choose where Veeam Agent should save the resulting log archive.


To export logs:

1. Launch the Veeam Agent control panel with the `veeam` or `veeamconfig ui` command.
2. In the Veeam Agent control panel, press the [M] key to open the **Miscellaneous** menu.
3. In the menu, select the **Export Logs** option and press [Enter].



4. In the **Choose logs directory** window, specify a directory to which you want to save the log archive:
 - a. In the **Choose logs directory** window, select the necessary directory and press [Enter].
 - b. Repeat the step 'a' until a path to the directory in which you want to save exported logs appears in the **Current directory** field.
 - c. To create a new directory, switch to the **Create Dir** button, press [Enter], then type a name for the new directory and press [Enter].

- d. Switch to the **Ok** button and press [Enter]. Veeam Agent will collect logs, export them to an archive file with the name `veeam_logs_<date>_<time>.tar.gz`, and save the archive to the specified directory.



Exporting Logs with Command Line Interface

You can use the Veeam Agent command line interface to collect and export product logs. To export logs, use the following command:

```
veeamconfig grablogs
```

Veeam Agent will collect logs, export them to an archive file with the name `veeam_logs_<date>_<time>.tar.gz`, and save the archive to the current working directory.

For example:

```
user@srv01:~$ veeamconfig grablogs
Logs have been exported successfully.
```

Getting Support

If you have any questions or want to share your feedback about Veeam Agent, you can use one of the following options:

- You can search for the information on the necessary subject in the current Veeam Agent for Linux User Guide.
- You can visit [Veeam R&D Forums](#) and share your opinion or ask a question.
- If you use Veeam Agent with an active license installed, you can visit [Veeam Customer Support Portal](#) and submit a support case to the Veeam Customer Support Team.

Using with Veeam Backup & Replication

If you have the Veeam backup infrastructure deployed in the production environment, you can use Veeam Agent together with Veeam Backup & Replication.

IMPORTANT

If you plan to use Veeam Agent for Linux 6.3.2 with Veeam Backup & Replication, you must install Veeam Backup & Replication 12.3 or later on the Veeam backup server.

For more information on managing connection to a Veeam backup server, see [Managing Veeam Backup & Replication Servers](#).

NOTE

This and subsequent sections describe tasks with Veeam Backup & Replication available for Veeam Agent operating in the standalone mode. For information about tasks available in Veeam Backup & Replication within the Veeam Agent management scenario, see the [Veeam Agent Management Guide](#).

Tasks with Veeam Backup & Replication

Veeam Backup & Replication lets you perform a number of additional data protection and disaster recovery tasks, as well as administrative actions with Veeam Agent backups. You can:

- [Grant access permissions on backup repositories](#).
- [Manage Veeam Agent licenses](#).

Data protection tasks

- [Create Veeam Agent backups on backup repositories](#).
- [Create Veeam Agent backups on Veeam Cloud Connect repositories](#).
- [Copy Veeam Agent backups to secondary backup repositories](#).
- [Archive Veeam Agent backups to tape](#).

Restore tasks

- [Restore Veeam Agent backups to Hyper-V VMs](#).
- [Restore Veeam Agent backups to VMware vSphere VMs](#).
- [Restore Veeam Agent backups to Nutanix VMs](#).
- [Restore files and folders from Veeam Agent backups](#).
- [Restore application items from Veeam Agent backups](#).
- [Restore disks from Veeam Agent backups](#).
- [Publish disks to analyze backup content](#).
- [Restore data from Veeam Agent backups to Amazon EC2](#).
- [Restore data from Veeam Agent backups to Microsoft Azure](#).

- [Restore data from Veeam Agent backups to Google Compute Engine.](#)
- [Export restore points of Veeam Agent backups to standalone full backup files.](#)

Administrative tasks

- [Import Veeam Agent backups.](#)
- [Enable and disable Veeam Agent backup jobs.](#)
- [Delete Veeam Agent backup jobs.](#)
- [View Veeam Agent backup properties.](#)
- [Create recovery tokens.](#)
- [Remove Veeam Agent backups.](#)
- [Delete Veeam Agent backups.](#)
- [Configure global settings.](#)
- [Assign roles to users.](#)

Setting Up User Permissions on Backup Repositories

To be able to store backups in a backup repository managed by a Veeam backup server, the user must have access permissions on this backup repository.

IMPORTANT

Veeam Agent for Linux does not support Veeam backup repositories with enabled KMS encryption. To learn more about KMS encryption for Veeam backup repositories, see the [Key Management System Keys](#) section in the Veeam Backup & Replication User Guide.

NOTE

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.

Access permissions are granted to security principals such as users and AD groups by the backup administrator working with Veeam Backup & Replication. Users with granted access permissions can target Veeam Agent backup jobs at this backup repository and perform restore from backups located in this backup repository.


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.


To grant access permissions to a security principal:

1. In Veeam Backup & Replication, 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 to Veeam Agent users.
 - The **Scale-out Repositories** node – if you want to grant access permissions on a scale-out backup repository to Veeam Agent users.

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 **Access Permissions** window, in the **Standalone applications** tab, 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 on this backup repository. Setting access permissions to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). However, we recommend this scenario 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 on this backup repository. Click **Add** to add the necessary users and groups to the list.
5. If you want to encrypt Veeam Agent backup files stored in the backup repository, select the **Encrypt backups stored in this repository** check box and choose the necessary password from the field below. If you have not specified a password beforehand, click **Add** on the right or the **Manage passwords** link to add a new password. Veeam Backup & Replication will encrypt files at the backup repository side using its built-in encryption mechanism. To learn more, see [Veeam Backup & Replication Documentation](#).



Managing License

If you plan to use Veeam Agent with Veeam Backup & Replication, you must install a license in Veeam Backup & Replication or Veeam Backup Enterprise Manager. The license must have a total number of instances that is sufficient to protect machines (servers and workstations) on which you plan to install Veeam Agent. For more information, see [Veeam Licensing Policy](#).

After Veeam Agent connects to Veeam Backup & Replication, Veeam Agent automatically starts consuming instances in the license. The product edition for Veeam Agent is selected depending on the OS running on the Veeam Agent computer. You can switch to another commercial edition of Veeam Agent manually if needed. If you do not want Veeam Agents to consume instances, you can restrict instance consumption. For more information, see [Managing Instance Consumption by Veeam Agents](#).

The number of backup jobs configured in Veeam Agent does not impact instance consumption. For example, if 2 backup jobs are configured in Veeam Agent that operates in the Server edition, this Veeam Agent will consume instances required for 1 server.

Veeam Agent obtains information about the license from Veeam Backup & Replication and keeps it locally on the Veeam Agent computer. Information about the license is valid for 32 days. If Veeam Agent does not connect to Veeam Backup & Replication during this period, Veeam Backup & Replication will revoke its license.

NOTE

In addition to managing Veeam Agent licenses, you can use the Veeam Backup & Replication console to manage Veeam Agent backup jobs and perform operations with backups created by these jobs.

If your backup server is connected to Veeam Backup Enterprise Manager, you can use Veeam Backup Enterprise Manager to manage licenses and perform restore tasks with Veeam Agent backups. You cannot manage Veeam Agent backup jobs with Veeam Backup Enterprise Manager.

For more information on Veeam Backup & Replication licensing, see the [Licensing](#) section in the Veeam Backup & Replication User Guide.


Managing Instance Consumption by Veeam Agents

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, you can restrict instance consumption.

If you restrict instance consumption, Veeam Backup & Replication will switch all Veeam Agents connected to this Veeam backup server to the free edition that offers limited capabilities. For information about Veeam Agent editions, see [Product Editions](#).

To restrict instance consumption by Veeam Agents:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, click the **Instances** tab.
3. On the **Instances** tab, clear the **Allow unlicensed agents to consume instances** check box.
4. Click **Close**.




Assigning License to Veeam Agent

After Veeam Agent connects to Veeam Backup & Replication, Veeam Agent automatically starts consuming the license. The product edition for Veeam Agent is selected depending on the type of the OS running on the protected computer.

You can also assign a license to Veeam Agent manually if needed. When you assign a license, you can select the product edition, too.

To assign a license:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.
3. In the **Licensed Instances** window, select the Veeam Agent to which you want to assign the license, click **Assign** and select the desired product edition: *Workstation* or *Server*.



Viewing Licensed Veeam Agents and Revoking License

When Veeam Agent connects to the backup server, Veeam Backup & Replication applies a license to the Veeam Agent. You can view to which Veeam Agents the license is currently applied.

To view a list of licensed Veeam Agents:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.


In the list of licensed instances, Veeam Backup & Replication displays Veeam Agents that have established a connection with the backup server when you created the backup job.

Revoking License from Veeam Agents


You can revoke the license from some Veeam Agents and re-apply it to other protected workloads. License revoking can be helpful, for example, if you do not want to use some Veeam Agents with Veeam Backup & Replication anymore.

To revoke a license from the Veeam Agent:

1. In Veeam Backup & Replication, from the main menu, select **License**.
2. In the **License Information** window, select the **Instances** tab and click **Manage**.
3. In the Licensed Instances window, select a Veeam Agent and click **Revoke**. Veeam Backup & Replication will revoke the license from the Veeam Agent, and the license will be freed for other workloads that you want to protect with Veeam products.



The Veeam Agent from which you have revoked the license will become unable to connect to the Veeam backup server but will remain in the **Licensed Instances** list. To allow this Veeam Agent to create backups in the Veeam backup repository, select the Veeam Agent and click **Remove**. During the next backup job session, the Veeam Agent will connect to the Veeam backup server and start consuming the license.



Performing Data Protection Tasks

You can perform the following data protection tasks:

- Back up your data and store the resulting backup files in one of the following types of Veeam backup repositories:
 - [In a backup repository managed by a Veeam backup server](#)
 - [In a Veeam Cloud Connect repository](#)
- [Copy Veeam Agent backups from the backup repository to a secondary backup repository with backup copy jobs.](#)
- [Use SureBackup.](#)
- [Archive Veeam Agent backups to tapes with backup to tape jobs.](#)

Backing Up to Backup Repositories

You can store backups created with Veeam Agent in backup repositories connected to Veeam backup servers. To do this, you must perform the following actions:

1. [Set up user permissions at the backup repository side.](#)
2. [Point the Veeam Agent backup job to the backup repository.](#)


NOTE

Consider the following:

- A Veeam Agent backup job can be started automatically upon the defined schedule or manually from the Veeam Agent computer. You cannot start, stop, retry or edit Veeam Agent backup jobs in the Veeam Backup & Replication console.
- If the user is granted restore permissions on the Veeam backup server, the user will be able to see all backups in the backup repository.
- The user who creates a Veeam Agent backup in the backup repository is set as the owner of the backup file. The backup file owner can access this file and restore data from it. If the user who is not the backup file owner needs to perform operations with the backup file, the user must have the Veeam Backup & Replication role that allows to perform these operations. To learn more about roles, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.

Backup jobs targeted at the backup repository become visible in Veeam Backup & Replication under the **Jobs > Backup** node in the **Home** view. Backups created with Veeam Agent are available under the **Backups > Disk** node in the **Home** view.

The Veeam Backup Administrator working with Veeam Backup & Replication can manage Veeam Agent backup jobs and restore data from Veeam Agent backups. To learn more, see [Restoring Data from Veeam Agent Backups](#) and [Performing Administration Tasks](#).



Backing Up to Cloud Repositories


You can store backups created with Veeam Agent in cloud repositories provided to you by a Veeam Cloud Connect service provider. To do this, you must connect to the service provider and point the backup job to the cloud repository. To learn more, see [Specify Service Provider Settings](#).

Veeam Agent Backups on Tenant Side

Backups created with Veeam Agent are available under the **Cloud** node in the **Home** view of the Veeam Backup & Replication console deployed on the tenant side.

The backup administrator working with Veeam Backup & Replication on the tenant side can manage Veeam Agent backups created in the cloud repository and restore data from such backups. To recover data from a Veeam Agent backup, you can perform the following operations:

- [Export computer disks as virtual disks](#).
- [Restore guest OS files](#).
- [Export restore points to standalone full backup files](#).



Veeam Agent Backups on Service Provider Side

The service provider can view information about backup and restore sessions performed by Veeam Agent users. The full list of sessions is available in the **History** view of the Veeam backup console deployed on the service provider side. The list of sessions performed within the last 24 hours is available under the **Last 24 hours** node in the **Cloud Connect** view of the Veeam backup console on the service provider side. The service provider cannot view detailed statistics about individual sessions in the list.

The service provider cannot perform restore tasks with Veeam Agent backups that are stored in the cloud repository. The service provider can perform the following restore tasks with unencrypted Veeam Agent backups stored in the cloud repository:

- Instant recovery
- Restore to public cloud repositories
- Disk restore
- Disk publish

To learn more, see the [Restoring Data from Tenant Backups](#) section in the Veeam Cloud Connect Guide.

The screenshot shows the Veeam Backup and Replication console. The top navigation bar includes 'Home' and 'View' tabs. Below the navigation bar is a toolbar with icons for Backup Job, Replication Job, CDP Policy, Tape Job, Backup Copy, SureBackup Job, Restore, Failover Plan, Import Backup, Export Backup, and Security & Compliance. The main area is titled 'Cloud Connect' and features a search bar. On the left, a sidebar shows a tree view with 'Cloud Connect' expanded, containing 'Cloud Gateways', 'Gateway Pools', 'Tenants', 'Backup Storage', 'Replica Resources', and 'Last 24 Hours'. The main pane displays a table of backup jobs.

Job Name	Session Type	Status	Start Time	End Time	Tenant	Data Sent	Data Received
srv001_daily_backup	Cloud Backup	Success	11/30/2023 11:24 AM	11/30/2023 11:25 AM	ABC Company	54.1 KB	655 MB
srv001_volume_backup	Cloud Backup	Success	11/30/2023 11:26 AM	11/30/2023 11:27 AM	ABC Company	15.6 KB	15.8 KB
srv002_system_backup	Cloud Backup	Success	11/30/2023 11:25 AM	11/30/2023 11:26 AM	ABC Company	97.8 KB	2.2 MB
srv003_system_backup	Cloud Backup	Success	11/30/2023 11:19 AM	11/30/2023 11:20 AM	ABC Company	40.6 KB	20.5 MB


At the bottom left, there is a status bar showing '5 sessions'.

Performing Backup Copy for Veeam Agent Backups

You can configure backup copy jobs that will copy backups created with Veeam Agent to a secondary backup repository.

Backup copy jobs treat Veeam Agent backups as usual backup files. The backup copy job setup and processing procedures practically do not differ from the same procedures for a backup copy job that processes VM backups. To learn more about backup copy jobs, see the [Backup Copy](#) section in the Veeam Backup & Replication User Guide.

When mapping a backup copy job to a Veeam Agent backup, consider the limitations listed in the [Map Backup File](#) section in the Veeam Backup & Replication User Guide.



Restoring Data from Copies of Veeam Agent Backups

Backups copied to the secondary backup repository do not preserve user access permissions. At the same time, users who created backups do not have access permissions on these secondary repositories. For this reason, users cannot restore data from their backups residing in the secondary site.

To overcome this limitation, you can delegate the restore task to backup administrators who work with Veeam Backup & Replication. Backup administrators can use Veeam Backup & Replication options to recover data from such backups: for example, perform file-level restore or retrieve necessary application items with Veeam Explorers.

You can also restore data from the copied backup stored in the target repository using Veeam Agent.

Using SureBackup

Veeam Backup & Replication offers the SureBackup technology to test backups and check if you can recover data from them. You can verify any restore point of a backed-up computer protected with Veeam Agent for Linux.

To learn more about the logic behind SureBackup, see the [How SureBackup Works](#) section in the Veeam Backup & Replication User Guide.

Before creating the SureBackup job, check limitations for Veeam Agent backups below. Then launch the **New SureBackup Job** wizard to create the SureBackup job. To learn more, see the [Creating SureBackup Job](#) section in the Veeam Backup & Replication User Guide.

Limitations

For backups created with Veeam Agent for Linux, SureBackup has the following limitations:

- SureBackup job in the **Backup verification and content scan only** mode is not supported.
- SureBackup is not supported for backups stored in the Veeam Cloud Connect repository.
- SureBackup is not supported for backups stored in the archive tier of the the scale-out backup repository.
- For backups created with Veeam Agents, you cannot exclude objects from the scope of the SureBackup job at the **Linked Jobs** step of the **New SureBackup Job** wizard.
- SureBackup is not supported for backups containing drives greater than 64 TB.
- If you plan to verify computer recovery with VMware vSphere, consider the following:
 - SureBackup is not supported for backups of 4 KB sector drives.
 - SureBackup is not supported for backups of storage spaces.
 - SureBackup is not supported for backups containing more than 54 drives.
- When Veeam Backup & Replication publishes virtual machines based on backed-up Veeam Agent computers in the isolated virtual environment, all these virtual machines are included in the first isolated network added during the virtual lab configuration. To learn more, see the [Create Isolated Networks](#) section in the Veeam Backup & Replication User Guide.
- You cannot use SureBackup with backup files created with Veeam Agent for Linux on Power.
- The successful recovery verification is not guaranteed for the following Linux distributions:
 - Amazon Linux 2
 - Amazon Linux 2023
 - openSUSE Tumbleweed
- The successful recovery verification is not guaranteed for backups of Linux-based systems that contain encrypted devices.

- If you want Veeam Backup & Replication to connect the recovered VM to the virtual network, one of the following configuration utilities must be installed on the protected computer:
 - Netplan
 - NetworkManager
 - sysconfig
 - systemd-networkd
 - ifupdown/ifupdown2
- SureBackup is not supported for file-level backups. You must use volume-level backup of the protected computer. The backup must include the `root` file system (`/`) and all partitions specified in the `/etc/fstab` file. To learn more about backup types, see [Backup Types](#).

Archiving Veeam Agent Backups to Tape

You can configure backup to tape jobs to archive Veeam Agent backups to tape.

Backup to tape jobs treat Veeam Agent backups as usual backup files. The archiving job setup and processing procedures practically do not differ from the regular ones. To learn more about backup to tape jobs, see the [Backup to Tape](#) section in the Veeam Backup & Replication User Guide.

NOTE

For the **After this job** option in the backup to tape job schedule settings, you cannot select a backup job managed by Veeam Agent or a standalone Veeam Agent backup job as the preceding backup job.

New Backup to Tape Job



Backups
Specify objects to be processed by this tape job.

Name

Backups

Media Pool

Incremental Backup

Options

Schedule

Summary

Backups:

Name	Type	Size	
Daily Backup	Linux Agent Ba...	3.17 GB	
System Backup	Linux Agent Ba...	333 GB	

Add...

Remove

Up

Down

Full:
335 GB
Incremental:
1.58 GB

Restoring Data from Veeam Agent Backups

You can perform the following restore operations:

- [Restore Veeam Agent backups to VMware vSphere VMs](#)
- [Restore Veeam Agent backups to Hyper-V VMs](#)
- [Restore Veeam Agent backups to Nutanix AHV VMs](#)
- [Restore Veeam Agent backups to Proxmox VE VMs](#)
- [Restore Veeam Agent backups to oVirt KVM VM](#)
- [Restore disks from Veeam Agent backups to oVirt KVM VM](#)
- [Restore data from Veeam Agent backups to Microsoft Azure](#)
- [Restore data from Veeam Agent backups to Amazon EC2](#)
- [Restore data from Veeam Agent backups to Google Compute Engine](#)
- [Restore individual files and folders from Veeam Agent backups](#)
- [Restore application items from Veeam Agent backups with Veeam Explorers](#)
- [Export computer disks as VMDK, VHD or VHDX disks](#)
- [Publish disks to analyze backup content](#)
- [Export restore points of Veeam Agent backups to standalone full backup files](#)

Restoring Veeam Agent Backup to vSphere VM

In the Veeam Backup & Replication console, you can use Instant Recovery to restore a Veeam Agent computer as a VMware vSphere VM in your virtualization environment.

A restored VMware vSphere VM will have the same settings as the backed-up Veeam Agent computer. During the restore process, Veeam Backup & Replication retrieves the settings of the Veeam Agent computer from the backup and applies them to the target VM. These settings include:

- Amount of RAM.
- Number of CPU cores.
- Number of network adapters.
- Network adapter settings.
- BIOS UUID.

If you do not want to preserve the backed-up machine UUID for a VMware vSphere VM, you can create a new UUID during the Instant Recovery configuration process.

- Number of disks and volumes.
- Size of volumes.

Considerations and Limitations

If you restore a Veeam Agent computer to a VMware vSphere VM, consider the following:


- You can use entire machine or volume-level backups of Linux computers. Volume-level backups must include the `root` file system (`/`) and all partitions specified in the `/etc/fstab` file.
- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- Make sure that the target host has enough resources for a new VM. Otherwise, your VM will reduce the target host performance.
- If you restore a workload to the production network, make sure that the original workload is powered off.
- If the disk you want to restore contains a Logical Volume Manager (LVM) volume group, consider the following:
 - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Backup & Replication will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
 - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Backup & Replication will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

TIP

After restore, you can remove unnecessary disks from the machine. To learn more, see [this Veeam KB article](#).

Restore to vSphere VM

The procedure of Instant Recovery for a Veeam Agent computer practically does not differ from the same procedure for a VM. The main difference from Instant Recovery is that you do not need to select the recovery mode, because Veeam Agent computers are always restored to a new location. To learn more, see the [Performing Instant Recovery of Workloads to VMware vSphere](#) section in the Veeam Backup & Replication User



Guide.

Restoring Veeam Agent Backup to Hyper-V VM

In the Veeam Backup & Replication console, you can use Instant Recovery to restore a Veeam Agent computer as a Hyper-V VM in your virtualization environment.

A restored Hyper-V VM will have the same settings as the backed-up Veeam Agent computer. During the restore process, Veeam Backup & Replication retrieves settings of the Veeam Agent computer from the backup and applies them to the target VM.

Considerations and Limitations

If you restore a Veeam Agent computer to a Hyper-V VM, consider the following:


- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot use backups stored in a Veeam Cloud Connect repository for this operation.
- To restore to a Hyper-V VM from a backup of a Linux computer, you must consider the Hyper-V limitations. To learn more, see [this Microsoft article](#).
-
- Make sure that the target host has enough resources for a new VM. Otherwise, your VM will reduce the target host performance.
- Veeam Agent computer disks are recovered as dynamically expanding virtual disks.
- By default, Veeam Backup & Replication automatically powers on a VM after restore. If you do not want to power on a VM after restore, you can change this setting during the Instant Recovery configuration process.
- If the disk you want to restore contains a Logical Volume Manager (LVM) volume group, consider the following:
 - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Backup & Replication will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
 - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Backup & Replication will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

TIP

After restore, you can remove unnecessary disks from the machine. To learn more, see [this Veeam KB article](#).

Restore to Hyper-V VM

The procedure of Instant Recovery for a Veeam Agent computer practically does not differ from the same procedure for a VM. The main difference from Instant Recovery is that you do not need to select the recovery mode, because Veeam Agent computers are always restored to a new location. To learn more, see the [Performing Instant Recovery of Workloads to Hyper-V](#) section in the Veeam Backup & Replication User Guide.



Restoring Veeam Agent Backup to Nutanix VM

You can use the Veeam Backup & Replication console to restore a Veeam Agent computer as a Nutanix AHV VM in your virtualization environment.

Considerations and Limitations

If you restore a Veeam Agent computer to a Nutanix AHV VM, consider the following:


- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- If the disk you want to restore contains a Logical Volume Manager (LVM) volume group, consider the following:
 - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
 - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

TIP

After restore, you can remove unnecessary disks from the machine. To learn more, see [this Veeam KB article](#).

Restore to Nutanix AHV

The procedure of restore to Nutanix AHV for a Veeam Agent computer practically does not differ from the same procedure for a VM. To learn more about restore to Nutanix AHV, see the [Restoring VMs Using Veeam Backup & Replication Console](#) section in the Veeam Backup for Nutanix AHV User Guide.



Restoring Veeam Agent Backup to Proxmox VM

You can use the Veeam Backup & Replication console to restore a Veeam Agent computer as a Proxmox VE VM in your virtualization environment.

Considerations and Limitations

If you restore a Veeam Agent computer to a Proxmox VE VM, consider the following:


- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- If the disk you want to restore contains a Logical Volume Manager (LVM) volume group, consider the following:
 - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
 - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

TIP

After restore, you can remove unnecessary disks from the machine. To learn more, see [this Veeam KB article](#).

Restore to Proxmox VE

The procedure of restore to Proxmox VE for a Veeam Agent computer practically does not differ from the same procedure for a VM. To learn more about restore to Proxmox VE, see the [Performing VM Restore](#) section in the Veeam Backup for Proxmox VE User Guide.



Restoring Veeam Agent Backup to oVirt KVM VM

You can use the Veeam Backup & Replication console to restore a Veeam Agent computer as an oVirt KVM VM in your virtualization environment.

Considerations and Limitations

If you restore a Veeam Agent computer to an oVirt KVM VM, consider the following:


- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- If the disk you want to restore contains a Logical Volume Manager (LVM) volume group, consider the following:
 - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
 - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

TIP

After restore, you can remove unnecessary disks from the machine. To learn more, see [this Veeam KB article](#).

Restore to oVirt KVM VM

The procedure of restore to an oVirt KVM VM for a Veeam Agent computer practically does not differ from the same procedure for a VM. To learn more, see the [Performing VM Restore](#) section in the Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization User Guide.



Restoring Disk from Veeam Agent Backup to oVirt KVM

You can use the Veeam Backup & Replication console to restore disks from a Veeam Agent computer backup to an oVirt KVM VM in your virtualization environment.

Considerations and Limitations

If you restore disks to an oVirt KVM VM, consider the following:


- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- If the disk you want to restore contains a Logical Volume Manager (LVM) volume group, consider the following:
 - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
 - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

TIP

After restore, you can remove unnecessary disks from the machine. To learn more, see [this Veeam KB article](#).

Restore to oVirt KVM VM

The procedure of restoring disks to an oVirt KVM VM for a Veeam Agent computer practically does not differ from the same procedure for a VM. To learn more, see the [Performing Disk Restore](#) section in the Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization User Guide.



Restoring to Microsoft Azure

You can use the Veeam Backup & Replication console to restore computers from Veeam Agent backups to Microsoft Azure.

Considerations and Limitations


If you restore a Veeam Agent computer to Microsoft Azure, consider the following:

- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- Veeam Agent backups must be created at the entire computer level or volume level.
- If you recover an EFI-based system to Microsoft Azure, Veeam Agent will restore a BIOS-based Generation 1 VM.
- Veeam Backup & Replication offers experimental support for generation 2 VMs within restore to Microsoft Azure feature. To learn more, see the [Generation 2 VM Support](#) section in the Veeam Backup & Replication User Guide.

After restore, you can remove unnecessary disks from the machine. To learn more, see [this Veeam KB article](#).

Restore to Microsoft Azure

The procedure of restore to Microsoft Azure from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about restore to Microsoft Azure, see the [Restoring to Microsoft Azure](#) section in the Veeam Backup & Replication User Guide.



Restoring to Amazon EC2

You can use the Veeam Backup & Replication console to restore computers from Veeam Agent backups to Amazon EC2.

Considerations and Limitations

If you restore a Veeam Agent computer to Amazon EC2, consider the following:


- You can use backups of Linux computers stored in a Veeam backup repository only. You cannot perform this operation with Veeam Agent backups stored in a Veeam Cloud Connect repository.
- Veeam Agent backups must be created at the entire computer level or volume level. If the disk you want to restore contains a Logical Volume Manager (LVM) volume group, consider the following:
 - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
 - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

TIP

After restore, you can remove unnecessary disks from the machine. To learn more, see [this Veeam KB article](#).

Restore to Amazon EC2

The procedure of restore to Amazon EC2 from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about restore to Amazon EC2, see the [Restoring to Amazon EC2](#) section in the Veeam Backup & Replication User Guide.



Restoring to Google Compute Engine

You can use the Veeam Backup & Replication console to restore computers from Veeam Agent backups to Google Compute Engine.

Considerations and Limitations

If you restore a Veeam Agent computer to Google Compute Engine, consider the following:


- You can use backups of Linux computers stored in a Veeam backup repository. You cannot perform this operation with Veeam Agent backups created on the Veeam Cloud Connect repository.
- Veeam Agent backups must be created at the entire computer level or volume level. If the disk you want to restore contains a Logical Volume Manager (LVM) volume group, consider the following:
 - Since LVM volume group is a logical entity that spans across the physical disks, Veeam Agent treats the original disk and the LVM volume group as separate entities. Therefore, Veeam Agent will restore the original disk and the LVM volume group as 2 separate disks. This way, all data, including the data within the LVM volume group, is accurately restored.
 - Root file system partition and boot partition must not be on LVM logical volumes. For more information on this limitation, see [Google documentation](#).
 - Restoring the original disk and the LVM volume groups as 2 separate disks requires an increased amount of storage space. For example, you restore a machine with 2 disks, and a separate LVM volume group is configured on each of these disks. In this case, Veeam Agent will restore 4 disks. The restored disks will consume the storage space equal to the size of the 2 original disks and the 2 LVM volume groups from these disks.

TIP

After restore, you can remove unnecessary disks from the machine. To learn more, see [this Veeam KB article](#).

Restore to Google Compute Engine

The procedure of restore to Google Compute Engine from a Veeam Agent backup practically does not differ from the same procedure for a VM backup. To learn more about restore to Google Compute Engine, see the [Restoring to Google Compute Engine](#) section in the Veeam Backup & Replication User Guide.



Restoring Files and Folders

You can use the Veeam Backup & Replication console to restore individual files and folders from Veeam Agent backups.

The procedure of file-level restore from a Veeam Agent backup is similar to the same procedure for a VM backup. To learn more about file-level restore, see the [Restore from Linux, Unix and Other File Systems](#) section in the Veeam Backup & Replication User Guide.

When you perform the file-level restore procedure, Veeam Backup & Replication provides the following options for mounting disks of the machine from the backup or replica:

- Mounting disks to the original host.

IMPORTANT

Veeam backup server must be able to resolve the domain name of the original host into IP address.

- Mounting disks to a helper host — any Linux host in your infrastructure with a [supported operating system](#).
- Mounting disks to a temporary helper appliance — a helper VM required to mount Linux computer disks from the backup.


If you have selected to mount disks to a temporary helper appliance, it is recommended that you add a vCenter Server and not a standalone ESXi host in the Veeam backup console. If Veeam Backup & Replication is set up to deploy a helper appliance on a standalone ESXi host, after Veeam Backup & Replication removes the helper appliance, the helper VM will be displayed in vCenter as orphaned.

Restoring Application Items

You can use Veeam Explorers to restore application items from backups created using Veeam Agent for Linux. Veeam Backup & Replication lets you restore items and objects from the following applications:

- Oracle
- PostgreSQL

The procedure of application item restore from a Veeam Agent backup does not differ from the same procedure for a VM backup. To learn more, see the [Application Item Restore](#) section in the Veeam Backup & Replication User Guide.



Exporting Disks

You can restore computer disks from Veeam Agent backups created using Veeam Agent for Linux and convert them to disks of the VMDK, VHD or VHDX format.

During disks restore, Veeam Backup & Replication creates standard virtual disks that can be used by VMware vSphere and Microsoft Hyper-V VMs.

- When you restore 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 restore a disk in the VHD/VHDX format, Veeam Backup & Replication creates a file of the VHD or VHDX format.

You can save converted disks locally on any server or SMB share added to the backup infrastructure or place disks on a datastore connected to an ESXi host (for VMDK disk format only). VMDK disks can be restored as thin provision and thick disks:

- Disks restored to a datastore are saved in the thin provisioned format.
- Disks restored to a server are saved in the thick provisioned format.

Veeam Backup & Replication supports batch disk restore. For example, if you choose to restore 2 computer disks, Veeam Backup & Replication will convert them to 2 virtual disks and store these disks in the specified location.

IMPORTANT

Consider the following:

- If the backup from which you export disks contains a Btrfs storage pool, during the disk export process Veeam Backup & Replication will create a separate disk and restore the Btrfs pool to this disk.
- If the disk you want to export contains a Logical Volume Manager (LVM) volume group, Veeam Agent will export the original disk and the LVM volume group as 2 separate disks. Among other things, this leads to the increase of the required storage space. For example, if you export all disks from a backup of a machine that has 2 disks and a separate LVM volume group on each disk, Veeam Agent will export 4 disks. The exported disks will consume storage space equal to the size of 2 original disks and 2 LVM volume groups from these disks.


To restore disks and convert them to the VMDK, VHD or VHDX format, perform the following steps in the **Export Disk** wizard:

Step 1. Launch Export Disk Wizard

To launch the **Export Disk** wizard, do either of the following:


- Open the **Home** tab and click **Restore > Agent > Disk restore > Export disk**. In this case, you will be able to select a backup of the necessary Veeam Agent computer at the **Backup** step of the wizard.
- Open the **Home** view. In the inventory pane, click the **Backups** node. In the working area, expand the necessary Veeam Agent backup, select the necessary computer in the backup and click **Export Disks** on the ribbon or right-click a computer in the backup and select **Export content as virtual disks**.

In this case, you will pass immediately to the **Restore Point** step of the wizard.




Step 2. Select Backup

At the **Backup** step of the wizard, select a backup from which you want to restore disks. In the list of backups, Veeam Backup & Replication displays all backups that are currently hosted on the Veeam backup repository and Veeam Cloud Connect repository.




Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select the necessary restore point from which you want to restore disks. In the list of points, Veeam Backup & Replication displays all restore points that have been created. Make sure that you select a restore point that relates to the selected backup.



Step 4. Select Disks

At the **Disks** step of the wizard, select check boxes next to those 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 disk.

1. From the **Server** list, select a server on which the resulting virtual disks must be saved. If you plan to save the disks in the VMDK format on a datastore, select an ESXi host to which this datastore is connected.
2. In the **Path to folder** field, specify a folder on the server or datastore where the virtual disks must be placed.
3. Select the export format for disks:
 - **VMDK** – select this option if you want to save the resulting virtual disk in the VMware VMDK format.
 - **VHD** – select this option if you want to save resulting virtual disk in the Microsoft Hyper-V VHD format.
 - **VHDX** – select this option if you want to save resulting virtual disk in the Microsoft Hyper-V VHDX format (supported by Microsoft Windows Server 2012 and later).
4. Click **Disk type** to specify how the resulting disk must be saved:
 - [For VMDK disk format] in the thin provisioned, lazy zeroed thick provisioned, or eagerly zeroed thick provisioned format
 - [For VHD and VHDX disk formats] in the dynamic or fixed format
5. [For export of a VMDK disk to an ESXi host] Click the **Pick proxy to use** link to select backup proxies over which backup data must be transported to the target datastore.

NOTE

Consider the following:

- If you have selected to store the resulting virtual disk in a datastore, you will be able to save the virtual disk in the VMDK format only. Other options will be disabled.
- If you have selected to store the resulting virtual disk on the server running Microsoft Windows Server OS and in the VMDK format, you will be able to save the virtual disk in the lazy zeroed thick provisioned format only.

The screenshot shows the 'Export Disk' dialog box with the 'Target' tab selected. The dialog has a sidebar on the left with options: Backup, Restore Point, Disks, Target (selected), Reason, and Summary. The main area contains the following fields and options:

- Server:** A dropdown menu showing 'winsrv0042019.tech.local'.
- Path to folder:** A text box containing 'C:\WeeamBackup' and a 'Browse...' button.
- Export format:** Three radio button options:
 - VMDK** (selected): This virtual disk type is used by VMware products such as VMware Workstation, or VMware vSphere. Maximum VMDK disk size is 62TB. Pick proxy to use.
 - VHD**: This virtual disk type is used by Microsoft products such as Microsoft Hyper-V or Microsoft Azure. Maximum VHD disk size is 2TB.
 - VHDX**: This virtual disk type is used by more recent versions of Microsoft products such as Microsoft Hyper-V. Maximum VHDX disk size is 64TB.
- Disk type:** A dropdown menu showing 'Thick (lazy)'.


At the bottom of the dialog are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Step 6. Specify Restore Reason

At the **Reason** step of the wizard, enter a reason for restoring the computer volume.

TIP

If you do not want to display the **Restore Reason** step of the wizard in future, select the **Do not show me this page again** check box.




The screenshot shows the 'Export Disk' wizard window, specifically the 'Reason' step. The window has a title bar 'Export Disk' and a close button. On the left is a sidebar with steps: Backup, Restore Point, Disks, Target, Reason (selected), and Summary. The main area is titled 'Reason' with a subtitle 'Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.' Below this is a text box labeled 'Restore reason:' containing the text 'Converting backup to virtual disk.' At the bottom left of the main area is a checkbox labeled 'Do not show me this page again' which is checked. At the bottom right are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

Step 7. Complete Restore Process

At the **Summary** step of the wizard, complete the disk restore procedure.

1. Review details for the disk to be restored.
2. Click **Finish** to start the restore procedure and exit the wizard.



Publishing Disks

You can use the Veeam backup console to publish disks from Veeam Agent backups.

TIP

You can publish disks using the PowerShell console. To learn more, see the [Disk Publishing \(Data Integration API\)](#) section in the Veeam PowerShell Reference.

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.

To publish disks from a backup of a Linux-based computer to a target server, Veeam Backup & Replication uses the FUSE protocol. After publishing, the target server can access the backup content and read the necessary data from the disk.

To learn more, see the [Disk Publishing](#) section in the Veeam Backup & Replication User Guide.

Performing Disk Publish

Before you publish disks, [check prerequisites](#). Then use the **Publish Disks** wizard.

1. [Launch the wizard](#).
2. [Select a Veeam Agent computer whose disks you want to publish](#).
3. [Select a restore point](#).
4. [Select disks](#).
5. [Specify the target server](#).
6. [Specify a reason for disk publishing](#).
7. [Finish working with the wizard](#).

Before You Begin

Before you publish disks, check the following requirements and limitations:

- The necessary ports must be opened on the target server. For more information, see [Ports](#).
- The target server must support the file system of the disk that you plan to publish.
- If data deduplication is enabled for some disks in a backup, data deduplication must be enabled on the target server.
- The 32-bit version of a Linux server is not supported as the target server.
- You cannot publish disks from backups stored in the Veeam Cloud Connect repository.


- If the disk you want to publish contains a Logical Volume Manager (LVM) volume group, Veeam Agent will publish the original disk and the LVM volume group as 2 separate disks. This may lead to the increase of the required storage space. For example, if you publish all disks from a backup of a machine that has 2 disks and a separate LVM volume group on each disk, Veeam Agent will publish 4 disks. The published disks will consume storage space equal to the size of 2 original disks and 2 LVM volume groups from these disks.

For the full list of limitations, see the [Considerations and Limitations](#) section in the Veeam Backup & Replication User Guide.

Step 1. Launch Publish Disks Wizard


To launch the **Publish Disks** wizard, do either of the following:

- On the **Home** tab, click **Restore > Agent > Disk Restore > Publish disk**.
- Open the **Home** view. In the inventory pane, click **Backups**. In the working area, expand the necessary Veeam Agent backup, select a computer whose disks you want to publish and click **Publish Disks** on the ribbon. Alternatively, you can right-click the computer and select **Publish disks**. In this case, you will proceed to the [Restore point](#) step of the wizard.



Step 2. Select Computer

At the **Machine** step of the wizard, expand a backup and select a Veeam Agent computer 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.

Machine

Restore point

Disks

Target

Reason

Summary

Computer name: **srv02 SystemBackup**

Data size: **208 GB**

Available restore points:

Created	Type	Backup
5 days ago (5:30 PM Thursda...	Increment	srv02 SystemBackup
8 days ago (6:08 PM Monday...	Full	srv02 SystemBackup

< Previous

Next >

Finish

Cancel

Step 4. Select Disks

At the **Disks** step of the wizard, select a check box next to the disks that you want to publish. Click **Select All** if you want to select all disks from the backup.

Machine

Restore point

Disks

Target

Reason

Summary

Disks:

Disk name	Size	Volumes
☐ Disk 0	160 GB	sda1
☐ Disk 1	158 GB	home; root; swap_1
☒ Disk 2	200 GB	sdb

Select All

Clear All

< Previous

Next >

Finish

Cancel

Step 5. Select Target Server

At the **Target** step of the wizard, select a Linux server that will have access to disk content.


You can select one of the following types of servers:

- A server added to the backup infrastructure.

If you want to add a new backup server to the backup infrastructure at this step, click **Add**. In this case, you will be able to add a new Linux server. To learn more, see the [Adding Linux Servers](#) section in the Veeam Backup & Replication User Guide.

- A temporary server. In this case, select *Specify a different host* from the drop-down list. In the **Target Server** window, specify the following settings:
 - a. In the **Host name** field, specify a server name or IP address of the server.
 - b. Select the account 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 a new account in the Credentials Manager. To learn more, see the [Credentials Manager](#) section in the Veeam Backup & Replication User Guide.
 - c. Click **Advanced** and customize connection settings in the **Network Settings window**. To learn more, see [Customizing Connection Settings](#).
- The original server. In this case, select *Original server* from the drop-down list.

If prompted, specify credentials for the target server.




The screenshot shows the 'Publish Disks' wizard window, specifically the 'Target' step. The window has a title bar with 'Publish Disks' and a close button. On the left, there is a sidebar with icons and labels: 'Machine', 'Restore point', 'Disks', 'Target' (highlighted), 'Reason', and 'Summary'. The main area is titled 'Target' with the instruction 'Select a target server to mount disks to.' Below this, there is a 'Target server:' label followed by a dropdown menu showing 'VBR12-01.tech.local (temporary added host)' and an 'Add...' button. At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.

Customizing Connection Settings

If necessary, you can customize connection settings for a target Linux server at the **Target** step of the **Publish Disks** wizard. To do so, click **Advanced** in the **Target Server** window and specify settings in the **Network Settings** window:

1. In the **Service console connection** section, specify an SSH timeout.
2. In the **Data transfer options** section, specify connection settings for file copy operations.
3. [For Linux server deployed outside NAT] In the **Preferred TCP connection role** section, select the **Run server on this side** check box.

To learn more about these settings, see the [Specify Credentials and SSH Settings](#) section in the Veeam Backup & Replication User Guide.




Step 6. Specify Restore 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 further will want to return this page, follow the instructions described in [this Veeam KB article](#).



The screenshot shows the 'Publish Disks' wizard window. The title bar says 'Publish Disks' with a close button. On the left is a sidebar with steps: Machine, Restore point, Disks, Target, Reason (selected), and Summary. The main area has a 'Reason' section with a sub-header 'Reason' and a description: 'Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.' Below this is a large text box labeled 'Restore reason:' containing the text 'Compare disk content'. At the bottom left of the main area is a checkbox labeled 'Do not show me this page again'. At the bottom right are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review the configured settings and click **Finish**.



What You Do Next

After the disks are published, go to the following locations on the target server to browse disks content:

- Go to the `/tmp/Veeam.Mount.Disks` location to browse disks images.
- Go to the `/tmp/Veeam.Mount.FS` location to browse disks content.

After you started a disks publishing session, you can view the session statistics or stop the session from the Veeam backup console. To learn more, see [Managing Publishing Disks Session](#).

Managing Publishing Disks Session

After you started a publishing session, you can check details about the 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 publishing 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 Veeam Agent computer 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 Veeam Agent computer 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, double-click the necessary publishing session and click **Cancel restore task** in the **Restore Session** window. Alternatively, you can select the necessary publishing session and click **Stop Publishing** on the ribbon or right-click the session and click **Stop Publishing**.
- Open the **Home** view. In the inventory pane select **Last 24 hours**. In the working area, double-click the necessary publishing session and click **Cancel restore task** in the **Restore Session** window. Alternatively, you can select the necessary publishing session and click **Stop** on the ribbon or right-click the session and click **Stop session**.

- Open the **History** view. In the inventory pane select **Restore**. In the working area, select the necessary publishing session and double-click it. In the **Restore Session** window, click **Cancel restore task**. Alternatively, you can right-click the publishing session and click **Stop session**.



Exporting Restore Point to Full Backup File

You can restore data from a specific restore point in a Veeam Agent backup and export this data to a standalone full backup file. The procedure of Veeam Agent backup export does not differ from the same procedure for a VM. To learn more, see the [Exporting Backups](#) section in the Veeam Backup & Replication User Guide.



Performing Administration Tasks

You can manage Veeam Agent backup jobs and backups created with these jobs. Veeam Backup & Replication allows you to perform the following administration tasks:

- [Import Veeam Agent backups.](#)
- [Enable and disable Veeam Agent backup jobs.](#)
- [Delete Veeam Agent backup jobs.](#)
- [View Veeam Agent backup properties.](#)
- [Create recovery token.](#)
- [Remove Veeam Agent backups.](#)
- [Delete Veeam Agent backups.](#)
- [Configure global settings.](#)
- [Assign roles to users.](#)

Importing Veeam Agent Backups

You may need to import a Veeam Agent backup in the Veeam Backup & Replication console in the following situations:

- The Veeam Agent backup is stored on a drive managed by another computer (not the Veeam backup server).
- The Veeam Agent backup is stored in a backup repository managed by another Veeam backup server.
- The Veeam Agent backup has been removed in the Veeam Backup & Replication console.


After importing, the Veeam Agent backup becomes available in the Veeam Backup & Replication console. You can restore data from such backup in a regular manner.

Before importing a backup, check the following prerequisites:

- The computer or server from which you plan to import the backup must be added to Veeam Backup & Replication. Otherwise you will not be able to access backup files.
- To be able to restore data from previous backup restore points, make sure that you have all incremental restore points in the same folder where the full backup file resides.

To import a Veeam Agent backup:

1. In Veeam Backup & Replication, click **Import Backup** on the **Home** tab.
2. From the **Computer** list, select the computer or 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. We recommend that you use the VBK files for import only if a corresponding VBM file is not available.
4. Click **OK**. The imported backup will become available in the **Home** view, under the **Backups > Disk (imported)** node in the inventory pane.




Importing Encrypted Backups

You can import Veeam Agent backups that were encrypted by Veeam Backup & Replication or Veeam Agent for Microsoft Windows.

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 need to specify the latest password. For Veeam Agent backups, you can use the latest password to restore data from all restore points in the backup chain, including those restore points that were encrypted with an old password.

If you enter correct password, Veeam Backup & Replication will decrypt the backup file. The backup will be moved under the **Backups > Disk (imported)** node in the inventory pane.




Enabling and Disabling Veeam Agent Backup Jobs

You can disable and enable Veeam Agent jobs in Veeam Backup & Replication.

When you disable the job, you prohibit the user to store the resulting backup in the backup repository. If the user starts a disabled job manually or the job starts by schedule, the job session will fail and report the "*Job is disabled on backup server*" error. To let Veeam Agent store backups in the backup repository again, you must enable the disabled job.

To disable or enable the scheduled backup job in Veeam Backup & Replication:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. Select the necessary job in the working area and click **Disable** on the ribbon, or right-click the necessary job in the working area and select **Disable**. To enable the disabled job, click **Disable** on the toolbar, or right-click the job and select **Disable** once again.




Viewing Veeam Agent Backup Job Statistics

You can view statistics about Veeam Agent backup jobs in the Veeam Backup & Replication console. Veeam Backup & Replication displays statistics for Veeam Agent backup jobs in the similar way as for regular backup jobs. The difference is that the list of objects included in the job contains a Veeam Agent machine instead of one or several VMs.

To view Veeam Agent backup job statistics:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. In the working area, select the necessary Veeam Agent backup job and click **Statistics** on the ribbon, or right-click the job and select **Statistics**.



Deleting Veeam Agent Backup Jobs

You can delete Veeam Agent backup jobs.

When you delete a Veeam Agent backup job, Veeam Backup & Replication removes all records about the job from its database and console. When the user starts a new Veeam Agent backup job session manually or the job starts automatically by schedule, the job will appear in the Veeam Backup & Replication console again, and records about a new job session will be stored in the Veeam Backup & Replication database.


NOTE

When you delete a Veeam Agent backup job, the backup files become orphaned and can be deleted by the background retention. For more information about the background retention, see the [Background Retention](#) section in the Veeam Backup & Replication User Guide.

To prevent the job from starting permanently, you must delete the job and unassign access rights permissions for this user from the backup repository. To completely delete the job, you must perform this operation in Veeam Agent on the Veeam Agent computer.

To remove a job:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click the **Jobs** node.
3. Select the necessary job in the working area and click **Delete** on the ribbon, or right-click the necessary job in the working area and select **Delete**.




Viewing Veeam Agent Backup Properties

You can view statistics about Veeam Agent backups.

To view Veeam Agent backup statistics:

1. In Veeam Backup & Replication, open the **Home** view.
2. In the inventory pane, click **Disk** under the **Backups** node.
3. In the working area, expand the **Agents** node, select the necessary backup and click **Properties** on the ribbon, or right-click the backup and select **Properties**.



Creating Recovery Token

You can create a recovery token for one or multiple Veeam Agent backups stored in a Veeam backup repository. A recovery token provides a limited-time access to the backup for which it was generated and can be used on the Veeam Agent computer side to connect to Veeam backup server during bare metal recovery. Using recovery tokens also eliminates the necessity to share Veeam backup server credentials.

To learn more about specifying a recovery token on the Veeam Agent computer side during bare metal recovery, see [Veeam Backup Repository Settings](#).

Considerations and Limitations

Before creating a recovery token, consider the following prerequisites and limitations:

- After generation, a recovery token remains valid for 24 hours.
- You can use a recovery token to restore files and folders only from the backups for which the token was generated. To view all backups in the repository and additional backup import options, use the password authentication instead.
- During recovery, Veeam Backup & Replication does not stop backup operations.
- You cannot create a recovery token for backups stored in a Veeam Cloud Connect repository.

Generating Recovery Token

To create a recovery token on the Veeam Backup & Replication side:

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. In the working area, right-click the backup and select **Create recovery token**.


You can create a recovery token for several backups. To do this, press and hold [Ctrl], select multiple backups, right-click one of the selected backups and select **Create recovery token**.

4. In the **Create Recovery Token** window, click **Create**.

You can modify the existing recovery token using the PowerShell console. To learn more, see the [Working with Tokens](#) section in the Veeam PowerShell Reference.

TIP

Alternatively, you can get access to the backup using user credentials. To learn more, see [Veeam Backup Repository Settings](#).




Removing Veeam Agent Backups

If you want to remove records about Veeam Agent backups from the Veeam Backup & Replication console and configuration database, you can use the **Remove from configuration** operation. When you remove a Veeam Agent backup from configuration, the actual backup files remain in the backup repository. You can import the backup to the Veeam Backup & Replication at any time later and restore data from it.

IMPORTANT

Removing backups from configuration is designed for experienced users only. Consider using the [Delete from disk](#) operation instead.

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. Press and hold the [Ctrl] key, select the backup, right-click the backup and select **Remove from configuration**.




Deleting Veeam Agent Backups from Disk

If you want to delete records about backups from the Veeam Backup & Replication console and configuration database and, additionally, delete backup files from the backup repository, you can use the **Delete from disk** operation.

To remove a Veeam Agent backup from the backup repository:

1. Open the **Home** view.
2. In the inventory pane, click **Backups**.
3. Select the necessary computer backup and click **Delete from > Disk** on the ribbon or right-click the computer and select **Delete from disk**.



Configuring Global Settings

Global settings configured on the Veeam backup server apply to Veeam Agent backup jobs as well. You can:

- Configure network throttling settings so that Veeam Agent backup job does not consume all network resources. To learn more, see the [Specifying I/O Settings](#) topic in the Veeam Backup & Replication User Guide.
- Configure the following global notification settings to get alerted about the Veeam Agent backup job results:
 - Email notifications. To learn more, see the [Specifying Email Notification Settings](#) section in the Veeam Backup & Replication User Guide.
 - SNMP notifications. To learn more, see the [Specifying SNMP Settings](#) section in the Veeam Backup & Replication User Guide.

Assigning Roles to Users

User roles configured on the Veeam backup server apply to Veeam Agent backup jobs as well.

To learn more, see the [Users and Roles](#) section in the Veeam Backup & Replication User Guide.