



Red Hat Enterprise Linux 7

Linux Domain Identity, Authentication, and Policy Guide

Using Red Hat Identity Management in Linux Environments

Red Hat Enterprise Linux 7 Linux Domain Identity, Authentication, and Policy Guide

Using Red Hat Identity Management in Linux Environments

Filip Hanzelka
Red Hat Customer Content Services
fhanzelk@redhat.com

Lucie Maňásková
Red Hat Customer Content Services
lmanasko@redhat.com

Aneta Šteflová Petrová
Red Hat Customer Content Services

Marc Muehlfeld
Red Hat Customer Content Services

Tomáš Čapek
Red Hat Customer Content Services

Ella Deon Ballard
Red Hat Customer Content Services

Legal Notice

Copyright © 2018 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](#). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Keywords

1. FreeIPA. 2. Identity Management. 3. IdM. 4. IPA.

Abstract

Identity and policy management, for both users and machines, is a core function for most enterprise environments. Identity Management provides a way to create an identity domain that allows machines to enroll to a domain and immediately access identity information required for single sign-on and authentication services, as well as policy settings that govern authorization and access. In addition to this guide, you can find documentation on other features and services related to Red Hat Enterprise Linux Identity Management in the following guides: The System-Level Authentication Guide documents different applications and services available to configure authentication on local systems, including the authconfig utility, the System Security Services Daemon (SSSD) service, the Pluggable Authentication Module (PAM) framework, Kerberos, the certmonger utility, and single sign-on (SSO) for applications. The Windows Integration Guide documents how to integrate Linux domains with Microsoft Windows Active

Directory (AD) using Identity Management. Among other topics, the guide covers various aspects of direct and indirect AD integration, using SSSD to access a Common Internet File System (CIFS), and the realmd system.

Table of Contents

PART I. OVERVIEW OF RED HAT IDENTITY MANAGEMENT	8
CHAPTER 1. INTRODUCTION TO RED HAT IDENTITY MANAGEMENT	9
1.1. THE GOAL OF RED HAT IDENTITY MANAGEMENT	9
1.2. THE IDENTITY MANAGEMENT DOMAIN	11
PART II. INSTALLING IDENTITY MANAGEMENT	16
CHAPTER 2. INSTALLING AND UNINSTALLING AN IDENTITY MANAGEMENT SERVER	17
2.1. PREREQUISITES FOR INSTALLING A SERVER	17
2.2. PACKAGES REQUIRED TO INSTALL AN IDM SERVER	24
2.3. INSTALLING AN IDM SERVER: INTRODUCTION	25
2.4. UNINSTALLING AN IDM SERVER	39
2.5. RENAMING A SERVER	40
CHAPTER 3. INSTALLING AND UNINSTALLING IDENTITY MANAGEMENT CLIENTS	41
3.1. PREREQUISITES FOR INSTALLING A CLIENT	41
3.2. PACKAGES REQUIRED TO INSTALL A CLIENT	42
3.3. INSTALLING A CLIENT	42
3.4. SETTING UP AN IDM CLIENT THROUGH KICKSTART	46
3.5. POST-INSTALLATION CONSIDERATIONS FOR CLIENTS	47
3.6. TESTING THE NEW CLIENT	48
3.7. UNINSTALLING A CLIENT	48
3.8. RE-ENROLLING A CLIENT INTO THE IDM DOMAIN	48
3.9. RENAMING CLIENT MACHINES	49
CHAPTER 4. INSTALLING AND UNINSTALLING IDENTITY MANAGEMENT REPLICAS	52
4.1. EXPLAINING IDM REPLICAS	52
4.2. DEPLOYMENT CONSIDERATIONS FOR REPLICAS	52
4.3. PREREQUISITES FOR INSTALLING A REPLICA	56
4.4. PACKAGES REQUIRED TO INSTALL A REPLICA	56
4.5. CREATING THE REPLICA: INTRODUCTION	57
4.6. TESTING THE NEW REPLICA	63
4.7. UNINSTALLING A REPLICA	63
PART III. ADMINISTRATION: MANAGING SERVERS	64
CHAPTER 5. THE BASICS OF MANAGING THE IDM SERVER AND SERVICES	65
5.1. STARTING AND STOPPING THE IDM SERVER	65
5.2. LOGGING INTO IDM USING KERBEROS	65
5.3. THE IDM COMMAND-LINE UTILITIES	67
5.4. THE IDM WEB UI	71
CHAPTER 6. MANAGING REPLICATION TOPOLOGY	77
6.1. EXPLAINING REPLICATION AGREEMENTS, TOPOLOGY SUFFIXES, AND TOPOLOGY SEGMENTS	77
6.2. WEB UI: USING THE TOPOLOGY GRAPH TO MANAGE REPLICATION TOPOLOGY	79
6.3. COMMAND LINE: MANAGING TOPOLOGY USING THE IPA TOPOLOGY* COMMANDS	84
6.4. REMOVING A SERVER FROM THE TOPOLOGY	86
6.5. MANAGING SERVER ROLES	88
CHAPTER 7. DISPLAYING AND RAISING THE DOMAIN LEVEL	93
7.1. DISPLAYING THE CURRENT DOMAIN LEVEL	93
7.2. RAISING THE DOMAIN LEVEL	94

CHAPTER 8. UPDATING AND MIGRATING IDENTITY MANAGEMENT	95
8.1. UPDATING IDENTITY MANAGEMENT	95
8.2. MIGRATING IDENTITY MANAGEMENT FROM RED HAT ENTERPRISE LINUX 6 TO VERSION 7	96
CHAPTER 9. BACKING UP AND RESTORING IDENTITY MANAGEMENT	104
9.1. FULL-SERVER BACKUP AND DATA-ONLY BACKUP	105
9.2. RESTORING A BACKUP	108
CHAPTER 10. DEFINING ACCESS CONTROL FOR IDM USERS	111
10.1. ACCESS CONTROLS FOR IDM ENTRIES	111
10.2. DEFINING SELF-SERVICE SETTINGS	112
10.3. DELEGATING PERMISSIONS OVER USERS	115
10.4. DEFINING ROLE-BASED ACCESS CONTROLS	117
PART IV. ADMINISTRATION: MANAGING IDENTITIES	135
CHAPTER 11. MANAGING USER ACCOUNTS	136
11.1. SETTING UP USER HOME DIRECTORIES	136
11.2. USER LIFE CYCLE	137
11.3. EDITING USERS	147
11.4. ENABLING AND DISABLING USER ACCOUNTS	149
11.5. ALLOWING NON-ADMIN USERS TO MANAGE USER ENTRIES	150
11.6. USING AN EXTERNAL PROVISIONING SYSTEM FOR USERS AND GROUPS	153
CHAPTER 12. MANAGING HOSTS	162
12.1. ABOUT HOSTS, SERVICES, AND MACHINE IDENTITY AND AUTHENTICATION	162
12.2. ABOUT HOST ENTRY CONFIGURATION PROPERTIES	163
12.3. ADDING HOST ENTRIES	164
12.4. DISABLING AND RE-ENABLING HOST ENTRIES	167
12.5. MANAGING PUBLIC SSH KEYS FOR HOSTS	168
12.6. SETTING ETHERS INFORMATION FOR A HOST	174
CHAPTER 13. MANAGING USER AND HOST GROUPS	175
13.1. HOW USER AND HOST GROUPS WORK IN IDM	175
13.2. ADDING AND REMOVING USER OR HOST GROUPS	178
13.3. ADDING AND REMOVING USER OR HOST GROUP MEMBERS	180
13.4. DISABLING USER PRIVATE GROUPS	182
13.5. SETTING SEARCH ATTRIBUTES FOR USERS AND USER GROUPS	184
13.6. DEFINING AUTOMATIC GROUP MEMBERSHIP FOR USERS AND HOSTS	185
CHAPTER 14. UNIQUE UID AND GID NUMBER ASSIGNMENTS	193
14.1. ID RANGES	193
14.2. ID RANGE ASSIGNMENTS DURING INSTALLATION	193
14.3. DISPLAYING CURRENTLY ASSIGNED ID RANGES	194
14.4. AUTOMATIC ID RANGE EXTENSION AFTER DELETING A REPLICA	194
14.5. MANUAL ID RANGE EXTENSION AND ASSIGNING A NEW ID RANGE	194
14.6. ENSURING THAT ID VALUES ARE UNIQUE	196
14.7. REPAIRING CHANGED UID AND GID NUMBERS	196
CHAPTER 15. USER AND GROUP SCHEMA	198
15.1. ABOUT CHANGING THE DEFAULT USER AND GROUP SCHEMA	200
15.2. APPLYING CUSTOM OBJECT CLASSES TO NEW USER ENTRIES	201
15.3. APPLYING CUSTOM OBJECT CLASSES TO NEW GROUP ENTRIES	203
15.4. SPECIFYING DEFAULT USER AND GROUP ATTRIBUTES	204
CHAPTER 16. MANAGING SERVICES	209

16.1. ADDING AND EDITING SERVICE ENTRIES AND KEYTABS	209
16.2. CONFIGURING CLUSTERED SERVICES	211
16.3. USING THE SAME SERVICE PRINCIPAL FOR MULTIPLE SERVICES	212
16.4. RETRIEVE EXISTING KEYTABS FOR MULTIPLE SERVERS	212
16.5. DISABLING AND RE-ENABLING SERVICE ENTRIES	214
CHAPTER 17. DELEGATING USER ACCESS TO HOSTS AND SERVICES	216
17.1. DELEGATING SERVICE MANAGEMENT	216
17.2. DELEGATING HOST MANAGEMENT	217
17.3. DELEGATING HOST OR SERVICE MANAGEMENT IN THE WEB UI	217
17.4. ACCESSING DELEGATED SERVICES	218
CHAPTER 18. ID VIEWS	220
Potential Negative Impact on SSSD Performance	220
Additional Resources	220
18.1. ATTRIBUTES AN ID VIEW CAN OVERRIDE	220
18.2. GETTING HELP FOR ID VIEW COMMANDS	221
18.3. DEFINING A DIFFERENT ATTRIBUTE VALUE FOR A USER ACCOUNT ON DIFFERENT HOSTS	221
CHAPTER 19. DEFINING ACCESS CONTROL FOR IDM USERS	227
CHAPTER 20. MANAGING KERBEROS FLAGS AND PRINCIPAL ALIASES	228
20.1. KERBEROS FLAGS FOR SERVICES AND HOSTS	228
20.2. MANAGING KERBEROS PRINCIPAL ALIASES FOR USERS, HOSTS, AND SERVICES	230
CHAPTER 21. INTEGRATING WITH NIS DOMAINS AND NETGROUPS	233
21.1. ABOUT NIS AND IDENTITY MANAGEMENT	233
21.2. ENABLING NIS IN IDENTITY MANAGEMENT	235
21.3. CREATING NETGROUPS	235
21.4. EXPOSING AUTOMOUNT MAPS TO NIS CLIENTS	239
21.5. MIGRATING FROM NIS TO IDM	240
PART V. ADMINISTRATION: MANAGING AUTHENTICATION	247
CHAPTER 22. USER AUTHENTICATION	248
22.1. USER PASSWORDS	248
22.2. ONE-TIME PASSWORDS	252
22.3. RESTRICTING ACCESS TO SERVICES AND HOSTS BASED ON HOW USERS AUTHENTICATE	260
22.4. MANAGING PUBLIC SSH KEYS FOR USERS	261
22.5. CONFIGURING SSSD TO PROVIDE A CACHE FOR THE OPENSSSH SERVICES	265
22.6. SMART-CARD AUTHENTICATION IN IDENTITY MANAGEMENT	267
22.7. USER CERTIFICATES	267
CHAPTER 23. SMART-CARD AUTHENTICATION IN IDENTITY MANAGEMENT	268
23.1. MANAGING SMART CARD LINKS IN THE IDENTITY MANAGEMENT SERVER	268
23.2. AUTHENTICATING TO AN IDENTITY MANAGEMENT CLIENT WITH A SMART CARD	279
23.3. AUTHENTICATING TO AN IDENTITY MANAGEMENT SYSTEM REMOTELY WITH A SMART CARD	281
23.4. CONFIGURING A USER NAME HINT POLICY FOR SMART-CARD AUTHENTICATION	284
23.5. PKINIT SMART-CARD AUTHENTICATION IN IDENTITY MANAGEMENT	285
23.6. AUTHENTICATING TO THE IDENTITY MANAGEMENT WEB UI WITH A SMART CARD	288
23.7. INTEGRATING IDENTITY MANAGEMENT SMART-CARD AUTHENTICATION WITH WEB APPLICATIONS	292
CHAPTER 24. MANAGING CERTIFICATES FOR USERS, HOSTS, AND SERVICES	295
24.1. MANAGING CERTIFICATES WITH THE INTEGRATED IDM CAS	295

24.2. MANAGING CERTIFICATES ISSUED BY EXTERNAL CAS	299
24.3. LISTING AND DISPLAYING CERTIFICATES	301
24.4. CERTIFICATE PROFILES	303
24.5. CERTIFICATE AUTHORITY ACL RULES	309
24.6. USING CERTIFICATE PROFILES AND ACLS TO ISSUE USER CERTIFICATES WITH THE IDM CAS	314
CHAPTER 25. STORING AUTHENTICATION SECRETS WITH VAULTS	321
25.1. HOW VAULTS WORK	321
25.2. PREREQUISITES FOR USING VAULTS	323
25.3. GETTING HELP FOR VAULT COMMANDS	323
25.4. STORING A USER'S PERSONAL SECRET	324
25.5. STORING A SERVICE SECRET IN A VAULT	326
25.6. STORING A COMMON SECRET FOR MULTIPLE USERS	329
CHAPTER 26. MANAGING CERTIFICATES AND CERTIFICATE AUTHORITIES	332
26.1. LIGHTWEIGHT SUB-CAS	332
26.2. RENEWING CERTIFICATES	334
26.3. INSTALLING A CA CERTIFICATE MANUALLY	336
26.4. CHANGING THE CERTIFICATE CHAIN	337
26.5. ALLOWING IDM TO START WITH EXPIRED CERTIFICATES	337
26.6. INSTALLING THIRD-PARTY CERTIFICATES FOR HTTP OR LDAP	338
26.7. CONFIGURING OCSP RESPONDERS	339
26.8. INSTALLING A CA INTO AN EXISTING IDM DOMAIN	340
26.9. REPLACING THE WEB SERVER'S AND LDAP SERVER'S CERTIFICATE	341
CHAPTER 27. KERBEROS PKINIT AUTHENTICATION IN IDM	343
27.1. DEFAULT PKINIT STATUS IN DIFFERENT IDM VERSIONS	343
27.2. DISPLAYING THE CURRENT PKINIT CONFIGURATION	343
27.3. CONFIGURING PKINIT IN IDM	344
27.4. ADDITIONAL RESOURCES	345
PART VI. ADMINISTRATION: MANAGING POLICIES	346
CHAPTER 28. DEFINING PASSWORD POLICIES	347
28.1. WHAT ARE PASSWORD POLICIES AND WHY ARE THEY USEFUL	347
28.2. HOW PASSWORD POLICIES WORK IN IDM	347
28.3. ADDING A NEW PASSWORD POLICY	350
28.4. MODIFYING PASSWORD POLICY ATTRIBUTES	351
28.5. CHANGING PASSWORD EXPIRATION DATE WITH IMMEDIATE EFFECT	352
CHAPTER 29. MANAGING THE KERBEROS DOMAIN	353
29.1. MANAGING KERBEROS TICKET POLICIES	353
29.2. REKEYING KERBEROS PRINCIPALS	355
29.3. PROTECTING KEYTABS	357
29.4. REMOVING KEYTABS	357
29.5. ADDITIONAL RESOURCES	358
CHAPTER 30. USING SUDO	359
30.1. THE SUDO UTILITY IN IDENTITY MANAGEMENT	359
30.2. SUDO RULES IN IDENTITY MANAGEMENT	359
30.3. CONFIGURING THE LOCATION FOR LOOKING UP SUDO POLICIES	360
30.4. ADDING SUDO COMMANDS, COMMAND GROUPS, AND RULES	362
30.5. MODIFYING SUDO COMMANDS AND COMMAND GROUPS	366
30.6. MODIFYING SUDO RULES	366

30.7. LISTING AND DISPLAYING SUDO COMMANDS, COMMAND GROUPS, AND RULES	377
30.8. DISABLING AND ENABLING SUDO RULES	377
30.9. REMOVING SUDO COMMANDS, COMMAND GROUPS, AND RULES	378
CHAPTER 31. CONFIGURING HOST-BASED ACCESS CONTROL	380
31.1. HOW HOST-BASED ACCESS CONTROL WORKS IN IDM	380
31.2. CONFIGURING HOST-BASED ACCESS CONTROL IN AN IDM DOMAIN	380
31.3. ADDING HBAC SERVICE ENTRIES FOR CUSTOM HBAC SERVICES	390
31.4. ADDING HBAC SERVICE GROUPS	391
CHAPTER 32. DEFINING SELINUX USER MAPS	393
32.1. ABOUT IDENTITY MANAGEMENT, SELINUX, AND MAPPING USERS	393
32.2. CONFIGURING SELINUX USER MAP ORDER AND DEFAULTS	395
32.3. MAPPING SELINUX USERS AND IDM USERS	397
PART VII. ADMINISTRATION: MANAGING NETWORK SERVICES	402
CHAPTER 33. MANAGING DNS	403
33.1. BIND IN IDENTITY MANAGEMENT	403
33.2. SUPPORTED DNS ZONE TYPES	404
33.3. DNS CONFIGURATION PRIORITIES	404
33.4. MANAGING MASTER DNS ZONES	405
33.5. MANAGING DYNAMIC DNS UPDATES	420
33.6. MANAGING DNS FORWARDING	427
33.7. MANAGING REVERSE DNS ZONES	433
33.8. DEFINING DNS QUERY POLICY	436
33.9. DNS LOCATIONS	436
33.10. UPDATING DNS RECORDS SYSTEMATICALLY WHEN USING EXTERNAL DNS	440
33.11. INSTALLING DNS SERVICES INTO AN EXISTING SERVER	443
CHAPTER 34. USING AUTOMOUNT	444
34.1. ABOUT AUTOMOUNT AND IDM	444
34.2. CONFIGURING AUTOMOUNT	444
34.3. SETTING UP A KERBEROS-AWARE NFS SERVER	450
34.4. CONFIGURING LOCATIONS	453
34.5. CONFIGURING MAPS	455
PART VIII. SECURITY HARDENING	462
CHAPTER 35. CONFIGURING TLS FOR IDENTITY MANAGEMENT	463
35.1. CONFIGURING THE HTTPD DAEMON	463
35.2. CONFIGURING THE DIRECTORY SERVER COMPONENT	463
35.3. CONFIGURING THE CERTIFICATE SERVER COMPONENT	464
35.4. RESULT	464
CHAPTER 36. DISABLING ANONYMOUS BINDS	465
PART IX. PERFORMANCE TUNING	466
CHAPTER 37. PERFORMANCE TUNING FOR BULK PROVISIONING OF ENTRIES	467
Recommendations and Prerequisites for Bulk Provisioning	467
Backing up the Current DS Tuning Parameter Values	468
Adjusting the Database, Domain Entry, and DN Cache Size	468
Disabling Unnecessary Services and Adjusting Database Locks	470
Importing the Entries	471
Re-enabling the Disabled Services and Restoring the Original Attribute Values	471

PART X. MIGRATION	474
CHAPTER 38. MIGRATING FROM AN LDAP DIRECTORY TO IDM	475
38.1. AN OVERVIEW OF AN LDAP TO IDM MIGRATION	475
38.2. EXAMPLES FOR USING IPA MIGRATE-DS	483
38.3. MIGRATING AN LDAP SERVER TO IDENTITY MANAGEMENT	486
38.4. MIGRATING OVER SSL	488
APPENDIX A. TROUBLESHOOTING: GENERAL GUIDELINES	489
A.1. INVESTIGATING FAILURES WHEN EXECUTING THE IPA UTILITY	489
A.2. INVESTIGATING KINIT AUTHENTICATION FAILURES	491
A.3. INVESTIGATING IDM WEB UI AUTHENTICATION FAILURES	493
A.4. INVESTIGATING SMART CARD AUTHENTICATION FAILURES	494
A.5. INVESTIGATING WHY A SERVICE FAILS TO START	494
A.6. TROUBLESHOOTING DNS	495
A.7. TROUBLESHOOTING REPLICATION	497
APPENDIX B. TROUBLESHOOTING: SOLUTIONS TO SPECIFIC PROBLEMS	498
B.1. IDENTITY MANAGEMENT SERVERS	498
B.2. IDENTITY MANAGEMENT REPLICAS	499
B.3. IDENTITY MANAGEMENT CLIENTS	504
B.4. LOGGING IN AND AUTHENTICATION PROBLEMS	506
B.5. VAULTS	508
APPENDIX C. A REFERENCE OF IDENTITY MANAGEMENT FILES AND LOGS	510
C.1. IDENTITY MANAGEMENT CONFIGURATION FILES AND DIRECTORIES	510
C.2. IDENTITY MANAGEMENT LOG FILES AND DIRECTORIES	512
C.3. IDM DOMAIN SERVICES AND LOG ROTATION	515
APPENDIX D. MANAGING REPLICAS AT DOMAIN LEVEL 0	517
D.1. REPLICA INFORMATION FILE	517
D.2. CREATING REPLICAS	517
D.3. MANAGING REPLICAS AND REPLICATION AGREEMENTS	521
D.4. PROMOTING A REPLICA TO A MASTER CA SERVER	525
APPENDIX E. REVISION HISTORY	527

PART I. OVERVIEW OF RED HAT IDENTITY MANAGEMENT

CHAPTER 1. INTRODUCTION TO RED HAT IDENTITY MANAGEMENT

This chapter explains the purpose of Red Hat Identity Management. It also provides basic information about the Identity Management domain, including the client and server machines that are part of the domain.

1.1. THE GOAL OF RED HAT IDENTITY MANAGEMENT

Red Hat Identity Management (IdM) provides a centralized and unified way to manage identity stores, authentication, policies, and authorization policies in a Linux-based domain. IdM significantly reduces the administrative overhead of managing different services individually and using different tools on different machines.

IdM is one of the few centralized identity, policy, and authorization software solutions that support:

- Advanced features of Linux operating system environments
- Unifying large groups of Linux machines
- Native integration with Active Directory

IdM creates a Linux-based and Linux-controlled domain:

- IdM builds on existing, native Linux tools and protocols. It has its own processes and configuration, but its underlying technologies are well-established on Linux systems and trusted by Linux administrators.
- IdM servers and clients are Red Hat Enterprise Linux machines. However, even though IdM does not support Windows clients directly, it allows integration with Active Directory environment.



NOTE

This guide describes using IdM in Linux environments only. For more information on integration with Active Directory, see the [Windows Integration Guide](#).

For information on the Samba suite, which allows integrating Linux machines into Active Directory environment, see the [Using Samba, Kerberos, and Winbind](#) chapter in the *Windows Integration Guide*.

1.1.1. Examples of Benefits Brought by IdM

Managing identities and policies with several Linux servers

Without IdM: Each server is administered separately. All passwords are saved on the local machines. The IT administrator manages users on every machine, sets authentication and authorization policies separately, and maintains local passwords.

With IdM: The IT administrator can:

- Maintain the identities in one central place: the IdM server

- Apply policies uniformly to multiples of machines at the same time
- Set different access levels for users by using host-based access control, delegation, and other rules
- Centrally manage privilege escalation rules
- Define how home directories are mounted

Enterprise single sign-on

Without IdM: Users log in to the system and are prompted for a password every single time they access a service or application. These passwords might be different, and the users have to remember which credential to use for which application.

With IdM: After users log in to the system, they can access multiple services and applications without being repeatedly asked for their credentials. This helps:

- Improve usability
- Reduce the security risk of passwords being written down or stored insecurely
- Boost user productivity

Managing a mixed Linux and Windows environment

Without IdM: Windows systems are managed in an Active Directory forest, but development, production, and other teams have many Linux systems. The Linux systems are excluded from the Active Directory environment.

With IdM: The IT administrator can:

- Manage the Linux systems using native Linux tools
- Integrate the Linux systems with the Windows systems, thus preserving a centralized user store
- Expand the Linux base easily
- Separate management of Linux and Active Directory machines and enable Linux and Windows admins to control their environment directly

1.1.2. Contrasting Identity Management with a Standard LDAP Directory

A standard LDAP directory, such as Red Hat Directory Server, is a general-purpose directory: it can be customized to fit a broad range of use cases.

- Schema: a flexible schema that can be customized for a vast array of entries, such as users, machines, network entities, physical equipment, or buildings.
- Typically used as: a back-end directory to store data for other applications, such as business applications that provide services on the Internet.

Identity Management (IdM) has a specific purpose: managing identities as well as authentication and authorization policies that relate to these identities.

- Schema: a specific schema that defines a particular set of entries relevant to its purpose, such as entries for user or machine identities.
- Typically used as: the identity and authentication server to manage identities within the boundaries of an enterprise or a project.

The underlying directory server technology is the same for both Red Hat Directory Server and IdM. However, IdM is optimized to manage identities. This limits its general extensibility, but also brings certain benefits: simpler configuration, better automation of resource management, and increased efficiency in managing identities.

Additional Resources

- [Identity Management or Red Hat Directory Server – Which One Should I Use?](#) on the *Red Hat Enterprise Linux Blog*.

1.2. THE IDENTITY MANAGEMENT DOMAIN

The Identity Management (IdM) domain consists of a group of machines that share the same configuration, policies, and identity stores. The shared properties allow the machines within the domain to be aware of each other and operate together.

From the perspective of IdM, the domain includes the following types of machines:

- IdM servers, which work as domain controllers
- IdM clients, which are enrolled with the servers

IdM servers are also IdM clients enrolled with themselves: server machines provide the same functionality as clients.

IdM supports Red Hat Enterprise Linux machines as the IdM servers and clients.



NOTE

This guide describes using IdM in Linux environments. For more information on integration with Active Directory, see the [Windows Integration Guide](#).

1.2.1. Identity Management Servers

The IdM servers act as central repositories for identity and policy information. They also host the services used by domain members. IdM provides a set of management tools to manage all the IdM-associated services centrally: the IdM web UI and command-line utilities.

For information on installing IdM servers, see [Chapter 2, Installing and Uninstalling an Identity Management Server](#).

To support redundancy and load balancing, the data and configuration can be replicated from one IdM server to another: a *replica* of the initial server. You can configure servers and their replicas to provide different services to clients. For more details on IdM replicas, see [Chapter 4, Installing and Uninstalling Identity Management Replicas](#).

1.2.1.1. Services Hosted by IdM Servers

Most of the following services are not strictly required to be installed on the IdM server. For example, services such as a certificate authority (CA), a DNS server, or a Network Time Protocol (NTP) server can be installed on an external server outside the IdM domain.

Kerberos KDC

IdM uses the Kerberos protocol to support single sign-on. With Kerberos, the user only needs to present the correct user name and password once. Then the user can access IdM services without the system prompting for the credentials again.

- For details on how Kerberos works, see the [System-Level Authentication Guide](#).
- For information on how to authenticate using Kerberos in IdM, see [Section 5.2, “Logging into IdM Using Kerberos”](#).
- For information on managing Kerberos in IdM, see [Chapter 29, Managing the Kerberos Domain](#).

LDAP directory server

IdM includes an internal LDAP directory server instance where it stores all the IdM information, such as information related to Kerberos, user accounts, host entries, services, policies, DNS, and others.

The LDAP directory server instance is based on the same technology as Red Hat Directory Server. However, it is tuned to IdM-specific tasks.




NOTE

This guide refers to this component as Directory Server.

Certificate authority

In most deployments, an integrated certificate authority (CA) is installed with the IdM server. You can also install the server without the integrated CA if you create and provide all required certificates independently.

- For more details on installing an IdM server with the different CA configurations, see [Section 2.3.2, “Determining What CA Configuration to Use”](#).



NOTE

This guide refers to this component as Certificate System when addressing the implementation and as certificate authority when addressing the services provided by the implementation.

Domain Name System (DNS)

IdM uses DNS for dynamic service discovery. The IdM client installation utility can use information from DNS to automatically configure the client machine. After the client is enrolled in the IdM domain, it uses DNS to locate IdM servers and services within the domain.


- For more information about service discovery, see the [System-Level Authentication Guide](#).

- For information on using DNS with IdM and important prerequisites, see [Section 2.1.3, “Host Name and DNS Configuration”](#).
- For details on installing an IdM server with or without integrated DNS, see [Section 2.3.1, “Determining Whether to Use Integrated DNS”](#).

Network Time Protocol


Many services require that servers and clients have the same system time, within a certain variance. For example, Kerberos tickets use time stamps to determine their validity and to prevent replay attacks. If the times between the server and client skew outside the allowed range, the Kerberos tickets are invalidated.

By default, IdM uses the Network Time Protocol (NTP) to synchronize clocks over a network. With NTP, a central server acts as an authoritative clock and the clients synchronize their times to match the server clock. The IdM server is configured as the NTP server for the IdM domain during the server installation process.



NOTE

Running an NTP server on an IdM server installed on a virtual machine can lead to inaccurate time synchronization in some environments. To avoid potential problems, do not run NTP on IdM servers installed on virtual machines. For more information on the reliability of an NTP server on a virtual machine, see [this Knowledgebase solution](#).



RHEL_467514_0318

Figure 1.1. The Identity Management Server: Unifying Services

1.2.2. Identity Management Clients

IdM clients are machines configured to operate within the IdM domain. They interact with the IdM servers to access domain resources. For example, they belong to the Kerberos domains configured on the servers, receive certificates and tickets issued by the servers,

and use other centralized services for authentication and authorization.

An IdM client does not require dedicated client software to interact as a part of the domain. It only requires proper system configuration of certain services and libraries, such as Kerberos or DNS. This configuration directs the client machine to use IdM services.

For information on installing IdM clients, see [Chapter 3, *Installing and Uninstalling Identity Management Clients*](#).

1.2.2.1. Services Hosted by IdM Clients

System Security Services Daemon

The System Security Services Daemon (SSSD) is a client-side application for caching credentials. Using SSSD on client machines is recommended because it simplifies the required client configuration. SSSD also provides additional features, for example:

- Offline client authentication, ensured by caching credentials from centralized identity and authentication stores locally
- Improved consistency of the authentication process, because it is not necessary to maintain both a central account and a local user account for offline authentication
- Integration with other services, such as **sudo**
- Host-based access control (HBAC) authorization

With SSSD, the IdM administrators can define all identity configuration centrally in the IdM server. Caching enables the local system to continue normal authentication operations if the IdM server becomes unavailable or if the client becomes offline.

For more information about SSSD, see the [System-Level Authentication Guide](#). SSSD also supports Windows Active Directory (AD). For more information about using SSSD with AD, see the [Windows Integration Guide](#).

certmonger

The **certmonger** service monitors and renews the certificates on the client. It can request new certificates for the services on the system.

For more information about **certmonger**, see the [System-Level Authentication Guide](#).




Figure 1.2. Interactions Between IdM Services

PART II. INSTALLING IDENTITY MANAGEMENT

CHAPTER 2. INSTALLING AND UNINSTALLING AN IDENTITY MANAGEMENT SERVER

An *Identity Management (IdM) server* is a domain controller: it defines and manages the IdM domain. To set up an IdM server, you must:

1. Install the necessary packages
2. Configure the machine using setup scripts

Red Hat strongly recommends to set up multiple domain controllers within your domain for load balancing and redundancy. These additional servers are *replicas* of the initial master IdM server.

This chapter describes installing the first, initial IdM server. For information on installing a replica from the initial server, see [Chapter 4, *Installing and Uninstalling Identity Management Replicas*](#).

2.1. PREREQUISITES FOR INSTALLING A SERVER

2.1.1. Hardware Recommendations

RAM is the most important hardware feature to size properly. To determine how much RAM you require, consider these recommendations:

- For 10,000 users and 100 groups: at least 3 GB of RAM and 1 GB swap space
- For 100,000 users and 50,000 groups: at least 16 GB of RAM and 4 GB of swap space



NOTE

A basic user entry or a simple host entry with a certificate is approximately 5 - 10 KiB in size.

For larger deployments, it is more effective to increase the RAM than to increase disk space because much of the data is stored in cache.

To increase performance, you can tune the underlying Directory Server to increase performance. For details, see [Optimizing System Performance](#) in the Directory Server Performance Tuning Guide.

2.1.2. System Requirements

Identity Management 4.4 is supported on Red Hat Enterprise Linux 7. Install an IdM server on a clean system without any custom configuration for services such as DNS, Kerberos, or Directory Server.

The IdM server installation overwrites system files to set up the IdM domain. IdM backs up the original system files to `/var/lib/ipa/sysrestore/`.

Federal Information Processing Standard (FIPS) support

In environments set up using Red Hat Enterprise Linux 7.4 and later:

- You can configure a new IdM server, replica, or client on a system with the FIPS mode enabled. The installation script automatically detects a system with FIPS enabled and configures IdM without the administrator's intervention.

To enable FIPS in the operating system, see [Enabling FIPS Mode](#) in the *Security Guide*.



IMPORTANT

You cannot:

- Enable FIPS mode on existing IdM servers previously installed with FIPS mode disabled.
- Install a replica in FIPS mode when using an existing IdM server with FIPS mode disabled.

In environments set up using Red Hat Enterprise Linux 7.3 and earlier:

- IdM does not support the FIPS mode. Disable FIPS on your system before installing an IdM server, replica, or client, and do not enable it after the installation.

For further details about FIPS mode, see [Federal Information Processing Standard \(FIPS\)](#) in the *Security Guide*.

Name Service Cache Daemon (NSCD) requirements

Red Hat recommends to disable NSCD on Identity Management machines. Alternatively, if disabling NSCD is not possible, only enable NSCD for maps that SSSD does not cache.

Both NSCD and the SSSD service perform caching, and problems can occur when systems use both services simultaneously. See the [System-Level Authentication Guide](#) for information on how to avoid conflicts between NSCD and SSSD.

IPv6 must be enabled on the system

Installing and running an IdM server requires IPv6 to be enabled on the network. Note that IPv6 is enabled by default on Red Hat Enterprise Linux 7 systems.

If you disabled IPv6 before, re-enable it as described in [How do I disable or enable the IPv6 protocol in Red Hat Enterprise Linux?](#) in Red Hat Knowledgebase.

2.1.3. Host Name and DNS Configuration

**WARNING**

Be extremely cautious and ensure that:

- you have a tested and functional DNS service available
- the service is properly configured

This requirement applies to IdM servers with integrated DNS services as well as to IdM servers installed without DNS. DNS records are vital for nearly all IdM domain functions, including running LDAP directory services, Kerberos, and Active Directory integration.

Note that the primary DNS domain and Kerberos realm cannot be changed after the installation.

The server host must have DNS properly configured regardless of whether the DNS server is integrated within IdM or hosted externally.

Identity Management requires one separate DNS domain to be used for service records. To avoid conflicts on the DNS level, the *primary DNS domain* used for IdM cannot be shared with any other system.

Note that host names of IdM clients are not required to be part of the primary DNS domain.

**NOTE**

For information on configuring users to access an IdM client using a host name from the Active Directory DNS domain, while the client itself is joined to IdM, see [IdM clients in an Active Directory DNS Domain](#) in the *Windows Integration Guide*.

Verifying the Server Host Name

The host name must be a fully qualified domain name, such as **server.example.com**. To verify your machine's host name, use the **hostname** utility:

```
[root@server ~]# hostname  
server.example.com
```

The output of **hostname** must not be **localhost** or **localhost6**.



IMPORTANT

The fully qualified domain name must be a valid DNS name, which means only numbers, alphabetic characters, and hyphens (-) are allowed. Other characters, like underscores, in the host name cause DNS failures. Additionally, the host name must be all lower-case; no capital letters are allowed.

For other recommended naming practices, see the [Red Hat Enterprise Linux Security Guide](#).

The fully qualified domain name must not resolve to the loopback address. It must resolve to the machine's public IP address, not to **127.0.0.1**.

Verifying the Forward and Reverse DNS Configuration

1. Obtain the IP address of the server. The **ip addr show** command displays both the IPv4 and IPv6 addresses:
 - The IPv4 address is displayed on the line starting with **inet**. In the following example, the configured IPv4 address is **192.0.2.1**.
 - The IPv6 address is displayed on the line starting with **inet6**. Only IPv6 addresses with **scope global** are relevant for this procedure. In the following example, the returned IPv6 address is **2001:DB8::1111**.

```
[root@server ~]# ip addr show
...
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast
state UP group default qlen 1000
    link/ether 00:1a:4a:10:4e:33 brd ff:ff:ff:ff:ff:ff
    inet 192.0.2.1/24 brd 192.0.2.255 scope global dynamic eth0
        valid_lft 106694sec preferred_lft 106694sec
    inet6 2001:DB8::1111/32 scope global dynamic
        valid_lft 2591521sec preferred_lft 604321sec
    inet6 fe80::56ee:75ff:fe2b:def6/64 scope link
        valid_lft forever preferred_lft forever
```

2. Verify the forward DNS configuration by using the **dig** utility and adding the host name.
 1. Run the **dig +short server.example.com A** command. The returned IPv4 address must match the IP address returned by **ip addr show**:

```
[root@server ~]# dig +short server.example.com A
192.0.2.1
```

2. Run the **dig +short server.example.com AAAA** command. If the command returns an address, it must match the IPv6 address returned by **ip addr show**:

```
[root@server ~]# dig +short server.example.com AAAA
2001:DB8::1111
```

**NOTE**

If no output is returned for the AAAA record, it does not indicate incorrect configuration; no output only means that no IPv6 address is configured in DNS for the server machine. If you do not intend to use the IPv6 protocol in your network, you can proceed with the installation in this situation.

3. Verify the reverse DNS configuration (PTR records) by using the **dig** utility and adding the IP address.

1. Run the **dig +short -x IPv4 address** command. The server host name must be displayed in the command output. For example:

```
[root@server ~]# dig +short -x 192.0.2.1
server.example.com
```

2. Use **dig** to query the IPv6 address as well if the **dig +short -x server.example.com AAAA** command in the previous step returned an IPv6 address. Again, the server host name must be displayed in the command output. For example:

```
[root@server ~]# dig +short -x 2001:DB8::1111
server.example.com
```

**NOTE**

If **dig +short server.example.com AAAA** in the previous step did not display any IPv6 address, querying the AAAA record does not output anything. In this case, this is normal behavior and does not indicate incorrect configuration.

If a different host name or no host name is displayed, even though **dig +short server.example.com** in the previous step returned an IP address, it indicates that the reverse DNS configuration is incorrect.

Verifying the Standards-compliance of DNS Forwarders

When configuring IdM with integrated DNS, verify that all DNS forwarders you want to use with the IdM DNS server comply with the [Extension Mechanisms for DNS \(EDNS0\)](#) and [DNS Security Extensions \(DNSSEC\)](#) standards. To do this, inspect the output of the following command for each forwarder separately:

```
$ dig +dnssec @IP_address_of_the_DNS_forwarder . SOA
```

The expected output displayed by the command contains the following information:

- status: **NOERROR**
- flags: **ra**
- EDNS flags: **do**
- The **RRSIG** record must be present in the **ANSWER** section

If any of these items is missing from the output, inspect the documentation of your DNS forwarder and verify that EDNS0 and DNSSEC are supported and enabled. In latest versions of the BIND server, the **dnssec-enable yes**; option must be set in the **/etc/named.conf** file.

For example, the expected output can look like this:

```
;; ->HEADER<- opcode: QUERY, status: NOERROR, id: 48655
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 2, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags: do; udp: 4096

;; ANSWER SECTION:
. 31679 IN SOA a.root-servers.net. nstld.verisign-grs.com. 2015100701 1800
900 604800 86400
. 31679 IN RRSIG SOA 8 0 86400 20151017170000 20151007160000 62530 .
GNVz7SQs [...]
```

The /etc/hosts File



IMPORTANT

Do not modify the **/etc/hosts** file manually. If **/etc/hosts** has been modified, make sure its contents conform to the following rules.

The following is an example of a correctly configured **/etc/hosts** file. It properly lists the IPv4 and IPv6 localhost entries for the host, followed by the IdM server IP address and host name as the first entry. Note that the IdM server host name cannot be part of the **localhost** entry.

```
127.0.0.1 localhost.localdomain localhost
::1 localhost6.localdomain6 localhost6
192.0.2.1 server.example.com server
2001:DB8::1111 server.example.com server
```

2.1.4. Port Requirements

IdM uses a number of ports to communicate with its services. These ports must be open and available for IdM to work. They cannot be in use by another service or blocked by a firewall.

- For a list of the required ports, see [the section called “List of Required Ports”](#).
- For a list of **firewalld** services that correspond to the required ports, see [the section called “List of firewalld Services”](#).

List of Required Ports

Table 2.1. Identity Management Ports

Service	Ports	Protocol
HTTP/HTTPS	80, 443	TCP
LDAP/LDAPS	389, 636	TCP
Kerberos	88, 464	TCP and UDP
DNS	53	TCP and UDP
NTP	123	UDP

NOTE

Do not be concerned that IdM uses ports 80 and 389.

- Port 80 (HTTP) is used to provide Online Certificate Status Protocol (OCSP) responses and Certificate Revocation Lists (CRL). Both are digitally signed and therefore secured against man-in-the-middle attacks.
- Port 389 (LDAP) uses STARTTLS and GSSAPI for encryption.

In addition, IdM can listen on port 8080 and in some installations also on ports 8443 and 749. However, these three ports are only used internally: even though IdM keeps them open, they are not required to be accessible from outside. It is recommended that you do not open ports 8080, 8443, and 749 and instead leave them blocked by a firewall.

List of firewalld Services**Table 2.2. firewalld Services**

Service name	For details, see:
freeipa-ldap	<code>/usr/lib/firewalld/services/freeipa-ldap.xml</code>
freeipa-ldaps	<code>/usr/lib/firewalld/services/freeipa-ldaps.xml</code>
dns	<code>/usr/lib/firewalld/services/dns.xml</code>

Opening the Required Ports

1. Make sure the **firewalld** service is running.

- To find out if **firewalld** is currently running:

```
# systemctl status firewalld.service
```

- To start **firewalld** and configure it to start automatically when the system boots:

```
# systemctl start firewalld.service
# systemctl enable firewalld.service
```

2. Open the required ports using the **firewall-cmd** utility. Choose one of the following options:

- a. Add the individual ports to the firewall by using the **firewall-cmd --add-port** command. For example, to open the ports in the default zone:

```
# firewall-cmd --permanent --add-port=
{80/tcp,443/tcp,list_of_ports}
```

- b. Add the **firewalld** services to the firewall by using the **firewall-cmd --add-service** command. For example, to open the ports in the default zone:

```
# firewall-cmd --permanent --add-service={freeipa-
ldap,list_of_services}
```

For details on using **firewall-cmd** to open ports on a system, see the [Security Guide](#) or the `firewall-cmd(1)` man page.

3. Reload the **firewall-cmd** configuration to ensure that the change takes place immediately:

```
# firewall-cmd --reload
```

Note that reloading **firewalld** on a system in production can cause DNS connection time outs. See also [Reloading the Firewall Using the Command-Line Interface](#) in the *Security Guide*. If required, to avoid the risk of time outs, repeat the commands without the **--permanent** option to apply the changes to the running system.

4. *Optional.* To verify that the ports are available now, use the **nc**, **telnet**, or **nmap** utilities to connect to a port or run a port scan.

2.2. PACKAGES REQUIRED TO INSTALL AN IDM SERVER

To install the packages required for a server without integrated DNS services:

```
# yum install ipa-server
```

To install the packages required for a server with integrated DNS services:

```
# yum install ipa-server ipa-server-dns
```




NOTE

To determine whether DNS is right for your use case, see [Section 2.3.1, “Determining Whether to Use Integrated DNS”](#).

The `ipa-server` package automatically installs other required packages as dependencies, such as:

- 389-ds-base for the Directory Server LDAP service
- krb5-server package for the Kerberos service
- various IdM-specific tools

2.3. INSTALLING AN IDM SERVER: INTRODUCTION



NOTE

The installation procedures and examples in the following sections are not mutually exclusive: you can combine them to achieve the required result. For example, you can install a server with integrated DNS and with an externally hosted root CA.

The **ipa-server-install** utility installs and configures an IdM server.

Before installing a server, see these sections:

- [Section 2.3.1, “Determining Whether to Use Integrated DNS”](#)
- [Section 2.3.2, “Determining What CA Configuration to Use”](#)

The **ipa-server-install** utility provides a non-interactive installation mode which allows automated and unattended server setup. For details, see [Section 2.3.7, “Installing a Server Non-Interactively”](#)

The **ipa-server-install** installation script creates a log file at **/var/log/ipaserver-install.log**. If the installation fails, the log can help you identify the problem.

2.3.1. Determining Whether to Use Integrated DNS

IdM supports installing a server with integrated DNS or without integrated DNS.

An IdM server with integrated DNS services

The integrated DNS server provided by IdM is not designed to be used as a general-purpose DNS server. It only supports features related to IdM deployment and maintenance. It does not support some of the advanced DNS features.

Red Hat strongly recommends IdM-integrated DNS for basic usage within the IdM deployment: When the IdM server also manages DNS, there is tight integration between DNS and native IdM tools which enables automating some of the DNS record management.

Note that even if an IdM server is used as a master DNS server, other external DNS servers can still be used as slave servers.

For example, if your environment is already using another DNS server, such as an Active Directory-integrated DNS server, you can delegate only the IdM primary domain to the IdM-integrated DNS. You are not required to migrate DNS zones over to the IdM-integrated DNS.


To install a server with integrated DNS, see [Section 2.3.3, “Installing a Server with Integrated DNS”](#)

An IdM server without integrated DNS services

An external DNS server is used to provide the DNS services. Consider installing an IdM server without DNS in these situations:

- If you require advanced DNS features beyond the scope of the IdM DNS
- In environments with a well-established DNS infrastructure which allows you to use external DNS servers

To install a server without integrated DNS, see [Section 2.3.4, “Installing a Server Without Integrated DNS”](#)



IMPORTANT

Make sure your system meets the DNS requirements described in [Section 2.1.3, “Host Name and DNS Configuration”](#).

Maintenance Requirements for Integrated or External DNS

When using an integrated DNS server, most of the DNS record maintenance is automated. You only must:

- set up correct delegation from the parent domain to the IdM servers

For example, if the IdM domain name is **ipa.example.com**, it must be properly delegated from the **example.com** domain.



NOTE

You can verify the delegation using the following command:

```
# dig @IP_address +norecurse +short ipa.example.com. NS
```

IP_address is the IP address of the server that manages the **example.com** DNS domain. If the delegation is correct, the command lists the IdM servers that have a DNS server installed.

When using an external DNS server, you must:

- manually create the new domain on the DNS server
- fill the new domain manually with records from the zone file that is generated by the IdM installer
- manually update the records after installing or removing a replica, as well as after any changes in the service configuration, such as after an Active Directory trust is configured

Preventing DNS Amplification Attacks

The default configuration of the IdM-integrated DNS server allows all clients to issue recursive queries to the DNS server. If your server is deployed in a network with an untrusted client, change the server's configuration to limit recursion to authorized clients only. ^[1]

To ensure that only authorized clients are allowed to issue recursive queries, add the appropriate access control list (ACL) statements to the `/etc/named.conf` file on your server. For example:

```
acl authorized { 192.0.2.0/24; 198.51.100.0/24; };
options {
    allow-query { any; };
    allow-recursion { authorized; };
};
```

2.3.2. Determining What CA Configuration to Use

IdM supports installing a server with an integrated IdM certificate authority (CA) or without a CA.

Server with an integrated IdM CA

This is the default configuration suitable for most deployments. Certificate System uses a *CA signing certificate* to create and sign the certificates in the IdM domain.



WARNING

Red Hat strongly recommends to keep the CA services installed on more than one server. For information on installing a replica of the initial server including the CA services, see [Section 4.5.4, “Installing a Replica with a CA”](#).

If you install the CA on only one server, you risk losing the CA configuration without a chance of recovery if the CA server fails. See [Section B.2.6, “Recovering a Lost CA Server”](#) for details.

The CA signing certificate must be signed by a *root CA*, which is the highest CA in the CA hierarchy. The root CA can be the IdM CA itself or an externally-hosted CA.

The IdM CA is the root CA

This is the default configuration.

To install a server with this configuration, see [Section 2.3.3, “Installing a Server with Integrated DNS”](#) and [Section 2.3.4, “Installing a Server Without Integrated DNS”](#).

An external CA is the root CA

The IdM CA is subordinate to an external CA. However, all certificates for the IdM domain are still issued by the Certificate System instance.

The external CA can be a corporate CA or a third-party CA, such as Verisign or Thawte. The certificates issued within the IdM domain are potentially subject to restrictions set by the external root CA for attributes like the validity period.

To install a server with an externally-hosted root CA, see [Section 2.3.5, “Installing a Server with an External CA as the Root CA”](#)

Server without a CA

This configuration option is suitable for very rare cases when restrictions within the infrastructure do not allow to install certificate services with the server.

You must request these certificates from a third-party authority prior to the installation:

- An LDAP server certificate and a private key
- An Apache server certificate and a private key
- Full CA certificate chain of the CA that issued the LDAP and Apache server certificates

Managing certificates without the integrated IdM CA presents a significant maintenance burden. Most notably:

- Creating, uploading, and renewing certificates is a manual process.
- The **certmonger** service is not used to track certificates. Therefore, it does not warn you of impending certificate expiration.

To install a server without an integrated CA, see [Section 2.3.6, “Installing Without a CA”](#)

2.3.3. Installing a Server with Integrated DNS



NOTE

If you are unsure what DNS or CA configuration is appropriate for you, see [Section 2.3.1, “Determining Whether to Use Integrated DNS”](#) and [Section 2.3.2, “Determining What CA Configuration to Use”](#).

To install a server with integrated DNS, provide the following information during the installation process:

DNS forwarders

The following DNS forwarder settings are supported:

- one or more forwarders (the **--forwarder** option in non-interactive installation)
- no forwarders (the **--no-forwarders** option in non-interactive installation)

If you are unsure whether to use DNS forwarding, see [Section 33.6, “Managing DNS Forwarding”](#).

Reverse DNS zones

The following reverse DNS zone settings are supported:

- automatic detection of the reverse zones that need to be created in IdM DNS (the default setting in interactive installation, the **--auto-reverse** option in non-interactive installation)
- no reverse zone auto-detection (the **--no-reverse** option in interactive installation)

Note that the **--allow-zone-overlap** option is ignored if the **--auto-reverse** option is set. Using the combination of options:

```
$ ipa-server-install --auto-reverse --allow-zone-overlap
```

thus does *not* create reverse zones which would overlap with already existing DNS zones, for example on another DNS server.

For non-interactive installation, add the **--setup-dns** option as well.

Example 2.1. Installing a Server with Integrated DNS

This procedure installs a server:

- with integrated DNS
- with integrated IdM CA as the root CA, which is the default CA configuration

1. Run the **ipa-server-install** utility.

```
# ipa-server-install
```

2. The script prompts to configure an integrated DNS service. Enter **yes**.

```
Do you want to configure integrated DNS (BIND)? [no]: yes
```

3. The script prompts for several required settings.
 - To accept the default values in brackets, press **Enter**.
 - To provide a value different than the proposed default value, enter the required value.

```
Server host name [server.example.com]:
Please confirm the domain name [example.com]:
Please provide a realm name [EXAMPLE.COM]:
```



WARNING

Red Hat strongly recommends that the Kerberos realm name is the same as the primary DNS domain name, with all letters uppercase. For example, if the primary DNS domain is **ipa.example.com**, use **IPA.EXAMPLE.COM** for the Kerberos realm name.

Different naming practices will prevent you from using Active Directory trusts and can have other negative consequences.

4. Enter the passwords for the Directory Server superuser, **cn=Directory Manager**, and for the **admin** IdM system user account.

```
Directory Manager password:
IPA admin password:
```

5. The script prompts for DNS forwarders.

```
Do you want to configure DNS forwarders? [yes]:
```

- To configure DNS forwarders, enter **yes**, and then follow the instructions on the command line.

The installation process will add the forwarder IP addresses to the **/etc/named.conf** file on the installed IdM server.

- For the forwarding policy default settings, see the **--forward-policy** description in the `ipa-dns-install(1)` man page.
 - See also [the section called “Forward Policies”](#) for details.
 - If you do not want to use DNS forwarding, enter **no**.
6. The script prompts to check if any DNS reverse (PTR) records for the IP addresses associated with the server need to be configured.

```
Do you want to search for missing reverse zones? [yes]:
```

If you run the search and missing reverse zones are discovered, the script asks you whether to create the reverse zones along with the PTR records.

```
Do you want to create reverse zone for IP 192.0.2.1 [yes]:
Please specify the reverse zone name [2.0.192.in-addr.arpa.]:
Using reverse zone(s) 2.0.192.in-addr.arpa.
```



NOTE

Using IdM to manage reverse zones is optional. You can use an external DNS service for this purpose instead.

7. Enter **yes** to confirm the server configuration.

```
Continue to configure the system with these values? [no]: yes
```

8. The installation script now configures the server. Wait for the operation to complete.
9. Add DNS delegation from the parent domain to the IdM DNS domain. For example, if the IdM DNS domain is **ipa.example.com**, add a name server (NS) record to the **example.com** parent domain.

**IMPORTANT**

This step must be repeated each time an IdM DNS server is installed.

The script recommends you to back up the CA certificate and to make sure the required network ports are open. For information about IdM port requirements and instructions on how to open these ports, see [Section 2.1.4, “Port Requirements”](#).

To test the new server:

1. Authenticate to the Kerberos realm using the admin credentials. This verifies that **admin** is properly configured and the Kerberos realm is accessible.

```
# kinit admin
```

2. Run a command such as **ipa user-find**. On a new server, the command prints the only configured user: **admin**.

```
# ipa user-find admin
-----
1 user matched
-----
User login: admin
Last name: Administrator
Home directory: /home/admin
Login shell: /bin/bash
UID: 939000000
GID: 939000000
Account disabled: False
Password: True
Kerberos keys available: True
-----
Number of entries returned 1
-----
```

2.3.4. Installing a Server Without Integrated DNS

**NOTE**

If you are unsure what DNS or CA configuration is appropriate for you, see [Section 2.3.1, “Determining Whether to Use Integrated DNS”](#) and [Section 2.3.2, “Determining What CA Configuration to Use”](#).

To install a server without integrated DNS, run the **ipa-server-install** utility without any DNS-related options.

Example 2.2. Installing a Server Without Integrated DNS

This procedure installs a server:

- without integrated DNS

- with integrated IdM CA as the root CA, which is the default CA configuration

1. Run the **ipa-server-install** utility.

```
# ipa-server-install
```

2. The script prompts to configure an integrated DNS service. Press **Enter** to select the default **no** option.

```
Do you want to configure integrated DNS (BIND)? [no]:
```

3. The script prompts for several required settings.

- To accept the default values in brackets, press **Enter**.
- To provide a value different than the proposed default value, enter the required value.

```
Server host name [server.example.com]:
Please confirm the domain name [example.com]:
Please provide a realm name [EXAMPLE.COM]:
```



WARNING

Red Hat strongly recommends that the Kerberos realm name is the same as the primary DNS domain name, with all letters uppercase. For example, if the primary DNS domain is **ipa.example.com**, use **IPA.EXAMPLE.COM** for the Kerberos realm name.

Different naming practices will prevent you from using Active Directory trusts and can have other negative consequences.

4. Enter the passwords for the Directory Server superuser, **cn=Directory Manager**, and for the **admin** IdM system user account.

```
Directory Manager password:
IPA admin password:
```

5. Enter **yes** to confirm the server configuration.

```
Continue to configure the system with these values? [no]: yes
```

6. The installation script now configures the server. Wait for the operation to complete.

7. The installation script produces a file with DNS resource records: the

`/tmp/ipa.system.records.UFRPto.db` file in the example output below. Add these records to the existing external DNS servers. The process of updating the DNS records varies depending on the particular DNS solution.

```
...
Restarting the KDC
Please add records in this file to your DNS system:
/tmp/ipa.system.records.UFRBto.db
Restarting the web server
...
```



IMPORTANT

The server installation is not complete until you add the DNS records to the existing DNS servers.

The script recommends you to back up the CA certificate and to make sure the required network ports are open. For information about IdM port requirements and instructions on how to open these ports, see [Section 2.1.4, “Port Requirements”](#).

To test the new server:

1. Authenticate to the Kerberos realm using the admin credentials. This verifies that **admin** is properly configured and the Kerberos realm is accessible.

```
# kinit admin
```

2. Run a command such as **ipa user-find**. On a new server, the command prints the only configured user: **admin**.

```
# ipa user-find admin
-----
1 user matched
-----
User login: admin
Last name: Administrator
Home directory: /home/admin
Login shell: /bin/bash
UID: 939000000
GID: 939000000
Account disabled: False
Password: True
Kerberos keys available: True
-----
Number of entries returned 1
-----
```

2.3.5. Installing a Server with an External CA as the Root CA

**NOTE**

If you are unsure what DNS or CA configuration is appropriate for you, see [Section 2.3.1, “Determining Whether to Use Integrated DNS”](#) and [Section 2.3.2, “Determining What CA Configuration to Use”](#).

To install a server and chain it with an external CA as the root CA, pass these options with the **ipa-server-install** utility:

- **--external-ca** specifies that you want to use an external CA.
- **--external-ca-type** specifies the type of the external CA. See the `ipa-server-install(1)` man page for details.

Otherwise, most of the installation procedure is the same as in [Section 2.3.3, “Installing a Server with Integrated DNS”](#) or [Section 2.3.4, “Installing a Server Without Integrated DNS”](#).

During the configuration of the Certificate System instance, the utility prints the location of the certificate signing request (CSR): **/root/ipa.csr**:

```
...
```

```
Configuring certificate server (pki-tomcatd): Estimated time 3 minutes 30 seconds
```

```
[1/8]: creating certificate server user
```

```
[2/8]: configuring certificate server instance
```

```
The next step is to get /root/ipa.csr signed by your CA and re-run  
/sbin/ipa-server-install as: /sbin/ipa-server-install --external-cert-  
file=/path/to/signed_certificate --external-cert-  
file=/path/to/external_ca_certificate
```


When this happens:

1. Submit the CSR located in **/root/ipa.csr** to the external CA. The process differs depending on the service to be used as the external CA.

**IMPORTANT**

It might be necessary to request the appropriate extensions for the certificate. The CA signing certificate generated for the Identity Management server must be a valid CA certificate. This requires either that the Basic Constraint be set to `CA=true` or that the Key Usage Extension be set on the signing certificate to allow it to sign certificates.

2. Retrieve the issued certificate and the CA certificate chain for the issuing CA in a base 64-encoded blob (either a PEM file or a Base_64 certificate from a Windows CA). Again, the process differs for every certificate service. Usually, a download link on a web page or in the notification email allows the administrator to download all the required certificates.

**IMPORTANT**

Be sure to get the full certificate chain for the CA, not just the CA certificate.

3. Run **ipa-server-install** again, this time specifying the locations and names of the newly-issued CA certificate and the CA chain files. For example:

```
# ipa-server-install --external-cert-
file=/tmp/servercert20110601.pem --external-cert-
file=/tmp/cacert.pem
```

**NOTE**

The **ipa-server-install --external-ca** command can sometimes fail with the following error:

```
ipa          : CRITICAL failed to configure ca instance Command
'/usr/sbin/pkispawn -s CA -f /tmp/configuration_file' returned
non-zero exit status 1
Configuration of CA failed
```


This failure occurs when the ***_proxy** environmental variables are set. For a solution on how to fix this problem, see [Section B.1.1, “External CA Installation Fails”](#)

2.3.6. Installing Without a CA

**NOTE**

If you are unsure what DNS or CA configuration is appropriate for you, see [Section 2.3.1, “Determining Whether to Use Integrated DNS”](#) and [Section 2.3.2, “Determining What CA Configuration to Use”](#).

To install a server without a CA, you must provide the required certificates manually by adding options to the **ipa-server-install** utility. Other than that, most of the installation procedure is the same as in [Section 2.3.3, “Installing a Server with Integrated DNS”](#) or [Section 2.3.4, “Installing a Server Without Integrated DNS”](#).

**IMPORTANT**

You cannot install a server or replica using self-signed third-party server certificates.

Certificates Required to Install an IdM Server without a CA

For a successful CA-less IdM server installation, you must provide the following certificates:

- The LDAP server certificate and private key, supplied using these options:
 - **--dirsrv-cert-file** for the certificate and private key files for the LDAP server certificate

- **--dirsrv-pin** for the password to access the private key in the files specified in **--dirsrv-cert-file**
- The Apache server certificate and private key, supplied using these options:
 - **--http-cert-file** for the certificate and private key files for the Apache server certificate
 - **--http-pin** for the password to access the private key in the files specified in **--http-cert-file**
- The full CA certificate chain of the CA that issued the LDAP and Apache server certificates, supplied using these options:
 - **--dirsrv-cert-file** and **--http-cert-file** for the certificate files with the full CA certificate chain or a part of it

The files provided using **--dirsrv-cert-file** and **--http-cert-file** must contain exactly one server certificate and exactly one private key. The contents of the files provided using **--dirsrv-cert-file** and **--http-cert-file** are often identical.

- If necessary, the certificate files to complete the full CA certificate chain, supplied using this option:
 - **--ca-cert-file**, which you can add this option multiple times
- Optionally, the certificate files to provide an external Kerberos key distribution center (KDC) PKINIT certificate, supplied using these options:
 - **--pkinit-cert-file** for the Kerberos KDC SSL certificate and private key
 - **--pkinit-pin** for the password to unlock the Kerberos KDC private key

If you do not provide the PKINIT certificate, **ipa-server-install** configures the IdM server with a local KDC with a self-signed certificate. For details, see [Chapter 27, Kerberos PKINIT Authentication in IdM](#).

The files provided using **--dirsrv-cert-file** and **--http-cert-file** combined with the files provided using **--ca-cert-file** must contain the full CA certificate chain of the CA that issued the LDAP and Apache server certificates.

For details on what the certificate file formats these options accept, see the `ipa-server-install(1)` man page.



NOTE

The listed command-line options are incompatible with the **--external-ca** option.

**NOTE**

Earlier versions of Identity Management used the **--root-ca-file** option to specify the PEM file of the root CA certificate. This is no longer necessary because the trusted CA is always the issuer of the DS and HTTP server certificates. IdM now automatically recognizes the root CA certificate from the certificates specified by **--dirsrv-cert-file**, **--http-cert-file**, and **--ca-cert-file**.

Example 2.3. Command example for installing an IdM server without a CA

```
[root@server ~]# ipa-server-install \
--http-cert-file /tmp/server.crt \
--http-cert-file /tmp/server.key \
--http-pin secret \
--dirsrv-cert-file /tmp/server.crt \
--dirsrv-cert-file /tmp/server.key \
--dirsrv-pin secret \
--ca-cert-file ca.crt
```

2.3.7. Installing a Server Non-Interactively**NOTE**

If you are unsure what DNS or CA configuration is appropriate for you, see [Section 2.3.1, “Determining Whether to Use Integrated DNS”](#) and [Section 2.3.2, “Determining What CA Configuration to Use”](#).

The minimum required options for a non-interactive installation are:

- **--ds-password** to provide the password for the Directory Manager (DM), the Directory Server super user
- **--admin-password** to provide the password for **admin**, the IdM administrator
- **--realm** to provide the Kerberos realm name
- **--unattended** to let the installation process select default options for the host name and domain name

Optionally, you can provide custom values for these settings:

- **--hostname** for the server host name
- **--domain** for the domain name

**WARNING**

Red Hat strongly recommends that the Kerberos realm name is the same as the primary DNS domain name, with all letters uppercase. For example, if the primary DNS domain is **ipa.example.com**, use **IPA.EXAMPLE.COM** for the Kerberos realm name.

Different naming practices will prevent you from using Active Directory trusts and can have other negative consequences.

For a complete list of options accepted by **ipa-server-install**, run the **ipa-server-install --help** command.

Example 2.4. Basic Installation without Interaction

1. Run the **ipa-server-install** utility, providing the required settings. For example, the following installs a server without integrated DNS and with an integrated CA:

```
# ipa-server-install --realm EXAMPLE.COM --ds-password
DM_password --admin-password admin_password --unattended
```

2. The setup script now configures the server. Wait for the operation to complete.
3. The installation script produces a file with DNS resource records: the **/tmp/ipa.system.records.UFRBto.db** file in the example output below. Add these records to the existing external DNS servers. The process of updating the DNS records varies depending on the particular DNS solution.

```
...
Restarting the KDC
Please add records in this file to your DNS system:
/tmp/ipa.system.records.UFRBto.db
Restarting the web server
...
```

**IMPORTANT**

The server installation is not complete until you add the DNS records to the existing DNS servers.

The script recommends you to back up the CA certificate and to make sure the required network ports are open. For information about IdM port requirements and instructions on how to open these ports, see [Section 2.1.4, “Port Requirements”](#).

To test the new server:

1. Authenticate to the Kerberos realm using the admin credentials. This verifies that **admin** is properly configured and the Kerberos realm is accessible.

```
# kinit admin
```

2. Run a command such as **ipa user-find**. On a new server, the command prints the only configured user: **admin**.

```
# ipa user-find admin
-----
1 user matched
-----
User login: admin
Last name: Administrator
Home directory: /home/admin
Login shell: /bin/bash
UID: 939000000
GID: 939000000
Account disabled: False
Password: True
Kerberos keys available: True
-----
Number of entries returned 1
-----
```

2.4. UNINSTALLING AN IDM SERVER



NOTE

At domain level **0**, the procedure is different. See [Section D.3.6, “Removing a Replica”](#).

Prerequisites

- Before uninstalling a server that serves as a certificate authority (CA), key recovery authority (KRA), or DNS Security Extensions (DNSSEC) server, make sure these services are running on another server in the domain.



WARNING

Removing the last replica that serves as a CA, KRA, or DNSSEC server can seriously disrupt the Identity Management functionality.

Procedure

To uninstall **server.example.com**:

1. On another server, use the **ipa server-del** command to delete **server.example.com** from the topology:

—

```
[root@another_server ~]# ipa server-del server.example.com
```

2. On **server.example.com**, use the **ipa-server-install --uninstall** command:

```
[root@server ~]# ipa-server-install --uninstall
```

3. Make sure all name server (NS) DNS records pointing to **server.example.com** are deleted from your DNS zones. This applies regardless of whether you use integrated DNS managed by IdM or external DNS.

2.5. RENAMING A SERVER

It is not possible to change the host name of an IdM server after it was set up. However, you can replace the server with a replica with a different name.

1. Create a new replica of the server, with a CA and with the new required host name or IP address. This is described in [Chapter 4, *Installing and Uninstalling Identity Management Replicas*](#).
2. Stop the initial IdM server instance.

```
[root@old_server ~]# ipactl stop
```

3. Verify that all other replicas and clients are working as before.
4. Uninstall the initial IdM server, as described in [Section 2.4, “Uninstalling an IdM Server”](#)

[1] For details, see the [DNS Amplification Attacks](#) page.

CHAPTER 3. INSTALLING AND UNINSTALLING IDENTITY MANAGEMENT CLIENTS

This chapter explains how to configure a system to join an Identity Management (IdM) domain as a client machine enrolled with a server.



NOTE

See [Section 1.2, “The Identity Management Domain”](#) for details on clients and servers in the IdM domain.

3.1. PREREQUISITES FOR INSTALLING A CLIENT

DNS requirements

Employ proper DNS delegation. For details on DNS requirements in IdM, see [Section 2.1.3, “Host Name and DNS Configuration”](#).

Do not alter the `resolv.conf` file on clients.

Port requirements

IdM clients connect to a number of ports on IdM servers to communicate with their services. These ports must be open *on the IdM servers* to work. For more information on which ports IdM requires, see [Section 2.1.4, “Port Requirements”](#).

On a client, open these ports in the outgoing direction. If you are using a firewall that does not filter outgoing packets, such as **firewalld**, the ports are already available in the outgoing direction.

Federal Information Processing Standard (FIPS) support

In environments set up using Red Hat Enterprise Linux 7.4 and later:

- You can configure a new IdM server, replica, or client on a system with the FIPS mode enabled. The installation script automatically detects a system with FIPS enabled and configures IdM without the administrator's intervention.

To enable FIPS in the operating system, see [Enabling FIPS Mode](#) in the *Security Guide*.



IMPORTANT

You cannot:

- Enable FIPS mode on existing IdM servers previously installed with FIPS mode disabled.
- Install a replica in FIPS mode when using an existing IdM server with FIPS mode disabled.

In environments set up using Red Hat Enterprise Linux 7.3 and earlier:

- IdM does not support the FIPS mode. Disable FIPS on your system before installing an IdM server, replica, or client, and do not enable it after the

installation.

For further details about FIPS mode, see [Federal Information Processing Standard \(FIPS\)](#) in the *Security Guide*.

Name Service Cache Daemon (NSCD) requirements

Red Hat recommends to disable NSCD on Identity Management machines. Alternatively, if disabling NSCD is not possible, only enable NSCD for maps that SSSD does not cache.

Both NSCD and the SSSD service perform caching, and problems can occur when systems use both services simultaneously. See the [System-Level Authentication Guide](#) for information on how to avoid conflicts between NSCD and SSSD.

3.2. PACKAGES REQUIRED TO INSTALL A CLIENT

Install the ipa-client package:

```
# yum install ipa-client
```

The ipa-client package automatically installs other required packages as dependencies, such as the System Security Services Daemon (SSSD) packages.

3.3. INSTALLING A CLIENT

The **ipa-client-install** utility installs and configures an IdM client. The installation process requires you to provide credentials that can be used to enroll the client. The following authentication methods are supported:

Credentials of a user authorized to enroll clients, such as admin

By default, **ipa-client-install** expects this option. See [Section 3.3.1, “Installing a Client Interactively”](#) for an example.

To provide the user credentials directly to **ipa-client-install**, use the **--principal** and **--password** options.

A random, one-time password pre-generated on the server

To use this authentication method, add the **--random** option to **ipa-client-install** option. See [Example 3.1, “Installing a Client Non-interactively Using a Random Password”](#).

A principal from a previous enrollment

To use this authentication method, add the **--keytab** option to **ipa-client-install**. See [Section 3.8, “Re-enrolling a Client into the IdM Domain”](#) for details.

See the ipa-client-install(1) man page for details.

The following sections document basic installation scenarios. For more details on using **ipa-client-install** and a complete list of the accepted options, see the ipa-client-install(1) man page.

3.3.1. Installing a Client Interactively

The following procedure installs a client while prompting the user for input when required. The user provides credentials of a user authorized to enroll clients into the domain, such as the **admin** user.

1. Run the **ipa-client-install** utility.

Add the **--enable-dns-updates** option to update the DNS records with the client machine's IP address if one of the following applies:

- the IdM server the client will be enrolled with was installed with integrated DNS
- the DNS server on the network accepts DNS entry updates with the GSS-TSIG protocol

Add the **--no-krb5-offline-passwords** option to disable storing Kerberos passwords in the SSSD cache.

2. The installation script attempts to obtain all the required settings automatically.
 - a. If your DNS zone and SRV records are set properly on your system, the script automatically discovers all the required values and prints them. Enter **yes** to confirm.

```
Client hostname: client.example.com
Realm: EXAMPLE.COM
DNS Domain: example.com
IPA Server: server.example.com
BaseDN: dc=example,dc=com
```

```
Continue to configure the system with these values? [no]: yes
```

If you want to install the system with different values, cancel the current installation. Then run **ipa-client-install** again, and specify the required values using command-line options.

For details, see the **DNS Autodiscovery** section in the `ipa-client-install(1)` man page.

- b. If the script fails to obtain some settings automatically, it prompts you for the values.



IMPORTANT

The fully qualified domain name must be a valid DNS name, which means only numbers, alphabetic characters, and hyphens (-) are allowed. Other characters, like underscores, in the host name cause DNS failures. Additionally, the host name must be all lower-case; no capital letters are allowed.

For other recommended naming practices, see the [Red Hat Enterprise Linux Security Guide](#).

3. The script prompts for a user whose identity will be used to enroll the client. By default, this is the **admin** user:


```
User authorized to enroll computers: admin
Password for admin@EXAMPLE.COM
```

4. The installation script now configures the client. Wait for the operation to complete.

```
Client configuration complete.
```

5. Run the **ipa-client-automount** utility, which automatically configures NFS for IdM. See [Section 34.2.1, “Configuring NFS Automatically”](#) for details.


3.3.2. Installing a Client Non-interactively

For a non-interactive installation, provide all required information to the **ipa-client-install** utility using command-line options. The minimum required options for a non-interactive installation are:

- options for specifying the credentials that will be used to enroll the client; see [Section 3.3, “Installing a Client”](#) for details
- **--unattended** to let the installation run without requiring user confirmation

If your DNS zone and SRV records are set properly on your system, the script automatically discovers all the other required values. If the script cannot discover the values automatically, provide them using command-line options.

- **--hostname** to specify a static host name for the client machine



IMPORTANT

The fully qualified domain name must be a valid DNS name, which means only numbers, alphabetic characters, and hyphens (-) are allowed. Other characters, like underscores, in the host name cause DNS failures. Additionally, the host name must be all lower-case; no capital letters are allowed.

For other recommended naming practices, see the [Red Hat Enterprise Linux Security Guide](#).

- **--server** to specify the host name of the IdM server the client will be enrolled with
- **--domain** to specify the DNS domain name of the IdM server the client will be enrolled with
- **--realm** to specify the Kerberos realm name

Add the **--enable-dns-updates** option to update the DNS records with the client machine's IP address if one of the following applies:

- the IdM server the client will be enrolled with was installed with integrated DNS
- the DNS server on the network accepts DNS entry updates with the GSS-TSIG protocol

Add the **--no-krb5-offline-passwords** option to disable storing Kerberos passwords in the SSSD cache.

For a complete list of options accepted by **ipa-client-install**, see the `ipa-client-install(1)` man page.

Example 3.1. Installing a Client Non-interactively Using a Random Password

This procedure installs a client without prompting the user for any input. The process includes pre-generating a random one-time password on the server that is used to authorize the enrollment.

1. On an existing server:

- a. Log in as the administrator:

```
$ kinit admin
```

- b. Add the new machine as an IdM host. Use the **--random** option with the **ipa host-add** command to generate the random password.

```
$ ipa host-add client.example.com --random
```

```
-----  
Added host "client.example.com"
```

```
-----  
Host name: client.example.com  
Random password: W5YpARl=7M.n  
Password: True  
Keytab: False  
Managed by: server.example.com
```

The generated password will become invalid after you use it to enroll the machine into the IdM domain. It will be replaced with a proper host keytab after the enrollment is finished.

2. On the machine where you want to install the client, run **ipa-client-install**, and use these options:
 - **--password** for the random password from the **ipa host-add** output



NOTE

The password often contains special characters. Therefore, enclose it in single quotes (').

- **--unattended** to let the installation run without requiring user confirmation

If your DNS zone and SRV records are set properly on your system, the script automatically discovers all the other required values. If the script cannot discover the values automatically, provide them using command-line options.

For example:

```
# ipa-client-install --password 'W5YpARl=7M.n' --domain  
example.com --server server.example.com --unattended
```

3. Run the **ipa-client-automount** utility, which automatically configures NFS for IdM. See [Section 34.2.1, “Configuring NFS Automatically”](#) for details.

3.4. SETTING UP AN IDM CLIENT THROUGH KICKSTART

A Kickstart enrollment automatically adds a new system to the IdM domain at the time Red Hat Enterprise Linux is installed. For details on Kickstart, see [Kickstart Installations](#) in the *Installation Guide*.

Preparing for a Kickstart client installation includes these steps:

1. [Section 3.4.1, “Pre-creating a Client Host Entry on the IdM Server”](#)
2. [Section 3.4.2, “Creating a Kickstart File for the Client”](#)

3.4.1. Pre-creating a Client Host Entry on the IdM Server

1. Log in as admin:

```
$ kinit admin
```

2. Create the host entry on the IdM server, and set a temporary password for the entry:

```
$ ipa host-add client.example.com --password=secret
```

The password is used by Kickstart to authenticate during the client installation and expires after the first authentication attempt. After the client is successfully installed, it authenticates using its keytab.

3.4.2. Creating a Kickstart File for the Client

A Kickstart file used to set up an IdM client must include the following:

- The **ipa-client** package in the list of packages to be installed:

```
%packages
@ X Window System
@ Desktop
@ Sound and Video
ipa-client
...
```

See [Package Selection](#) in the *Installation Guide* for details.

- Post-installation instructions that:
 - ensure SSH keys are generated before enrollment
 - runs the **ipa-client-install** utility, specifying:
 - all required information to access and configure the IdM domain services

- the password which you set when pre-creating the client host on the IdM server, in [Section 3.4.1, “Pre-creating a Client Host Entry on the IdM Server”](#)

For example:

```
%post --log=/root/ks-post.log

# Generate SSH keys to ensure that ipa-client-install uploads
them to the IdM server
/usr/sbin/sshd-keygen

# Run the client install script
/usr/sbin/ipa-client-install --hostname=client.example.com --
domain=EXAMPLE.COM --enable-dns-updates --mkhomedir -w secret --
realm=EXAMPLE.COM --server=server.example.com
```

For a non-interactive installation, add also the **--unattended** option.

To let the client installation script request a certificate for the machine:

- Add the **--request-cert** option to **ipa-client-install**.
- Set the system bus address to **/dev/null** for both the **getcert** and **ipa-client-install** utility in the kickstart **chroot** environment. To do this, add these lines to the post-installation instruction file before the **ipa-client-install** instruction:

```
# env DBUS_SYSTEM_BUS_ADDRESS=unix:path=/dev/null getcert list
# env DBUS_SYSTEM_BUS_ADDRESS=unix:path=/dev/null ipa-client-
install
```



NOTE

Red Hat recommends not to start the **sshd** service prior to the kickstart enrollment. While starting **sshd** before enrolling the client generates the SSH keys automatically, using the above script is the preferred solution.

See [Post-installation Script](#) in the *Installation Guide* for details.

For details on using Kickstart, see [How Do You Perform a Kickstart Installation?](#) in the *Installation Guide*. For examples of Kickstart files, see [Sample Kickstart Configurations](#).

3.5. POST-INSTALLATION CONSIDERATIONS FOR CLIENTS

3.5.1. Removing Pre-Identity Management Configuration

The **ipa-client-install** script does not remove any previous LDAP and SSSD configuration from the **/etc/openldap/ldap.conf** and **/etc/sss/sss.conf** files. If you modified the configuration in these files before installing the client, the script adds the new client values, but comments them out. For example:

```
BASE    dc=example,dc=com
```

```
URI      ldap://ldap.example.com

#URI ldaps://server.example.com # modified by IPA
#BASE dc=ipa,dc=example,dc=com # modified by IPA
```

To apply the new Identity Management configuration values:

1. Open **/etc/openldap/ldap.conf** and **/etc/sss/sss.conf**.
2. Delete the previous configuration.
3. Uncomment the new Identity Management configuration.
4. Server processes that rely on system-wide LDAP configuration might require a restart to apply the changes. Applications that use **openldap** libraries typically import the configuration when started.

3.6. TESTING THE NEW CLIENT

Check that the client can obtain information about users defined on the server. For example, to check the default **admin** user:

```
[user@client ~]$ id admin
uid=1254400000(admin) gid=1254400000(admins) groups=1254400000(admins)
```

3.7. UNINSTALLING A CLIENT

Uninstalling a client removes the client from the IdM domain, along with all of the IdM-specific configuration for system services, such as SSSD. This restores the client machine's previous configuration.

1. Run the **ipa-client-install --uninstall** command:

```
# ipa-client-install --uninstall
```

2. Remove the DNS entries for the client host manually from the server. See [Section 33.4.6, “Deleting Records from DNS Zones”](#).

3.8. RE-ENROLLING A CLIENT INTO THE IDM DOMAIN

If a client virtual machine has been destroyed and you still have its keytab, you can re-enroll the client:

- Interactively, using administrator credentials. See [Section 3.8.1, “Re-enrolling a Client Interactively Using the Administrator Account”](#).
- Non-interactively, using a previously backed-up keytab file. See [Section 3.8.2, “Re-enrolling a Client Non-interactively Using the Client Keytab”](#).

**NOTE**

You can only re-enroll clients whose domain entry is still active. If you uninstalled a client (using **ipa-client-install --uninstall**) or disabled its host entry (using **ipa host-disable**), you cannot re-enroll it.

During re-enrollment, IdM performs the following:

- Revokes the original host certificate
- Generates a new host certificate
- Creates new SSH keys
- Generates a new keytab

3.8.1. Re-enrolling a Client Interactively Using the Administrator Account

1. Re-create the client machine with the same host name.
2. Run the **ipa-client-install --force-join** command on the client machine:

```
# ipa-client-install --force-join
```

3. The script prompts for a user whose identity will be used to enroll the client. By default, this is the **admin** user:

```
User authorized to enroll computers: admin
Password for admin@EXAMPLE.COM
```

3.8.2. Re-enrolling a Client Non-interactively Using the Client Keytab

Re-enrollment using the client keytab is appropriate for automated installation or in other situations when using the administrator password is not feasible.

1. Back up the original client's keytab file, for example in the **/tmp** or **/root** directory.
2. Re-create the client machine with the same host name.
3. Re-enroll the client, and specify the keytab location using the **--keytab** option:

```
# ipa-client-install --keytab /tmp/krb5.keytab
```

**NOTE**

The keytab specified in the **--keytab** option is only used when authenticating to initiate the enrollment. During the re-enrollment, IdM generates a new keytab for the client.

3.9. RENAMING CLIENT MACHINES

This section explains how to rename an IdM client. The process involves:

- [the section called “Identifying Current Service and Keytab Configuration”](#).
- [the section called “Removing the Client Machine from the IdM Domain”](#).
- [the section called “Re-enrolling the Client with a New Host Name”](#).



WARNING

Renaming a client is a manual procedure. Red Hat does not recommend it unless changing the host name is absolutely required.

Identifying Current Service and Keytab Configuration

Before uninstalling the current client, make note of certain settings for the client. You will apply this configuration after re-enrolling the machine with a new host name.

1. Identify which services are running on the machine:
 - a. Use the **ipa service-find** command, and identify services with certificates in the output:

```
$ ipa service-find client.example.com
```

- b. In addition, each host has a default *host service* which does not appear in the **ipa service-find** output. The service principal for the host service, also called a *host principal*, is **host/client.example.com**.
2. Identify all host groups to which the machine belongs.

```
# ipa hostgroup-find client.example.com
```

3. For all service principals displayed by **ipa service-find client.example.com**, determine the location of the corresponding keytabs on **client.example.com**.

Each service on the client system has a Kerberos principal in the form *service_name/hostname@REALM*, such as **ldap/client.example.com@EXAMPLE.COM**.

Removing the Client Machine from the IdM Domain

1. Unenroll the client machine from the IdM domain. See [Section 3.7, “Uninstalling a Client”](#).
2. For each identified keytab other than **/etc/krb5.keytab**, remove the old principals:

```
[root@client ~]# ipa-rmkeytab -k /path/to/keytab -r EXAMPLE.COM
```

See [Section 29.4, “Removing Keytabs”](#).

3. On an IdM server, remove the host entry. This removes all services and revokes all certificates issued for that host:

```
[root@server ~]# ipa host-del client.example.com
```

At this point, the host is completely removed from IdM.

Re-enrolling the Client with a New Host Name

1. Rename the machine as required.
2. Re-enroll the machine as an IdM client. See [Section 3.8, “Re-enrolling a Client into the IdM Domain”](#).
3. On an IdM server, add a new keytab for every service identified in [the section called “Identifying Current Service and Keytab Configuration”](#).

```
[root@server ~]# ipa service-add service_name/new_host_name
```

4. Generate certificates for services that had a certificate assigned in [the section called “Identifying Current Service and Keytab Configuration”](#). You can do this:
 - Using the IdM administration tools. See [Chapter 24, Managing Certificates for Users, Hosts, and Services](#).
 - Using the **certmonger** utility. See [Working with certmonger](#) in the *System-Level Authentication Guide* or the `certmonger(8)` man page.
5. Re-add the client to the host groups identified in [the section called “Identifying Current Service and Keytab Configuration”](#). See [Section 13.3, “Adding and Removing User or Host Group Members”](#).

CHAPTER 4. INSTALLING AND UNINSTALLING IDENTITY MANAGEMENT REPLICAS

Replicas are created by cloning the configuration of existing Identity Management servers. Therefore, servers and their replicas share identical core configuration. The replica installation process copies the existing server configuration and installs the replica based on that configuration.

Maintaining several server replicas is a recommended backup solution to avoid data loss, as described in the ["Backup and Restore in IdM/IPA" Knowledgebase solution](#).



NOTE

Another backup solution, recommended primarily for situations when rebuilding the IdM deployment from replicas is not possible, is the **ipa-backup** utility, as described in [Chapter 9, Backing Up and Restoring Identity Management](#).

4.1. EXPLAINING IDM REPLICAS


Replicas are created as clones of the initial master servers. Once a replica is created, it is functionally identical to the master server: servers and replicas created from these servers share the same internal information about users, machines, certificates, and configured policies.



NOTE

For more information on the types of machines in the IdM topology, see [Section 1.2, "The Identity Management Domain"](#).

Replication is the process of copying data between replicas. The information between replicas is shared using *multi-master replication*: all replicas joined through a replication agreement receive updates and are therefore considered data masters.



RHEL_404973_0516

Figure 4.1. Server and Replica Agreements

4.2. DEPLOYMENT CONSIDERATIONS FOR REPLICAS

4.2.1. Distribution of Server Services in the Topology

IdM servers can run a number of services, such as a certificate authority (CA) or DNS. A replica can run the same services as the server it was created from, but it is not necessary.

For example, you can install a replica without DNS services, even if the initial server runs DNS. Similarly, you can set up a replica as a DNS server even if the initial server was installed without DNS.




Figure 4.2. Replicas with Different Services

CA Services on Replicas

If you set up a replica without a CA, it will forward all requests for certificate operations to the CA server in your topology.



WARNING

Red Hat strongly recommends to keep the CA services installed on more than one server. For information on installing a replica of the initial server including the CA services, see [Section 4.5.4, “Installing a Replica with a CA”](#).

If you install the CA on only one server, you risk losing the CA configuration without a chance of recovery if the CA server fails. See [Section B.2.6, “Recovering a Lost CA Server”](#) for details.

If you set up a CA on the replica, its configuration must mirror the CA configuration of the initial server.

- For example, if the server includes an integrated IdM CA as the root CA, the replica must also be installed with an integrated CA as the root CA.
- See [Section 2.3.2, “Determining What CA Configuration to Use”](#) for the supported CA configuration options.

4.2.2. Replica Topology Recommendations

Red Hat recommends to follow these guidelines:

Configure no more than 60 replicas in a single IdM domain

Red Hat guarantees to support environments with 60 replicas or less.

Configure *at least two*, but *no more than four* replication agreements per each replica

Configuring additional replication agreements ensures that information is replicated not just between the initial replica and the master server, but between other replicas as well.

- If you create replica B from server A and then replica C from server A, replicas B and C are not directly joined, so data from replica B must first be replicated to server A before propagating to replica C.




Figure 4.3. Replicas B and C Are Not Joined in a Replication Agreement

Setting up an additional replication agreement between replica B and replica C ensures the data is replicated directly, which improves data availability, consistency, failover tolerance, and performance.




Figure 4.4. Replicas B and C Are Joined in a Replication Agreement

See [Chapter 6, Managing Replication Topology](#) for details on managing replication agreements.

Configuring more than four replication agreements per replica is unnecessary. A large number of replication agreements per server does not bring significant additional benefits, because one consumer server can only be updated by one master at a time, so

the other agreements are meanwhile idle and waiting. Additionally, configuring too many replication agreements can have a negative impact on overall performance.



NOTE

The **ipa topologysuffix-verify** command checks if your topology meets the most important recommendations. Run **ipa topologysuffix-verify -help** for details.

The command requires you to specify the topology suffix. See [Section 6.1, “Explaining Replication Agreements, Topology Suffixes, and Topology Segments”](#) for details.




Figure 4.5. Topology Example

4.2.2.1. Tight Cell Topology

One of the most resilient topologies is to create a cell configuration for the servers and replicas with a small number of servers in a cell:

- Each of the cells is a *tight cell*, where all servers have replication agreements with each other.
- Each server has one replication agreement with another server *outside* the cell. This ensures that every cell is loosely coupled to every other cell in the domain.

To accomplish a tight cell topology:

- Have at least one IdM server in each main office, data center, or locality. Preferably, have two IdM servers.
- Do not have more than four servers per data center.
- In small offices, rather than using a replica, use SSSD to cache credentials and an off-site IdM server as the data back end.

4.3. PREREQUISITES FOR INSTALLING A REPLICA

The installation requirements for replicas are the same as for IdM servers. Make sure that the replica machine meets all of the prerequisites listed in [Section 2.1, “Prerequisites for Installing a Server”](#).

In addition to the general server requirements, you must also meet the following conditions:

The replica must be running the same or later version of IdM

For example, if the master server is running on Red Hat Enterprise Linux 7 and uses the IdM 4.4 packages, then the replica must also run on Red Hat Enterprise Linux 7 or later and use IdM version 4.4 or later. This ensures that configuration can be properly copied from the server to the replica.



IMPORTANT

IdM does not support creating a replica of an earlier version than the version of the master. If you try to create a replica using an earlier version, the installation fails.

The replica needs additional ports to be open

In addition to the standard IdM server port requirements described in [Section 2.1.4, “Port Requirements”](#), make sure you also meet the following:

- At domain level 0, keep the *TCP port 22* open during the replica setup process. This port is required in order to use SSH to connect to the master server.



NOTE

For details on domain levels, see [Chapter 7, Displaying and Raising the Domain Level](#).

- If one of the servers is running Red Hat Enterprise Linux 6 and has a CA installed, keep also *TCP port 7389* open during and after the replica configuration. In a purely Red Hat Enterprise Linux 7 environment, port 7389 is not required.

For information on how to open ports using the **firewall-cmd** utility, see [Section 2.1.4, “Port Requirements”](#).

4.4. PACKAGES REQUIRED TO INSTALL A REPLICA

Replica package requirements are the same as server package requirements. See [Section 2.2, “Packages Required to Install an IdM Server”](#).

4.5. CREATING THE REPLICA: INTRODUCTION

The **ipa-replica-install** utility is used to install a new replica from an existing IdM server.



NOTE

This chapter describes the simplified replica installation introduced in Red Hat Enterprise Linux 7.3. The procedures require domain level 1 (see [Chapter 7, Displaying and Raising the Domain Level](#)).

For documentation on installing a replica at domain level 0, see [Appendix D, Managing Replicas at Domain Level 0](#).

You can install a new replica:

- on an existing IdM client by *promoting* the client to a replica: see [the section called “Promoting an Existing Client to a Replica”](#)
- on a machine that has not yet been enrolled in the IdM domain: see [the section called “Installing a Replica on a Machine That Is Not a Client”](#)

In both of these situations, you can customize your replica by adding options to **ipa-replica-install**: see [the section called “Using ipa-replica-install to Configure the Replica for Your Use Case”](#).



IMPORTANT

If the IdM server you are replicating has a trust with Active Directory, set up the replica as a trust agent after running **ipa-replica-install**. See [Trust Controllers and Trust Agents](#) in the *Windows Integration Guide*.

Promoting an Existing Client to a Replica

To install the replica on an existing client, you must make sure the client is authorized to be promoted. To achieve this, choose one of the following:

Provide a privileged user's credentials

The default privileged user is **admin**. There are multiple ways to provide the user's credentials. You can:

- let IdM prompt you to get the credentials interactively



NOTE

This is the default way to provide the privileged user's credentials. If no credentials are available when **ipa-replica-install** runs, the installation automatically prompts you.

- log in as the user before running **ipa-replica-install** on the client:

```
$ kinit admin
```

- add the user's principal name and password to **ipa-replica-install** directly:

```
# ipa-replica-install --principal admin --admin-password  
admin_password
```

Add the client to the **ipaservers** host group

Membership in **ipaservers** grants the machine elevated privileges analogous to a privileged user's credentials. You will not be required to provide the user's credentials.

Example: [Section 4.5.1, “Promoting a Client to a Replica Using a Host Keytab”](#)

Installing a Replica on a Machine That Is Not a Client

When run on a machine that has not yet been enrolled in the IdM domain, **ipa-replica-install** first enrolls the machine as a client and then installs the replica components.

To install a replica in this situation, choose one of the following:

Provide a privileged user's credentials

The default privileged user is **admin**. To provide the credentials, add the principal name and password to **ipa-replica-install** directly:

```
# ipa-replica-install --principal admin --admin-password admin_password
```

Provide a random password for the client

You must generate the random password on a server before installing the replica. You will not be required to provide the user's credentials during the installation.

Example: [Section 4.5.2, “Installing a Replica Using a Random Password”](#)

By default, the replica is installed against the first IdM server discovered by the client installer. To install the replica against a particular server, add the following options to **ipa-replica-install**:

- **--server** for the server's fully qualified domain name (FQDN)
- **--domain** for the IdM DNS domain

Using **ipa-replica-install** to Configure the Replica for Your Use Case

When run without any options, **ipa-replica-install** only sets up basic server services. To install additional services, such as DNS or a certificate authority (CA), add options to **ipa-replica-install**.



WARNING

Red Hat strongly recommends to keep the CA services installed on more than one server. For information on installing a replica of the initial server including the CA services, see [Section 4.5.4, “Installing a Replica with a CA”](#).

If you install the CA on only one server, you risk losing the CA configuration without a chance of recovery if the CA server fails. See [Section B.2.6, “Recovering a Lost CA Server”](#) for details.

For example scenarios of installing a replica with the most notable options, see:

- [Section 4.5.3, “Installing a Replica with DNS”](#), using `--setup-dns` and `--forwarder`
- [Section 4.5.4, “Installing a Replica with a CA”](#), using `--setup-ca`
- [Section 4.5.5, “Installing a Replica from a Server without a CA”](#), using `--dirsrv-cert-file`, `--dirsrv-pin`, `--http-cert-file`, and `--http-pin`

For a complete list of the options used to configure the replica, see the `ipa-replica-install(1)` man page.

4.5.1. Promoting a Client to a Replica Using a Host Keytab

In this procedure, an existing IdM client is promoted to a replica using its own host keytab to authorize the promotion.

The procedure does not require you to provide the administrator or Directory Manager (DM) credentials. It is therefore more secure because no sensitive information is exposed on the command line.

1. On an existing server:

- a. Log in as the administrator.

```
$ kinit admin
```

- b. Add the client machine to the **ipaservers** host group.

```
$ ipa hostgroup-add-member ipaservers --hosts client.example.com
Host-group: ipaservers
Description: IPA server hosts
Member hosts: server.example.com, client.example.com
-----
Number of members added 1
-----
```

Membership in **ipaservers** grants the machine elevated privileges analogous to the administrator's credentials.

2. On the client, run the **ipa-replica-install** utility.

```
# ipa-replica-install
```

4.5.2. Installing a Replica Using a Random Password

In this procedure, a replica is installed from scratch on a machine that is not yet an IdM client. To authorize the enrollment, a client-specific random password valid for one client enrollment only is used.

The procedure does not require you to provide the administrator or Directory Manager (DM) credentials. It is therefore more secure because no sensitive information is exposed on the command line.

1. On an existing server:

- a. Log in as the administrator.

```
$ kinit admin
```

- b. Add the new machine as an IdM host. Use the **--random** option with the **ipa host-add** command to generate a random one-time password to be used for the replica installation.

```
$ ipa host-add client.example.com --random
```

```
-----  
Added host "client.example.com"  
-----
```

```
Host name: client.example.com  
Random password: W5YpARl=7M.n  
Password: True  
Keytab: False  
Managed by: server.example.com
```

The generated password will become invalid after you use it to enroll the machine into the IdM domain. It will be replaced with a proper host keytab after the enrollment is finished.

- c. Add the machine to the **ipaservers** host group.

```
$ ipa hostgroup-add-member ipaservers --hosts client.example.com  
Host-group: ipaservers  
Description: IPA server hosts  
Member hosts: server.example.com, client.example.com  
-----
```

```
Number of members added 1  
-----
```

Membership in **ipaservers** grants the machine elevated privileges required to set up the necessary server services.

2. On the machine where you want to install the replica, run **ipa-replica-install**, and provide the random password using the **--password** option. Enclose the password in single quotes (') because it often contains special characters:

```
# ipa-replica-install --password 'W5YpARl=7M.n'
```

4.5.3. Installing a Replica with DNS

This procedure works for installing a replica on a client as well as on a machine that is not part of the IdM domain yet. See [Section 4.5, “Creating the Replica: Introduction”](#) for details.

1. Run **ipa-replica-install** with these options:

- **--setup-dns** to create a DNS zone if it does not exist already and configure the replica as the DNS server
- **--forwarder** to specify a forwarder, or **--no-forwarder** if you do not want to use any forwarders

To specify multiple forwarders for failover reasons, use **--forwarder** multiple times.

For example:

```
# ipa-replica-install --setup-dns --forwarder 192.0.2.1
```



NOTE

The **ipa-replica-install** utility accepts a number of other options related to DNS settings, such as **--no-reverse** or **--no-host-dns**. For more information about them, see the `ipa-replica-install(1)` man page.

2. If the initial server was created with DNS enabled, the replica is automatically created with the proper DNS entries. The entries ensure that IdM clients will be able to discover the new server.

If the initial server did not have DNS enabled, add the DNS records manually. The following DNS records are necessary for the domain services:

- **_ldap._tcp**
- **_kerberos._tcp**
- **_kerberos._udp**
- **_kerberos-master._tcp**
- **_kerberos-master._udp**
- **_ntp._udp**
- **_kpasswd._tcp**
- **_kpasswd._udp**

This example shows how to verify that the entries are present:

a. Set the appropriate values for the **DOMAIN** and **NAMESERVER** variables:

■

```
# DOMAIN=example.com
# NAMESERVER=replica
```

- b. Use the following command to check for the DNS entries:

```
# for i in _ldap._tcp._kerberos._tcp._kerberos._udp._kerberos-
master._tcp._kerberos-master._udp._ntp._udp ; do
dig @${NAMESERVER} ${i}.${DOMAIN} srv +nocmd +noquestion
+nocomments +nostats +noaa +noadditional +noauthority
done | egrep "^_"

_ldap._tcp.example.com. 86400      IN      SRV      0 100 389
server1.example.com.
_ldap._tcp.example.com. 86400      IN      SRV      0 100 389
server2.example.com.
_kerberos._tcp.example.com. 86400 IN      SRV      0 100 88
server1.example.com.
...
```

3. *Optional, but recommended.* Manually add other DNS servers as backup servers in case the replica becomes unavailable. See [Section 33.11.1, “Setting up Additional Name Servers”](#). This is recommended especially for situations when the new replica is your first DNS server in the IdM domain.

4.5.4. Installing a Replica with a CA

This procedure works for installing a replica on a client as well as on a machine that is not part of the IdM domain yet. See [Section 4.5, “Creating the Replica: Introduction”](#) for details.

1. Run **ipa-replica-install** with the **--setup-ca** option.

```
[root@replica ~]# ipa-replica-install --setup-ca
```

2. The **--setup-ca** option copies the CA configuration from the initial server's configuration, regardless of whether the IdM CA on the server is a root CA or whether it is subordinated to an external CA.



NOTE

For details on the supported CA configurations, see [Section 2.3.2, “Determining What CA Configuration to Use”](#).

4.5.5. Installing a Replica from a Server without a CA

This procedure works for installing a replica on a client as well as on a machine that is not part of the IdM domain yet. See [Section 4.5, “Creating the Replica: Introduction”](#) for details.



IMPORTANT

You cannot install a server or replica using self-signed third-party server certificates.

- Run **ipa-replica-install**, and provide the required certificate files by adding

these options:

- **--dirsrv-cert-file**
- **--dirsrv-pin**
- **--http-cert-file**
- **--http-pin**

For details about the files that are provided using these options, see [Section 2.3.6, “Installing Without a CA”](#).

For example:

```
[root@replica ~]# ipa-replica-install \
--dirsrv-cert-file /tmp/server.crt \
--dirsrv-cert-file /tmp/server.key \
--dirsrv-pin secret \
--http-cert-file /tmp/server.crt \
--http-cert-file /tmp/server.key \
--http-pin secret
```



NOTE

Do not add the **--ca-cert-file** option. The **ipa-replica-install** utility takes this part of the certificate information automatically from the master server.

4.6. TESTING THE NEW REPLICA

To check if replication works as expected after creating a replica:

1. Create a user on one of the servers:

```
[admin@server1 ~]$ ipa user-add test_user --first=Test --last=User
```

2. Make sure the user is visible on the other server:

```
[admin@server2 ~]$ ipa user-show test_user
```

4.7. UNINSTALLING A REPLICA

See [Section 2.4, “Uninstalling an IdM Server”](#).

PART III. ADMINISTRATION: MANAGING SERVERS

CHAPTER 5. THE BASICS OF MANAGING THE IDM SERVER AND SERVICES

This chapter describes the Identity Management command-line and UI tools that are available to manage the IdM server and services, including methods for authenticating to IdM.

5.1. STARTING AND STOPPING THE IDM SERVER

A number of different services are installed together with an IdM server, including Directory Server, Certificate Authority (CA), DNS, Kerberos, and others. Use the **ipactl** utility to stop, start, or restart the entire IdM server along with all the installed services.

To start the entire IdM server:

```
# ipactl start
```

To stop the entire IdM server:

```
# ipactl stop
```

To restart the entire IdM server:

```
# ipactl restart
```

If you only want to stop, start, or restart an individual service, use the **systemctl** utility, described in the [System Administrator's Guide](#). For example, using **systemctl** to manage individual services is useful when customizing the Directory Server behavior: the configuration changes require restarting the Directory Server instance, but it is not necessary to restart all the IdM services.



IMPORTANT

To restart multiple IdM domain services, Red Hat always recommends to use **ipactl**. Because of dependencies between the services installed with the IdM server, the order in which they are started and stopped is critical. The **ipactl** utility ensures that the services are started and stopped in the appropriate order.


5.2. LOGGING INTO IDM USING KERBEROS

IdM uses the Kerberos protocol to support single sign-on. With Kerberos, the user only needs to present the correct user name and password once. Then the user can access IdM services without the system prompting for the credentials again.

By default, only machines that are members of the IdM domain can use Kerberos to authenticate to IdM. However, it is possible to configure external systems for Kerberos authentication as well; for more information, see [Section 5.4.4, “Configuring an External System for Kerberos Authentication to the Web UI”](#).

Using kinit

To log in to IdM from the command line, use the **kinit** utility.

**NOTE**

To use **kinit**, the **krb5-workstation** package must be installed.

When run without specifying a user name, **kinit** logs into IdM under the user name of the user that is currently logged-in on the local system. For example, if you are logged-in as **local_user** on the local system, running **kinit** attempts to authenticate you as the **local_user** IdM user:

```
[local_user@server ~]$ kinit
Password for local_user@EXAMPLE.COM:
```

**NOTE**

If the user name of the local user does not match any user entry in IdM, the authentication attempt fails.

To log in as a different IdM user, pass the required user name as a parameter to the **kinit** utility. For example, to log in as the **admin** user:

```
[local_user@server ~]$ kinit admin
Password for admin@EXAMPLE.COM:
```

Obtaining Kerberos Tickets Automatically

The **pam_krb5** pluggable authentication module (PAM) and SSSD can be configured to automatically obtain a TGT for a user after a successful login in to the desktop environment on an IdM client machine. This ensures that after logging in, the user is not required to run **kinit**.

On IdM systems that have IdM configured in SSSD as the identity and authentication provider, SSSD obtains the TGT automatically after the user logs in with the corresponding Kerberos principal name.

For information on configuring **pam_krb5**, see the **pam_krb5(8)** man page. For general information about PAM, see the [System-Level Authentication Guide](#).

Storing Multiple Kerberos Tickets

By default, Kerberos only stores one ticket per logged-in user in the credential cache. Whenever a user runs **kinit**, Kerberos overwrites the currently-stored ticket with the new ticket. For example, if you use **kinit** to authenticate as **user_A**, the ticket for **user_A** will be lost after you authenticate again as **user_B**.

To obtain and store another TGT for a user, set a different credential cache, which ensures the contents of the previous cache are not overwritten. You can do this in one of the following two ways:

- Run the **export KRB5CCNAME=path_to_different_cache** command, and then use **kinit** to obtain the ticket.
- Run the **kinit -c path_to_different_cache** command, and then reset the **KRB5CCNAME** variable.

To restore the original TGT stored in the default credential cache:

1. Run the **kdestroy** command.
2. Restore the default credential cache location using the **unset \$KRB5CCNAME** command.

Checking the Current Logged-in User

To verify what TGT is currently stored and used for authentication, use the **klist** utility to list cached tickets. In the following example, the cache contains a ticket for **user_A**, which means that only **user_A** is currently allowed to access IdM services:

```
$ klist
Ticket cache: KEYRING:persistent:0:0
Default principal: user_A@EXAMPLE.COM


Valid starting          Expires                Service principal
11/10/2015 08:35:45    11/10/2015 18:35:45    krbtgt/EXAMPLE.COM@EXAMPLE.COM
```

5.3. THE IDM COMMAND-LINE UTILITIES

The basic command-line script for IdM is named **ipa**. The **ipa** script is a parent script for a number of subcommands. These subcommands are then used to manage IdM. For example, the **ipa user-add** command adds a new user:

```
$ ipa user-add user_name
```

Command-line management has certain benefits over management in UI; for example, the command-line utilities allow management tasks to be automated and performed repeatedly in a consistent way without manual intervention. Additionally, while most management operations are available both from the command line and in the web UI, some tasks can only be performed from the command line.



NOTE

This section only provides a general overview of the **ipa** subcommands. More information is available in the other sections dedicated to specific areas of managing IdM. For example, for information about managing user entries using the **ipa** subcommands, see [Chapter 11, Managing User Accounts](#).

5.3.1. Getting Help for ipa Commands

The **ipa** script can display help about a particular set of subcommands: *atopic*. To display the list of available topics, use the **ipa help topics** command:

```
$ ipa help topics

automember          Auto Membership Rule.
automount           Automount
caacl               Manage CA ACL rules.
...
```

To display help for a particular topic, use the **ipa help topic_name** command. For example, to display information about the **automember** topic:


```
$ ipa help automember
```

Auto Membership Rule.

Bring clarity to the membership of hosts and users by configuring inclusive or exclusive regex patterns, you can automatically assign a new entries into a group or hostgroup based upon attribute information.

...

EXAMPLES:

```
Add the initial group or hostgroup:
  ipa hostgroup-add --desc="Web Servers" webservers
  ipa group-add --desc="Developers" devel
...
```

The **ipa** script can also display a list of available **ipa** commands. To do this, use the **ipa help commands** command:

```
$ ipa help commands
automember-add                Add an automember rule.
automember-add-condition      Add conditions to an automember
rule.
...
```

For detailed help on the individual **ipa** commands, add the **--help** option to a command. For example:

```
$ ipa automember-add --help

Usage: ipa [global-options] automember-add AUTOMEMBER-RULE [options]

Add an automember rule.
Options:
  -h, --help                show this help message and exit
  --desc=STR                 A description of this auto member rule
...
```

For more information about the **ipa** utility, see the **ipa(1)** man page.

5.3.2. Setting a List of Values

IdM stores entry attributes in lists. For example:

```
ipaUserSearchFields: uid,givenname,sn,telephonenumber,ou,title
```

Any update to a list of attributes overwrites the previous list. For example, an attempt to add a single attribute by only specifying this attribute replaces the whole previously-defined list with the single new attribute. Therefore, when changing a list of attributes, you must specify the whole updated list.

IdM supports the following methods of supplying a list of attributes:

- Using the same command-line argument multiple times within the same command invocation. For example:

```
$ ipa permission-add --permissions=read --permissions=write --
permissions=delete
```

- Enclosing the list in curly braces, which allows the shell to do the expansion. For example:

```
$ ipa permission-add --permissions={read,write,delete}
```

5.3.3. Using Special Characters

When passing command-line arguments in **ipa** commands that include special characters, such as angle brackets (< and >), ampersand (&), asterisk (*), or vertical bar (|), you must escape these characters by using a backslash (\). For example, to escape an asterisk (*):

```
$ ipa certprofile-show certificate_profile --out=exported\*profile.cfg
```

Commands containing unescaped special characters do not work as expected because the shell cannot properly parse such characters.

5.3.4. Searching IdM Entries

Listing IdM Entries

Use the **ipa *-find** commands to search for a particular type of IdM entries. For example:

- To list all users:

```
$ ipa user-find
-----
4 users matched
-----
...
```

- To list user groups whose specified attributes contain **keyword**:

```
$ ipa group-find keyword
-----
2 groups matched
-----
...
```

To configure the attributes IdM searches for users and user groups, see [Section 13.5, “Setting Search Attributes for Users and User Groups”](#).

When searching user groups, you can also limit the search results to groups that contain a particular user:

```
$ ipa group-find --user=user_name
```

You can also search for groups that do not contain a particular user:

```
$ ipa group-find --no-user=user_name
```

Showing Details for a Particular Entry

Use the **ipa *-show** command to display details about a particular IdM entry. For example:

```
$ ipa host-show server.example.com
Host name: server.example.com
Principal name: host/server.example.com@EXAMPLE.COM
...
```

5.3.4.1. Adjusting the Search Size and Time Limit

Some search results, such as viewing lists of users, can return a very large number of entries. By tuning these search operations, you can improve overall server performance when running the **ipa *-find** commands, such as **ipa user-find**, and when displaying corresponding lists in the web UI.

The search size limit:

- Defines the maximum number of entries returned for a request sent to the server from a client, the IdM command-line tools, or the IdM web UI.
- Default value: 100 entries.

The search time limit:

- Defines the maximum time that the server waits for searches to run. Once the search reaches this limit, the server stops the search and returns the entries that discovered in that time.
- Default value: 2 seconds.

If you set the values to **-1**, IdM will not apply any limits when searching.



IMPORTANT

Setting search size or time limits too high can negatively affect server performance.

Web UI: Adjusting the Search Size and Time Limit

To adjust the limits globally for all queries:

1. Select **IPA Server** → **Configuration**.
2. Set the required values in the **Search Options** area.
3. Click **Save** at the top of the page.

Command Line: Adjusting the Search Size and Time Limit

To adjust the limits globally for all queries, use the **ipa config-mod** command and add the **--searchrecordslimit** and **--searchtimelimit** options. For example:

```
$ ipa config-mod --searchrecordslimit=500 --searchtimelimit=5
```

From the command line, you can also adjust the limits only for a specific query. To do this, add the **--sizelimit** or **--timelimit** options to the command. For example:

```
$ ipa user-find --sizelimit=200 --timelimit=120
```

5.4. THE IDM WEB UI

The Identity Management web UI is a web application for IdM administration. It has most of the capabilities of the **ipa** command-line utility. Therefore, the users can choose whether they want to manage IdM from the UI or from the command line.



NOTE

Management operations available to the logged-in user depend on the user's access rights. For the **admin** user and other users with administrative privileges, all management tasks are available. For regular users, only a limited set of operations related to their own user account is available.

5.4.1. Supported Web Browsers

Identity Management supports the following browsers for connecting to the web UI:

- Mozilla Firefox 38 and later
- Google Chrome 46 and later

5.4.2. Accessing the Web UI and Authenticating

The web UI can be accessed both from IdM server and client machines, as well as from machines outside of the IdM domain. However, to access the UI from a non-domain machine, you must first configure the non-IdM system to be able to connect to the IdM Kerberos domain; see [Section 5.4.4, “Configuring an External System for Kerberos Authentication to the Web UI”](#) for more details.

5.4.2.1. Accessing the Web UI

To access the web UI, type the IdM server URL into the browser address bar:

```
https://server.example.com
```

This opens the IdM web UI login screen in your browser.




Figure 5.1. Web UI Login Screen

5.4.2.2. Available Login Methods

The user can authenticate to the web UI in the following ways:

With an active Kerberos ticket

If the user has a valid TGT obtained with the **kinit** utility, clicking **Login** automatically authenticates the user. Note that the browser must be configured properly to support Kerberos authentication.

For information on obtaining a Kerberos TGT, see [Section 5.2, “Logging into IdM Using Kerberos”](#). For information on configuring the browser, see [Section 5.4.3, “Configuring the Browser for Kerberos Authentication”](#).

By providing user name and password

To authenticate using a user name and password, enter the user name and password on the web UI login screen.

IdM also supports one-time password (OTP) authentication. For more information, see [Section 22.2, “One-Time Passwords”](#).

With a smart card

For more information, see [Section 23.6, “Authenticating to the Identity Management Web UI with a Smart Card”](#).

After the user authenticates successfully, the IdM management window opens.




Figure 5.2. The IdM Web UI Layout

5.4.2.3. Authenticating to the IdM Web UI as an AD User

Active Directory (AD) users can log in to the IdM web UI with their user name and password. In the web UI, AD users can perform only a limited set of operations related to their own user account, unlike IdM users who can perform management operations related to their administrative privileges.

To enable web UI login for AD users, the IdM administrator must define an ID override for each AD user in the Default Trust View. For example:

```
[admin@server ~]$ ipa idoverrideuser-add 'Default Trust View'
ad_user@ad.example.com
```

For details on ID views in AD, see [Using ID Views in Active Directory Environments](#) in the *Windows Integration Guide*.

5.4.3. Configuring the Browser for Kerberos Authentication

To enable authentication with Kerberos credentials, you must configure your browser to support Kerberos negotiation for accessing the IdM domain. Note that if your browser is not configured properly for Kerberos authentication, an error message appears after clicking **Login** on the IdM web UI login screen.




Figure 5.3. Kerberos Authentication Error

You can configure your browser for Kerberos authentication in three ways:

- Automatically from the IdM web UI. This option is only available for Firefox. See [the section called “Automatic Firefox Configuration in the Web UI”](#) for details.
- Automatically from the command line during the IdM client installation. This option is only available for Firefox. See [the section called “Automatic Firefox Configuration from the Command Line”](#) for details.
- Manually in the Firefox configuration settings. This option is available for all supported browsers. See [the section called “Manual Browser Configuration”](#) for details.

**NOTE**

The System-Level Authentication Guide includes a [troubleshooting guide for Kerberos authentication in Firefox](#). If Kerberos authentication is not working as expected, see this troubleshooting guide for more advice.

Automatic Firefox Configuration in the Web UI

To automatically configure Firefox from the IdM web UI:

1. Click the link for browser configuration on the web UI login screen.




Figure 5.4. Link to Configuring the Browser in the Web UI

2. Choose the link for Firefox configuration to open the Firefox configuration page.




Figure 5.5. Link to the Firefox Configuration Page

3. Follow the steps on the Firefox configuration page.

Automatic Firefox Configuration from the Command Line

Firefox can be configured from the command line during IdM client installation. To do this, use the **--configure-firefox** option when installing the IdM client with the **ipa-client-install** utility:

```
# ipa-client-install --configure-firefox
```

The **--configure-firefox** option creates a global configuration file with default Firefox settings that enable Kerberos for single sign-on (SSO).

Manual Browser Configuration

To manually configure your browser:

1. Click the link for browser configuration on the web UI login screen.




Figure 5.6. Link to Configuring the Browser in the Web UI

2. Choose the link for manual browser configuration.




Figure 5.7. Link to the Manual Configuration Page

3. Look for the instructions to configure your browser and follow the steps.

5.4.4. Configuring an External System for Kerberos Authentication to the Web UI

To enable Kerberos authentication to the web UI from a system that is not a member of the IdM domain, you must define an IdM-specific Kerberos configuration file on the external machine. Enabling Kerberos authentication on external systems is especially useful when your infrastructure includes multiple realms or overlapping domains.

To create the Kerberos configuration file:

1. Copy the `/etc/krb5.conf` file from the IdM server to the external machine. For example:

```
# scp /etc/krb5.conf
root@externalmachine.example.com:/etc/krb5_ipa.conf
```


**WARNING**

Do not overwrite the existing **krb5.conf** file on the external machine.

2. On the external machine, set the terminal session to use the copied IdM Kerberos configuration file:

```
$ export KRB5_CONFIG=/etc/krb5_ipa.conf
```

3. Configure the browser on the external machine as described in [Section 5.4.3, “Configuring the Browser for Kerberos Authentication”](#).

Users on the external system can now use the **kinit** utility to authenticate against the IdM server domain.

5.4.5. Proxy Servers and Port Forwarding in the Web UI

Using proxy servers to access the web UI does not require any additional configuration in IdM.

Port forwarding is not supported with the IdM server. However, because it is possible to use proxy servers, an operation similar to port forwarding can be configured using proxy forwarding with OpenSSH and the SOCKS option. This can be configured using the **-D** option of the **ssh** utility; for more information on using **-D**, see the `ssh(1)` man page.

CHAPTER 6. MANAGING REPLICATION TOPOLOGY

This chapter describes how to manage replication between servers in an Identity Management (IdM) domain.



NOTE

This chapter describes simplified topology management introduced in Red Hat Enterprise Linux 7.3. The procedures require domain level 1 (see [Chapter 7, *Displaying and Raising the Domain Level*](#)).

For documentation on managing topology at domain level 0, see [Section D.3, “Managing Replicas and Replication Agreements”](#).

For details on installing an initial replica and basic information on replication, see [Chapter 4, *Installing and Uninstalling Identity Management Replicas*](#).

6.1. EXPLAINING REPLICATION AGREEMENTS, TOPOLOGY SUFFIXES, AND TOPOLOGY SEGMENTS

Replication Agreements

Data stored on an IdM server is replicated based on replication agreements: when two servers have a replication agreement configured, they share their data.

Replication agreements are always bilateral: the data is replicated from the first replica to the other one as well as from the other replica to the first one.



NOTE

For additional details, see [Section 4.1, “Explaining IdM Replicas”](#).

Topology Suffixes

Topology suffixes store the data that is replicated. IdM supports two types of topology suffixes: **domain** and **ca**. Each suffix represents a separate back end, a separate replication topology.

When a replication agreement is configured, it joins two topology suffixes of the same type on two different servers.

The domain suffix: **dc=example,dc=com**

The **domain** suffix contains all domain-related data.

When two replicas have a replication agreement between their **domain** suffixes, they share directory data, such as users, groups, and policies.

The ca suffix: **o=ipaca**

The **ca** suffix contains data for the Certificate System component. It is only present on servers with a certificate authority (CA) installed.

When two replicas have a replication agreement between their **ca** suffixes, they share certificate data.




Figure 6.1. Topology Suffixes

An initial topology segment is set up between two servers by the **ipa-replica-install** script when installing a new replica.

Example 6.1. Viewing Topology Suffixes

The **ipa topologysuffix-find** command displays a list of topology suffixes:


```
$ ipa topologysuffix-find
-----
2 topology suffixes matched
-----
  Suffix name: ca
  Managed LDAP suffix DN: o=ipaca

  Suffix name: domain
  Managed LDAP suffix DN: dc=example,dc=com
-----
Number of entries returned 2
-----
```

Topology Segments

When two replicas have a replication agreement between their suffixes, the suffixes form a *topology segment*. Each topology segment consists of a *left node* and a *right node*. The nodes represent the servers joined in the replication agreement.

Topology segments in IdM are always bidirectional. Each segment represents two replication agreements: from server A to server B, and from server B to server A. The data is therefore replicated in both directions.



RHEL_404973_0916

Figure 6.2. Topology Segments**Example 6.2. Viewing Topology Segments**

The **ipa topologysegment-find** command shows the current topology segments configured for the domain or CA suffixes. For example, for the domain suffix:

```
$ ipa topologysegment-find
Suffix name: domain
-----
1 segment matched
-----
Segment name: server1.example.com-to-server2.example.com
Left node: server1.example.com
Right node: server2.example.com
Connectivity: both
-----
Number of entries returned 1
-----
```

In this example, domain-related data is only replicated between two servers: **server1.example.com** and **server1.example.com**.

To display details for a particular segment only, use the **ipa topologysegment-show** command:

```
$ ipa topologysegment-show
Suffix name: domain
Segment name: server1.example.com-to-server2.example.com
Segment name: server1.example.com-to-server2.example.com
Left node: server1.example.com
Right node: server2.example.com
Connectivity: both
```

6.2. WEB UI: USING THE TOPOLOGY GRAPH TO MANAGE REPLICATION TOPOLOGY

Accessing the Topology Graph

The topology graph in the web UI shows the relationships between the servers in the domain:

1. Select **IPA Server** → **Topology** → **Topology Graph**.
2. If you make any changes to the topology that are not immediately reflected in the graph, click **Refresh**.

Customizing the Topology View

You can move individual topology nodes by dragging the mouse:




Figure 6.3. Moving Topology Graph Nodes

You can zoom in and zoom out the topology graph using the mouse wheel:




Figure 6.4. Zooming the Topology Graph

You can move the canvas of the topology graph by holding the left mouse button:




Figure 6.5. Moving the Topology Graph Canvas

Interpreting the Topology Graph

Servers joined in a domain replication agreement are connected by an orange arrow. Servers joined in a CA replication agreement are connected by a blue arrow.

Topology graph example: recommended topology

Figure 6.6, “Recommended Topology Example” shows one of the possible recommended topologies for four servers: each server is connected to at least two other servers, and more than one server is a CA master.




Figure 6.6. Recommended Topology Example

Topology graph example: discouraged topology

In Figure 6.7, “Discouraged Topology Example: Single Point of Failure”, **server1** is a single point of failure. All the other servers have replication agreements with this server, but not with any of the other servers. Therefore, if **server1** fails, all the other servers will become isolated.

Avoid creating topologies like this.




Figure 6.7. Discouraged Topology Example: Single Point of Failure

For details on topology recommendations, see [Section 4.2, “Deployment Considerations for Replicas”](#).

6.2.1. Setting up Replication Between Two Servers

1. In the topology graph, hover your mouse over one of the server nodes.




Figure 6.8. Domain or CA Options

2. Click on the **domain** or the **ca** part of the circle depending on what type of topology segment you want to create.
3. A new arrow representing the new replication agreement appears under your mouse pointer. Move your mouse to the other server node, and click on it.




Figure 6.9. Creating a New Segment

4. In the **Add Topology Segment** window, click **Add** to confirm the properties of the new segment.

IdM creates a new topology segment between the two servers, which joins them in a replication agreement. The topology graph now shows the updated replication topology:




Figure 6.10. New Segment Created

6.2.2. Stopping Replication Between Two Servers

1. Click on an arrow representing the replication agreement you want to remove. This highlights the arrow.




Figure 6.11. Topology Segment Highlighted

2. Click **Delete**.
3. In the **Confirmation** window, click **OK**.

IdM removes the topology segment between the two servers, which deletes their replication agreement. The topology graph now shows the updated replication topology:




Figure 6.12. Topology Segment Deleted

6.3. COMMAND LINE: MANAGING TOPOLOGY USING THE IPA TOPOLOGY* COMMANDS

6.3.1. Getting Help for Topology Management Commands

To display all commands used to manage replication topology:

```
$ ipa help topology
```

To display detailed help for a particular command, run it with the **--help** option:

```
$ ipa topologysuffix-show --help
```

6.3.2. Setting up Replication Between Two Servers

1. Use the **ipa topologysegment-add** command to create a topology segment for the two servers. When prompted, provide:

- the required topology suffix: **domain** or **ca**



NOTE

If you want to create a segment between **ca** suffixes, both servers must have a CA installed. See [Section 26.8, “Installing a CA Into an Existing IdM Domain”](#).

- the left node and the right node, representing the two servers
- optionally, a custom name for the segment

For example:

```
$ ipa topologysegment-add
Suffix name: domain
Left node: server1.example.com
Right node: server2.example.com
Segment name [server1.example.com-to-server2.example.com]:
new_segment
-----
Added segment "new_segment"
-----
Segment name: new_segment
Left node: server1.example.com
Right node: server2.example.com
Connectivity: both
```

Adding the new segment joins the servers in a replication agreement.

2. *Optional.* Use the **ipa topologysegment-show** command to verify that the new segment is configured.

```
$ ipa topologysegment-show
Suffix name: domain
Segment name: new_segment
Segment name: new_segment
Left node: server1.example.com
Right node: server2.example.com
Connectivity: both
```

6.3.3. Stopping Replication Between Two Servers

1. To stop replication, you must delete the corresponding replication segment between the servers. To do that, you need to know the segment name.

If you do not know the name, use the **ipa topologysegment-find** command to display all segments, and locate the required segment in the output. When prompted, provide the required topology suffix: **domain** or **ca**. For example:

```
$ ipa topologysegment-find
Suffix name: domain
-----
8 segments matched
-----
Segment name: new_segment
Left node: server1.example.com
Right node: server2.example.com
Connectivity: both

...

-----
Number of entries returned 8
-----
```

2. Use the **ipa topologysegment-del** command to remove the topology segment joining the two servers.

```
$ ipa topologysegment-del
Suffix name: domain
Segment name: new_segment
-----
Deleted segment "new_segment"
-----
```

Deleting the segment removes the replication agreement.

3. *Optional.* Use the **ipa topologysegment-find** command to verify that the segment is no longer listed.

```
$ ipa topologysegment-find
Suffix name: domain
-----
7 segments matched
-----
Segment name: server2.example.com-to-server3.example.com
Left node: server2.example.com
Right node: server3.example.com
Connectivity: both

...

-----
Number of entries returned 7
-----
```

6.4. REMOVING A SERVER FROM THE TOPOLOGY

IdM does not allow removing a server from the topology if one of the following applies:

- the server being removed is the only server connecting other servers with the rest of the topology; this would cause the other servers to become isolated, which is not allowed

- the server being removed is your last CA or DNS server

In these situations, the attempt fails with an error. For example, on the command line:


```
$ ipa server-del
Server name: server1.example.com
Removing server1.example.com from replication topology, please wait...
ipa: ERROR: Server removal aborted:

Removal of 'server1.example.com' leads to disconnected topology in suffix
'domain':
Topology does not allow server server2.example.com to replicate with
servers:
    server3.example.com
    server4.example.com
...
```

6.4.1. Web UI: Removing a Server from the Topology

To remove a server from the topology without uninstalling the server components from the machine:

1. Select **IPA Server** → **Topology** → **IPA Servers**.
2. Click on the name of the server you want to delete.




☐	Server name	Min domain level	Max domain level	Managed suffixes
☐	server1.example.com	0	1	domain, ca
☐	server2.example.com	0	1	domain
☐	server3.example.com	0	1	domain, ca

Showing 1 to 3 of 3 entries.

Figure 6.13. Selecting a Server

3. Click **Delete Server**.

6.4.2. Command Line: Removing a Server from the Topology



IMPORTANT

Removing a server is an irreversible action. If you remove a server, the only way to introduce it back into the topology is to install a new replica on the machine.

To remove **server1.example.com**:

1. On another server, run the `ipa server-del` command to remove `server1.example.com`. The command removes all topology segments pointing to the server:

```
[user@server2 ~]$ ipa server-del
Server name: server1.example.com
Removing server1.example.com from replication topology, please
wait...
-----
Deleted IPA server "server1.example.com"
-----
```

2. On `server1.example.com`, run the `ipa server-install --uninstall` command to uninstall the server components from the machine.

```
[root@server1 ~]# ipa server-install --uninstall
```

6.5. MANAGING SERVER ROLES

Based on the services installed on an IdM server, it can perform various *server roles*. For example: CA server, DNS server, or key recovery authority (KRA) server.

6.5.1. Viewing Server Roles

Web UI: Viewing Server Roles

For a complete list of the supported server roles, see **IPA Server → Topology → Server Roles**.

- Role status **absent** means that no server in the topology is performing the role.
- Role status **enabled** means that one or more servers in the topology are performing the role.

Server Roles

Refresh

Role name	Role status
AD trust agent	absent
AD trust controller	absent
CA server	enabled

Figure 6.14. Server Roles in the Web UI

Command Line: Viewing Server Roles

The `ipa config-show` command displays all CA servers, NTP servers, and the current CA renewal master:

```
$ ipa config-show
...
```

```

IPA masters: server1.example.com, server2.example.com,
server3.example.com
IPA CA servers: server1.example.com, server2.example.com
IPA NTP servers: server1.example.com, server2.example.com,
server3.example.com
IPA CA renewal master: server1.example.com

```

The **ipa server-show** command displays a list of roles enabled on a particular server. For example, for a list of roles enabled on *server.example.com*:

```

$ ipa server-show
Server name: server.example.com
...
Enabled server roles: CA server, DNS server, NTP server, KRA server

```

The **ipa server-find --servrole** searches for all servers with a particular server role enabled. For example, to search for all CA servers:

```

$ ipa server-find --servrole "CA server"
-----
2 IPA servers matched
-----
Server name: server1.example.com
...
Server name: server2.example.com
...
-----
Number of entries returned 2
-----

```

6.5.2. Promoting a Replica to a Master CA Server



NOTE

This section describes changing the CA renewal master at domain level 1 (see [Chapter 7, Displaying and Raising the Domain Level](#)). For documentation on changing the CA renewal master at domain level 0, see [Section D.4, “Promoting a Replica to a Master CA Server”](#).

In a topology that includes multiple replicas, one of them acts as the master CA server: it manages the renewal of CA subsystem certificates and generates certificate revocation lists (CRLs). By default, the master CA is the initial server from which replicas were created.

If you plan to take the master CA server offline or decommission it, promote another CA server to take its place as the new CA renewal master:

1. Configure the replica to handle CA subsystem certificate renewal.
 - See [Section 6.5.2.1, “Changing the Current CA Renewal Master”](#) for domain level 1.
 - See [Section D.4.1, “Changing Which Server Handles Certificate Renewal”](#) for domain level 0.

2. Configure the replica to generate CRLs. See [Section 6.5.2.2, “Changing Which Server Generates CRLs”](#).
3. Before decommissioning the previous master CA server, make sure the new master works properly. See [Section 6.5.2.3, “Verifying That the New Master CA Server Is Configured Correctly”](#).

6.5.2.1. Changing the Current CA Renewal Master

Web UI: Changing the Current CA Renewal Master

1. Select **IPA Server** → **Configuration**.
2. In the **IPA CA renewal master** field, select the new CA renewal master.

Command Line: Changing the Current CA Renewal Master

Use the `ipa config-mod --ca-renewal-master-server` command:

```
$ ipa config-mod --ca-renewal-master-server
new_ca_renewal_master.example.com
...
IPA masters: old_ca_renewal_master.example.com,
new_ca_renewal_master.example.com
IPA CA servers: old_ca_renewal_master.example.com,
new_ca_renewal_master.example.com
IPA NTP servers: old_ca_renewal_master.example.com,
new_ca_renewal_master.example.com
IPA CA renewal master: new_ca_renewal_master.example.com
```

The output confirms that the update was successful.

6.5.2.2. Changing Which Server Generates CRLs

To change which server generates CRLs, stop CRL generation on the current CRL generation master, and then enable it on the other server.

Identifying the Current CRL Generation Master

Examine the `/etc/pki/pki-tomcat/ca/CS.cfg` file on each server with a CA installed:

- On the CRL generation master, the `ca.crl.MasterCRL.enableCRLUpdates` parameter is set to **true**:

```
# grep ca.crl.MasterCRL.enableCRLUpdates /etc/pki/pki-
tomcat/ca/CS.cfg
ca.crl.MasterCRL.enableCRLUpdates=true
```

- On CRL generation clones, the parameter is set to **false**.

Stopping CRL Generation on the Current CRL Generation Master

1. Stop the CA service:

```
# systemctl stop pki-tomcatd@pki-tomcat.service
```

2. Disable CRL generation on the server. Open the `/etc/pki/pki-tomcat/ca/CS.cfg` file, and set the values of the `ca.crl.MasterCRL.enableCRLCache` and `ca.crl.MasterCRL.enableCRLUpdates` parameters to **false**:

```
ca.crl.MasterCRL.enableCRLCache=false
ca.crl.MasterCRL.enableCRLUpdates=false
```

3. Start the CA service:

```
# systemctl start pki-tomcatd@pki-tomcat.service
```

4. Configure Apache to redirect CRL requests to the new master. Open the `/etc/httpd/conf.d/ipa-pki-proxy.conf` file, and uncomment the **RewriteRule** argument:

```
# Only enable this on servers that are not generating a CRL
RewriteRule ^/ipa/crl/MasterCRL.bin
https://server.example.com/ca/ee/ca/getCRL?
op=getCRL&crlIssuingPoint=MasterCRL [L,R=301,NC]
```

5. Restart Apache:

```
# systemctl restart httpd.service
```

Before, this server responded to CRL requests. Now, all CRL requests are routed to the previous CA master.

Configure a Server to Generate CRLs

1. Stop the CA service:

```
# systemctl stop pki-tomcatd@pki-tomcat.service
```

2. Enable CRL generation on the server. Set the values of the `ca.crl.MasterCRL.enableCRLCache` and `ca.crl.MasterCRL.enableCRLUpdates` parameters to **true**:

```
ca.crl.MasterCRL.enableCRLCache=true
ca.crl.MasterCRL.enableCRLUpdates=true
```

3. Start the CA service:

```
# systemctl start pki-tomcatd@pki-tomcat.service
```

4. Configure Apache to disable redirecting CRL requests. Open the `/etc/httpd/conf.d/ipa-pki-proxy.conf` file, and comment out the **RewriteRule** argument:

```
#RewriteRule ^/ipa/crl/MasterCRL.bin
https://server.example.com/ca/ee/ca/getCRL?
op=getCRL&crlIssuingPoint=MasterCRL [L,R=301,NC]
```


Before, all CRL requests were routed to the previous CA master. Now, this server will respond to CRL requests.

5. Restart Apache:

```
# systemctl restart httpd.service
```

6.5.2.3. Verifying That the New Master CA Server Is Configured Correctly

Make sure the `/var/lib/ipa/pki-ca/publish/MasterCRL.bin` file exists on the new master CA server.

The file is generated based on the time interval defined in the `/etc/pki/pki-tomcat/ca/CS.cfg` file using the `ca.crl.MasterCRL.autoUpdateInterval` parameter. The default value is 240 minutes (4 hours).

If the file exists, the new master CA server is configured correctly, and you can safely dismiss the previous CA master system.

CHAPTER 7. DISPLAYING AND RAISING THE DOMAIN LEVEL

The domain level indicates what operations and capabilities are available in the IdM topology.

Domain level 1

Examples of available functionality:

- simplified **ipa-replica-install** (see [Section 4.5, “Creating the Replica: Introduction”](#))
- enhanced topology management (see [Chapter 6, Managing Replication Topology](#))



IMPORTANT

Domain level 1 was introduced in Red Hat Enterprise Linux 7.3 with IdM version 4.4. To use the domain level 1 features, all your replicas must be running Red Hat Enterprise Linux 7.3 or later.

If your first server was installed with Red Hat Enterprise Linux 7.3, the domain level for your domain is automatically set to 1.

If you upgrade all servers to IdM version 4.4 from earlier versions, the domain level is not raised automatically. If you want to use domain level 1 features, raise the domain level manually, as described in [Section 7.2, “Raising the Domain Level”](#).

Domain level 0

Examples of available functionality:

- **ipa-replica-install** requires a more complicated process of creating a replica information file on the initial server and copying it to the replica (see [Section D.2, “Creating Replicas”](#))
- more complicated and error-prone topology management using **ipa-replica-manage** and **ipa-csreplica-manage** (see [Section D.3, “Managing Replicas and Replication Agreements”](#))

7.1. DISPLAYING THE CURRENT DOMAIN LEVEL

Command Line: Displaying the Current Domain Level

1. Log in as the administrator:

```
$ kinit admin
```

2. Run the **ipa domainlevel-get** command:

```
$ ipa domainlevel-get
-----
Current domain level: 0
-----
```

■

Web UI: Displaying the Current Domain Level

Select **IPA Server** → **Domain Level**.

7.2. RAISING THE DOMAIN LEVEL

**IMPORTANT**

This is a non-reversible operation. If you raise the domain level from **0** to **1**, you cannot downgrade from **1** to **0** again.

Command Line: Raising the Domain Level

1. Log in as the administrator:

```
$ kinit admin
```

2. Run the **ipa domainlevel-set** command and provide the required level:

```
$ ipa domainlevel-set 1
-----
Current domain level: 1
-----
```

Web UI: Raising the Domain Level

1. Select **IPA Server** → **Domain Level**.
2. Click **Set Domain Level**.

CHAPTER 8. UPDATING AND MIGRATING IDENTITY MANAGEMENT

8.1. UPDATING IDENTITY MANAGEMENT

You can use the **yum** utility to update the Identity Management packages on the system.

Additionally, if a new minor Red Hat Enterprise Linux version is available, such as 7.3, **yum** upgrades the Identity Management server or client to this version.




NOTE

This section does not describe migrating Identity Management from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7. If you want to migrate, see [Section 8.2, “Migrating Identity Management from Red Hat Enterprise Linux 6 to Version 7”](#).

8.1.1. Considerations for Updating Identity Management

- After you update the Identity Management packages on at least one server, all other servers in the topology receive the updated schema, even if you do not update their packages. This ensures that any new entries which use the new schema can be replicated among the other servers.
- Downgrading Identity Management packages is not supported.



IMPORTANT

Do not run the **yum downgrade** command on any of the `ipa-*` packages.

- Red Hat recommends upgrading to the next version only. For example, if you want to upgrade to Identity Management for Red Hat Enterprise Linux 7.4, we recommend upgrading from Identity Management for Red Hat Enterprise Linux 7.3. Upgrading from earlier versions can cause problems.

8.1.2. Using yum to Update the Identity Management Packages

To update all Identity Management packages on a server or client:

```
# yum update ipa-*
```



WARNING

When upgrading multiple Identity Management servers, wait at least 10 minutes between each upgrade.

When two or more servers are upgraded simultaneously or with only short intervals between the upgrades, there is not enough time to replicate the post-upgrade data changes throughout the topology, which can result in conflicting replication events.

Related Information

- For details on using the **yum** utility, see [Yum](#) in the *System Administrator's Guide*.

IMPORTANT

Due to [CVE-2014-3566](#), the Secure Socket Layer version 3 (SSLv3) protocol needs to be disabled in the **mod_nss** module. You can ensure that by following these steps:

1. Edit the **/etc/httpd/conf.d/nss.conf** file and set the **NSSProtocol** parameter to **TLSv1.0** (for backward compatibility), **TLSv1.1**, and **TLSv1.2**.

```
NSSProtocol TLSv1.0,TLSv1.1,TLSv1.2
```

2. Restart the **httpd** service.

```
# systemctl restart httpd.service
```

Note that Identity Management in Red Hat Enterprise Linux 7 automatically performs the above steps when the **yum update ipa-*** command is launched to upgrade the main packages.

8.2. MIGRATING IDENTITY MANAGEMENT FROM RED HAT ENTERPRISE LINUX 6 TO VERSION 7

This procedure describes how to migrate all data and configuration from Red Hat Enterprise Linux 6 Identity Management to Red Hat Enterprise Linux 7 servers. The migration procedure includes:

- Migrating the Red Hat Enterprise Linux 6-based certificate authority (CA) master server to Red Hat Enterprise Linux 7.
- Transitioning all services to the new Red Hat Enterprise Linux 7 server. These services include CRL and certificate creating, DNS management, or Kerberos KDC administration.
- Decommissioning the original Red Hat Enterprise Linux 6 CA master.

In the following procedures:

- **rhel7.example.com** is the Red Hat Enterprise Linux 7 system that will become the new CA master.
- **rhel6.example.com** is the original Red Hat Enterprise Linux 6 CA master.



NOTE

To identify which Red Hat Enterprise Linux 6 server is the master CA server, determine on which server the **certmonger** service tracks the **renew_ca_cert** command. Run this command on every Red Hat Enterprise Linux 6 server:

```
[root@rhel6 ~]# getcert list -d /var/lib/pki-ca/alias -n
"subsystemCert cert-pki-ca" | grep post-save
post-save command:
/usr/lib64/ipa/certmonger/renew_ca_cert "subsystemCert
cert-pki-ca"
```

The post-save action that executes **renew_ca_cert** is defined only for the CA master.

8.2.1. Prerequisites for Migrating Identity Management from Red Hat Enterprise Linux 6 to 7

- Update the **rhel6.example.com** system to the latest Red Hat Enterprise Linux 6 version.
- On the **rhel6.example.com** system, upgrade the ipa-* packages:

```
[root@rhel6 ~]# yum update ipa-*
```

This step also makes sure that you have applied the [RHBA-2015:0231-2](#) advisory, which provides the **2.3-6.el6_6** version of the bind-dyndb-ldap package and is available with the Red Hat Enterprise Linux 6.6 Extended Update Support (EUS).



WARNING

Using an earlier version of bind-dyndb-ldap results in inconsistent behavior in DNS forward zones serving between the Red Hat Enterprise Linux 6.6 DNS servers and Red Hat Enterprise Linux 7 DNS servers.

- Make sure the **rhel7.example.com** system meets the requirements in [Section 2.1, “Prerequisites for Installing a Server”](#) and [Section 4.3, “Prerequisites for Installing a Replica”](#).

- On the **rhel7.example.com** system, install the required packages. See [Section 2.2, “Packages Required to Install an IdM Server”](#).

8.2.2. Updating the Identity Management Schema on Red Hat Enterprise Linux 6

The **copy-schema-to-ca.py** schema update script prepares **rhel6.example.com** for the installation of the **rhel7.example.com** replica. Updating the schema is necessary due to schema changes between Identity Management version 3.1 and later versions.

1. Copy the **copy-schema-to-ca.py** schema update script from the **rhel7.example.com** system to the **rhel6.example.com** system. For example:

```
[root@rhel7 ~]# scp /usr/share/ipa/copy-schema-to-ca.py
root@rhel6:/root/
```

2. Run the updated **copy-schema-to-ca.py** script on **rhel6.example.com**.

```
[root@rhel6 ~]# python copy-schema-to-ca.py
ipa          : INFO      Installed /etc/dirsrv/slapd-PKI-
IPA//schema/60kerberos.ldif
[... output truncated ...]
ipa          : INFO      Schema updated successfully
```

8.2.3. Installing the Red Hat Enterprise Linux 7 Replica

1. On the **rhel6.example.com** system, create the replica file you will use to install the **rhel7.example.com** replica. For example, to create a replica file for **rhel7.example.com** whose IP address is **192.0.2.1**:

```
[root@rhel6 ~]# ipa-replica-prepare rhel7.example.com --ip-address
192.0.2.1

Directory Manager (existing master) password:
Preparing replica for rhel7.example.com from rhel6.example.com
[... output truncated ...]
The ipa-replica-prepare command was successful
```

See also [Section D.1, “Replica Information File”](#) and [Section D.2, “Creating Replicas”](#).

2. Copy the replica information file from **rhel6.example.com** to **rhel7.example.com**.

```
[root@rhel6 ~]# scp /var/lib/ipa/replica-info-
replica.example.com.gpg root@rhel7:/var/lib/ipa/
```

3. Install the **rhel7.example.com** replica using the replica file. For example, the following command uses these options:
 - **--setup-ca** to set up the Certificate System component
 - **--setup-dns** and **--forwarder** to configure an integrated DNS server and set a forwarder
 - **--ip-address** to specify the IP address of the **rhel7.example.com** system

```
[root@rhel7 ~]# ipa-replica-install /var/lib/ipa/replica-info-
rhel7.example.com.gpg --setup-ca --ip-address 192.0.2.1 --setup-dns
--forwarder 192.0.2.20
Directory Manager (existing master) password:

Checking DNS forwarders, please wait ...
Run connection check to master
[... output truncated ...]
Client configuration complete.
```

See also:

- [Section D.2, “Creating Replicas”](#), which describes creating replicas using replica information files
- [Section 2.3.1, “Determining Whether to Use Integrated DNS”](#) and [Section 2.3.2, “Determining What CA Configuration to Use”](#)

4. Verify that the Identity Management services are running on **rhel7.example.com**.

```
[root@rhel7 ~]# ipactl status
Directory Service: RUNNING
[... output truncated ...]
ipa: INFO: The ipactl command was successful
```

8.2.4. Transitioning the CA Services to the Red Hat Enterprise Linux 7 Server

Before you begin:

- Verify that **rhel6.example.com** and **rhel7.example.com** CAs are both configured as master servers.

```
[root@rhel7 ~]$ kinit admin
[root@rhel7 ~]$ ipa-csreplica-manage list
rhel6.example.com: master
rhel7.example.com: master
```

To display details about a replication agreement:

```
[root@rhel7 ~]# ipa-csreplica-manage list --verbose
rhel7.example.com
rhel7.example.com
last init status: None
last init ended: 1970-01-01 00:00:00+00:00
last update status: Error (0) Replica acquired successfully:
Incremental update succeeded
last update ended: 2017-02-13 13:55:13+00:00
```

On the **rhel6.example.com** original master CA, stop the CA subsystem certificate renewal:

1. Disable tracking for the original CA certificates.

```
[root@rhel6 ~]# getcert stop-tracking -d /var/lib/pki-ca/alias -n
```



```
"auditSigningCert cert-pki-ca"
Request "20181127184547" removed.
[root@rhel6 ~]# getcert stop-tracking -d /var/lib/pki-ca/alias -n
"ocspSigningCert cert-pki-ca"
Request "20181127184548" removed.
[root@rhel6 ~]# getcert stop-tracking -d /var/lib/pki-ca/alias -n
"subsystemCert cert-pki-ca"
Request "20181127184549" removed.
[root@rhel6 ~]# getcert stop-tracking -d /etc/httpd/alias -n ipaCert
Request "20181127184550" removed.
```

2. Reconfigure **rhel6.example.com** to retrieve renewed certificates from a new master CA.

- a. Copy the renewal helper script into the **certmonger** service directory, and set the appropriate permissions.

```
[root@rhel6 ~]# cp /usr/share/ipa/ca_renewal
/var/lib/certmonger/cas/
[root@rhel6 ~]# chmod 0600 /var/lib/certmonger/cas/ca_renewal
```

- b. Update the SELinux configuration.

```
[root@rhel6 ~]# restorecon /var/lib/certmonger/cas/ca_renewal
```

- c. Restart **certmonger**.

```
[root@rhel6 ~]# service certmonger restart
```

- d. Check that the CA is listed to retrieve certificates.

```
[root@rhel6 ~]# getcert list-cas
...
CA 'dogtag-ipa-retrieve-agent-submit':
    is-default: no
    ca-type: EXTERNAL
    helper-location: /usr/libexec/certmonger/dogtag-ipa-retrieve-
agent-submit
```

- e. Obtain the CA certificate database PIN.

```
[root@rhel6 ~]# grep internal= /var/lib/pki-ca/conf/password.conf
```

- f. Configure **certmonger** to track the certificates for external renewal. This requires the database PIN.

```
[root@rhel6 ~]# getcert start-tracking \
-c dogtag-ipa-retrieve-agent-submit \
-d /var/lib/pki-ca/alias \
-n "auditSigningCert cert-pki-ca" \
-B /usr/lib64/ipa/certmonger/stop_pkicad \
-C '/usr/lib64/ipa/certmonger/restart_pkicad \
"auditSigningCert cert-pki-ca"' \
```

```

-T "auditSigningCert cert-pki-ca" \
-P database_pin
New tracking request "20181127184743" added.
[root@rhel6 ~]# getcert start-tracking \
-c dogtag-ipa-retrieve-agent-submit \
-d /var/lib/pki-ca/alias \
-n "ocspSigningCert cert-pki-ca" \
-B /usr/lib64/ipa/certmonger/stop_pkicad \
-C '/usr/lib64/ipa/certmonger/restart_pkicad \
"ocspSigningCert cert-pki-ca"' \
-T "ocspSigningCert cert-pki-ca" \
-P database_pin
New tracking request "20181127184744" added.
[root@rhel6 ~]# getcert start-tracking \
-c dogtag-ipa-retrieve-agent-submit \
-d /var/lib/pki-ca/alias \
-n "subsystemCert cert-pki-ca" \
-B /usr/lib64/ipa/certmonger/stop_pkicad \
-C '/usr/lib64/ipa/certmonger/restart_pkicad \
"subsystemCert cert-pki-ca"' \
-T "subsystemCert cert-pki-ca" \
-P database_pin
New tracking request "20181127184745" added.
[root@rhel6 ~]# getcert start-tracking \
-c dogtag-ipa-retrieve-agent-submit \
-d /etc/httpd/alias \
-n ipaCert \
-C /usr/lib64/ipa/certmonger/restart_httpd \
-T ipaCert \
-p /etc/httpd/alias/pwdfile.txt
New tracking request "20181127184746" added.

```

Move CRL generation from the original **rhel6.example.com** CA master to **rhel7.example.com**.

1. On **rhel6.example.com**, stop CRL generation:

- a. Stop the CA service.

```
[root@rhel6 ~]# service pki-cad stop
```

- b. Disable CRL generation on **rhel6.example.com**. Open the **/var/lib/pki-ca/conf/CS.cfg** file, and set the values of the **ca.crl.MasterCRL.enableCRLCache** and **ca.crl.MasterCRL.enableCRLUpdates** parameters to **false**.

```
ca.crl.MasterCRL.enableCRLCache=false
ca.crl.MasterCRL.enableCRLUpdates=false
```

- c. Start the CA service.

```
[root@rhel6 ~]# service pki-cad start
```

2. On **rhel6.example.com**, configure Apache to redirect CRL requests to the new master, **rhel7.example.com**.
 - a. Open the **/etc/httpd/conf.d/ipa-pki-proxy.conf** file. Uncomment the **RewriteRule** argument, and replace the server host name with the **rhel7.example.com** host name in the server URL:

```
RewriteRule ^/ipa/crl/MasterCRL.bin
https://rhel7.example.com/ca/ee/ca/getCRL?
op=getCRL&crlIssuingPoint=MasterCRL [L,R=301,NC]
```

- b. Restart Apache.

```
[root@rhel6 ~]# service httpd restart
```

3. On **rhel7.example.com**, configure **rhel7.example.com** as the new CA master:
 - a. Configure **rhel7.example.com** to handle CA subsystem certificate renewal, as described in [Section D.4.1, “Changing Which Server Handles Certificate Renewal”](#).
 - b. Configure **rhel7.example.com** to general certificate revocation lists (CRLs), as described in [the section called “Configure a Server to Generate CRLs”](#).

Related Information

- See [Section 6.5.2, “Promoting a Replica to a Master CA Server”](#) for details on CA subsystem certificate renewal and CRLs.

8.2.5. Stop the Red Hat Enterprise Linux 6 Server

Stop all service on **rhel6.example.com** to force domain discovery to the new **rhel7.example.com** server.


```
[root@rhel6 ~]# ipactl stop
Stopping CA Service
Stopping pki-ca: [ OK ]
Stopping HTTP Service
Stopping httpd: [ OK ]
Stopping MEMCACHE Service
Stopping ipa_memcached: [ OK ]
Stopping DNS Service
Stopping named: . [ OK ]
Stopping KPASSWD Service
Stopping Kerberos 5 Admin Server: [ OK ]
Stopping KDC Service
Stopping Kerberos 5 KDC: [ OK ]
Stopping Directory Service
Shutting down dirsrv:
    EXAMPLE-COM... [ OK ]
    PKI-IPA... [ OK ]
```

After this, using the **ipa** utility will contact the new server through a remote procedure call (RPC).

8.2.6. Next Steps After Migrating the Master CA Server

For each Red Hat Enterprise Linux 6 server in your topology:

1. Create a replica file from **rhel7.example.com**.



NOTE

After installing a Red Hat Enterprise Linux 7 replica from a Red Hat Enterprise Linux 6 server, the domain level for the Identity Management domain is automatically set to 0.

Red Hat Enterprise Linux 7.3 introduced an easier way to install and manage replicas. To use these features, your topology must be at domain level 1. See [Chapter 7, *Displaying and Raising the Domain Level*](#).

2. Use the replica file to install a new replica on another Red Hat Enterprise Linux 7 system.

See [Chapter 4, *Installing and Uninstalling Identity Management Replicas*](#).

To decommission a Red Hat Enterprise Linux 6 server:

- Remove the server from the topology by executing the removal commands on a Red Hat Enterprise Linux 7 server.

See [Section 2.4, “Uninstalling an IdM Server”](#).

CHAPTER 9. BACKING UP AND RESTORING IDENTITY MANAGEMENT

Red Hat Enterprise Linux Identity Management provides a solution to manually back up and restore the IdM system, for example when a server stops performing correctly or data loss occurs. During backup, the system creates a directory containing information on your IdM setup and stores it. During restore, you can use this backup directory to bring your original IdM setup back.




IMPORTANT

Use the backup and restore procedures described in this chapter only if you cannot rebuild the lost part of the IdM server group from the remaining servers in the deployment, by reinstalling the lost replicas as replicas of the remaining ones.

The ["Backup and Restore in IdM/IPA" Knowledgebase solution](#) describes how to avoid losses by maintaining several server replicas. Rebuilding from an existing replica with the same data is preferable, because the backed-up version usually contains older, thus potentially outdated, information.

The potential threat scenarios that backup and restore can prevent include:

- Catastrophic hardware failure on a machine occurs and the machine becomes incapable of further functioning. In this situation, you can reinstall the operating system from scratch, configure the machine with the same fully qualified domain name (FQDN) and host name, install the IdM packages as well as all other optional packages relating to IdM that were present on the original system, and restore the full backup of the IdM server.
- An upgrade on an isolated machine fails. The operating system remains functional, but the IdM data is corrupted, which is why you want to restore the IdM system to a known good state.



IMPORTANT

In cases of hardware or upgrade failure, such as the two mentioned above, restore from backup only if all replicas or a replica with a special role, such as the only certificate authority (CA), were lost. If a replica with the same data still exists, it is recommended to delete the lost replica and then rebuild it from the remaining one.

- Undesirable changes were made to the LDAP content, for example entries were deleted, and you want to revert them. Restoring backed-up LDAP data returns the LDAP entries to the previous state without affecting the IdM system itself.

The restored server becomes the only source of information for IdM; other master servers are re-initialized from the restored server. Any data created after the last backup was made are lost. Therefore you should not use the backup and restore solution for normal system maintenance. If possible, always rebuild the lost server by reinstalling it as a replica.

The backup and restore features can be managed only from the command line and are not available in the IdM web UI.

9.1. FULL-SERVER BACKUP AND DATA-ONLY BACKUP

IdM offers two backup options:

Full-IdM server backup

Full-server backup creates a backup copy of all the IdM server files as well as LDAP data, which makes it a standalone backup. IdM affects hundreds of files; the files that the backup process copies is a mix of whole directories and specific files, such as configuration files or log files, and relate directly to IdM or to various services that IdM depends on. Because the full-server backup is a raw file backup, it is performed offline. The script that performs the full-server backup stops all IdM services to ensure a safe course of the backup process.

For the full list of files and directories that the full-server backup copies, see [Section 9.1.3, “List of Directories and Files Copied During Backup”](#).

Data-only Backup

The data-only backup only creates a backup copy of LDAP data and the changelog. The process backs up the **IPA-REALM** instance and can also back up multiple back ends or only a single back end; the back ends include the **IPA** back end and the **CA Dogtag** back end. This type of backup also backs up a record of the LDAP content stored in LDIF (LDAP Data Interchange Format). The data-only backup can be performed both online and offline.

By default, IdM stores the created backups in the `/var/lib/ipa/backup/` directory. The naming conventions for the subdirectories containing the backups are:

- **ipa-full-YEAR-MM-DD-HH-MM-SS** in the GMT time zone for the full-server backup
- **ipa-data-YEAR-MM-DD-HH-MM-SS** in the GMT time zone for the data-only backup

9.1.1. Creating a Backup

Both full-server and data-only backups are created using the **ipa-backup** utility which must always be run as root.

To create a full-server backup, run **ipa-backup**.



IMPORTANT

Performing a full-server backup stops all IdM services because the process must run offline. The IdM services will start again after the backup is finished.

To create a data-only backup, run the **ipa-backup --data** command.

You can add several additional options to **ipa-backup**:

- **--online** performs an online backup; this option is only available with data-only backups
- **--logs** includes the IdM service log files in the backup

For further information on using **ipa-backup**, see the `ipa-backup(1)` man page.

9.1.2. Encrypting Backup

You can encrypt the IdM backup using the GNU Privacy Guard (GPG) encryption.

To create a GPG key:

1. Create a **keygen** file containing the key details, for example, by running **cat >keygen <<EOF** and providing the required encryption details to the file from the command line:

```
[root@server ~]# cat >keygen <<EOF
> %echo Generating a standard key
> Key-Type: RSA
> Key-Length:2048
> Name-Real: IPA Backup
> Name-Comment: IPA Backup
> Name-Email: root@example.com
> Expire-Date: 0
> %pubring /root/backup.pub
> %secring /root/backup.sec
> %commit
> %echo done
> EOF
[root@server ~]#
```

2. Generate a new key pair called **backup** and feed the contents of **keygen** to the command. The following example generates a key pair with the path names **/root/backup.sec** and **/root/backup.pub**:

```
[root@server ~]# gpg --batch --gen-key keygen
[root@server ~]# gpg --no-default-keyring --secret-keyring
/root/backup.sec \
    --keyring /root/backup.pub --list-secret-keys
```

To create a GPG-encrypted backup, pass the generated **backup** key to **ipa-backup** by supplying the following options:

- **--gpg**, which instructs **ipa-backup** to perform the encrypted backup
- **--gpg-keyring=GPG_KEYRING**, which provides the full path to the GPG keyring without the file extension.

For example:

```
[root@server ~]# ipa-backup --gpg --gpg-keyring=/root/backup
```



NOTE

You might experience problems if your system uses the **gpg2** utility to generate GPG keys because **gpg2** requires an external program to function. To generate the key purely from console in this situation, add the **pinentry-program /usr/bin/pinentry-curses** line to the **.gnupg/gpg-agent.conf** file before generating a key.

9.1.3. List of Directories and Files Copied During Backup

Directories:

```
/usr/share/ipa/html
/root/.pki
/etc/pki-ca
/etc/pki/pki-tomcat
/etc/sysconfig/pki
/etc/httpd/alias
/var/lib/pki
/var/lib/pki-ca
/var/lib/ipa/sysrestore
/var/lib/ipa-client/sysrestore
/var/lib/ipa/dnssec
/var/lib/sss/pubconf/krb5.include.d/
/var/lib/authconfig/last
/var/lib/certmonger
/var/lib/ipa
/var/run/dirsrv
/var/lock/dirsrv
```

Files:

```
/etc/named.conf
/etc/named.keytab
/etc/resolv.conf
/etc/sysconfig/pki-ca
/etc/sysconfig/pki-tomcat
/etc/sysconfig/dirsrv
/etc/sysconfig/ntpd
/etc/sysconfig/krb5kdc
/etc/sysconfig/pki/ca/pki-ca
/etc/sysconfig/ipa-dnskeysyncd
/etc/sysconfig/ipa-ods-exporter
/etc/sysconfig/named
/etc/sysconfig/ods
/etc/sysconfig/authconfig
/etc/ipa/nssdb/pwdfile.txt
/etc/pki/ca-trust/source/ipa.p11-kit
/etc/pki/ca-trust/source/anchors/ipa-ca.crt
/etc/nsswitch.conf
/etc/krb5.keytab
/etc/sss/sss.conf
/etc/openldap/ldap.conf
/etc/security/limits.conf
/etc/httpd/conf/password.conf
/etc/httpd/conf/ipa.keytab
/etc/httpd/conf.d/ipa-pki-proxy.conf
/etc/httpd/conf.d/ipa-rewrite.conf
/etc/httpd/conf.d/nss.conf
/etc/httpd/conf.d/ipa.conf
/etc/ssh/sshd_config
/etc/ssh/ssh_config
/etc/krb5.conf
/etc/ipa/ca.crt
```



```
/etc/ipa/default.conf
/etc/dirsrv/ds.keytab
/etc/ntp.conf
/etc/samba/smb.conf
/etc/samba/samba.keytab
/root/ca-agent.p12
/root/cacert.p12
/var/kerberos/krb5kdc/kdc.conf
/etc/systemd/system/multi-user.target.wants/ipa.service
/etc/systemd/system/multi-user.target.wants/sss.service
/etc/systemd/system/multi-user.target.wants/certmonger.service
/etc/systemd/system/pki-tomcatd.target.wants/pki-tomcatd@pki-
tomcat.service
/var/run/ipa/services.list
/etc/openssl/conf.xml
/etc/openssl/kasp.xml
/etc/ipa/dnssec/softhsm2.conf
/etc/ipa/dnssec/softhsm_pin_so
/etc/ipa/dnssec/ipa-ods-exporter.keytab
/etc/ipa/dnssec/ipa-dnskeysyncd.keytab
/etc/idm/nssdb/cert8.db
/etc/idm/nssdb/key3.db
/etc/idm/nssdb/secmod.db
/etc/ipa/nssdb/cert8.db
/etc/ipa/nssdb/key3.db
/etc/ipa/nssdb/secmod.db
```

Log files and directories:

```
/var/log/pki-ca
/var/log/pki/
/var/log/dirsrv/slapd-PKI-IPA
/var/log/httpd
/var/log/ipaserver-install.log
/var/log/kadmind.log
/var/log/pki-ca-install.log
/var/log/messages
/var/log/ipaclient-install.log
/var/log/secure
/var/log/ipaserver-uninstall.log
/var/log/pki-ca-uninstall.log
/var/log/ipaclient-uninstall.log
/var/named/data/named.run
```

9.2. RESTORING A BACKUP

If you have a directory with a backup created using **ipa-backup**, you can restore your IdM server or the LDAP content to the state in which they were when the backup was performed. You cannot restore a backup on a host different from the host on which the backup was originally created.

**NOTE**

Uninstalling an IdM server does not automatically remove the backup of this server.

9.2.1. Restoring from the Full-Server or Data-Only Backup

**IMPORTANT**

It is recommended that you uninstall a server before performing a full-server restore on it.

Both full-server and data-only backups are restored using the **ipa-restore** utility which must always be run as root. Pass the backup to the command:

- Pass only the name of the directory with the backup if it is located in the default **/var/lib/ipa/backup/** directory.
- Pass the full path to the backup if the directory containing the backup is not located in the default directory. For example:

```
[root@server ~]# ipa-restore /path/to/backup
```

The **ipa-restore** utility automatically detects what type of backup the backup directory contains and by default performs the same type of restore.

You can add the following options to **ipa-restore**:

- **--data** performs a data-only restore from a full-server backup, that is, restores only the LDAP data component from a backup directory containing the full-server backup
- **--online** restores the LDAP data in a data-only restore online
- **--instance** specifies which 389 DS instance is restored. IdM in Red Hat Enterprise Linux 7 only uses the **IPA-REALM** instance, but it might be possible, for example, to create a backup on a system with separate instances; in such cases, **--instance** allows you to restore only **IPA-REALM**. For example:

```
[root@server ~]# ipa-restore --instance=IPA-REALM /path/to/backup
```

You can use this option only when performing a data-only restore.

- **--backend** specifies which back end is restored; without this option, **ipa-restore** restores all back ends it discovers. The arguments defining the possible back ends are **userRoot**, which restores the IPA data back end, and **ipaca**, which restores the CA back end.

You can use this option only when performing a data-only restore.

- **--no-logs** restores the backup without restoring the log files

To avoid authentication problems on an IdM master, clear the SSSD cache after a restore:

1. Stop the SSSD service:

```
[root@server ~]# systemctl stop sssd
```

2. Remove all cached content from SSSD:

```
[root@server ~]# find /var/lib/sss/ ! -type d | xargs rm -f
```

3. Start the SSSD service:

```
[root@server ~]# systemctl start sssd
```

**NOTE**

It is recommended that you reboot your system after restoring from backup.

For further information on using **ipa-restore**, see the `ipa-restore(1)` man page.

9.2.2. Restoring with Multiple Master Servers

Restoring from backup sets the restored server as the new data master, and you will be required to reinitialize all other masters after the restore. To reinitialize the other masters, run the **ipa-replica-manage** command and, on masters that have a CA installed, the **ipa-csreplica-manage** command. For example:

```
[root@server ~]# ipa-replica-manage re-initialize --  
from=restored_master_FQDN
```

For further information on replication during restore and on restoration on other masters, see the `ipa-restore(1)` man page.

9.2.3. Restoring from an Encrypted Backup

If you want to restore from a backup encrypted with GPG, provide the full path to the private and public keys using the **--gpg-keyring** option. For example:

```
[root@server ~]# ipa-restore --gpg-keyring=/root/backup /path/to/backup
```

CHAPTER 10. DEFINING ACCESS CONTROL FOR IDM USERS

Access control is a set of security features which defines who can access certain resources, such as machines, services or entries, and what kinds of operations they are allowed to perform. Identity Management provides several access control areas to make it clear what kind of access is being granted and to whom it is granted. As part of this, Identity Management draws a distinction between access controls to resources within the domain and access control to the IdM configuration itself.

This chapter details the different internal access control mechanisms that are available for users within IdM to the IdM server and other IdM users.

10.1. ACCESS CONTROLS FOR IDM ENTRIES

Access control defines the rights or permissions users have been granted to perform operations on other users or objects.

The Identity Management access control structure is based on standard LDAP access controls. Access within the IdM server is based on the IdM users, stored in the back end Directory Server instance, who are allowed to access other IdM entities, also stored as LDAP entries in the Directory Server instance.

An access control instruction (ACI) has three parts:

Actor

This is the entity who is being granted permission to do something. In LDAP access control models, this is called the *bind rule* because it defines who the user is and can optionally require other limits on the bind attempt, such as restricting attempts to a certain time of day or a certain machine.

Target

This defines the entry which the actor is allowed to perform operations on.

Operation type

Operation type — the last part determines what kinds of actions the user is allowed to perform. The most common operations are add, delete, write, read, and search. In Identity Management, all users are implicitly granted read and search rights to all entries in the IdM domain, with restrictions only for sensitive attributes like passwords and Kerberos keys. Anonymous users are restricted from seeing security-related configuration, like **sudo** rules and host-based access control.

When any operation is attempted, the first thing that the IdM client does is send user credentials as part of the bind operation. The back end Directory Server checks those user credentials and then checks the user account to see if the user has permission to perform the requested operation.

10.1.1. Access Control Methods in Identity Management

To make access control rules simple and clear to implement, Identity Management divides access control definitions into three categories:

Self-service rules

Self-service rules, which define what operations a user can perform on his own personal entry. The access control type only allows write permissions to attributes within the entry; it does not allow add or delete operations for the entry itself.

Delegation rules

Delegation rules, which allow a specific user group to perform write (edit) operations on specific attributes for users in another user group. Like self-service rules, this form of access control rule is limited to editing the values of specific attributes; it does not grant the ability to add or remove whole entries or control over unspecified attributes.

Role-based access control

Role-based access control, which creates special access control groups which are then granted much broader authority over all types of entities in the IdM domain. Roles can be granted edit, add, and delete rights, meaning they can be granted complete control over entire entries, not just selected attributes.

Some roles are already created and available within Identity Management. Special roles can be created to manage any type of entry in specific ways, such as hosts, automount configuration, netgroups, DNS settings, and IdM configuration.

10.2. DEFINING SELF-SERVICE SETTINGS

Self-service access control rules define the operations that an entity can perform on itself. These rules define only what attributes a user (or other IdM entity) can edit on their personal entries.

10.2.1. Creating Self-Service Rules from the Web UI

1. On the **IPA Server** tab in the top menu, select the **Role-Based Access Control** → **Self Service Permissions** subtab.
2. Click **Add** at the top of the list of the self-service access control instructions.




Figure 10.1. Adding a Current Self-Service Rule

3. Enter the name of the rule in the pop-up window. Spaces are allowed.

Add Self Service Permission

Self-service *

Attributes *

☐ audio	☐ businesscategory
☐ carlicense	☐ cn
☐ departmentnumber	☐ description
☐ homedirectory	☐ homephone
☐ homepostaladdress	☐ inetuserhttpurl
☐ inetuserstatus	☒ initials
☐ internationalisdnnnumber	☐ ipasshpubkey
☐ ipatokenradiusconfiglink	☐ ipatokenradiususername
☐ ipauniqueid	☐ ipauserauthtype
☒ jpegphoto	☐ krbcanonicalname

* Required field

Figure 10.2. Form for Adding a Self-Service Rule

4. Select the check boxes by the attributes which this ACI will permit users to edit.
5. Click the **Add** button to save the new self-service ACI.

10.2.2. Creating Self-Service Rules from the Command Line

A new self-service rule can be added using the **selfservice-add** command. These two options are required:

- **--permissions** to set which permissions – such as write, add, or delete – the ACI grants
- **--attrs** to give the full list of attributes which this ACI grants permission to.

```
[jsmith@server ~]$ ipa selfservice-add "Users can manage their own name
details" --permissions=write --attrs=givenname --attrs=displayname --
attrs=title --attrs=initials
```

```
-----
Added selfservice "Users can manage their own name details"
-----
```

```
Self-service name: Users can manage their own name details
Permissions: write
Attributes: givenname, displayname, title, initials
```

10.2.3. Editing Self-Service Rules

In the self-service entry in the web UI, the only element that can be edited is the list of attributes that are included in the ACI. The check boxes can be selected or deselected.

Self Service Permissions > User Self service

Self Service Permission: User Self service

Settings

Refresh Reset Update

General

Self-service name User Self service

Attributes *

☐ audio	☒ businesscategory
☒ carlicense	☒ cn
☐ departmentnumber	☒ description
☐ destinationindicator	☒ displayname
☐ employeenumber	☒ employeetype
☒ facsimiletelephonenumber	☒ gecos
☐ gidnumber	☒ givenname
☐ homedirectory	☒ homephone
☐ homepostaladdress	☒ inetuserhttpurl
☐ inetuserstatus	☒ initials

Figure 10.3. Self-Service Edit Page

With the command line, self-service rules are edited using the **ipa selfservice-mod** command. The **--attrs** option overwrites whatever the previous list of supported attributes was, so always include the complete list of attributes along with any new attributes.

```
[jsmith@server ~]$ ipa selfservice-mod "Users can manage their own name
details" --attrs=givenname --attrs=displayname --attrs=title --
attrs=initials --attrs=surname
-----
Modified selfservice "Users can manage their own name details"
-----
Self-service name: Users can manage their own name details
Permissions: write
Attributes: givenname, displayname, title, initials
```



IMPORTANT

Include all of the attributes when modifying a self-service rule, including existing ones.

10.3. DELEGATING PERMISSIONS OVER USERS

Delegation is very similar to roles in that one group of users is assigned permission to manage the entries for another group of users. However, the delegated authority is much more similar to self-service rules in that complete access is granted but only to specific user attributes, not to the entire entry. Also, the groups in delegated authority are existing IdM user groups instead of roles specifically created for access controls.

10.3.1. Delegating Access to User Groups in the Web UI

1. On the **IPA Server** tab in the top menu, select the **Role-Based Access Control** → **Delegations** subtab.
2. Click the **Add** link at the top of the list of the delegation access control instructions.




Figure 10.4. Adding a New Delegation

3. Name the new delegation ACI.
4. Set the permissions by selecting the check boxes whether users will have the right to view the given attributes (read) and add or change the given attributes (write).

Some users may have a need to see information, but should not be able to edit it.

5. In the **User group** drop-down menu, select the group *who is being granted permissions* to the entries of users in the user group.

Add Delegation [X]

Delegation name *

Permissions ☒ read ☒ write

User group *

Member user group *

Attributes *

☐ audio	☒ businesscategory
☒ carlicense	☐ cn
☒ departmentnumber	☐ description
☐ destinationindicator	☐ displayname
☐ employeenumber	☒ employeetype
☒ facsimiletelephonenumber	☐ gecos
☐ gidnumber	☒ givenname
☐ homedirectory	☐ homephone
☐ homepostaladdress	☐ inetuserhttpurl

* Required field

Figure 10.5. Form for Adding a Delegation

6. In the **Member user group** drop-down menu, select the group *whose entries can be edited* by members of the delegation group.
7. In the attributes box, select the check boxes by the attributes to which the member user group is being granted permission.
8. Click the **Add** button to save the new delegation ACI.

10.3.2. Delegating Access to User Groups in the Command Line

A new delegation access control rule is added using the **delegation-add** command. There are three required arguments:

- **--group**, the group *who is being granted permissions* to the entries of users in the user group.
- **--membergroup**, the group *whose entries can be edited* by members of the delegation group.
- **--attrs**, the attributes which users in the member group are allowed to view or edit.

For example:

```
$ ipa delegation-add "basic manager attrs" --attrs=manager --attrs=title -
--attrs=employeetype --attrs=employeeenumber --group=engineering_managers --
memberof=engineering
-----
Added delegation "basic manager attrs"
-----
Delegation name: basic manager attrs
Permissions: write
Attributes: manager, title, employeetype, employeeenumber
Member user group: engineering
User group: engineering_managers
```

Delegation rules are edited using the **delegation-mod** command. The **--attrs** option overwrites whatever the previous list of supported attributes was, so always include the complete list of attributes along with any new attributes.

```
[jsmith@server ~]$ ipa delegation-mod "basic manager attrs" --
attrs=manager --attrs=title --attrs=employeetype --attrs=employeeenumber --
attrs=displayname
-----
Modified delegation "basic manager attrs"
-----
Delegation name: basic manager attrs
Permissions: write
Attributes: manager, title, employeetype, employeeenumber, displayname
Member user group: engineering
User group: engineering_managers
```



IMPORTANT

Include all of the attributes when modifying a delegation rule, including existing ones.

10.4. DEFINING ROLE-BASED ACCESS CONTROLS

Role-based access control grants a very different kind of authority to users compared to self-service and delegation access controls. Role-based access controls are fundamentally administrative, providing the ability to modify entries.

There are three parts to role-based access controls: the *permission*, the *privilege* and the *role*. A privilege consists of one or more permissions, and a role consists of one or more privileges.

- A *permission* defines a specific operation or set of operations (such as read, write, add, or delete) and the target entries within the IdM LDAP directory to which those operations apply. Permissions are building blocks; they can be assigned to multiple privileges as needed.

With IdM permissions, you can control which users have access to which objects and even which attributes of these objects. IdM enables you to whitelist or blacklist individual attributes or change the entire visibility of a specific IdM function, such as users, groups, or sudo, to all anonymous users, all authenticated users, or just a

certain group of privileged users. This flexible approach to permissions is useful in scenarios when, for example, the administrator wants to limit access of users or groups only to the specific sections these users or groups need to access and to make the other sections completely hidden to them.

- A *privilege* is a group of permissions that can be applied to a role. For example, a permission can be created to add, edit, and delete automount locations. Then that permission can be combined with another permission relating to managing FTP services, and they can be used to create a single privilege that relates to managing filesystems.



NOTE

A privilege, in the context of Red Hat Identity Management, has a very specific meaning of an atomic unit of access control on which permissions and then roles are created. *Privilege escalation* as a concept of regular users temporarily gaining additional privileges does not exist in Red Hat Identity Management. Privileges are assigned to users by using Role-Based Access Controls (RBAC). Users either have the role that grants access, or they do not.

Apart from users, privileges are also assigned to user groups, hosts, host groups and network services. This practice permits a fine-grained control of operations by a set of users on a set of hosts via specific network services.

- A *role* is a list of privileges which users specified for the role possess.

It is possible to create entirely new permissions, as well as to create new privileges based on existing permissions or new permissions. Red Hat Identity Management provides the following range of pre-defined roles.

Table 10.1. Predefined Roles in Red Hat Identity Management

Role	Privilege	Description
Helpdesk	Modify Users and Reset passwords, Modify Group membership	Responsible for performing simple user administration tasks
IT Security Specialist	Netgroups Administrators, HBAC Administrator, Sudo Administrator	Responsible for managing security policy such as host-based access controls, sudo rules
IT Specialist	Host Administrators, Host Group Administrators, Service Administrators, Automount Administrators	Responsible for managing hosts

Role	Privilege	Description
Security Architect	Delegation Administrator, Replication Administrators, Write IPA Configuration, Password Policy Administrator	Responsible for managing the Identity Management environment, creating trusts, creating replication agreements
User Administrator	User Administrators, Group Administrators, Stage User Administrators	Responsible for creating users and groups

10.4.1. Roles

10.4.1.1. Creating Roles in the Web UI

1. Open the **IPA Server** tab in the top menu, and select the **Role-Based Access Control** subtab.
2. Click the **Add** link at the top of the list of the role-based access control instructions.





Figure 10.6. Adding a New Role

3. Enter the role name and a description.



Add Role ✕


Role name *

Description

* Required field

Figure 10.7. Form for Adding a Role

4. Click the **Add and Edit** button to save the new role and go to the configuration page.
5. At the top of the **Users** tab, or in the **Users Groups** tab when adding groups, click **Add**.



Identity Policy Authentication Network Services **IPA Server**

Role Based Access Control ▾ ID Ranges ID Views Realm Domains

Roles » Example Role

Role: Example Role

Example Role members:

Users (1)	User Groups	Hosts	Host Groups	Services	Privileges
Refresh Delete +Add					
☐ User login					
☐ employee					

Showing 1 to 1 of 1 entries.

Figure 10.8. Adding Users

6. Select the users on the left and use the > button to move them to the **Prospective** column.

Add Users into Role Example Role

Filter available Users Filter

Available			Prospective	
☐	User login	>	☐	User login
☒	admin		☐	manager
☐	helpdesk			

<

Add Cancel

Figure 10.9. Selecting Users

- At the top of the **Privileges** tab, click **Add**.

Identity Policy Authentication Network Services IPA Server Config

Role Based Access Control ID Ranges ID Views Realm Domains Config

Roles » Example Role

Role: Example Role

Example Role members:

Users (1)	User Groups	Hosts	Host Groups	Services	Privileges (4)	Settings
Refresh Delete Add						
☐	Privilege name					
☐	Automount Administrators					
☐	Certificate Administrators					
☐	DNS Administrators					

Figure 10.10. Adding Privileges

- Select the privileges on the left and use the > button to move them to the **Prospective** column.




Figure 10.11. Selecting Privileges

9. Click the **Add** button to save.

10.4.1.2. Creating Roles in the Command Line

1. Add the new role:

```
[root@server ~]# kinit admin
[root@server ~]# ipa role-add --desc="User Administrator" useradmin
-----
Added role "useradmin"
-----
Role name: useradmin
Description: User Administrator
```

2. Add the required privileges to the role:

```
[root@server ~]# ipa role-add-privilege --privileges="User
Administrators" useradmin
Role name: useradmin
Description: User Administrator
```

```
Privileges: user administrators
```

```
-----  
Number of privileges added 1
```

3. Add the required groups to the role. In this case, we are adding only a single group, **useradmins**, which already exists.

```
[root@server ~]# ipa role-add-member --groups=useradmins useradmin  
Role name: useradmin  
Description: User Administrator  
Member groups: useradmins  
Privileges: user administrators  
-----  
Number of members added 1
```

10.4.2. Permissions

10.4.2.1. Creating New Permissions from the Web UI

1. Open the **IPA Server** tab in the top menu, and select the **Role-Based Access Control** subtab.
2. Select the **Permissions** task link.




Figure 10.12. Permissions Task

3. Click the **Add** button at the top of the list of the permissions.




Figure 10.13. Adding a New Permission

4. Define the properties for the new permission in the form that shows up.

Add Permission

Permission name *

Bind rule type
☒ permission
☐ all
☐ anonymous

Granted rights *

☒ read
☐ search
☐ compare

☐ write
☐ add
☐ delete

☐ all

Type

Subtree *

Extra target filter

Target DN

Member of group

Effective attributes

* Required field

Figure 10.14. Form for Adding a Permission

5. Click the **Add** button under the form to save the permission.

You can specify the following permission properties:

1. Enter the name of the new permission.
2. Select the appropriate **Bind rule type**:
 - **permission** is the default permission type, granting access through privileges and roles

- **all** specifies that the permission applies to all authenticated users
- **anonymous** specifies that the permission applies to all users, including unauthenticated users

**NOTE**

It is not possible to add permissions with a non-default bind rule type to privileges. You also cannot set a permission that is already present in a privilege to a non-default bind rule type.

3. Choose the rights that the permission grants in **Granted rights**.
4. Define the method to identify the target entries for the permission:
 - **Type** specifies an entry type, such as user, host, or service. If you choose a value for the **Type** setting, a list of all possible attributes which will be accessible through this ACI for that entry type appears under **Effective Attributes**.

Defining **Type** sets **Subtree** and **Target DN** to one of the predefined values.

- **Subtree** specifies a subtree entry; every entry beneath this subtree entry is then targeted. Provide an existing subtree entry, as **Subtree** does not accept wildcards or non-existent domain names (DNs). For example:

```
cn=automount,dc=example,dc=com
```

- **Extra target filter** uses an LDAP filter to identify which entries the permission applies to. The filter can be any valid LDAP filter, for example:

```
(!(objectclass=posixgroup))
```

IdM automatically checks the validity of the given filter. If you enter an invalid filter, IdM warns you about this after you attempt to save the permission.


- **Target DN** specifies the domain name (DN) and accepts wildcards. For example:

```
uid=*,cn=users,cn=accounts,dc=com
```

- **Member of group** sets the target filter to members of the given group.

After you fill out the filter settings and click **Add**, IdM validates the filter. If all the permission settings are correct, IdM will perform the search. If some of the permissions settings are incorrect, IdM will display a message informing you about which setting is set incorrectly.

5. If you set **Type**, choose the **Effective attributes** from the list of available ACI attributes. If you did not use **Type**, add the attributes manually by writing them into the **Effective attributes** field. Add a single attribute at a time; to add multiple attributes, click **Add** to add another input field.



IMPORTANT

If you do not set any attributes for the permission, then all attributes are included by default.

10.4.2.2. Creating New Permissions from the Command Line

To add a new permission, issue the **ipa permission-add** command. Specify the properties of the permission by supplying the corresponding options:

- Supply the name of the permission. For example:

```
[root@server ~]# ipa permission-add "dns admin permission"
```

- **--bindtype** specifies the bind rule type. This options accepts the **all**, **anonymous**, and **permission** arguments. For example:

```
--bindtype=all
```

If you do not use **--bindtype**, the type is automatically set to the default **permission** value.



NOTE

It is not possible to add permissions with a non-default bind rule type to privileges. You also cannot set a permission that is already present in a privilege to a non-default bind rule type.

- **--permissions** lists the rights granted by the permission. You can set multiple attributes by using multiple **--permissions** options or by listing the options in a comma-separated list inside curly braces. For example:

```
--permissions=read --permissions=write
--permissions={read,write}
```

- **--attrs** gives the list of attributes over which the permission is granted. You can set multiple attributes by using multiple **--attrs** options or by listing the options in a comma-separated list inside curly braces. For example:

```
--attrs=description --attrs=automountKey
--attrs={description,automountKey}
```

The attributes provided with **--attrs** must exist and be allowed attributes for the given object type, otherwise the command fails with schema syntax errors.

- **--type** defines the entry object type, such as user, host, or service. Each type has its own set of allowed attributes. For example:

```
[root@server ~]# ipa permission-add "manage service" --
permissions=all --type=service --attrs=krbprincipalkey --
attrs=krbprincipalname --attrs=managedby
```

- **--subtree** gives a subtree entry; the filter then targets every entry beneath this subtree entry. Provide an existing subtree entry; **--subtree** does not accept wildcards or non-existent domain names (DNs). Include a DN within the directory.

Because IdM uses a simplified, flat directory tree structure, **--subtree** can be used to target some types of entries, like automount locations, which are containers or parent entries for other configuration. For example:

```
[root@server ~]# ipa permission-add "manage automount locations" --  
subtree="ldap://ldap.example.com:389/cn=automount,dc=example,dc=com"  
--permissions=write --attrs=automountmapname --attrs=automountkey --  
attrs=automountInformation
```

The **--type** and **--subtree** options are mutually exclusive.

- **--filter** uses an LDAP filter to identify which entries the permission applies to. IdM automatically checks the validity of the given filter. The filter can be any valid LDAP filter, for example:

```
[root@server ~]# ipa permission-add "manage Windows groups" --  
filter="(!(objectclass=posixgroup))" --permissions=write --  
attrs=description
```

- **--memberof** sets the target filter to members of the given group after checking that the group exists. For example:

```
[root@server ~]# ipa permission-add ManageHost --permissions="write"  
--subtree=cn=computers,cn=accounts,dc=testrelm,dc=com --  
attr=nshostlocation --memberof=admins
```

- **--targetgroup** sets target to the specified user group after checking that the group exists.

The **Target DN** setting, available in the web UI, is not available on the command line.



NOTE

For information about modifying and deleting permissions, run the **ipa permission-mod --help** and **ipa permission-del --help** commands.

10.4.2.3. Default Managed Permissions

Managed permissions are permissions that come preinstalled with Identity Management. They behave like other permissions created by the user, with the following differences:

- You cannot modify their name, location, and target attributes.
- You cannot delete them.
- They have three sets of attributes:
 - *default* attributes, which are managed by IdM and the user cannot modify them

- *included* attributes, which are additional attributes added by the user; to add an included attribute to a managed permission, specify the attribute by supplying the **--includedattrs** option with the **ipa permission-mod** command
- *excluded* attributes, which are attributes removed by the user; to add an excluded attribute to a managed permission, specify the attribute by supplying the **--excludedattrs** option with the **ipa permission-mod** command

A managed permission applies to all attributes that appear in the default and included attribute sets but not in the excluded set.

If you use the **--attrs** option when modifying a managed permission, the included and excluded attribute sets automatically adjust, so that only the attributes supplied with **--attrs** are enabled.



NOTE

While you cannot delete a managed permission, setting its bind type to **permission** and removing the managed permission from all privileges effectively disables it.

Names of all managed permissions start with **System:**, for example *System: Add Sudo rule* or *System: Modify Services*.

Earlier versions of IdM used a different scheme for default permissions, which, for example, forbade the user from modifying the default permissions and the user could only assign them to privileges. Most of these default permissions have been turned into managed permissions, however, the following permissions still use the previous scheme:

- Add Automember Rebuild Membership Task
- Add Replication Agreements
- Certificate Remove Hold
- Get Certificates status from the CA
- Modify DNA Range
- Modify Replication Agreements
- Remove Replication Agreements
- Request Certificate
- Request Certificates from a different host
- Retrieve Certificates from the CA
- Revoke Certificate
- Write IPA Configuration

If you attempt to modify a managed permission from the web UI, the attributes that you cannot modify will be disabled.

Permission: System: Modify Users

Settings Privileges (2)

Refresh Reset Update

Permission settings

Permission name

System: Modify Users

Bind rule type

☒ permission ☐ all ☐ anonymous

Granted rights

☐ read ☐ search ☐ compare ☒ write

☐ add ☐ delete ☐ all

Figure 10.15. Disabled Attributes

If you attempt to modify a managed permission from the command line, the system will not allow you to change the attributes that you cannot modify. For example, attempting to change a default **System: Modify Users** permission to apply to groups fails:

```
$ ipa permission-mod 'System: Modify Users' --type=group
ipa: ERROR: invalid 'ipapermlocation': not modifiable on managed
permissions
```

You can, however, make the **System: Modify Users** permission not to apply to the **GECOS** attribute:

```
$ ipa permission-mod 'System: Modify Users' --excludedattrs=gecos
-----
Modified permission "System: Modify Users"
```

10.4.2.4. Permissions in Earlier Versions of Identity Management

Earlier versions of Identity Management handled permissions differently, for example:

- The global IdM ACI granted read access to all users of the server, even anonymous ones – that is, not authenticated – users.
- Only write, add, and delete permission types were available. The read permission was available too, but it was of little practical value because all users, including unauthenticated ones, had read access by default.

The current version of Identity Management contains options for setting permissions which are much more fine-grained:

- The global IdM ACI does not grant read access to unauthenticated users.
- It is now possible to, for example, add both a filter and a subtree in the same permission.
- It is possible to add search and compare rights.

The new way of handling permissions has significantly improved the IdM capabilities for controlling user or group access, while retaining backward compatibility with the earlier versions. Upgrading from an earlier version of IdM deletes the global IdM ACI on all servers and replaces it with *managed permissions*.

Permissions created in the previous way are automatically converted to the current style whenever you modify them. If you do not attempt to change them, the permissions of the previous type stay unconverted. Once a permission uses the current style, it can never downgrade to the previous style.



NOTE

It is still possible to assign permissions to privileges on servers running an earlier version of IdM.

The **ipa permission-show** and **ipa permission-find** commands recognize both the current permissions and the permissions of the previous style. While the outputs from both of these commands display permissions in the current style, the permissions themselves remain unchanged; the commands upgrade the permission entries before outputting the data only in memory, without committing the changes to LDAP.

Permissions with both the previous and the current characteristics have effect on all servers – those running previous versions of IdM, as well as those running the current IdM version. However, you cannot create or modify permissions with the current permissions on servers running previous versions of IdM.

10.4.3. Privileges

10.4.3.1. Creating New Privileges from the Web UI

1. Open the **IPA Server** tab in the top menu, and select the **Role-Based Access Control** subtab.
2. Select the **Privileges** task link.




Figure 10.16. Privileges Task

- Click the **Add** link at the top of the list of the privileges.




Figure 10.17. Adding a New Privilege

- Enter the name and a description of the privilege.

Add Privilege

Privilege name *

New Privilege

Description

For employees

* Required field

Add

Add and Add Another

Add and Edit

Cancel

Figure 10.18. Form for Adding a Privilege

- Click the **Add and Edit** button to go to the privilege configuration page to add permissions.
- Select the **Permissions** tab.

- Click **Add** at the top of the list of the permissions to add permission to the privilege.




Figure 10.19. Adding Permissions

- Click the check box by the names of the permissions to add, and use the > button to move the permissions to the **Prospective** column.




Figure 10.20. Selecting Permissions

9. Click the **Add** button to save.

10.4.3.2. Creating New Privileges from the Command Line

Privilege entries are created using the **privilege-add** command, and then permissions are added to the privilege group using the **privilege-add-permission** command.

1. Create the privilege entry.

```
[jsmith@server ~]$ ipa privilege-add "managing filesystems" --
desc="for filesystems"
```

2. Assign the required permissions. For example:

```
[jsmith@server ~]$ ipa privilege-add-permission "managing
filesystems" --permissions="managing automount" --
permissions="managing ftp services"
```

PART IV. ADMINISTRATION: MANAGING IDENTITIES

CHAPTER 11. MANAGING USER ACCOUNTS

This chapter covers general management and configuration of user accounts.

11.1. SETTING UP USER HOME DIRECTORIES

It is recommended that every user has a home directory configured. The default expected location for user home directories is in the `/home/` directory. For example, IdM expects a user with the `user_login` login to have a home directory set up at `/home/user_login`.



NOTE

You can change the default expected location for user home directories using the **`ipa config-mod`** command.

IdM does not automatically create home directories for users. However, you can configure a PAM home directory module to create a home directory automatically when a user logs in. Alternatively, you can add home directories manually using NFS shares and the **`automount`** utility.

11.1.1. Mounting Home Directories Automatically Using the PAM Home Directory Module

Supported PAM Home Directory Modules

To configure a PAM home directory module to create home directories for users automatically when they log in to the IdM domain, use one of the following PAM modules:

- **`pam_oddjob_mkhomedir`**
- **`pam_mkhomedir`**

IdM first attempts to use **`pam_oddjob_mkhomedir`**. If this module is not installed, IdM attempts to use **`pam_mkhomedir`** instead.

Configuring the PAM Home Directory Module

Enabling the PAM home directory module has local effect. Therefore, you must enable the module individually on each client and server where it is required.

To configure the module during the installation of the server or client, use the **`--mkhomedir`** option with the **`ipa-server-install`** or **`ipa-client-install`** utility when installing the machine.

To configure the module on an already installed server or client, use the **`authconfig`** utility. For example:

```
# authconfig --enablemkhomedir --update
```

For more information on using **`authconfig`** to create home directories, see the [System-Level Authentication Guide](#).

11.1.2. Mounting Home Directories Manually

You can use an NFS file server to provide a **/home/** directory that will be available to all machines in the IdM domain, and then mount the directory on an IdM machine using the **automount** utility.

Potential Problems When Using NFS

Using NFS can potentially have negative impact on performance and security. For example, using NFS can lead to security vulnerabilities resulting from granting root access to the NFS user, performance issues with loading the entire **/home/** directory tree, or network performance issues for using remote servers for home directories.

To reduce the effect of these problems, it is recommended to follow these guidelines:

- Use **automount** to mount only the user's home directory and only when the user logs in. Do not use it to load the entire **/home/** tree.
- Use a remote user who has limited permissions to create home directories, and mount the share on the IdM server as this user. Because the IdM server runs as an **httpd** process, it is possible to use **sudo** or a similar program to grant limited access to the IdM server to create home directories on the NFS server.

Configuring Home Directories Using NFS and automount

To manually add home directories to the IdM server from separate locations using NFS shares and **automount**:

1. Create a new location for the user directory maps.

```
$ ipa automountlocation-add userdirs
Location: userdirs
```

2. Add a direct mapping to the new location's **auto.direct** file. The **auto.direct** file is the **automount** map automatically created by the **ipa-server-install** utility. In the following example, the mount point is **/share**:

```
$ ipa automountkey-add userdirs auto.direct --key=/share --info="-ro,soft, server.example.com:/home/share"

Key: /share
Mount information: -ro,soft, server.example.com:/home/share
```

For more details on using **automount** with IdM, see [Chapter 34, Using Automount](#).

11.2. USER LIFE CYCLE


Identity Management supports three user account states: *stage*, *active*, and *preserved*.

- **Stage** users are not allowed to authenticate. This is an initial state. Some of the user account properties required for active users might not yet be set.
- **Active** users are allowed to authenticate. All required user account properties must be set in this state.
- **Preserved** users are former **active** users. They are considered inactive and cannot authenticate to IdM. Preserved users retain most of the account properties they had as active users, but they are not part of any user groups.

**NOTE**

The list of users in the **preserved** state can provide a history of past user accounts.

User entries can also be permanently deleted from the IdM database. Deleting a user entry permanently removes the entry itself and all its information from IdM, including group memberships and passwords. Any external configuration for the user, such as the system account and home directory, is not deleted, but is no longer accessible through IdM.

**IMPORTANT**

Deleted user accounts cannot be restored. When you delete a user account, all the information associated with the account is lost permanently.

A new administrator user can only be created by another administrator, such as the default **admin** user. If you accidentally delete all administrator accounts, the Directory Manager must create a new administrator manually in the Directory Server.

User Life Cycle Management Operations

To manage user provisioning, the administrator can move user accounts from one state to another. New user accounts can be added as either **active** or **stage**, but not as **preserved**.

IdM supports the following operations for user life cycle management:

stage → active

When an account in the **stage** state is ready to be properly activated, the administrator moves it to the **active** state.

active → preserved

After the user leaves the company, the administrator moves the account to the **preserved** state.

preserved → active

A former user joins the company again. The administrator restores the user account by moving it from the **preserved** state back to the **active** state.

preserved → stage

A former user is planning to join the company again. The administrator moves the account from the **preserved** state to the **stage** state to prepare the account for later reactivation.

You can also permanently delete active, stage, and preserved users from IdM. Note that you cannot move stage users to the **preserved** state, you can only delete them permanently.




Figure 11.1. User Life Cycle Operations

11.2.1. Adding Stage or Active Users

Adding Users in the Web UI

1. Select the **Identity** → **Users** tab.
2. Select the **Active users** or **Stage users** category, depending on whether you want to add a user in the **active** or **stage** state.




Figure 11.2. Selecting User Category

For more information about the **active** or **stage** user life cycle states, see [Section 11.2, “User Life Cycle”](#).

3. Click **Add** at the top of the users list.




Figure 11.3. Adding a User

4. Fill in the **Add User** form.

Note that if you do not set a user login manually, IdM generates the login automatically based on the specified first name and last name.

5. Click **Add**.

Alternatively, click **Add and Add Another** to start adding another user or **Add and Edit** to start editing the new user entry. For information on editing user entries, see [Section 11.3, “Editing Users”](#).

Adding Users from the Command Line

To add a new user in the **active** state, use the **ipa user-add** command. To add a new user in the **stage** state, use the **ipa stageuser-add** command.



NOTE

For more information about the **active** or **stage** user life cycle states, see [Section 11.2, “User Life Cycle”](#).

When run without any options, **ipa user-add** and **ipa stageuser-add** prompt you for the minimum required user attributes and use default values for the other attributes. Alternatively, you can add options specifying various attributes directly to the commands.

In the interactive session, after you run the command without any options, IdM proposes an automatically generated user login based on the provided first name and last name and displays it in brackets ([]). To accept the default login, confirm by pressing **Enter**. To specify a custom login, do not confirm the default and specify the custom login instead.

```
$ ipa user-add
First name: first_name
Last name: last_name
User login [default_login]: custom_login
```

Adding options to **ipa user-add** and **ipa stageuser-add** enables you to define custom values for many of the user attributes. This means that you can specify more information than in the interactive session. For example, to add a stage user:


```
$ ipa stageuser-add stage_user_login --first=first_name --last=last_name -
-email=email_address
```

For a complete list of options accepted by **ipa user-add** and **ipa stageuser-add**, run the commands with the **--help** option added.

11.2.1.1. User Name Requirements

IdM supports user names that can be described by the following regular expression:

```
[a-zA-Z0-9_.] [a-zA-Z0-9_.-]{0,252} [a-zA-Z0-9_.$-]?
```



NOTE

User names ending with the trailing dollar sign (\$) are supported to enable Samba 3.x machine support.

If you add a user whose user name contains uppercase characters, IdM automatically converts the name to lowercase when saving it. Therefore, IdM always requires users to enter their user names all lowercase when logging in. Additionally, it is not possible to add users whose user names only differ in letter casing, such as **user** and **User**.

The default maximum length for user names is 32 characters. To change it, use the **ipa config-mod --maxusername** command. For example, to increase the maximum user name length to 64 characters:

```
$ ipa config-mod --maxusername=64
Maximum username length: 64
...
```

11.2.1.2. Defining a Custom UID or GID Number

If you add a new user entry without specifying a custom UID or GID number, IdM automatically assigns an ID number that is next available in the ID range. This means that users' ID numbers are always unique. For more information about ID ranges, see [Chapter 14, Unique UID and GID Number Assignments](#)

When you specify a custom ID number, the server does not validate whether the custom ID number is unique. Due to this, multiple user entries might have the same ID number assigned. Red Hat recommends to prevent having multiple entries with the same ID number.

11.2.2. Listing Users and Searching for Users

Listing Users in the Web UI

1. Select the **Identity** → **Users** tab.
2. Select the **Active users**, **Stage users**, or **Preserved users** category.




Figure 11.4. Listing Users

Displaying Information About a User in the Web UI

To display detailed information about a user, click the name of the user in the list of users:




Figure 11.5. Displaying User Information

Listing Users from the Command Line

To list all active users run the `ipa user-find` command. To list all stage users, use the `ipa stageuser-find` command. To list preserved users, run the `ipa user-find --preserved=true` command.

For example:

```
$ ipa user-find
-----
23 users matched
-----
User login: admin
Last name: Administrator
Home directory: /home/admin
Login shell: /bin/bash
UID: 1453200000
GID: 1453200000
Account disabled: False
Password: True
Kerberos keys available: True
```

```
User login: user
...
```

By adding options and arguments to **ipa user-find** and **ipa stageuser-find**, you can define the search criteria and filter the search results. For example, to display all active users with a specific title defined:

```
$ ipa user-find --title=user_title
-----
2 users matched
-----
    User login: user
    ...
    Job Title: Title
    ...

    User login: user2
    ...
    Job Title: Title
    ...
```

Similarly, to display all stage users whose login contains **user**:

```
$ ipa user-find user
-----
3 users matched
-----
User login: user
...

User login: user2
...

User login: user3
...
```

For a complete list of options accepted by **ipa user-find** and **ipa stageuser-find**, run the commands with the **--help** option added.

Displaying Information about a User from the Command Line

To display information about an active or preserved user, use the **ipa user-show** command:

```
$ ipa user-show user_login
    User login: user_login
    First name: first_name
    Last name: last_name
    ...
```

To display information about a stage user, use the **ipa stageuser-show** command:

11.2.3. Activating, Preserving, Deleting, and Restoring Users

This section describes moving user accounts between different user life cycle states. For

details on the life cycle states in IdM, see [Section 11.2, “User Life Cycle”](#).

Managing User Life Cycle in the Web UI

To activate a stage user:

- In the **Stage users** list, select the user to activate, and click **Activate**.




Figure 11.6. Activating a User

To preserve or delete a user:

1. In the **Active users** or **Stage users** lists, select the user. Click **Delete**.




Figure 11.7. Deleting a User

2. If you selected an active user, select **delete** or **preserve**. If you selected a stage user, you can only delete the user. The default UI option is **delete**.

For example, to preserve an active user:




Figure 11.8. Selecting the Delete Mode in the Web UI

To confirm, click the **Delete** button.

To restore a preserved user:

- In the **Preserved users** list, select the user to restore, and click **Restore**.



Figure 11.9. Restoring a User



NOTE

Restoring a user does not restore all of the account's previous attributes. For example, the user's password is not restored and must be defined again.

Note that in the web UI, it is not possible to move a user from the **preserved** state to the **stage** state.

Managing User Life Cycle from the Command Line

To activate a user account by moving it from **stage** to **active**, use the **ipa stageuser-activate** command.

```
$ ipa stageuser-activate user_login
-----
Stage user user_login activated
-----
...
```

To preserve or delete a user account, use the **ipa user-del** or **ipa stageuser-del** commands.

- To remove an active user permanently from the IdM database, run **ipa user-del** without any options.

```
$ ipa user-del user_login
-----
Deleted user "user3"
-----
```

- To preserve an active user account, run **ipa user-del** with the **--preserve** option.

```
$ ipa user-del --preserve user_login
-----
Deleted user "user_login"
-----
```

- To remove a stage user permanently from the IdM database, run **ipa stageuser-del**.

```
$ ipa stageuser-del user_login
-----
Deleted stage user "user_login"
-----
```

NOTE

When deleting multiple users, use the **--continue** option to force the command to continue regardless of errors. A summary of the successful and failed operations is printed to the **stdout** standard output stream when the command completes.

```
$ ipa user-del --continue user1 user2 user3
```

If **--continue** is not used, the command proceeds with deleting users until it encounters an error, after which it stops and exits.

To restore a preserved user account by moving it from **preserved** to **active**, use the **ipa user-undel** command.

```
$ ipa user-undel user_login
-----
Undeleted user account "user_login"
-----
```

To restore a preserved user account by moving it from **preserved** to **stage**, use the **ipa user-stage** command.

```
$ ipa user-stage user_login
-----
Staged user account "user_login"
-----
```

**NOTE**

Restoring a user account does not restore all of the account's previous attributes. For example, the user's password is not restored and must be defined again.

For more information about these commands and the options they accept, run them with the **--help** option added.

11.3. EDITING USERS

Editing Users in the Web UI

1. Select the **Identity** → **Users** tab.
2. Search the **Active users**, **Stage users**, or **Preserved users** category to find the user to edit.
3. Click the name of the user to edit.

User categories

Active users >

Stage users

Preserved users

Active users

Search

☐	User login	First name	Last name	Status	UID
☐	admin		Administrator	✓ Enabled	14532
☐	user	User	User	✓ Enabled	14532

Figure 11.10. Selecting a User to Edit

4. Edit the user attribute fields as required.
5. Click **Save** at the top of the page.

✓ User: user

user is a member of:

Settings

User Groups

Netgroups

Roles

HBAC Rules

Sudo R

Refresh

Revert

Save

Actions ▾

Identity Settings

Job Title

First name

User

Figure 11.11. Save Modified User Attributes

After you update user details in the web UI, the new values are not synchronized immediately. It might take up to approximately 5 minutes before the new values are reflected at the client system.

Editing Users from the Command Line

To modify a user in the **active** or **preserved** states, use the **ipa user-mod** command. To modify a user in the **stage** state, use the **ipa stageuser-mod** command.

The **ipa user-mod** and **ipa stageuser-mod** commands accept the following options:

- the user login, which identifies the user account to be modified
- options specifying the new attribute values

For a complete list of user entry attributes that can be modified from the command line, see the list of options accepted by **ipa user-mod** and **ipa stageuser-mod**. To display the list of options, run the commands with the **--help** option added.

Simply adding an attribute option to **ipa user-mod** or **ipa stageuser-mod** overwrites the current attribute value. For example, the following changes a user's title or adds a new title if the user did not yet have a title specified:

```
$ ipa user-mod user_login --title=new_title
```

For LDAP attributes that are allowed to have multiple values, IdM also accepts multiple values. For example, a user can have two email addresses saved in their user account. To add an additional attribute value without overwriting the existing value, use the **--addattr** option together with the option to specify the new attribute value. For example, to add a new email address to a user account that already has an email address specified:

```
$ ipa user-mod user --addattr=mobile=new_mobile_number
-----
Modified user "user"
-----
  User login: user
...
  Mobile Telephone Number: mobile_number, new_mobile_number
...
```

To set two attribute values at the same time, use the **--addattr** option twice:

```
$ ipa user-mod user --addattr=mobile=mobile_number_1 --
addattr=mobile=mobile_number_2
```

The **ipa user-mod** command also accepts the **--setattr** option for setting attribute values and the **--delattr** option for deleting attribute values. These options are used in a way similar to using **--addattr**. For details, see the output of the **ipa user-mod --help** command.

**NOTE**

To overwrite the current email address for a user, use the **--email** option. However, to add an additional email address, use the **mail** option with the **--addattr** option:

```
$ ipa user-mod user --email=email@example.com
$ ipa user-mod user --addattr=mail=another_email@example.com
```


11.4. ENABLING AND DISABLING USER ACCOUNTS

The administrator can disable and enable active user accounts. Disabling a user account deactivates the account. Disabled user accounts cannot be used to authenticate. A user whose account has been disabled cannot log into IdM and cannot use IdM services, such as Kerberos, or perform any tasks.

Disabled user accounts still exist within IdM and all of the associated information remains unchanged. Unlike preserved user accounts, disabled user accounts remain in the **active** state. Therefore, they are displayed in the output of the **ipa user-find** command. For example:

```
$ ipa user-find
...
  User login: user
  First name: User
  Last name: User
  Home directory: /home/user
  Login shell: /bin/sh
  UID: 1453200009
  GID: 1453200009
  Account disabled: True
  Password: False
  Kerberos keys available: False
...
```


Any disabled user account can be enabled again.

**NOTE**

After disabling a user account, existing connections remain valid until the user's Kerberos TGT and other tickets expire. After the ticket expires, the user will not be able to renew it.

Enabling and Disabling User Accounts in the Web UI

1. Select the **Identity** → **Users** tab.
2. From the **Active users** list, select the required user or users, and then click **Disable** or **Enable**.



Active users						
Search		Refresh Delete + Add - Disable ✓ Enable Act				
☐	User login	First name	Last name	Status	UID	Email address
☐	admin		Administrator	✓ Enabled	1453200000	
☒	user	User	User	✓ Enabled	1453200009	
☐	user2	User2	User2	✓ Enabled	1453200007	

Figure 11.12. Disabling or Enabling a User Account

Disabling and Enabling User Accounts from the Command Line

To disable a user account, use the **ipa user-disable** command.

```
$ ipa user-disable user_login
-----
Disabled user account "user_login"
-----
```

To enable a user account, use the **ipa user-enable** command.

```
$ ipa user-enable user_login
-----
Enabled user account "user_login"
-----
```

11.5. ALLOWING NON-ADMIN USERS TO MANAGE USER ENTRIES

By default, only the **admin** user is allowed to manage user life cycle and disable or enable user accounts. To allow another, non-admin user to do this, create a new role, add the relevant permissions to this role, and assign the non-admin user to the role.

By default, IdM includes the following privileges related to managing user accounts:

Modify Users and Reset passwords

This privilege includes permissions to modify various user attributes.

User Administrators

This privilege includes permissions to add active users, activate non-active users, remove users, modify user attributes, and other permissions.

Stage User Provisioning

This privilege includes a permission to add stage users.

Stage User Administrator

This privilege includes permissions to perform a number of life cycle operations, such as adding stage users or moving users between life cycle states. However, it does not include permissions to move users to the active state.

For information on defining roles, permissions, and privileges, see [Section 10.4, “Defining Role-Based Access Controls”](#).

Allowing Different Users to Perform Different User Management Operations

The different privileges related to managing user accounts can be added to different users. For example, you can separate privileges for employee account entry and activation by:

- Configuring one user as a *stage user administrator*, who is allowed to add future employees to IdM as stage users, but not to activate them.
- Configuring another user as a *security administrator*, who is allowed to activate the stage users after their employee credentials are verified on the first day of employment.

To allow a user to perform certain user management operations, create a new role with the required privilege or privileges, and assign the user to that role.

Example 11.1. Allowing a Non-admin User to Add Stage Users

This example shows how to create a user who is only allowed to add new stage users, but not to perform any other stage user management operations.

1. Log in as the **admin** user or another user allowed to manage role-based access control.

```
$ kinit admin
```

2. Create a new custom role to manage adding stage users.

- a. Create the **System Provisioning** role.

```
$ ipa role-add --desc "Responsible for provisioning stage
users" "System Provisioning"
-----
Added role "System Provisioning"
-----
Role name: System Provisioning
Description: Responsible for provisioning stage users
```

- b. Add the **Stage User Provisioning** privilege to the role. This privilege provides the ability to add stage users.

```
$ ipa role-add-privilege "System Provisioning" --
privileges="Stage User Provisioning"
Role name: System Provisioning
Description: Responsible for provisioning stage users
Privileges: Stage User Provisioning
-----
Number of privileges added 1
-----
```

3. Grant a non-admin user the rights to add stage users.

- a. If the non-admin user does not yet exist, create a new user. In this example, the user is named **stage_user_admin**.

```
$ ipa user-add stage_user_admin --password
First name: first_name
Last name: last_name
Password:
Enter password again to verify:
...
```

- b. Assign the **stage_user_admin** user to the **System Provisioning** role.

```
$ ipa role-add-member "System Provisioning" --
users=stage_user_admin
Role name: System Provisioning
Description: Responsible for provisioning stage users
Member users: stage_user_admin
Privileges: Stage User Provisioning
-----
Number of members added 1
-----
```

- c. To make sure the **System Provisioning** role is configured correctly, you can use the **ipa role-show** command to display the role settings.

```
$ ipa role-show "System Provisioning"
-----
1 role matched
-----
Role name: System provisioning
Description: Responsible for provisioning stage users
Member users: stage_user_admin
Privileges: Stage User Provisioning
-----
Number of entries returned 1
-----
```

4. Test adding a new stage user as the **stage_user_admin** user.

- a. Log in as **stage_user_admin**. Note that if you created **stage_user_admin** as a new user in one of the previous steps, IdM will ask you to change the initial password set by **admin**.

```
$ kinit stage_user_admin
Password for stage_user_admin@EXAMPLE.COM:
Password expired. You must change it now.
Enter new password:
Enter it again:
```

- b. To make sure your Kerberos ticket for **admin** has been replaced with a Kerberos ticket for **stage_user_admin**, you can use the **klist** utility.

```
$ klist
Ticket cache: KEYRING:persistent:0:krb_ccache_xIlCQDW
```

```
Default principal: stage_user_admin@EXAMPLE.COM
```

```
Valid starting      Expires      Service principal
02/25/2016 11:42:20 02/26/2016 11:42:20  krbtgt/EXAMPLE.COM
```

c. Add a new stage user.

```
$ ipa stageuser-add stage_user
First name: first_name
Last name: last_name
ipa: ERROR: stage_user: stage user not found
```



NOTE

The error that IdM reports after adding a stage user is expected. The **stage_user_admin** is only allowed to add stage users, not to display information about them. Therefore, instead of displaying a summary of the newly added **stage_user** settings, IdM displays the error.

The **stage_user_admin** user is not allowed to display information about stage users. Therefore, an attempt to display information about the new **stage_user** user while logged in as **stage_user_admin** fails:

```
$ ipa stageuser-show stage_user
ipa: ERROR: stage_user: stage user not found
```

To display information about **stage_user**, you can log in as **admin**:

```
$ kinit admin
Password for admin@EXAMPLE.COM:
$ ipa stageuser-show stage_user
User login: stage_user
First name: Stage
Last name: User
...
```

11.6. USING AN EXTERNAL PROVISIONING SYSTEM FOR USERS AND GROUPS

Identity Management supports configuring your environment, so that an external solution for managing identities is used to provision user and group identities in IdM. This section describes an example of such configuration. The example includes:

- [Section 11.6.1, “Configuring User Accounts to Be Used by the External Provisioning System”](#)
- [Section 11.6.2, “Configuring IdM to Automatically Activate Stage User Accounts”](#)
- [Section 11.6.3, “Configuring the LDAP Provider of the External Provisioning System to Manage the IdM Identities”](#)

11.6.1. Configuring User Accounts to Be Used by the External Provisioning System

This procedure shows how to configure two IdM user accounts to be used by the external provisioning system. By adding the accounts to a group with an appropriate password policy, you enable the external provisioning system to manage user provisioning in IdM.

1. Create a user, **provisionator**, with the privileges to add stage users. The user account will be used by the external provisioning system to add new stage users.

- a. Add the **provisionator** user account:

```
$ ipa user-add provisionator --first=provisioning --last=account --password
```

- b. Grant the **provisionator** user the required privileges.

Create a custom role, **System Provisioning**, to manage adding stage users:

```
$ ipa role-add --desc "Responsible for provisioning stage users" "System Provisioning"
```

Add the **Stage User Provisioning** privilege to the role. This privilege provides the ability to add stage users:

```
$ ipa role-add-privilege "System Provisioning" --privileges="Stage User Provisioning"
```

Add the **provisionator** user to the role:

```
$ ipa role-add-member --users=provisionator "System Provisioning"
```

2. Create a user, **activator**, with the privileges to manage user accounts. The user account will be used to automatically activate stage users added by the external provisioning system.

- a. Add the **activator** user account:

```
$ ipa user-add activator --first=activation --last=account --password
```

- b. Grant the **activator** user the required privileges.

Add the user to the default **User Administrator** role:

```
$ ipa role-add-member --users=activator "User Administrator"
```

3. Create a user group for service and application accounts:

```
$ ipa group-add service-accounts
```

4. Update the password policy for the group. The following policy prevents password expiration and lockout for the account but compensates the potential risks by requiring complex passwords:

```
$ ipa pwpolicy-add service-accounts --maxlife=10000 --minlife=0 --
history=0 --minclasses=4 --minlength=20 --priority=1 --maxfail=0 --
failinterval=1 --lockouttime=0
```

5. Add the provisioning and activation accounts to the group for service and application accounts:

```
$ ipa group-add-member service-accounts --users=
{provisionator,activator}
```

6. Change the passwords for the user accounts:

```
$ kpasswd provisionator
$ kpasswd activator
```


Changing the passwords is necessary because passwords of new IdM users expire immediately.

Additional resources:

- For details on adding new users, see [Section 11.2.1, “Adding Stage or Active Users”](#).
- For details on granting users the privileges required to manage other user accounts, see [Section 11.5, “Allowing Non-admin Users to Manage User Entries”](#).
- For details on managing IdM password policies, see [Chapter 28, Defining Password Policies](#).

11.6.2. Configuring IdM to Automatically Activate Stage User Accounts

This procedure shows how to create a script for activating stage users. The system runs the script automatically at specified time intervals. This ensures that new user accounts are automatically activated and available for use shortly after they are created.



IMPORTANT

The procedure assumes that the new user accounts do not require validation before the script adds them to IdM. For example, validation is not required when the users have already been validated by the owner of the external provisioning system.

It is sufficient to enable the activation process on only one of your IdM servers.

1. Generate a keytab file for the activation account:

```
# ipa-getkeytab -s example.com -p "activator" -k /etc/krb5.ipa-
activation.keytab
```


If you want to enable the activation process on more than one IdM server, generate the keytab file on one server only. Then copy the keytab file to the other servers.

2. Create a script, **/usr/local/sbin/ipa-activate-all**, with the following contents to activate all users:

```
#!/bin/bash

kinit -k -i activator

ipa stageuser-find --all --raw | grep " uid:" | cut -d ":" -f 2 |
while read uid; do ipa stageuser-activate ${uid}; done
```

3. Edit the permissions and ownership for the **ipa-activate-all** script to make it executable:

```
# chmod 755 /usr/local/sbin/ipa-activate-all
# chown root:root /usr/local/sbin/ipa-activate-all
```

4. Create a **systemd** unit file, **/etc/systemd/system/ipa-activate-all.service**, with the following contents:

```
[Unit]
Description=Scan IdM every minute for any stage users that must be
activated

[Service]
Environment=KRB5_CLIENT_KTNAME=/etc/krb5.ipa-activation.keytab
Environment=KRB5CCNAME=FILE:/tmp/krb5cc_ipa-activate-all
ExecStart=/usr/local/sbin/ipa-activate-all
```

5. Create a **systemd** timer, **/etc/systemd/system/ipa-activate-all.timer**, with the following contents:

```
[Unit]
Description=Scan IdM every minute for any stage users that must be
activated

[Timer]
OnBootSec=15min
OnUnitActiveSec=1min

[Install]
WantedBy=multi-user.target
```

6. Enable **ipa-activate-all.timer**:

```
# systemctl enable ipa-activate-all.timer
```

Additional resources:

- For more information on **systemd** unit files, see the [Managing Services with systemd Unit Files](#) chapter of the *System Administrator's Guide*.

11.6.3. Configuring the LDAP Provider of the External Provisioning System to Manage the IdM Identities

This section shows templates for various user and group management operations. Using these templates, you can configure the LDAP provider of your provisioning system to manage IdM user accounts. For example, you can configure the system to inactivate a user account after the employee has left the company.

Managing User Accounts Using LDAP

You can add new user entries, modify existing entries, move users between different life cycle states, or delete users by editing the underlying Directory Server database. To edit the database, use the **ldapmodify** utility.

The following LDIF-formatted templates provide information on what attributes to modify using **ldapmodify**. For detailed example procedures, see [Example 11.2, “Adding a Stage User with **ldapmodify**”](#) and [Example 11.3, “Preserving a User with **ldapmodify**”](#).

Adding a new stage user

Adding a user with UID and GID automatically assigned:

```
dn: uid=user_login,cn=staged
users,cn=accounts,cn=provisioning,dc=example,dc=com
changetype: add
objectClass: top
objectClass: inetorgperson
uid: user_login
sn: surname
givenName: first_name
cn: full_name
```

Adding a user with UID and GID statically assigned:

```
dn: uid=user_login,cn=staged
users,cn=accounts,cn=provisioning,dc=example,dc=com
changetype: add
objectClass: top
objectClass: person
objectClass: inetorgperson
objectClass: organizationalperson
objectClass: posixaccount
uid: user_login
uidNumber: UID_number
gidNumber: GID_number
sn: surname
givenName: first_name
cn: full_name
homeDirectory: /home/user_login
```

You are not required to specify any IdM object classes when adding stage users. IdM adds these classes automatically after the users are activated.

Note that the distinguished name (DN) of the created entry must start with **uid=user_login**.

Modifying existing users

Before modifying a user, obtain the user's distinguished name (DN) by searching by the user's login. In the following example, the *user_allowed_to_read* user in the following example is a user allowed to read user and group information, and *password* is this user's password:

```
# ldapsearch -LLL -x -D
"uid=user_allowed_to_read,cn=users,cn=accounts,dc=example,dc=com" -w
"password" -H ldap://server.example.com -b "cn=users, cn=accounts,
dc=example, dc=com" uid=user_login
```

To modify a user's attribute:

```
dn: distinguished_name
changetype: modify
replace: attribute_to_modify
attribute_to_modify: new_value
```

To disable a user:

```
dn: distinguished_name
changetype: modify
replace: nsAccountLock
nsAccountLock: TRUE
```

To enable a user:

```
dn: distinguished_name
changetype: modify
replace: nsAccountLock
nsAccountLock: FALSE
```

To preserve a user:

```
dn: distinguished_name
changetype: modrdn
newrdn: uid=user_login
deleteoldrdn: 0
newsuperior: cn=deleted users,cn=accounts,cn=provisioning,dc=example
```

Updating the **nssAccountLock** attribute has no effect on stage and preserved users. Even though the update operation completes successfully, the attribute value remains **nssAccountLock: TRUE**.

Creating a new group

To create a new group:

```
dn: cn=group_distinguished_name,cn=groups,cn=accounts,dc=example,dc=com
changetype: add
objectClass: top
objectClass: ipaobject
objectClass: ipausergroup
objectClass: groupofnames
objectClass: nestedgroup
```

```
objectClass: posixgroup
cn: group_name
gidNumber: GID_number
```

Modifying groups

Before modifying a group, obtain the group's distinguished name (DN) by searching by the group's name.

```
# ldapsearch -Y GSSAPI -H ldap://server.example.com -b
"cn=groups,cn=accounts,dc=example,dc=com" "cn=group_name"
```

To delete an existing group:

```
dn: group_distinguished_name
changetype: delete
```

To add a member to a group:

```
dn: group_distinguished_name
changetype: modify
add: member
member: uid=user_login,cn=users,cn=accounts,dc=example,dc=com
```

To remove a member from a group:

```
dn: distinguished_name
changetype: modify
delete: member
member: uid=user_login,cn=users,cn=accounts,dc=example,dc=com
```

Do not add stage or preserved users to groups. Even though the update operation completes successfully, the users will not be updated as members of the group. Only active users can belong to groups.

Example 11.2. Adding a Stage User with `ldapmodify`

To add a new **stageuser** user using the standard **interorgperson** object class:

1. Use **ldapmodify** to add the user.

```
# ldapmodify -Y GSSAPI
SASL/GSSAPI authentication started
SASL username: admin@EXAMPLE
SASL SSF: 56
SASL data security layer installed.
dn: uid=stageuser,cn=staged
users,cn=accounts,cn=provisioning,dc=example
changetype: add
objectClass: top
objectClass: inetorgperson
cn: Stage
sn: User
```

```
adding new entry "uid=stageuser,cn=staged
users,cn=accounts,cn=provisioning,dc=example"
```

2. Consider validating the contents of the stage entry to make sure your provisioning system added all required POSIX attributes and the stage entry is ready to be activated. To display the new stage user's LDAP attributes using the **ipa stageuser-show --all --raw** command. Note that the user is explicitly disabled by the **nsaccountlock** attribute:

```
$ ipa stageuser-show stageuser --all --raw
dn: uid=stageuser,cn=staged
users,cn=accounts,cn=provisioning,dc=example
uid: stageuser
sn: User
cn: Stage
has_password: FALSE
has_keytab: FALSE
nsaccountlock: TRUE
objectClass: top
objectClass: inetorgperson
objectClass: organizationalPerson
objectClass: person
```

Example 11.3. Preserving a User with ldapmodify

To preserve **user** by using the **LDAPmodrdn** operation:

1. Use the **ldapmodify** utility to modify the user entry.

```
$ ldapmodify -Y GSSAPI
SASL/GSSAPI authentication started
SASL username: admin@EXAMPLE
SASL SSF: 56
SASL data security layer installed.
dn: uid=user1,cn=users,cn=accounts,dc=example
changetype: modrdn
newrdn: uid=user1
deleteoldrdn: 0
newsuperior: cn=deleted
users,cn=accounts,cn=provisioning,dc=example

modifying rdn of entry "uid=user1,cn=users,cn=accounts,dc=example"
```

2. Optionally, verify the user has been preserved by listing all preserved users.

```
$ ipa user-find --preserved=true
-----
1 user matched
-----
  User login: user1
  First name: first_name
  Last name: last_name
  ...
```

Number of entries returned 1

CHAPTER 12. MANAGING HOSTS

Both DNS and Kerberos are configured as part of the initial client configuration. This is required because these are the two services that bring the machine within the IdM domain and allow it to identify the IdM server it will connect with. After the initial configuration, IdM has tools to manage both of these services in response to changes in the domain services, changes to the IT environment, or changes on the machines themselves which affect Kerberos, certificate, and DNS services.

This chapter describes how to manage identity services that relate directly to the client machine:

- DNS entries and settings
- Machine authentication
- Host name changes (which affect domain services)

12.1. ABOUT HOSTS, SERVICES, AND MACHINE IDENTITY AND AUTHENTICATION

The basic function of an enrollment process is to create a *host* entry for the client machine in the IdM directory. This host entry is used to establish relationships between other hosts and even services within the domain (as described in [Chapter 1, Introduction to Red Hat Identity Management](#)). These relationships are part of *delegating* authorization and control to hosts within the domain.

A host entry contains all of the information about the client within IdM:

- Service entries associated with the host
- The host and service principal
- Access control rules
- Machine information, such as its physical location and operating system

Some services that run on a host can also belong to the IdM domain. Any service that can store a Kerberos principal or an SSL certificate (or both) can be configured as an IdM service. Adding a service to the IdM domain allows the service to request an SSL certificate or keytab from the domain. (Only the public key for the certificate is stored in the service record. The private key is local to the service.)

An IdM domain establishes a commonality between machines, with common identity information, common policies, and shared services. Any machine which belongs to a domain functions as a client of the domain, which means it uses the services that the domain provides. An IdM domain provides three main services specifically for machines:

- DNS
- Kerberos
- Certificate management

Like users, machines are an identity that is managed by IdM. Client machines use DNS to identify IdM servers, services, and domain members. These are, like user identities, stored

in the 389 Directory Server instance for the IdM server. Like users, machines can be authenticated to the domain using Kerberos or certificates.

From the machine perspective, there are several tasks that can be performed that access these domain services:

- Joining the DNS domain (*machine enrollment*)
- Managing DNS entries and zones
- Managing machine authentication

Authentication in IdM includes machines as well as users. Machine authentication is required for the IdM server to trust the machine and to accept IdM connections from the client software installed on that machine. After authenticating the client, the IdM server can respond to its requests. IdM supports three different approaches to machine authentication:

- SSH keys. The SSH public key for the host is created and uploaded to the host entry. From there, the System Security Services Daemon (SSSD) uses IdM as an identity provider and can work in conjunction with OpenSSH and other services to reference the public keys located centrally in Identity Management. This is described in [Section 12.5, “Managing Public SSH Keys for Hosts”](#).
- Key tables (or *keytabs*, a symmetric key resembling to some extent a user password) and machine certificates. Kerberos tickets are generated as part of the Kerberos services and policies defined by the server. Initially granting a Kerberos ticket, renewing the Kerberos credentials, and even destroying the Kerberos session are all handled by the IdM services. Managing Kerberos is covered in [Chapter 29, Managing the Kerberos Domain](#).
- Machine certificates. In this case, the machine uses an SSL certificate that is issued by the IdM server's certificate authority and then stored in IdM's Directory Server. The certificate is then sent to the machine to present when it authenticates to the server. On the client, certificates are managed by a service called *certmonger*.

12.2. ABOUT HOST ENTRY CONFIGURATION PROPERTIES

A host entry can contain information about the host that is outside its system configuration, such as its physical location, MAC address, keys, and certificates.

This information can be set when the host entry is created if it is created manually; otherwise, most of this information needs to be added to the host entry after the host is enrolled in the domain.

Table 12.1. Host Configuration Properties

UI Field	Command-Line Option	Description
Description	--desc = <i>description</i>	A description of the host.
Locality	--locality = <i>locality</i>	The geographic location of the host.

UI Field	Command-Line Option	Description
Location	--location = <i>location</i>	The physical location of the host, such as its data center rack.
Platform	--platform = <i>string</i>	The host hardware or architecture.
Operating system	--os = <i>string</i>	The operating system and version for the host.
MAC address	--macaddress = <i>address</i>	The MAC address for the host. This is a multi-valued attribute. The MAC address is used by the NIS plug-in to create a NIS ethers map for the host.
SSH public keys	--sshpubkey = <i>string</i>	The full SSH public key for the host. This is a multi-valued attribute, so multiple keys can be set.
Principal name (not editable)	--principalname = <i>principal</i>	The Kerberos principal name for the host. This defaults to the host name during the client installation, unless a different principal is explicitly set in the -p . This can be changed using the command-line tools, but cannot be changed in the UI.
Set One-Time Password	--password = <i>string</i>	Sets a password for the host which can be used in bulk enrollment.
-	--random	Generates a random password to be used in bulk enrollment.
-	--certificate = <i>string</i>	A certificate blob for the host.
-	--updatedns	This sets whether the host can dynamically update its DNS entries if its IP address changes.

12.3. ADDING HOST ENTRIES

12.3.1. Adding Host Entries from the Web UI

1. Open the **Identity** tab, and select the **Hosts** subtab.
2. Click **Add** at the top of the hosts list.




Host name	Description	Enrolled
server.example.com		True

Showing 1 to 1 of 1 entries.

Figure 12.1. Adding Host Entries

3. Fill in the machine name and select the domain from the configured zones in the drop-down list. If the host has already been assigned a static IP address, then include that with the host entry so that the DNS entry is fully created.

Optionally, to add an extra value to the host for some use cases, use the **Class** field. Semantics placed on this attribute are for local interpretation.



Add Host

Host Name *

DNS Zone *

Class

IP Address

Force ☒

* Required field

Figure 12.2. Add Host Wizard

DNS zones can be created in IdM, which is described in [Section 33.4.1, “Adding and Removing Master DNS Zones”](#). If the IdM server does not manage the DNS server, the zone can be entered manually in the menu area, like a regular text field.



NOTE

Select the **Force** check box if you want to skip checking whether the host is resolvable via DNS.

- Click the **Add and Edit** button to go directly to the expanded entry page and fill in more attribute information. Information about the host hardware and physical location can be included with the host entry.

Host: server.zone.example.com

server.zone.examp... is a member of:

Settings Host Groups Netgroups Roles HBAC Rules Sudo Rules

Refresh Revert Save Actions

Host Settings

Host name [server.zone.example.com](#)

Principal name host/server.zone.example.com@EXAMPLE.COM

Description

Class

Locality

Figure 12.3. Expanded Entry Page

12.3.2. Adding Host Entries from the Command Line

Host entries are created using the **host-add** command. This command adds the host entry to the IdM Directory Server. The full list of options with **host-add** are listed in the **ipa host** manpage. At its most basic, an add operation only requires the client host name to add the client to the Kerberos realm and to create an entry in the IdM LDAP server:

```
$ ipa host-add client1.example.com
```

If the IdM server is configured to manage DNS, then the host can also be added to the DNS resource records using the **--ip-address** and **--force** options.

Example 12.1. Creating Host Entries with Static IP Addresses

```
$ ipa host-add --force --ip-address=192.168.166.31 client1.example.com
```

Commonly, hosts may not have a static IP address or the IP address may not be known at the time the client is configured. For example, laptops may be preconfigured as Identity Management clients, but they do not have IP addresses at the time they are configured. Hosts which use DHCP can still be configured with a DNS entry by using **--force**. This essentially creates a placeholder entry in the IdM DNS service. When the DNS service dynamically updates its records, the host's current IP address is detected and its DNS record is updated.

Example 12.2. Creating Host Entries with DHCP

```
$ ipa host-add --force client1.example.com
```

Host records are deleted using the **host-del** command. If the IdM domain uses DNS, then the **--updatedns** option also removes the associated records of any kind for the host from the DNS.

```
$ ipa host-del --updatedns client1.example.com
```

12.4. DISABLING AND RE-ENABLING HOST ENTRIES

Active hosts can be accessed by other services, hosts, and users within the domain. There can be situations when it is necessary to remove a host from activity. However, deleting a host removes the entry and all the associated configuration, and it removes it permanently.

12.4.1. Disabling Host Entries

Disabling a host prevents domain users from access it without permanently removing it from the domain. This can be done by using the **host-disable** command.

For example:

```
[jsmith@ipaserver ~]$ kinit admin
[jsmith@ipaserver ~]$ ipa host-disable server.example.com
```

**IMPORTANT**

Disabling a host entry not only disables that host. It disables every configured service on that host as well.

12.4.2. Re-enabling Hosts

Disabling a host essentially kills its current, active keytabs. Removing the keytabs effectively removes the host from the IdM domain without otherwise touching its configuration entry.

To re-enable a host, simply use the **ipa-getkeytab** command. The **-s** option sets which IdM server to request the keytab, **-p** gives the principal name, and **-k** gives the file to which to save the keytab.

For example, requesting a new host keytab:

```
[jsmith@ipaserver ~]$ ipa-getkeytab -s ipaserver.example.com -p
host/server.example.com -k /etc/krb5.keytab -D
fqdn=server.example.com,cn=computers,cn=accounts,dc=example,dc=com -w
password
```

If the **ipa-getkeytab** command is run on an active IdM client or server, then it can be run without any LDAP credentials (**-D** and **-w**). The IdM user uses Kerberos credentials to authenticate to the domain. To run the command directly on the disabled host, then supply

LDAP credentials to authenticate to the IdM server. The credentials should correspond to the host or service which is being re-enabled.

12.5. MANAGING PUBLIC SSH KEYS FOR HOSTS

OpenSSH uses *public keys* to authenticate hosts. One machine attempts to access another machine and presents its key pair. The first time the host authenticates, the administrator on the target machine has to approve the request manually. The machine then stores the host's public key in a **known_hosts** file. Any time that the remote machine attempts to access the target machine again, the target machine simply checks its **known_hosts** file and then grants access automatically to approved hosts.

There are a few problems with this system:

- The **known_hosts** file stores host entries in a triplet of the host IP address, host name, and key. This file can rapidly become out of date if the IP address changes (which is common in virtual environments and data centers) or if the key is updated.
- SSH keys have to be distributed manually and separately to all machines in an environment.
- Administrators have to approve host keys to add them to the configuration, but it is difficult to verify either the host or key issuer properly, which can create security problems.

On Red Hat Enterprise Linux, the System Security Services Daemon (SSSD) can be configured to cache and retrieve host SSH keys so that applications and services only have to look in one location for host keys. Because SSSD can use Identity Management as one of its identity information providers, Identity Management provides a universal and centralized repository of keys. Administrators do not need to worry about distributing, updating, or verifying host SSH keys.

12.5.1. About the SSH Key Format

When keys are uploaded to the IdM entry, the key format can be either an [OpenSSH-style key](#) or a raw [RFC 4253-style blob](#). Any RFC 4253-style key is automatically converted into an OpenSSH-style key before it is imported and saved into the IdM LDAP server.

The IdM server can identify the type of key, such as an RSA or DSA key, from the uploaded key blob. However, in a key file such as `~/.ssh/known_hosts`, a key entry is identified by the host name and IP address of the server, its type, then lastly the key itself. For example:

```
host.example.com,1.2.3.4 ssh-rsa AAA...ZZZ==
```

This is slightly different than a user public key entry, which has the elements in the order *type key== comment*:

```
"ssh-rsa ABCD1234...== ipaclient.example.com"
```

All three parts from the key file can be uploaded to and viewed for the host entry. In that case, the host public key entry from the `~/.ssh/known_hosts` file needs to be reordered to match the format of a user key, *type key== comment*:

```
ssh-rsa AAA...ZZZ== host.example.com,1.2.3.4
```

The key type can be determined automatically from the content of the public key, and the comment is optional, to make identifying individual keys easier. The only required element is the public key blob itself.

12.5.2. About ipa-client-install and OpenSSH

The **ipa-client-install** script, by default, configures an OpenSSH server and client on the IdM client machine. It also configures SSSD to perform host and user key caching. Essentially, simply configuring the client does all of the configuration necessary for the host to use SSSD, OpenSSH, and Identity Management for key caching and retrieval.

If the SSH service is enabled with the client installation (which is the default), then an RSA key is created when the **ssh** service is first started.



NOTE

When the machine is added as an IdM client using **ipa-client-install**, the client is created with two SSH keys, RSA and DSS.

There is an additional client configuration option, **--ssh-trust-dns**, which can be run with **ipa-client-install** and automatically configures OpenSSH to trust the IdM DNS records, where the key fingerprints are stored.

Alternatively, it is possible to disable OpenSSH at the time the client is installed, using the **--no-sshd** option. This prevents the install script from configuring the OpenSSH server.

Another option, **--no-dns-sshfp**, prevents the host from creating DNS SSHFP records with its own DNS entries. This can be used with or without the **--no-sshd** option.

12.5.3. Uploading Host SSH Keys Through the Web UI

1. The key for a host can probably be retrieved from a `~/.ssh/known_hosts`. For example:

```
server.example.com,1.2.3.4 ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAQEApxjBvSFskTU0WQW4e0weeo0DZZ08F9Ud21xLLy6F
0hzwpxFGIyxvXZ52+siHBHbbqGL5+14N7UvElruyslIHx9LYUR/pPKSMXCGyboLy5aTN
l50Q5EHwrhVnFDIKXkvp45945R7SKYCUtRumm0Iw6wq0XD4o+ILeVbV3wmcB1bXs36Zv
C/M6riefn9PcJmh6vNCvIsbMY6S+FhkWUTTi0XJjUDYRLlwM273FfWhzHK+SSQXeBp/z
InlgFvJhSZMRi9HZpDoqxLbBB9QIdIw6U4MIjNmKsSI/ASpkFm2GuQ7ZK9KuMItY2AoC
uIRmRADf8iYNHBTXNfFurGogXwRDjQ==
```

If necessary, generate a host key. When using the OpenSSH tools, make sure to use a blank passphrase and to save the key to a different location than the user's `~/.ssh/` directory, so it will not overwrite any existing keys.

```
[jsmith@server ~]$ ssh-keygen -t rsa -C "server.example.com,1.2.3.4"
Generating public/private rsa key pair.
Enter file in which to save the key (/home/jsmith/.ssh/id_rsa):
/home/jsmith/.ssh/host_keys
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/jsmith/.ssh/host_keys.
Your public key has been saved in /home/jsmith/.ssh/host_keys.pub.
```

```

The key fingerprint is:
SHA256:GAUIDVVEgly7rs1lTWP6oguHz8BKvyZkpqCqVSsmi7c
server.example.com
The key's randomart image is:
+--[ RSA 2048 ]-----+
|           ..        |
|            .+       |
|             .*      |
|            0 . . . * |
|           S + . 0+  |
|            E . . . . |
|           . = . 0   |
|            0 . . . 0 |
|             .....   |
+-----+

```

2. Copy the public key from the key file. The full key entry has the form *host name,IP type key==*. Only the *key==* is required, but the entire entry can be stored. To use all elements in the entry, rearrange the entry so it has the order *type key== [host name,IP]*

```

[jsmith@server ~]$ cat /home/jsmith/.ssh/host_keys.pub

ssh-rsa AAAAB3NzaC1yc2E...tJG1PK2Mq++wQ== server.example.com,1.2.3.4

```

3. Open the **Identity** tab, and select the **Hosts** subtab.
4. Click the name of the host to edit.

Hosts			
Search		Refresh	Delete
+ Add		Actions v	
☐	Host name	Description	Enrolled
☐	server.example.com		True
Showing 1 to 1 of 1 entries.			

Figure 12.4. List of Hosts

5. In the **Host Settings** area of the **Settings** tab, click **Add** next to **SSH public keys**.

Host Settings

Host name `server.example.com`

Principal name `host/server.example.com@EXAMPLE.COM`

Description

SSH public keys

existing_ssh_key

Figure 12.5. Adding an SSH Key

- Paste in the public key for the host, and click **Set**.

Set SSH key

SSH public key:

```
ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQACro6NZA2FfjeSdMFLLNzw+KnUjksNqSGBSePpryTxfE0Xw9NS
0gQ7blzgopL4N/3f4g
/M5dik3GxkUX00gcM0CVFf961TETmNvYam6Sn7r++1IY2SSmK4GpIfKTr40vT+mnUeZs6aFEIiRKLNy
x5EgkyTXSu5QT
/AfcDluo9hdu42XvmU0ZCYE3460eaNQ5uVCTJmazhJScdFhwesruUtKCKcoIHSu6gZeoAr5PHuJfni+
XIVsLK5V/oRuc0sqAKpKVEF8U5DGANB6VdaQoqTQko5PS0q3HEzJ54DE5mLE3wqURNnrrfX
/R3+TF+b1GXpHs7pKD3Ugo08f0HNT8801
server.example.com
```

Figure 12.6. Setting an SSH Key

The **SSH public keys** area now shows the new key. Clicking **Show/Set key** opens the submitted key.

- To upload multiple keys, click the **Add** link below the list of public keys, and upload the other keys.

8. When all the keys have been submitted, click **Save** at the top of the host's page to save the changes.

When the public key is saved, the entry is displayed as the key fingerprint, the comment (if one was included), and the key type^[2].

After uploading the host keys, configure SSSD to use Identity Management as one of its identity domains and set up OpenSSH to use the SSSD tooling for managing host keys. This is covered in [the "Configuring Services: OpenSSH and Cached Keys" in the System-Level Authentication Guide](#).

12.5.4. Adding Host Keys from the Command Line

Host SSH keys are added to host entries in IdM, either when the host is created using **host-add** or by modifying the entry later.



NOTE

RSA and DSS host keys are created by the **ipa-client-install** command, unless the SSH service is explicitly disabled in the installation script.

1. Run the **host-mod** command with the **--sshpubkey** option to upload the base64-encoded public key to the host entry.

Adding a host key also changes the DNS SSHFP entry for the host, so also use the **-updatedns** option to update the host's DNS entry.

For example:

```
[jsmith@server ~]$ ipa host-mod --sshpubkey="ssh-rsa RjlzYQo==" --  
updatedns host1.example.com
```

A real key also usually ends with an equal sign (=) but is longer.

To upload more than one key, enter multiple **--sshpubkey** command-line parameters:

```
--sshpubkey="RjlzYQo==" --sshpubkey="ZEt0TAo=="
```



NOTE

A host can have multiple public keys.

2. After uploading the host keys, configure SSSD to use Identity Management as one of its identity domains and set up OpenSSH to use the SSSD tooling for managing host keys. This is covered in [the "Configuring Services: OpenSSH and Cached Keys" in the System-Level Authentication Guide](#).

12.5.5. Removing Host Keys

Host keys can be removed once they expire or are no longer valid.

To remove an individual host key, it is easiest to remove the key through the web UI:

1. Open the **Identity** tab, and select the **Hosts** subtab.
2. Click the name of the host to edit.




Figure 12.7. List of Hosts

3. In the **SSH public keys** area, click **Delete** by the fingerprint of the key to remove it.




Figure 12.8. Public Key Deletion

4. Click **Save** at the top of the host's page to save the changes.

The command-line tools can be used to remove all keys. This is done by running **ipa host-mod** with the **--sshpubkey=** set to a blank value; this removes *all* public keys for the host. Also, use the **--updatedns** option to update the host's DNS entry. For example:

```
[jsmith@server ~]$ kinit admin
[jsmith@server ~]$ ipa host-mod --sshpubkey= --updatedns host1.example.com
```

12.6. SETTING **ETHERS** INFORMATION FOR A HOST

NIS can host an **ethers** table which can be used to manage DHCP configuration files for systems based on their platform, operating system, DNS domain, and MAC address — all information stored in host entries in IdM.

In Identity Management, each system is created with a corresponding **ethers** entry in the directory, in the **ou=ethers** subtree.

```
cn=server,ou=ethers,dc=example,dc=com
```

This entry is used to create a NIS map for the **ethers** service which can be managed by the NIS compatibility plug-in in IdM.

To configure NIS maps for **ethers** entries:

1. Add the MAC address attribute to a host entry. For example:

```
[jsmith@server ~]$ kinit admin
[jsmith@server ~]$ ipa host-mod --macaddress=12:34:56:78:9A:BC
server.example.com
```

2. Open the **nsswitch.conf** file.
3. Add a line for the **ethers** service, and set it to use LDAP for its lookup.

```
ethers: ldap
```

4. Check that the **ethers** information is available for the client.

```
[root@server ~]# getent ethers server.example.com
```

[2] The key type is determined automatically from the key itself, if it is not included in the uploaded key.

CHAPTER 13. MANAGING USER AND HOST GROUPS

13.1. HOW USER AND HOST GROUPS WORK IN IDM

13.1.1. What User and Host Groups Are

A user group is a set of users with common privileges, password policies, and other characteristics.

A host group is a set of IdM hosts with common access control rules and other characteristics.

For example, you can define groups around company departments, physical locations, or access control requirements.

13.1.2. Supported Group Members

A user group in IdM can include:

- IdM users
- other IdM user groups
- external users, which are users that exist outside IdM

A host group in IdM can include:

- IdM servers and clients
- other IdM host groups

13.1.3. Direct and Indirect Group Members

User and host group attributes in IdM apply to both direct and indirect members: when group B is a member of group A, all users in group B are considered members of group A.

For example, in [Figure 13.1, “Direct and Indirect Group Membership”](#):

- User 1 and User 2 are *direct members* of group A.
- User 3, User 4, and User 5 are *indirect members* of group A.




Figure 13.1. Direct and Indirect Group Membership

If you set a password policy for user group A, the policy applies to all users in user group B as well.

Example 13.1. Viewing Direct and Indirect Group Members

1. Create two groups: **group_A** and **group_B**. See [Section 13.2, “Adding and Removing User or Host Groups”](#).
2. Add:
 - one user as a member of **group_A**
 - another user as a member of **group_B**
 - **group_B** as a member of **group_A**

See [Section 13.3, “Adding and Removing User or Host Group Members”](#).

3. In the web UI: Select **Identity** → **Groups**. From the individual group types which are listed in a side bar on the left, select **User Groups**, and click the name of **group_A**. Switch between **Direct Membership** and **Indirect Membership**.




Figure 13.2. Indirect and Direct Members

4. From the command line: Use the **ipa group-show** command:

```
$ ipa group-show group_A
...
Member users: user_1
Member groups: group_B
Indirect Member users: user_2
```

The list of indirect members does not include external users from trusted Active Directory domains. The Active Directory trust user objects are not visible in the IdM interface because they do not exist as LDAP objects within IdM.

13.1.4. User Group Types in IdM

POSIX groups (the default)

POSIX groups support POSIX attributes for their members. Note that groups that interact with Active Directory cannot use POSIX attributes.

Non-POSIX groups

All group members of this type of group must belong to the IdM domain.

External groups

External groups allow adding group members that exist in an identity store outside of the IdM domain. The external store can be a local system, an Active Directory domain, or a directory service.

Non-POSIX and external groups do not support POSIX attributes. For example, these groups do not have a GID defined.

Example 13.2. Searching for Different Types of User Groups

1. Run the **ipa group-find** command to display all user groups.
2. Run the **ipa group-find --posix** command to display all POSIX groups.
3. Run the **ipa group-find --nonposix** command to display all non-POSIX groups.
4. Run the **ipa group-find --external** command to display all external groups.

13.1.5. User and Host Groups Created by Default

Table 13.1. User and Host Groups Created by Default

Group Name	User or Host	Default Group Members
ipausers	User group	All IdM users
admins	User group	Users with administrative privileges, initially the default admin user
editors	User group	Users allowed to edit other IdM users in the web UI, without all the rights of an administrative user
trust admins	User group	Users with privileges to manage Active Directory trusts
ipaservers	Host group	All IdM server hosts

Adding a user to a user group applies the privileges and policies associated with the group. For example, adding a user to the **admins** group grants the user administrative privileges.

**WARNING**

Be careful when adding hosts to the **ipaservers** host group. All hosts in **ipaservers** have the ability to promote themselves to an IdM server.

In addition, IdM creates *user private groups* by default whenever a new user is created in IdM.

- The user private group has the same name as the user for which it was created.
- The user is the only member of the user private group.
- GID of the private groups matches the UID of the user.

Example 13.3. Viewing User Private Groups

Run the **ipa group-find --private** command to display all user private groups:

```
$ ipa group-find --private
-----
2 groups matched
-----
Group name: user1
Description: User private group for user1
GID: 830400006

Group name: user2
Description: User private group for user2
GID: 830400004
-----
Number of entries returned 2
-----
```

In some situations, it is better to avoid creating user private groups, such as when a NIS group or another system group already uses the GID that would be assigned to the user private group. See [Section 13.4, “Disabling User Private Groups”](#).

13.2. ADDING AND REMOVING USER OR HOST GROUPS

To add a group, you can use:

- The web UI (see [the section called “Web UI: Adding a User or Host Group”](#))
- The command line (see [the section called “Command Line: Adding a User or Host Group”](#))

IdM enables specifying a custom GID when creating a user group. If you do this, be careful to avoid ID conflicts. See [Section 14.6, “Ensuring That ID Values Are Unique”](#). If you do not specify a custom GID, IdM automatically assigns a GID from the available ID range.

To remove a group, you can use:

- The web UI (see [the section called “Web UI: Removing a User or Host Group”](#))
- The command line (see [the section called “Command Line: Removing a User or Host Group”](#))

Note that removing a group does not delete the group members from IdM.

Web UI: Adding a User or Host Group

1. Click **Identity** → **Groups**, and select **User Groups** or **Host Groups** in the left sidebar.
2. Click **Add** to start adding the group.
3. Fill out the information about the group.

For details on user group types, see [Section 13.1.4, “User Group Types in IdM”](#).

4. Click **Add** to confirm.

Command Line: Adding a User or Host Group

1. Log in as the administrator:

```
$ kinit admin
```

2. To add a user group, use the **ipa group-add** command. To add a host group, use the **ipa hostgroup-add** command.

```
$ ipa group-add group_name
-----
Added group "group_name"
-----
```

By default, **ipa group-add** adds a POSIX user group. To specify a different group type, add options to **ipa group-add**:

- **--nonposix** to create a non-POSIX group
- **--external** to create an external group

For details on group types, see [Section 13.1.4, “User Group Types in IdM”](#).

Web UI: Removing a User or Host Group

1. Click **Identity** → **Groups** and select **User Groups** or **Host Groups** in the left sidebar.
2. Select the group to remove, and click **Delete**.

Command Line: Removing a User or Host Group

1. Log in as the administrator:

```
$ kinit admin
```


2. To delete a user group, use the **ipa group-del *group_name*** command. To delete a host group, use the **ipa hostgroup-del *group_name*** command.

```
$ ipa group-del group_name
-----
Deleted group "group_name"
-----
```

13.3. ADDING AND REMOVING USER OR HOST GROUP MEMBERS

To add members to user groups, you can use:

- The IdM web UI (see [the section called “Web UI: Adding a Member to a User or Host Group”](#))
- The command line (see [the section called “Command Line: Adding a Member to a User Group”](#))



IMPORTANT

When adding another user group as a member, do not create recursive groups. For example, if Group A is a member of Group B, do not add Group B as a member of Group A. Recursive groups can cause unpredictable behavior.

To remove members from user groups, you can use:

- The IdM web UI (see [the section called “Web UI: Removing a Member from a User Group”](#))
- The command line (see [the section called “Command Line: Removing a Member from a User Group”](#))

Web UI: Adding a Member to a User or Host Group

1. Click **Identity** → **Groups** and select **User Groups** or **Host Groups** in the left sidebar.
2. Click the name of the group.
3. Select the type of group member you want to add. For example, **Users**, **User Groups**, or **External** for user groups.

User Group: group

group members:

Users User Groups External

Refresh Delete Add

group is a member of:

User Groups Netgroups Roles

☐ User login	UID	Email address
No entries.		

Figure 13.3. Adding User Group Members

4. Click **Add**.
5. Select the member you want to add, and click **Add** to confirm.

Command Line: Adding a Member to a User Group

1. *Optional.* Use the **ipa group-find** or **ipa hostgroup-find** command to find the group.
2. To add a member to a user group, use the **ipa group-add-member** command. To add a member to a host group, use the **ipa hostgroup-add-member** command.

When adding a user group member, specify the member using these options:

- **--users** adds an IdM user
- **--external** adds a user that exists outside the IdM domain, in the format of **DOMAIN\user_name** or **user_name@domain**
- **--groups** adds an IdM user group

When adding a host group member, specify the member using these options:

- **--hosts** adds an IdM host
- **--groups** adds an IdM host group

Example 13.4. Example commands for adding a member to a user group

To add *user1*, *user2*, and *group1* to a group named *group_name*:

```
$ ipa group-add-member group_name --users=user1 --users=user2 --groups=group1
```

To add *ad_user* from a domain named *ad_domain* to a group named *group_name*, you can choose how to specify the external user. For example:

```
$ ipa group-add-member group_name --external='AD_DOMAIN\ad_user'
$ ipa group-add-member group_name --external='ad_user@AD_DOMAIN'
$ ipa group-add-member group_name --external='ad_user@AD_DOMAIN.EXAMPLE.COM'
```

Web UI: Removing a Member from a User Group

1. Click **Identity** → **Groups** and select **User Groups** or **Host Groups** in the left sidebar.
2. Click the name of the group.
3. Select the type of group member you want to remove. For example, **Users**, **User Groups**, or **External** for user groups.




Figure 13.4. Removing User Group Members

4. Select the check box next to the required member.
5. Click **Delete**.

Command Line: Removing a Member from a User Group

1. *Optional.* Use the **ipa group-show** or **ipa hostgroup-show** command to confirm that the group includes the member you want to remove.
2. To remove a user group member, use the **ipa group-remove-member** command. To remove a host group member, use the **ipa hostgroup-remove-member** command.

When removing a user group member, specify the member using these options:

- **--users** removes an IdM user
- **--external** removes a user that exists outside the IdM domain, in the format of **DOMAIN\user_name** or **user_name@domain**
- **--groups** removes an IdM user group

When removing a host group member, specify the member using these options:

- **--hosts** removes an IdM host
- **--groups** removes an IdM host group

For example, to remove *user1*, *user2*, and *group1* from a group called *group_name*:

```
$ ipa group-remove-member group_name --users=user1 --users=user2 --groups=group1
```

13.4. DISABLING USER PRIVATE GROUPS

To ensure that IdM does not create a default user private group for a new user, choose one of the following:

- [Section 13.4.1, “Creating a User without a User Private Group”](#)
- [Section 13.4.2, “Disabling User Private Groups Globally for All Users”](#)

Even after you disable creating default user private groups, IdM will still require a GID when adding new users. To ensure that adding the user succeeds, see [Section 13.4.3, “Adding a User with User Private Groups Disabled”](#).



NOTE

If you want to disable creating default user private groups because of GID conflicts, consider changing the default UID and GID assignment ranges. See [Chapter 14, *Unique UID and GID Number Assignments*](#).

13.4.1. Creating a User without a User Private Group

Add the **--noprivate** option to the **ipa user-add** command. Note that for the command to succeed, you must specify a custom private group. See [Section 13.4.3, “Adding a User with User Private Groups Disabled”](#).

13.4.2. Disabling User Private Groups Globally for All Users

1. Log in as the administrator:

```
$ kinit admin
```

2. IdM uses the Directory Server Managed Entries Plug-in to manage user private groups. List the instances of the plug-in:

```
$ ipa-managed-entries --list
```

3. To ensure IdM does not create user private groups, disabling the plug-in instance responsible for managing user private groups:

```
$ ipa-managed-entries -e "UPG Definition" disable
Disabling Plugin
```



NOTE

To re-enable the **UPG Definition** instance later, use the **ipa-managed-entries -e "UPG Definition" enable** command.

4. Restart Directory Server to load the new configuration.

```
# systemctl restart dirsrv.target
```

13.4.3. Adding a User with User Private Groups Disabled

To make sure adding a new user succeeds when creating default user private groups is disabled, choose one of the following:

- Specify a custom GID when adding a new user. The GID does not have to correspond to an already existing user group.

For example, when adding a user from the command line, add the **--gid** option to the **ipa user-add** command.

- Use an automember rule to add the user to an existing group with a GID. See [Section 13.6, “Defining Automatic Group Membership for Users and Hosts”](#).

13.5. SETTING SEARCH ATTRIBUTES FOR USERS AND USER GROUPS

When searching entries for a specified keyword using the **ipa user-find *keyword*** and **ipa group-find *keyword*** commands, IdM only searches certain attributes. Most notably:

- In user searches: first name, last name, user name (login ID), job title, organization unit, phone number, UID, email address.
- In group searches: group name, description.

The following procedure shows how to configure IdM to search other attributes as well. Note that IdM always searches the default attributes. For example, even if you remove the job title attribute from the list of user search attributes, IdM will still search user titles.

Prerequisites

Before adding a new attribute, make sure that a corresponding index exists within the LDAP directory for this attribute. Most standard LDAP attributes have indexes in LDAP, but if you want to add a custom attribute, you must create an index manually. See [Creating Standard Indexes](#) in the *Directory Server Administration Guide*.

Web UI: Setting Search Attributes

1. Select **IPA Server** → **Configuration**.
2. In the **User Options** area, set the user search attributes in **User search fields**.
3. In the **Group Options** area, set the group search attributes in **Group search fields**.
4. Click **Save** at the top of the page.

Command Line: Setting Search Attributes

Use the **ipa config-mod** command with these options:

- **--usersearch** defines a new list of search attributes for users
- **--groupsearch** defines a new list of search attributes for groups

For example:

```
$ ipa config-mod --usersearch={uid,givenname,sn,telephonenumber,ou,title}
$ ipa config-mod --groupsearch={cn,description}
```

13.6. DEFINING AUTOMATIC GROUP MEMBERSHIP FOR USERS AND HOSTS

13.6.1. How Automatic Group Membership Works in IdM

13.6.1.1. What Automatic Group Membership Is

Using automatic group membership, you can assign users and hosts to groups automatically based on their attributes. For example, you can:

- Divide employees' user entries into groups based on the employees' manager, location, or any other attribute.
- Divide hosts based on their class, location, or any other attribute.
- Add all users or all hosts to a single global group.

13.6.1.2. Benefits of Automatic Group Membership

Reduced overhead of managing group membership manually

With automatic group membership, the administrator no longer assigns users and hosts to groups manually.

Improved consistency in user and host management

With automatic group membership, users and hosts are assigned to groups based on strictly defined and automatically evaluated criteria.

Easier management of group-based settings

Various settings are defined for groups and then applied to individual group members, for example **sudo** rules, **automount**, or access control. When using automatic group membership, users and hosts are automatically added to specified groups, which makes managing group-based settings easier.

13.6.1.3. Automember Rules

When configuring automatic group membership, the administrator defines *automember rules*. An automember rule applies to a specific user or host group. It includes *conditions* that the user or host must meet to be included or excluded from the group:

Inclusive conditions

When a user or host entry meets an inclusive condition, it will be included in the group.

Exclusive conditions

When a user or host entry meets an exclusive condition, it will *not* be included in the group.

The conditions are specified as regular expressions in the Perl-compatible regular expressions (PCRE) format. For more information on PCRE, see the `pcresyntax(3)` man page.

IdM evaluates exclusive conditions before inclusive conditions. In case of a conflict, exclusive conditions take precedence over inclusive conditions.

13.6.2. Adding an Automember Rule

To add an automember rule using:

- The IdM web UI, see [the section called “Web UI: Add an Automember Rule”](#)
- The command line, see [the section called “Command Line: Add an Automember Rule”](#)

After you add an automember rule:

- All entries created in the future will become members of the specified group. If an entry meets conditions specified in multiple automember rules, it will be added to all the corresponding groups.
- Existing entries will *not* become members of the specified group. See [Section 13.6.3, “Applying Automember Rules to Existing Users and Hosts”](#) for more information.

Web UI: Add an Automember Rule

1. Select **Identity** → **Automember** → **User group rules** or **Host group rules**.
2. Click **Add**.
3. In the **Automember rule** field, select the group to which the rule will apply. Click **Add and Edit**.
4. Define one or more inclusive and exclusive conditions. See [Section 13.6.1.3, “Automember Rules”](#) for details.
 - a. In the **Inclusive** or **Exclusive** sections, click **Add**.
 - b. In the **Attribute** field, select the required attribute.
 - c. In the **Expression** field, define the regular expression.
 - d. Click **Add**.

For example, the following condition targets all users with any value (.*) in their user login attribute (**uid**).

Figure 13.5. Adding Automember Rule Conditions

Command Line: Add an Automember Rule

1. Use the **ipa automember-add** command to add an automember rule. When prompted, specify:
 - **Automember rule**, which matches the target group name.
 - **Grouping Type**, which specifies whether the rule targets a user group or a host group. To target a user group, enter **group**. To target a host group, enter **hostgroup**.

For example, to add an automember rule for a user group named **user_group**:

```
$ ipa automember-add
Automember Rule: user_group
Grouping Type: group
-----
Added automember rule "user_group"
-----
Automember Rule: user_group
```

2. Define one or more inclusive and exclusive conditions. See [Section 13.6.1.3, “Automember Rules”](#) for details.
 - a. To add a condition, use the **ipa automember-add-condition** command. When prompted, specify:
 - **Automember rule**, which matches the target group name.
 - **Attribute Key**, which specifies the entry attribute to which the filter will apply. For example, **manager** for users.
 - **Grouping Type**, which specifies whether the rule targets a user group or a host group. To target a user group, enter **group**. To target a host group, enter **hostgroup**.

- **Inclusive regex** and **Exclusive regex**, which specify one or more conditions as regular expressions. If you only want to specify one condition, press **Enter** when prompted for the other.

For example, the following condition targets all users with any value (.*) in their user login attribute (**uid**).

```
$ ipa automember-add-condition
Automember Rule: user_group
Attribute Key: uid
Grouping Type: group
[Inclusive Regex]: .*
[Exclusive Regex]:
-----
Added condition(s) to "user_group"
-----
Automember Rule: user_group
Inclusive Regex: uid=.*
-----
Number of conditions added 1
-----
```

- To remove a condition, use the **ipa automember-remove-condition** command.

Example 13.5. Command Line: Creating an Automember Rule to Add All Entries to a Single Group

By creating an inclusive condition for an attribute that all user or host entries contain, such as **cn** or **fqdn**, you can ensure that all users or hosts created in the future will be added to a single group.

1. Create the group, such as a host group named **all_hosts**. See [Section 13.2, “Adding and Removing User or Host Groups”](#).
2. Add an automember rule for the new host group. For example:

```
$ ipa automember-add
Automember Rule: all_hosts
Grouping Type: hostgroup
-----
Added automember rule "all_hosts"
-----
Automember Rule: all_hosts
```

3. Add an inclusive condition that targets all hosts. In the following example, the inclusive condition targets hosts that have any value (.*) in the **fqdn** attribute:

```
$ ipa automember-add-condition
Automember Rule: all_hosts
Attribute Key: fqdn
Grouping Type: hostgroup
[Inclusive Regex]: .*
[Exclusive Regex]:
-----
Added condition(s) to "all_hosts"
```

```

-----
Automember Rule: all_hosts
Inclusive Regex: fqdn=.*
-----
Number of conditions added 1
-----

```

All hosts added in the future will automatically become members of the **all_hosts** group.

Example 13.6. Command Line: Creating an Automember Rule for Synchronized AD Users

Windows users synchronized from Active Directory (AD) share the **ntUser** object class. By creating an automember condition that targets all users with **ntUser** in their **objectclass** attribute, you can ensure that all synchronized AD users created in the future will be included in a common group for AD users.

1. Create a user group for the AD users, such as **ad_users**. See [Section 13.2, “Adding and Removing User or Host Groups”](#).
2. Add an automember rule for the new user group. For example:

```

$ ipa automember-add
Automember Rule: ad_users
Grouping Type: group
-----
Added automember rule "ad_users"
-----
Automember Rule: ad_users

```

3. Add an inclusive condition to filter the AD users. In the following example, the inclusive condition targets all users that have the **ntUser** value in the **objectclass** attribute:

```

$ ipa automember-add-condition
Automember Rule: ad_users
Attribute Key: objectclass
Grouping Type: group
[Inclusive Regex]: ntUser
[Exclusive Regex]:
-----
Added condition(s) to "ad_users"
-----
Automember Rule: ad_users
Inclusive Regex: objectclass=ntUser
-----
Number of conditions added 1
-----

```

All AD users added in the future will automatically become members of the **ad_users** user group.

13.6.3. Applying Automember Rules to Existing Users and Hosts

Automember rules apply automatically to user and hosts entries created after the rules were added. They are not applied retrospectively to entries that existed before the rules were added.

To apply automember rules to entries that existed before you added the rules, manually rebuild automatic membership. Rebuilding automatic membership re-evaluates all existing automember rules and applies them either to all entries or to specific entries.

Web UI: Rebuild Automatic Membership for Existing Entries

To rebuild automatic membership for all users or all hosts:

1. Select **Identity** → **Users** or **Hosts**.
2. Click **Actions** → **Rebuild auto membership**.




Figure 13.6. Rebuilding Automatic Membership for All Users or Hosts

To rebuild automatic membership for a single user or host only:

1. Select **Identity** → **Users** or **Hosts**, and click on the required user login or host name.
2. Click **Actions** → **Rebuild auto membership**.




Figure 13.7. Rebuilding Automatic Membership for a Single User or Host

Command Line: Rebuild Automatic Memberhips for Existing Entries

To rebuild automatic membership for all users, use the **ipa automember-rebuild --type=group** command:

```
$ ipa automember-rebuild --type=group
-----
Automember rebuild task finished. Processed (9) entries.
-----
```

To rebuild automatic membership for all users, use the **ipa automember-rebuild --type=hostgroup** command.

To rebuild automatic membership for a specified user or users, use the **ipa automember-rebuild --users=user** command:

```
$ ipa automember-rebuild --users=user1 --users=user2
-----
Automember rebuild task finished. Processed (2) entries.
-----
```

To rebuild automatic membership for a specified host or hosts, use the **ipa automember-rebuild --hosts=example.com** command.

13.6.4. Configuring a Default Automember Group

When a default automember group is configured, user or host entries that do not match any automember rule are automatically added to the default group.

1. Use the **ipa automember-default-group-set** command to configure a default automember group. When prompted, specify:
 - **Default (fallback) Group**, which specifies the target group name.

- **Grouping Type**, which specifies whether the target is a user group or a host group. To target a user group, enter **group**. To target a host group, enter **hostgroup**.

For example:

```
$ ipa automember-default-group-set
Default (fallback) Group: default_user_group
Grouping Type: group
-----
Set default (fallback) group for automember "default_user_group"
-----
Default (fallback) Group:
cn=default_user_group,cn=groups,cn=accounts,dc=example,dc=com
```

2. To verify that the group is set correctly, use the **ipa automember-default-group-show** command. The command displays the current default automember group. For example:

```
$ ipa automember-default-group-show
Grouping Type: group
Default (fallback) Group:
cn=default_user_group,cn=groups,cn=accounts,dc=example,dc=com
```

To remove the current default automember group, use the **ipa automember-default-group-remove** command.

CHAPTER 14. UNIQUE UID AND GID NUMBER ASSIGNMENTS

An IdM server generates user ID (UID) and group ID (GID) values and simultaneously ensures that replicas never generate the same IDs. The need for unique UIDs and GIDs might even be across IdM domains, if a single organization uses multiple separate domains.

14.1. ID RANGES

The UID and GID numbers are divided into *ID ranges*. By keeping separate numeric ranges for individual servers and replicas, the chances are minimal that an ID value issued for an entry is already used by another entry on another server or replica.

The Distributed Numeric Assignment (DNA) plug-in, as part of the back end 389 Directory Server instance for the domain, ensures that ranges are updated and shared between servers and replicas; the plug-in manages the ID ranges across all masters and replicas. Every server or replica has a current ID range and an additional **next** ID range that the server or replica uses after the current range has been depleted. For more information about the DNA Directory Server plug-in, see the [Red Hat Directory Server Deployment Guide](#).

14.2. ID RANGE ASSIGNMENTS DURING INSTALLATION

During server installation, the **ipa-server-install** command by default automatically assigns a random current ID range to the installed server. The setup script randomly selects a range of 200,000 IDs from a total of 10,000 possible ranges. Selecting a random range in this way significantly reduces the probability of conflicting IDs in case you decide to merge two separate IdM domains in the future.

However, you can define a current ID range manually during server installation by using the following two options with **ipa-server-install**:

- **--idstart** gives the starting value for UID and GID numbers; by default, the value is selected at random,
- **--idmax** gives the maximum UID and GID number; by default, the value is the **--idstart** starting value plus 199,999.

If you have a single IdM server installed, a new user or group entry receives a random ID from the whole range. When you install a new replica and the replica requests its own ID range, the initial ID range for the server splits and is distributed between the server and replica: the replica receives half of the remaining ID range that is available on the initial master. The server and replica then use their respective portions of the original ID range for new entries. Also, if less than 100 IDs from the ID range that was assigned to a replica remain, meaning the replica is close to depleting its allocated ID range, the replica contacts the other available servers with a request for a new ID range.

A server receives an ID range the first time the DNA plug-in is used; until then, the server has no ID range defined. For example, when you create a replica from a master server, the replica does not receive an ID range immediately. The replica requests an ID range from the initial master only when the first ID is about to be assigned on the replica.

**NOTE**

If the initial master stops functioning before the replica requests an ID range from it, the replica is unable to contact the master with a request for the ID range. An attempt to add a new user on the replica fails. In such situations, you can find out what ID range is assigned to the disabled master and assign an ID range to the replica manually, which is described in [Section 14.5, “Manual ID Range Extension and Assigning a New ID Range”](#).

14.3. DISPLAYING CURRENTLY ASSIGNED ID RANGES

To display which ID ranges are configured for a server, use the following commands:

- **ipa-replica-manage dnarange-show** displays the current ID range that is set on all servers or, if you specify a server, only on the specified server, for example:

```
# ipa-replica-manage dnarange-show
masterA.example.com: 1001-1500
masterB.example.com: 1501-2000
masterC.example.com: No range set

# ipa-replica-manage dnarange-show masterA.example.com
masterA.example.com: 1001-1500
```

- **ipa-replica-manage dnanextrange-show** displays the next ID range currently set on all servers or, if you specify a server, only on the specified server, for example:

```
# ipa-replica-manage dnanextrange-show
masterA.example.com: 1001-1500
masterB.example.com: No on-deck range set
masterC.example.com: No on-deck range set

# ipa-replica-manage dnanextrange-show masterA.example.com
masterA.example.com: 1001-1500
```

For more information about these two commands, see the `ipa-replica-manage(1)` man page.

14.4. AUTOMATIC ID RANGE EXTENSION AFTER DELETING A REPLICA

When you delete a functioning replica, the **ipa-replica-manage del** command retrieves the ID ranges that were assigned to the replica and adds them as a next range to other available IdM replicas. This ensures that ID ranges remain available to be used by other replicas.


After you delete a replica, you can verify which ID ranges are configured for other servers by using the **ipa-replica-manage dnarange-show** and **ipa-replica-manage dnanextrange-show** commands, described in [Section 14.3, “Displaying Currently Assigned ID Ranges”](#).

14.5. MANUAL ID RANGE EXTENSION AND ASSIGNING A NEW ID RANGE

In certain situations, it is necessary to manually adjust an ID range:

An assigned ID range has been depleted

A replica has exhausted the ID range that was assigned to it, and requesting additional IDs failed because no more free IDs are available in the ID ranges of other replicas. You want to extend the ID range assigned to the replica. This might involve splitting an existing ID range or extending it past the initial configured ID range for the server. Alternatively, you might want to assign a new ID range.



NOTE

If you assign a new ID range, the UIDs of the already existing entries on the server or replica stay the same. This does not pose a problem because even if you change the current ID range, IdM keeps a record of what ranges were assigned in the past.

A replica stopped functioning

ID range is not automatically retrieved when a replica dies and needs to be deleted, which means the ID range previously assigned to the replica becomes unavailable. You want to recover the ID range and make it available for other replicas.

If you want to recover the ID range belonging to a server that stopped functioning and assign it to another server, first find out what are the ID range values using the **ipa-replica-manage dnrange-show** command described in [Section 14.3, “Displaying Currently Assigned ID Ranges”](#), and then manually assign that ID range to the server. Also, to avoid duplicate UIDs or GIDs, make sure that no ID value from the recovered range was previously assigned to a user or group; you can do this by examining the UIDs and GIDs of existent users and groups.

To manually define the ID ranges, use the following two commands:

- **ipa-replica-manage dnrange-set** allows you to define the current ID range for a specified server:

```
# ipa-replica-manage dnrange-set masterA.example.com 1250-1499
```

- **ipa-replica-manage dnanextrange-set** allows you to define the next ID range for a specified server:

```
# ipa-replica-manage dnanextrange-set masterB.example.com 1001-5000
```

For more information about these commands, see the `ipa-replica-manage(1)` man page.



IMPORTANT

Be careful not to create overlapping ID ranges. If any of the ID ranges you assign to servers or replicas overlap, it could result in two different servers assigning the same ID value to different entries.

Do not set ID ranges that include UID values of 1000 and lower; these values are reserved for system use. Also, do not set an ID range that would include the 0 value; the SSSD

service does not handle the 0 ID value.

When extending an ID range manually, make sure that the newly extended range is included in the IdM ID range; you can check this using the **ipa idrange-find** command. Run the **ipa idrange-find -h** command to display help for how to use **ipa idrange-find**.

14.6. ENSURING THAT ID VALUES ARE UNIQUE

It is recommended to avoid conflicting UIDs or GIDs. UIDs and GIDs should always be unique: two users should not have the same UID, and two groups should not have the same GID.

Automatic ID assignment

When a user or a group is created interactively or without a manually specified ID number, the server assigns the next available ID number from the ID range to the user account. This ensures that the UID or GID is always unique.

Manual ID assignment

When you assign an ID to a user or a group entry manually, the server does not verify that the specified UID or GID is unique; it does not warn you of a conflict if you choose a value that is already used by another entry.

As explained in [Section 14.7, “Repairing Changed UID and GID Numbers”](#), the SSSD service does not handle entries with identical IDs. If two entries share the same ID number, a search for this ID only returns the first entry. However, if you search for other attributes or run the **ipa user-find --all** command, both entries are returned.

UIDs and GIDs are both selected from the same ID range. A user and a group can have the same ID; no conflict arises in this situation because the UID and the GID are set in two different attributes: **uidNumber** and **gidNumber**.



NOTE

Setting the same ID for both a user and a group allows you to configure user private groups. To create a unique system group for a user in this way, set the same ID value for a user and also for a group, in which the only member is the mentioned user.

14.7. REPAIRING CHANGED UID AND GID NUMBERS

When a user logs into an IdM system or service, SSSD on that system caches their user name together with the UID and GID of the user. SSSD then uses the UID as the identifying key for the user. If a user with the same user name but a different UID attempts to log into the system, SSSD registers two different UIDs and assumes that there are two different users with conflicting user names. This can pose a problem if a UID of a user changes. In such a situation, SSSD incorrectly interprets the user with a modified UID as a new user, instead of recognizing that it is the same user with a different UID. If the UID of an existing user changes, the user cannot log into SSSD and associated services and domains. This also affects client applications that use SSSD for identity information.

To work around this problem, if a UID or GID changes, clear the SSSD cache, which ensures that the user is able to log in again. For example, to clear the SSSD cache for a specified user, use the **sss_cache** utility as follows:

```
[root@server ~]# sss_cache -u user
```

CHAPTER 15. USER AND GROUP SCHEMA

When a user entry is created, it is automatically assigned certain LDAP object classes which, in turn, make available certain attributes. LDAP attributes are the way that information is stored in the directory. (This is discussed in detail in the *Directory Server Deployment Guide* and the *Directory Server Schema Reference*.)

Table 15.1. Default Identity Management User Object Classes

Object Classes	Description
ipaobject ipasshuser	IdM object classes
person organizationalperson inetorgperson inetuser posixAccount	Person object classes
krbprincipalaux krbticketpolicyaux	Kerberos object classes
mepOriginEntry	Managed entries (template) object classes

A number of attributes are available to user entries. Some are set manually and some are set based on defaults if a specific value is not set. There is also an option to add any attributes available in the object classes in [Table 15.1, “Default Identity Management User Object Classes”](#), even if there is not a UI or command-line argument for that attribute. Additionally, the values generated or used by the default attributes can be configured, as in [Section 15.4, “Specifying Default User and Group Attributes”](#).

Table 15.2. Default Identity Management User Attributes

UI Field	Command-Line Option	Required, Optional, or Default ^[a]
User login	<i>username</i>	Required
First name	--first	Required
Last name	--last	Required
Full name	--cn	Optional
Display name	--displayname	Optional

UI Field	Command-Line Option	Required, Optional, or Default ^[a]
Initials	--initials	Default
Home directory	--homedir	Default
GECOS field	--gecos	Default
Shell	--shell	Default
Kerberos principal	--principal	Default
Email address	--email	Optional
Password	--password ^[b]	Optional
User ID number	--uid	Default
Group ID number	--gidnumber	Default
Street address	--street	Optional
City	--city	Optional
State/Province	--state	Optional
Zip code	--postalcode	Optional
Telephone number	--phone	Optional
Mobile telephone number	--mobile	Optional
Pager number	--pager	Optional
Fax number	--fax	Optional
Organizational unit	--orgunit	Optional
Job title	--title	Optional
Manager	--manager	Optional
Car license	--carlicense	Optional
	--noprivate	Optional

UI Field	Command-Line Option	Required, Optional, or Default ^[a]
SSH Keys	--sshpubkey	Optional
Additional attributes	--addattr	Optional
Department Number	--departmentnumber	Optional
Employee Number	--employeenumber	Optional
Employee Type	--employeetype	Optional
Preferred Language	--preferredlanguage	Optional
[a] Required attributes must be set for every entry. Optional attributes may be set, while default attributes are automatically added with a predefined value unless a specific value is given. [b] The script prompts for the new password, rather than accepting a value with the argument.		

15.1. ABOUT CHANGING THE DEFAULT USER AND GROUP SCHEMA

It is possible to add or change the object classes and attributes used for user and group entries ([Chapter 15, User and Group Schema](#)).

The IdM configuration provides some validation when object classes are changed:

- All of the object classes and their specified attributes must be known to the LDAP server.
- All default attributes that are configured for the entry must be supported by the configured object classes.

There are limits to the IdM schema validation, however. Most important, the IdM server does not check that the defined user or group object classes contain all of the required object classes for IdM entries. For example, all IdM entries require the **ipaobject** object class. However, when the user or group schema is changed, the server does not check to make sure that this object class is included; if the object class is accidentally deleted, then future entry add operations will fail.

Also, all object class changes are atomic, not incremental. The entire list of default object classes has to be defined every time there is a change. For example, a company may create a custom object class to store employee information like birthdays and employment start dates. The administrator cannot simply add the custom object class to the list; he must set the entire list of current default object classes *plus* the new object class. The *existing* default object classes must always be included when the configuration is updated. Otherwise, the current settings will be overwritten, which causes serious performance problems.

15.2. APPLYING CUSTOM OBJECT CLASSES TO NEW USER ENTRIES

User and group accounts are created with a predefined set of LDAP object classes applied to the entry. Any attributes which belong to the object class can be added to the user entry.


While the standard and IdM-specific LDAP object classes will cover most deployment scenarios, administrators can create custom object classes with custom attributes. Note that after an administrator modifies the list of default object classes, new entries will contain the custom object classes but the old entries are not automatically modified.

15.2.1. From the Web UI

1. Add all of the custom schema elements to the 389 Directory Server instance used by Identity Management. Adding schema elements is described in [the schema chapter of the Directory Server Administrator's Guide](#).
2. Open the **IPA Server** tab.
3. Select the **Configuration** subtab.
4. Scroll to the **User Options** area.


Figure 15.1. User Options in Server Configuration

5. At the bottom of the users area, click **Add** to include a new field for another object class.



IMPORTANT

Always include the *existing* default object classes when the configuration is updated. Otherwise, the current settings will be overwritten. If any object classes required by Identity Management are not included, then subsequent attempts to add an entry will fail with object class violations.



Default user *
objectclasses


ipaobject	Delete
person	Delete
inetuser	Delete
posixaccount	Delete
Add	

Figure 15.2. Changing Default User Object Classes

6. When the changes are complete, click **Save** at the top of the **Configuration** page.

15.2.2. From the Command Line

1. Add all of the custom schema elements to the 389 Directory Server instance used by Identity Management. Adding schema elements is described in [the schema chapter of the Directory Server Administrator's Guide](#).
2. Add the new object class to the list of object classes added to entries. The option for user object classes is **--userobjectclasses**.



IMPORTANT

Always include the *existing* default object classes when the configuration is updated. Otherwise, the current settings will be overwritten. If any object classes required by Identity Management are not included, then subsequent attempts to add an entry will fail with object class violations.

All object classes must be included in the list of object classes. The information passed with the **config-mod** command overwrites the previous values. This can be done by specifying each object class with a **--userobjectclasses** argument or by listing all of the object classes in a comma-separated list inside curly braces with no spaces allowed, such as `{attr1,attr2,attr3}`. Especially for long lists, it can be easier to use the curly braces than multiple options. For example:

```
[bjensen@server ~]$ ipa config-mod --
userobjectclasses={top,person,organizationalperson,inetorgperson,ine
tuser,posixaccount,krbprincipalaux,krbticketpolicyaux,ipaobject,ipas
huser,employeeinfo}
```

**NOTE**

To use the curly braces option, the **brace expansion** feature must be switched on. To activate the feature, use the **set** command:

```
# set -o braceexpand
```

15.3. APPLYING CUSTOM OBJECT CLASSES TO NEW GROUP ENTRIES

As with user entries, administrators may create custom object classes with custom attributes. These can be added automatically by adding the object classes to the IdM server configuration. Note that after an administrator modifies the list of default object classes, new entries will contain the custom object classes but the old entries are not automatically modified.

15.3.1. From the Web UI

1. Add all of the custom schema elements to the 389 Directory Server instance used by Identity Management. Adding schema elements is described in [the schema chapter of the Directory Server Administrator's Guide](#).
2. Open the **IPA Server** tab.
3. Select the **Configuration** subtab.
4. Scroll to the **Group Options** area.

Group Options

Group search * fields


Default group * objectclasses

top	Delete
ipaobject	Delete
groupofnames	Delete
ipausergroup	Delete
nestedgroup	Delete

Add

Figure 15.3. Group Options in Server Configuration

5. Click **Add** to include a new field for another object class.




IMPORTANT

Always include the *existing* default object classes when the configuration is updated. Otherwise, the current settings will be overwritten. If any object classes required by Identity Management are not included, then subsequent attempts to add an entry will fail with object class violations.

6. When the changes are complete, click **Save** at the top of the **Configuration** page.

15.3.2. From the Command Line

1. Add all of the custom schema elements to the 389 Directory Server instance used by Identity Management. Adding schema elements is described in [the schema chapter of the Directory Server Administrator's Guide](#).
2. Add the new object class to the list of object classes added to entries. The option for group object classes is **--groupobjectclasses**.



IMPORTANT

Always include the *existing* default object classes when the configuration is updated. Otherwise, the current settings will be overwritten. If any object classes required by Identity Management are not included, then subsequent attempts to add an entry will fail with object class violations.

All object classes must be included in the list of object classes. The information passed with the **config-mod** command overwrites the previous values. This can be done by specifying each object class with a **--groupobjectclasses** argument or by listing all of the object classes in a comma-separated list inside curly braces with no spaces allowed, such as `{attr1,attr2,attr3}`. Especially for long lists, it can be easier to use the curly braces than multiple options. For example:

```
[bjensen@server ~]$ ipa config-mod --  
groupobjectclasses={top,groupofnames,nestedgroup,ipausergroup,ipaobj  
ect,ipasshuser,employee}
```

15.4. SPECIFYING DEFAULT USER AND GROUP ATTRIBUTES

Identity Management uses a template when it creates new entries.

For users, the template is very specific. Identity Management uses default values for several core attributes for IdM user accounts. These defaults can define actual values for user account attributes (such as the home directory location) or it can define the format of attribute values, such as the user name length. These settings also define the object classes assigned to users.

For groups, the template only defines the assigned object classes.

These default definitions are all contained in a single configuration entry for the IdM server, **cn=ipaconfig,cn=etc,dc=example,dc=com**.

The configuration can be changed using the **ipa config-mod** command.


Table 15.3. Default User Parameters

Field	Command-Line Option	Descriptions
Maximum user name length	--maxusername	Sets the maximum number of characters for user names. The default value is 32.
Root for home directories	--homedirectory	Sets the default directory to use for user home directories. The default value is /home .
Default shell	--defaultshell	Sets the default shell to use for users. The default value is /bin/sh .
Default user group	--defaultgroup	Sets the default group to which all newly created accounts are added. The default value is ipausers , which is automatically created during the IdM server installation process.
Default e-mail domain	--emaildomain	Sets the email domain to use to create email addresses based on the new accounts. The default is the IdM server domain.
Search time limit	--searchtimelimit	Sets the maximum amount of time, in seconds, to spend on a search before the server returns results.
Search size limit	--searchrecordslimit	Sets the maximum number of records to return in a search.
User search fields	--usersearch	Sets the fields in a user entry that can be used as a search string. Any attribute listed has an index kept for that attribute, so setting too many attributes could affect server performance.

Field	Command-Line Option	Descriptions
Group search fields	--groupsearch	Sets the fields in a group entry that can be used as a search string.
Certificate subject base		Sets the base DN to use when creating subject DN's for client certificates. This is configured when the server is set up.
Default user object classes	--userobjectclasses	Defines an object class that is used to create IdM user accounts. This can be invoked multiple times. The complete list of object classes must be given because the list is overwritten when the command is run.
Default group object classes	--groupobjectclasses	Defines an object class that is used to create IdM group accounts. This can be invoked multiple times. The complete list of object classes must be given because the list is overwritten when the command is run.
Password expiration notification	--pwdexpnotify	Sets how long, in days, before a password expires for the server to send a notification.
Password plug-in features		Sets the format of passwords that are allowed for users.

15.4.1. Viewing Attributes from the Web UI

1. Open the **IPA Server** tab.
2. Select the **Configuration** subtab.
3. The complete configuration entry is shown in three sections, one for all search limits, one for user templates, and one for group templates.



Configuration

Refresh Revert Save

Search Options

Search size limit *

100

Search time limit *

2

Figure 15.4. Setting Search Limits




User Options

User search fields *

uid,givenname,sn,telephonenumber,ou,1

Figure 15.5. User Attributes



Group Options

Group search fields *

cn,description

Figure 15.6. Group Attributes

15.4.2. Viewing Attributes from the Command Line

The **config-show** command shows the current configuration which applies to all new user accounts. By default, only the most common attributes are displayed; use the **--all** option to show the complete configuration.

```
[bjensen@server ~]$ kinit admin
[bjensen@server ~]$ ipa config-show --all
dn: cn=ipaConfig,cn=etc,dc=example,dc=com
Maximum username length: 32
Home directory base: /home
```

```
Default shell: /bin/sh
Default users group: ipausers
Default e-mail domain: example.com
Search time limit: 2
Search size limit: 100
User search fields: uid,givenname,sn,telephonenumber,ou,title
Group search fields: cn,description
Enable migration mode: FALSE
Certificate Subject base: O=EXAMPLE.COM
Default group objectclasses: top, groupofnames, nestedgroup, ipausergroup,
ipaobject
Default user objectclasses: top, person, organizationalperson,
inetorgperson, inetuser, posixaccount, krbprincipalaux,
krbticketpolicyaux, ipaobject, ipasshuser
Password Expiration Notification (days): 4
Password plugin features: AllowNThash
SELinux user map order: guest_u:s0$xguest_u:s0$user_u:s0$staff_u:s0-
s0:c0.c1023$unconfined_u:s0-s0:c0.c1023
Default SELinux user: unconfined_u:s0-s0:c0.c1023
Default PAC types: MS-PAC, nfs:NONE
cn: ipaConfig
objectclass: nsContainer, top, ipaGuiConfig, ipaConfigObject
```

CHAPTER 16. MANAGING SERVICES

Some services that run on a host can also belong to the IdM domain. Any service that can store a Kerberos principal or an SSL certificate (or both) can be configured as an IdM service. Adding a service to the IdM domain allows the service to request an SSL certificate or keytab from the domain. (Only the public key for the certificate is stored in the service record. The private key is local to the service.)

An IdM domain establishes a commonality between machines, with common identity information, common policies, and shared services. Any machine which belongs to a domain functions as a client of the domain, which means it uses the services that the domain provides. An IdM domain (as described in [Chapter 1, Introduction to Red Hat Identity Management](#)) provides three main services specifically for machines:

- DNS
- Kerberos
- Certificate management

16.1. ADDING AND EDITING SERVICE ENTRIES AND KEYTABS

As with host entries, service entries for the host (and any other services on that host which will belong to the domain) must be added manually to the IdM domain. This is a two step process. First, the service entry must be created, and then a keytab must be created for that service which it will use to access the domain.

By default, Identity Management saves its HTTP keytab to `/etc/httpd/conf/ipa.keytab`.



NOTE

This keytab is used for the web UI. If a key were stored in `ipa.keytab` and that keytab file is deleted, the IdM web UI will stop working, because the original key would also be deleted.

Similar locations can be specified for each service that needs to be made Kerberos aware. There is no specific location that must be used, but, when using `ipa-getkeytab`, you should avoid using `/etc/krb5.keytab`. This file should not contain service-specific keytabs; each service should have its keytab saved in a specific location and the access privileges (and possibly SELinux rules) should be configured so that only this service has access to the keytab.

16.1.1. Adding Services and Keytabs from the Web UI

1. Open the **Identity** tab, and select the **Services** subtab.
2. Click the **Add** button at the top of the services list.
3. Select the service type from the drop-down menu, and give it a name.
4. Select the host name of the IdM host on which the service is running. The host name is used to construct the full service principal name.
5. Click the **Add** button to save the new service principal.

6. Use the **ipa-getkeytab** command to generate and assign the new keytab for the service principal.

```
[root@ipaserver ~]# # ipa-getkeytab -s ipaserver.example.com -p
HTTP/server.example.com -k /etc/httpd/conf/krb5.keytab -e aes256-cts
```

- The realm name is optional. The IdM server automatically appends the Kerberos realm for which it is configured. You cannot specify a different realm.
- The host name must resolve to a DNS A record for it to work with Kerberos. You can use the **--force** flag to force the creation of a principal should this prove necessary.
- The **-e** argument can include a list of encryption types to include in the keytab. This supersedes any default encryption type. Lists of entries can be set by using the option multiple times with the same command invocation or by listing the options in a comma-separated list inside curly braces, such as **--option={val1,val2,val3}**.



WARNING

Creating a new key resets the secret for the specified principal. This means that all other keytabs for that principal are rendered invalid.

16.1.2. Adding Services and Keytabs from the Command Line

1. Create the service principal. The service is recognized through a name like *service/FQDN*:

```
# ipa service-add serviceName/hostname
```

For example:

```
$ ipa service-add HTTP/server.example.com
-----
Added service "HTTP/server.example.com@EXAMPLE.COM"
-----
Principal: HTTP/server.example.com@EXAMPLE.COM
Managed by: ipaserver.example.com
```

2. Create the service keytab file using the **ipa-getkeytab** command. This command is run on the client in the IdM domain. (Actually, it can be run on any IdM server or client, and then the keys copied to the appropriate machine. However, it is simplest to run the command on the machine with the service being created.)

The command requires the Kerberos service principal (**-p**), the IdM server name (**-s**), the file to write (**-k**), and the encryption method (**-e**). Be sure to copy the keytab to the appropriate directory for the service.

For example:

```
# ipa-getkeytab -s server.example.com -p HTTP/server.example.com -k
/etc/httpd/conf/krb5.keytab -e aes256-cts
```

- The realm name is optional. The IdM server automatically appends the Kerberos realm for which it is configured. You cannot specify a different realm.
- The host name must resolve to a DNS A record for it to work with Kerberos. You can use the **--force** flag to force the creation of a principal should this prove necessary.
- The **-e** argument can include a comma-separated list of encryption types to include in the keytab. This supersedes any default encryption type. Lists of entries can be set by using the option multiple times with the same command invocation or by listing the options in a comma-separated list inside curly braces, such as **--option={val1,val2,val3}**.



WARNING

The **ipa-getkeytab** command resets the secret for the specified principal. This means that all other keytabs for that principal are rendered invalid.

16.2. CONFIGURING CLUSTERED SERVICES

The IdM server is not *cluster aware*. However, it is possible to configure a clustered service to be part of IdM by synchronizing Kerberos keys across all of the participating hosts and configuring services running on the hosts to respond to whatever names the clients use.

1. Enroll all of the hosts in the cluster into the IdM domain.
2. Create any service principals and generate the required keytabs.
3. Collect any keytabs that have been set up for services on the host, including the host keytab at **/etc/krb5.keytab**.
4. Use the **ktutil** command to produce a single keytab file that contains the contents of all of the keytab files.
 1. For each file, use the **rkt** command to read the keys from that file.
 2. Use the **wkt** command to write all of the keys which have been read to a new keytab file.
5. Replace the keytab files on each host with the newly-created combined keytab file.
6. At this point, each host in this cluster can now impersonate any other host.
7. Some services require additional configuration to accommodate cluster members which do not reset host names when taking over a failed service.

- For `sshd`, set `GSSAPIStrictAcceptorCheck no` in `/etc/ssh/sshd_config`.
- For `mod_auth_kerb`, set `KrbServiceName Any` in `/etc/httpd/conf.d/auth_kerb.conf`.



NOTE

For SSL servers, the subject name or a subject alternative name for the server's certificate must appear correct when a client connects to the clustered host. If possible, share the private key among all of the hosts.

If each cluster member contains a subject alternative name which includes the names of all the other cluster members, that satisfies any client connection requirements.

16.3. USING THE SAME SERVICE PRINCIPAL FOR MULTIPLE SERVICES

Within a cluster, the same service principal can be used for multiple services, spread across different machines.

1. Retrieve a service principal using the `ipa-getkeytab` command.

```
# ipa-getkeytab -s kdc.example.com -p HTTP/server.example.com -k  
/etc/httpd/conf/krb5.keytab -e aes256-cts
```

2. Either direct multiple servers or services to use the same file, or copy the file to individual servers as required.

16.4. RETRIEVE EXISTING KEYTABS FOR MULTIPLE SERVERS

In some scenarios, like in a cluster environment, the same keytab file is required for a service represented on one common host name by different machines. IdM commands can be used to retrieve the same keytab on each of the hosts.

To prepare the common host name and the service principal, run the following commands on an IdM server:

1. Authenticate as **admin** user:

```
[root@ipaserver ~]# kinit admin
```

2. Add a common forward DNS record for all IP addresses that share this host name:

```
[root@ipaserver ~]# ipa dnsrecord-add idm.example.com cluster --a-  
rec={192.0.2.40,192.0.2.41}  
Record name: cluster  
A record: 192.0.2.40, 192.0.2.41
```

3. Create a new host entry object for the common DNS name:

```
[root@ipaserver ~]# ipa host-add cluster.idm.example.com  
-----
```

```
Added host "cluster.idm.example.com"
-----
Host name: cluster.idm.example.com
Principal name: host/cluster.idm.example.com@IDM.EXAMPLE.COM
Password: False
Keytab: False
Managed by: cluster.idm.example.com
```

4. Add the service principal for the host:

```
[root@ipaserver ~]# ipa service-add HTTP/cluster.idm.example.com
-----
Added service "HTTP/cluster.idm.example.com@IDM.EXAMPLE.COM"
-----
Principal: HTTP/cluster.idm.example.com@IDM.EXAMPLE.COM
Managed by: cluster.idm.example.com
```

5. Add the hosts to the service, that should be able to retrieve the keytab from IdM:

```
[root@ipaserver ~]# ipa service-allow-retrieve-keytab
HTTP/cluster.idm.example.com --hosts=
{node01.idm.example.com,node02.idm.example.com}
Principal: HTTP/cluster.idm.example.com@IDM.EXAMPLE.COM
Managed by: cluster.idm.example.com
Hosts allowed to retrieve keytab: node01.idm.example.com,
node02.idm.example.com
-----
Number of members added 2
-----
```

6. Grant permission to create a new keytab to one host:

```
[root@ipaserver ~]# ipa service-allow-create-keytab
HTTP/cluster.idm.example.com --hosts=node01.idm.example.com
Principal: HTTP/cluster.idm.example.com@IDM.EXAMPLE.COM
Managed by: cluster.idm.example.com
Hosts allowed to retrieve keytab: node01.idm.example.com,
node02.idm.example.com
Hosts allowed to create keytab: node01.idm.example.com
-----
Number of members added 1
-----
```

On the clients, follow these steps:

1. Authenticate with the hosts Kerberos keytab:

```
# kinit -kt /etc/krb5.keytab
```

2. 1. On the client you granted the respective permission to, generate a new keytab and store it in a file:

```
[root@node01 ~]# ipa-getkeytab -s ipaserver.idm.example.com -p
HTTP/cluster.idm.example.com -k /tmp/client.keytab
```

-
- 2. On all other clients, retrieve the existing keytab from the IdM server by adding the **-r** option to the command:

```
[root@node02 ~]# ipa-getkeytab -r -s ipaserver.idm.example.com -p  
HTTP/cluster.idm.example.com -k /tmp/client.keytab
```



WARNING

Be aware that if you omit the **-r** option, a new keytab will be generated. This invalidates all previously retrieved keytabs for this service principal.

16.5. DISABLING AND RE-ENABLING SERVICE ENTRIES

Active services can be accessed by other services, hosts, and users within the domain. There can be situations when it is necessary to remove a host or a service from activity. However, deleting a service or a host removes the entry and all the associated configuration, and it removes it permanently.

16.5.1. Disabling Service Entries

Disabling a service prevents domain users from access it without permanently removing it from the domain. This can be done by using the **service-disable** command.

For a service, specify the principal for the service. For example:

```
[jsmith@ipaserver ~]$ kinit admin  
[jsmith@ipaserver ~]$ ipa service-disable HTTP/server.example.com
```



IMPORTANT

Disabling a host entry not only disables that host. It disables every configured service on that host as well.

16.5.2. Re-enabling Services

Disabling a service essentially kills its current, active keytabs. Removing the keytabs effectively removes the service from the IdM domain without otherwise touching its configuration entry.

To re-enable a service, simply use the **ipa-getkeytab** command. The **-s** option sets which IdM server to request the keytab, **-p** gives the principal name, and **-k** gives the file to which to save the keytab.

For example, requesting a new HTTP keytab:

```
[root@ipaserver ~]# ipa-getkeytab -s ipaserver.example.com -p  
HTTP/server.example.com -k /etc/httpd/conf/krb5.keytab -e aes256-cts
```

CHAPTER 17. DELEGATING USER ACCESS TO HOSTS AND SERVICES

Manage means being able to retrieve a keytab and certificates for another host or service. Every host and service has a **managedby** entry which lists what hosts or services can manage it. By default, a host can manage itself and all of its services. It is also possible to allow a host to manage other hosts, or services on other hosts, by updating the appropriate delegations or providing a suitable **managedby** entry.

An IdM service can be managed from any IdM host, as long as that host has been granted, or *delegated*, permission to access the service. Likewise, hosts can be delegated permissions to other hosts within the domain.




Figure 17.1. Host and Service Delegation



NOTE

If a host is delegated authority to another host through a **managedBy** entry, it does not mean that the host has also been delegated management for all services on that host. Each delegation has to be performed independently.

17.1. DELEGATING SERVICE MANAGEMENT

A host is delegated control over a service using the **service-add-host** command. There are two parts to delegating the service: specifying the principal and identifying the hosts with the control:

```
# ipa service-add-host principal --hosts=hostnames
```

For example:

```
[root@server ~]# ipa service-add-host HTTP/web.example.com --
hosts=client1.example.com
```

Once the host is delegated authority, the host principal can be used to manage the service:

```
[root@server ~]# kinit -kt /etc/krb5.keytab host/`hostname`
# ipa-getkeytab -s `hostname` -k /tmp/test.keytab -p HTTP/web.example.com
Keytab successfully retrieved and stored in: /tmp/test.keytab
```

To create a ticket for this service, create a certificate request on the host with the delegated authority and use the **cert-request** command to create a service entry and load the certification information:

```
[root@server ~]# ipa cert-request --add --principal=HTTP/web.example.com
web.csr
Certificate: MIICETCCAXqgA...[snip]
Subject: CN=web.example.com,O=EXAMPLE.COM
Issuer: CN=EXAMPLE.COM Certificate Authority
Not Before: Tue Feb 08 18:51:51 2011 UTC
Not After: Mon Feb 08 18:51:51 2016 UTC
Serial number: 1005
```

For more information on creating certificate requests and using **ipa cert-request**, see [Section 24.1.1, “Requesting New Certificates for a User, Host, or Service”](#).

17.2. DELEGATING HOST MANAGEMENT

Hosts are delegated authority over other hosts through the **host-add-managedby** command. This creates a **managedby** entry. Once the **managedby** entry is created, then the host can retrieve a keytab for the host it has delegated authority over.

1. Log in as the admin user.

```
[root@server ~]# kinit admin
```

2. Add the **managedby** entry. For example, this delegates authority over *client2* to *client1*.

```
[root@server ~]# ipa host-add-managedby client2.example.com --
hosts=client1.example.com
```

3. Obtain a ticket as the host **client1** and then retrieve a keytab for **client2**:

```
[root@server ~]# kinit -kt /etc/krb5.keytab host/`hostname`
[root@server ~]# ipa-getkeytab -s `hostname` -k /tmp/client2.keytab
-p host/client2.example.com
Keytab successfully retrieved and stored in: /tmp/client2.keytab
```

17.3. DELEGATING HOST OR SERVICE MANAGEMENT IN THE WEB UI

Each host and service entry has a configuration tab that indicates what hosts have been delegated management control over that host or service.

1. Open the **Identity** tab, and select the **Hosts** or **Services** subtab.
2. Click the name of the host or service *that you are going to grant delegated management to*.

- Click the **Hosts** subtab on the far right of the host/service entry. This is the tab which lists hosts that *can manage* the selected host/service.

Host: example.server.com

example.server.com... is a member of:

example.server.com... is managed by:

Settings Host Groups Netgroups Roles HBAC Rules Sudo Rules **Hosts (1)**

Refresh Revert Save Actions

Host Settings

Host name example.server.com

Principal name host/example.server.com@REDHAT.COM

Figure 17.2. Host Subtab

- Click the **Add** link at the top of the list.
- Click the check box by the names of the hosts to which to delegate management for the host/service. Click the right arrow button, **>**, to move the hosts to the selection box.

Add Hosts Managing Host example.server.com

Filter available Hosts Filter

Available Prospective

☐ Host name ☐ Host name

☐ test.example.com

> <

Add Cancel

Figure 17.3. Host/Service Delegation Management

- Click the **Add** button to close the selection box and to save the delegation settings.

17.4. ACCESSING DELEGATED SERVICES

For both services and hosts, if a client has delegated authority, it can obtain a keytab for that principal on the local machine. For services, this has the format *service/hostname@REALM*. For hosts, the *service* is **host**.

With **kinit**, use the **-k** option to load a keytab and the **-t** option to specify the keytab.

For example, to access a host:

```
[root@server ~]# kinit -kt /etc/krb5.keytab  
host/ipa.example.com@EXAMPLE.COM
```

To access a service:


```
[root@server ~]# kinit -kt /etc/httpd/conf/krb5.keytab  
HTTP/ipa.example.com@EXAMPLE.COM
```


CHAPTER 18. ID VIEWS

ID views enable you to specify new values for POSIX user or group attributes, as well as to define on which client host or hosts the new values will apply.

For example, you can use ID views to:

- define different attribute values for different environments; see [Section 18.3, “Defining a Different Attribute Value for a User Account on Different Hosts”](#)
- replace a previously generated attribute value with a different value



IMPORTANT

You can apply ID views only to IdM clients, not to IdM servers.

Potential Negative Impact on SSSD Performance

Applying an ID view can have a negative impact on SSSD performance, because certain optimizations and ID views cannot run at the same time. For example, ID views prevent SSSD from optimizing the process of looking up groups on the server:

- With ID views, SSSD must check every member on the returned list of group member names if the group name is overridden.
- Without ID views, SSSD can only collect the user names from the member attribute of the group object.

This negative effect mostly becomes apparent when the SSSD cache is empty or after clearing the cache, which makes all entries invalid.

Additional Resources

ID views also have several use cases in environments involving Active Directory. For details, see the [ID Views and Migrating Existing Environments to Trust](#) chapter in the *Windows Integration Guide*.

18.1. ATTRIBUTES AN ID VIEW CAN OVERRIDE

ID views consist of user and group ID overrides. The overrides define the new attribute values.

User and group ID overrides can define new values for the following attributes:

User attributes

- Login name (**uid**)
- GECOS entry (**gecos**)
- UID number (**uidNumber**)
- GID number (**gidNumber**)
- Login shell (**loginShell**)
- Home directory (**homeDirectory**)

- SSH public keys (**ipaSshPubkey**)
- Certificate (**userCertificate**)

Group attributes

- Group name (**cn**)
- Group GID number (**gidNumber**)

18.2. GETTING HELP FOR ID VIEW COMMANDS

To display all commands used to manage ID views and overrides:

```
$ ipa help idviews
```

To display detailed help for a particular command, add the **--help** option to the command:

```
$ ipa idview-add --help
```

18.3. DEFINING A DIFFERENT ATTRIBUTE VALUE FOR A USER ACCOUNT ON DIFFERENT HOSTS

An administrator can create multiple ID views that override an attribute value used by a user account and apply these ID views to different client hosts. Example: A service account is configured to use different SSH public keys when authenticating on different hosts.


This section includes the following procedures:

- [Section 18.3.1, “Web UI: Overriding an Attribute Value for a Specific Host”](#)
- [Section 18.3.2, “Command Line: Overriding an Attribute Value for a Specific Host”](#)

The procedures show how to create an ID view for a client host named **host1.example.com**. To override the attribute values on the other hosts as well, use the procedures to create multiple ID views, one for each host.

In the following procedures:

- **user** is the user account whose attribute needs to be overridden
- **host1.example.com** is the host on which the ID view will be applied



IMPORTANT

After you create a new ID view, restart SSSD on all clients where the ID view is applied.

If the new ID view changes a UID or GID, clear the SSSD cache on these clients as well.

18.3.1. Web UI: Overriding an Attribute Value for a Specific Host

Before managing ID views, log in to the IdM web UI as a user with the required privileges, such as **admin**.

Creating a New ID View

- 1. Under the **Identity** tab, select the **ID Views** subtab.
- 2. Click **Add** and provide a name for the ID view.

Add ID View

ID View Name *

example_for_host1

Description

ID view to be applied on host1.example.com

* Required field

Add

Add and Add Another

Add and Edit

Cancel

Figure 18.1. Adding an ID View

- 3. Click **Add** to confirm.

The new ID view is now displayed in the list of ID views.

ID Views

Search

Refresh

Delete

Add

Un-apply from host

☐	ID View Name	Description
☐	example_for_host1	ID view to be applied on host1.example.com

Showing 1 to 1 of 1 entries.

Figure 18.2. List of ID Views

Adding a User Override to the ID View

- 1. In the list of ID views, click the name of the ID view.




Figure 18.3. Editing an ID View

2. Under the **Users** tab, click **Add** to add the user override.
3. Select the user account whose attribute value to override, and click **Add**.

The user override is now displayed on the **example_for_host1** ID view page.




Figure 18.4. List of Overrides

Specifying the Attribute to Override

1. Click the override that you want to use to change the attribute value.





Figure 18.5. Editing an Override

2. Define the new value for the attribute.

For example, to override the SSH public key used by the user account:

- a. Click **SSH public keys: Add**.



User ID override: user

Refresh Revert Save Actions ▾

User to override	user
Description	
User login	
GECOS	
UID	
GID	
Login shell	
Home directory	
SSH public keys	Add
Certificates	Add

Figure 18.6. Adding an SSH Public Key

- b. Paste in the public key.



NOTE

For details on adding SSH keys to IdM, see [Section 22.4, “Managing Public SSH Keys for Users”](#).

3. Click **Save** to update the override.

Applying the ID View to a Specific Host

1. In the list of ID views, click the name of the ID view.




Figure 18.7. Editing an ID View

2. Under the **Hosts** tab, click **Apply to hosts**.
3. Select the **host1.example.com** host, and move it to the **Prospective** column.
4. Click **Apply**.

The host is now displayed in the list of hosts to which the ID view applies.




Figure 18.8. Listing Hosts to Which an ID View Applies

18.3.2. Command Line: Overriding an Attribute Value for a Specific Host

Before managing ID views, request a ticket for a user with the required privileges. For example:

```
$ kinit admin
```

1. Create a new ID view. For example, the create an ID view named **example_for_host1**:

```
$ ipa idview-add example_for_host1
-----
Added ID View "example_for_host1"
-----
ID View Name: example_for_host1
```

2. Add a user override to the **example_for_host1** ID view. The **ipa idoverrideuser-add** command requires the name of the ID view and the user to override.
 - To specify the new attribute value, add the corresponding command-line option as well. For a list of the available options, run **ipa idoverrideuser-add --help**. For example, use the **--sshpubkey** option to override the SSH public key value:

```
$ ipa idoverrideuser-add example_for_host1 user --sshpubkey="ssh-
rsa AAAAB3NzaC1yrRqFE...gWRL71/miPIZ user@example.com"
-----
Added User ID override "user"
-----
Anchor to override: user
SSH public key: ssh-rsa
                  AAAAB3NzaC1yrRqFE...gWRL71/miPIZ
                  user@example.com
```



NOTE

For details on adding SSH keys to IdM, see [Section 22.4, “Managing Public SSH Keys for Users”](#).

- The **ipa idoverrideuser-add --certificate** command replaces all existing certificates for the account in the specified ID view. To append an additional certificate, use the **ipa idoverrideuser-add-cert** command instead:

```
$ ipa idoverrideuser-add-cert example_for_host1 user --
certificate="MIIEATCC..."
```

- Using the **ipa idoverrideuser-mod** command, you can also specify new attribute values for an existing user override.

3. Apply **example_for_host1** to the **host1.example.com** host:

```
$ ipa idview-apply example_for_host1 --hosts=host1.example.com
-----
Applied ID View "example_for_host1"
-----
hosts: host1.example.com
-----
Number of hosts the ID View was applied to: 1
-----
```



NOTE

The **ipa idview-apply** command also accepts the **--hostgroups** option. The option applies the ID view to hosts that belong to the specified host group, but does not associate the ID view with the host group itself. Instead, the **--hostgroups** option expands the members of the specified host group and applies the **--hosts** option individually to every one of them.

CHAPTER 19. DEFINING ACCESS CONTROL FOR IDM USERS

Access control is a set of security features which defines who can access certain resources, such as machines, services or entries, and what kinds of operations they are allowed to perform. Identity Management provides several access control areas to make it clear what kind of access is being granted and to whom it is granted. As part of this, Identity Management draws a distinction between access controls to resources within the domain and access control to the IdM configuration itself.

For details on the different internal access control mechanisms that are available for users within IdM to the IdM server and other IdM users, see [Chapter 10, *Defining Access Control for IdM Users*](#).

CHAPTER 20. MANAGING KERBEROS FLAGS AND PRINCIPAL ALIASES

20.1. KERBEROS FLAGS FOR SERVICES AND HOSTS

Various Kerberos flags can be used to define certain specific aspects of the Kerberos ticket behavior. You can add these flags to service and host Kerberos principals.

Principals in IdM accept the following two Kerberos flags:

OK_AS_DELEGATE

Use this flag to specify Kerberos tickets trusted for delegation.

AD clients check the **OK_AS_DELEGATE** flag on the Kerberos ticket to determine whether the user credentials can be forwarded or delegated to the specific server; AD forwards the TGT only to services or hosts with **OK_AS_DELEGATE** set. With this flag, SSSD can add the AD user TGT to the default Kerberos credentials cache on the IdM client machine.

REQUIRES_PRE_AUTH

Use this flag to specify that only pre-authenticated tickets are allowed to authenticate to the principal.

With the **REQUIRES_PRE_AUTH** flag set, the KDC requires additional authentication: the KDC issues the TGT for a principal with **REQUIRES_PRE_AUTH** only if the TGT has been pre-authenticated.

You can use **REQUIRES_PRE_AUTH** to disable pre-authentication for selected services or hosts, which lowers the load on the KDC but also slightly increases the possibility of a brute-force attack on a long-term key to succeed.

20.1.1. Setting Kerberos Flags from the Web UI

From the IdM web UI, you can currently only add the **OK_AS_DELEGATE** flag to a principal:

1. Select the **Services** subtab, accessible through the **Identity** main tab.




Figure 20.1. List of Services

2. Click on the service to which you want to add the flag.
3. Check the **Trusted for delegation** option.

Service: DNS/ipa.demo1.freeipa.org@DEMO1.FREEIPA.ORG

DNS/ipa.demo1.fre... is a member of: DNS/ipa.demo1.fre... is managed by:

Settings Roles Hosts (1)

Refresh Reset Update Actions

Service Settings

Principal DNS/ipa.demo1.freeipa.org@DEMO1.FREEIPA.ORG

Service DNS

Host Name ipa.demo1.freeipa.org

PAC type

- ☒ Inherited from server configuration
- ☐ Override inherited settings
 - ☐ MS-PAC
 - ☐ PAD

Trusted for delegation ☒ Undo

Figure 20.2. Adding the OK_AS_DELEGATE Flag

20.1.2. Setting Kerberos Flags from the Command Line

To add a flag to a principal from the command line or to remove a flag, add one of the following options to the **ipa service-mod** command:

- **--ok-as-delegate** for **OK_AS_DELEGATE**
- **--requires-pre-auth** for **REQUIRES_PRE_AUTH**

To add a flag, set the corresponding option to **1**. For example, to add the **OK_AS_DELEGATE** flag to the *test/ipa.example.com@EXAMPLE.COM* principal:

```
$ ipa service-mod test/ipa.example.com@EXAMPLE.COM --ok-as-delegate=1
```

To remove a flag or to disable it, set the corresponding option to **0**. For example, to disable the **REQUIRES_PRE_AUTH** flag for the *test/ipa.example.com@EXAMPLE.COM* principal:

```
$ ipa service-mod test/ipa.example.com@EXAMPLE.COM --requires-pre-auth=0
```

To find out if **OK_AS_DELEGATE** is currently set for a principal, run the **kvno** utility and then the **klist -f** command. **OK_AS_DELEGATE** is represented by the **0** character in the **klist -f** output:

```
$ kvno test/ipa.example.com@EXAMPLE.COM
$ klist -f
Ticket cache: KEYRING:persistent:0:0
Default principal: admin@EXAMPLE.COM

Valid starting Expires Service principal
02/19/2014 09:59:02 02/20/2014 08:21:33 test/ipa/example.com@EXAMPLE.COM
Flags: FATO
```

To find out what flags are currently set for a principal, use the **kadmin.local** utility. The current flags are displayed on the **Attributes** line of **kadmin.local** output, for example:

```
# kadmin.local
kadmin.local: getprinc test/ipa.example.com
Principal: test/ipa.example.com@EXAMPLE.COM
Expiration date: [never]
Last password change: Mon Sep 16 15:44:21 EDT 2013
Password expiration date: [none]
Maximum ticket life: 1 day 00:00:00
Maximum renewable life: 7 days 00:00:00
Last modified: Mon Oct 14 23:42:53 EDT 2013 (admin/admin@EXAMPLE.COM)
Last successful authentication: Wed Mar 11 08:01:03 EDT 2015
Last failed authentication: [never]
Failed password attempts: 0
Number of keys: 6
Key: vno 1, aes256-cts-hmac-sha1-96, no salt
Key: vno 1, aes128-cts-hmac-sha1-96, no salt
Key: vno 1, des3-cbc-sha1, no salt
Key: vno 1, arcfour-hmac, no salt
Key: vno 1, camellia128-cts-cmac, no salt
Key: vno 1, camellia256-cts-cmac, no salt
MKey: vno 1
Attributes: REQUIRES_PRE_AUTH OK_AS_DELEGATE OK_TO_AUTH_AS_DELEGATE
Policy: [none]
```

20.2. MANAGING KERBEROS PRINCIPAL ALIASES FOR USERS, HOSTS, AND SERVICES

When you create a new user, host, or service, a Kerberos principal in the following format is automatically added:

- *user_name@REALM*
- *host/host_name@REALM*
- *service_name/host_name@REALM*

In some scenarios, it is beneficial for the administrator to enable users, hosts, or services to authenticate against Kerberos applications using an alias, for example:

- The user name changed, but the user should be able to login using both the previous and new user name.
- The user needs to log in using the email address even if the IdM Kerberos realm differs from the email domain.

Note that if you rename a user, the object keeps the aliases and the previous canonical principal name.

20.2.1. Kerberos Principal Alias

Adding a Kerberos Principal Alias

To add the alias name **useralias** to the account **user**, enter:

```
[root@ipaserver ~]# ipa user-add-principal user useralias
-----
Added new aliases to user "user"
-----
      User login: user
      Principal alias: user@IDM.EXAMPLE.COM, useralias@IDM.EXAMPLE.COM
```

To add an alias to a host or service, use the **ipa host-add-principal** or **ipa service-add-principal** command respectively instead.

If you use an alias name to authenticate, pass the **-C** option to the **kinit** command:

```
[root@ipaserver ~]# kinit -C useralias
Password for user@IDM.EXAMPLE.COM:
```

Removing a Kerberos Principal Alias

To remove the alias **useralias** from the account **user**, enter:

```
[root@ipaserver ~]# ipa user-remove-principal user useralias
-----
Removed aliases from user "user"
-----
      User login: user
      Principal alias: user@IDM.EXAMPLE.COM
```

To remove an alias from a host or service, use the **ipa host-remove-principal** or **ipa service-remove-principal** command respectively instead.

Note that you cannot remove the canonical principal name:

```
[root@ipaserver ~]# ipa user-show user
      User login: user
      ...
      Principal name: user@IDM.EXAMPLE.COM
      ...

[root@ipaserver ~]# ipa user-remove-principal user user
ipa: ERROR: invalid 'krbprincipalname': at least one value equal to the
canonical principal name must be present
```

20.2.2. Kerberos Enterprise Principal Alias

Enterprise principal aliases can use any domain suffix except for user principal name (UPN) suffixes, NetBIOS names, or domain names of trusted Active Directory forest domains.

**NOTE**

When adding or removing enterprise principal aliases, escape the @ symbol using two backslashes (\\). Otherwise, the shell interprets the @ symbol as part of the Kerberos realm name and leads to the following error:

```
ipa: ERROR: The realm for the principal does not match the realm
for this IPA server
```

Adding a Kerberos Enterprise Principal Alias

To add the enterprise principal alias **user@example.com** to the *user* account:

```
[root@ipaserver ~]# ipa user-add-principal user user\\@example.com
-----
Added new aliases to user "user"
-----
      User login: user
      Principal alias: user@IDM.EXAMPLE.COM,
user\\@example.com@IDM.EXAMPLE.COM
```

To add an enterprise alias to a host or service, use the **ipa host-add-principal** or **ipa service-add-principal** command respectively instead.

If you use an enterprise principal name to authenticate, pass the **-E** option to the **kinit** command:

```
[root@ipaserver ~]# kinit -E user@example.com
Password for user\\@example.com@IDM.EXAMPLE.COM:
```

Removing a Kerberos Enterprise Principal Alias

To remove the enterprise principal alias **user@example.com** from the account **user**, enter:

```
[root@ipaserver ~]# ipa user-remove-principal user user\\@example.com
-----
Removed aliases from user "user"
-----
      User login: user
      Principal alias: user@IDM.EXAMPLE.COM
```

To remove an alias from a host or service, use the **ipa host-remove-principal** or **ipa service-remove-principal** command respectively instead.

CHAPTER 21. INTEGRATING WITH NIS DOMAINS AND NETGROUPS

21.1. ABOUT NIS AND IDENTITY MANAGEMENT

In UNIX environments, the network information service (NIS) is a common way to centrally manage identities and authentication. NIS, which was originally named *Yellow Pages* (YP), centrally manages authentication and identity information such as:

- Users and passwords
- Host names and IP addresses
- POSIX groups.

For modern network infrastructures, NIS is considered too insecure because, for example, it neither provides host authentication, nor is data sent encrypted over the network. To work around the problems, NIS is often integrated with other protocols to enhance security.

If you use Identity Management (IdM), you can use the NIS server plug-in to connect clients that cannot be fully migrated to IdM. IdM integrates netgroups and other NIS data into the IdM domain. Additionally, you can easily migrate user and host identities from a NIS domain to IdM.

NIS in Identity Management

NIS objects are integrated and stored in the Directory Server back end in compliance with [RFC 2307](#). IdM creates NIS objects in the LDAP directory and clients retrieve them through, for example, System Security Services Daemon (SSSD) or `nss_ldap` using an encrypted LDAP connection.

IdM manages netgroups, accounts, groups, hosts, and other data. IdM uses a NIS listener to map passwords, groups, and netgroups to IdM entries.

NIS Plug-ins in Identity Management

For NIS support, IdM uses the following plug-ins provided in the `slapi-nis` package:

NIS Server Plug-in

The NIS Server plug-in enables the IdM-integrated LDAP server to act as a NIS server for clients. In this role, Directory Server dynamically generates and updates NIS maps according to the configuration. Using the plug-in, IdM serves clients using the NIS protocol as an NIS server.

For further details, see [Section 21.2, “Enabling NIS in Identity Management”](#).

Schema Compatibility Plug-in

The Schema Compatibility plug-in enables the Directory Server back end to provide an alternate view of entries stored in part of the directory information tree (DIT). This includes adding, dropping, or renaming attribute values, and optionally retrieving values for attributes from multiple entries in the tree.

For further details, see the `/usr/share/doc/slapi-nis-version/sch-getting-started.txt` file.

21.1.1.1. NIS Netgroups in Identity Management

NIS entities can be stored in netgroups. Compared to UNIX groups, netgroups provide support for:

- Nested groups (groups as members of other groups).
- Grouping hosts.

A netgroup defines a set of the following information: host, user, and domain. This set is called a *triple*. These three fields can contain:

- A value.
- A dash (-), which specifies "no valid value"
- No value. An empty field specifies a wildcard.

```
(host.example.com,,nisdomain.example.com)
(-,user,nisdomain.example.com)
```

When a client requests a NIS netgroup, IdM translates the LDAP entry :

- to a traditional NIS map and sends it to the client over the NIS protocol by using the NIS plug-in.
- to an LDAP format that is compliant with [RFC 2307](#) or RFC 2307bis.

21.1.1.1.1. Displaying NIS Netgroup Entries

IdM stores users and groups in the ***memberUser*** attribute, and hosts and host groups in ***memberHost***. The following example shows a netgroup entry in Directory Server component of IdM:

Example 21.1. A NIS Entry in Directory Server

```
dn: ipaUniqueID=d4453480-cc53-11dd-ad8b-0800200c9a66,cn=ng,cn=alt,...
...
cn: netgroup1
memberHost: fqdn=host1.example.com,cn=computers,cn=accounts,...
memberHost: cn=VirtGuests,cn=hostgroups,cn=accounts,...
memberUser: cn=demo,cn=users,cn=accounts,...
memberUser: cn=Engineering,cn=groups,cn=accounts,...
nisDomainName: nisdomain.example.com
```

In IdM, you can manage netgroup entries using the **ipa netgroup-*** commands. For example, to display a netgroup entry:

Example 21.2. Displaying a Netgroup Entry

```
[root@server ~]# ipa netgroup-show netgroup1
Netgroup name: netgroup1
Description: my netgroup
NIS domain name: nisdomain.example.com
```

```
Member Host: VirtGuests
Member Host: host1.example.com
Member User: demo
Member User: Engineering
```

21.2. ENABLING NIS IN IDENTITY MANAGEMENT

To enable NIS in Identity Management:

1. Enable the NIS listener and compatibility plug-ins:

```
[root@ipaserver ~]# ipa-nis-manage enable
[root@ipaserver ~]# ipa-compat-manage enable
```

2. *Optional:* Set a fixed port for the NIS remote procedure calls (RPC).

When using NIS, the client must know to which port on the IdM server to use to establish the connection. Using the default settings, IdM binds to an unused random port when the server starts. This port is sent to the port mapper service the client uses to request the port number.

For a more strict firewall configuration, you can set a fixed port. For example, to set the port to **514**:

```
[root@ipaserver ~]# ldapmodify -x -D 'cn=directory manager' -W
dn: cn=NIS Server,cn=plugins,cn=config
changetype: modify
add: nsslapd-pluginarg0
nsslapd-pluginarg0: 514
```



NOTE

You can set any unused port number below 1024 for the setting.

3. Enable and start the port mapper service:

```
[root@ipaserver ~]# systemctl enable rpcbind.service
[root@ipaserver ~]# systemctl start rpcbind.service
```

4. Restart Directory Server:

```
[root@ipaserver ~]# systemctl restart dirsrv.target
```

21.3. CREATING NETGROUPS

21.3.1. Adding a Netgroup

To add a Netgroup, you can use:

- the IdM web UI (see [the section called “Web UI: Adding a Netgroup”](#))

- the command line (see [the section called “Command Line: Adding a Netgroup”](#))

Web UI: Adding a Netgroup

1. Select **Identity** → **Groups** → **Netgroups**
2. Click **Add**.
3. Enter a unique name and, optionally, a description. The group name is the identifier used for the netgroup in the IdM domain. You cannot change it later.
4. Click **Add and Edit** to save the changes and to start editing the entry.
5. The default NIS domain is set to the IdM domain name. Optionally, you can enter the name of the alternative NIS domain in the NIS domain name field.

The screenshot shows the 'Netgroup: server.example.com' configuration page. At the top, it displays 'server.example.com members: server.example.com is a member of:'. Below this are three tabs: 'Settings', 'Netgroups', and 'Netgroups', with the third 'Netgroups' tab selected. There are three buttons: 'Refresh', 'Revert', and 'Save'. The main section is titled 'General' and contains three fields: 'Netgroup name' with the value 'server.example.com', 'Description' with the value 'An example' and an 'Undo' button, and 'NIS domain name' with the value 'example.com' and an 'Undo' button.

Figure 21.1. Netgroup Tab

The **NIS domain name** field sets the domain that appears in the netgroup triple. It does not affect which NIS domain the Identity Management NIS listener responds to.

6. Add members, as described in [the section called “Web UI: Adding Members to a Netgroup”](#).
7. Click **Save**.

Command Line: Adding a Netgroup

You can add a new netgroup using the **ipa netgroup-add** command. Specify:

- the group name.
- optionally, a description.
- optionally, the NIS domain name if it is different than the IdM domain name.



NOTE

The **--nisdomain** option sets the domain that appears in the netgroup triple. It does not affect which NIS domain the Identity Management listener responds to.

For example:

```
[root@server ~]# ipa netgroup-add --desc="Netgroup description" --
nisdomain="example.com" example-netgroup
```

To add members to the netgroup, see [the section called “Command Line: Adding Members to a Netgroup”](#).

21.3.2. Adding Members to a Netgroup

Beside users and hosts, netgroups can contain user groups, host groups, and other netgroups (nested groups) as members. Depending on the size of a group, it can take up to several minutes after you create a nested groups for the members of the child group to show up as members of the parent group.

To add members to a Netgroup, you can use:

- the IdM web UI (see [the section called “Web UI: Adding Members to a Netgroup”](#))
- the command line (see [the section called “Command Line: Adding Members to a Netgroup”](#))



WARNING


Do not create recursive nested groups. For example, if *GroupA* is a member of *GroupB*, do not add *GroupB* as a member of *GroupA*. Recursive groups are not supported and can cause unpredictable behavior.

Web UI: Adding Members to a Netgroup

To add members to a netgroup using the Web UI:

1. Select **Identity** → **Groups** → **Netgroups**
2. Click the name of the netgroup to which to add members.

- Click **Add** next to the required member type.



Netgroup: server.example.com

server.example.com members: server.example.com is a member of:

Settings Netgroups Netgroups

Refresh Revert Save


User

User category the rule applies to: ☐ Anyone ☒ Specified Users and Groups

☐	Users	Delete	+Add
☐	User Groups	Delete	+Add

Figure 21.2. User Menu in the Netgroup Tab

- Select the members you want to add, and click > to confirm.



Add Users into Netgroup server.example.com

Filter available Users Filter

Available

☐	Users
☐	admin
☐	employee
☐	helpdesk
☐	manager
☐	test
☐	test_1

Prospective

☐	Users

Add Cancel

Figure 21.3. Add User Menu in the Netgroup Tab

- Click **Add**.

Command Line: Adding Members to a Netgroup

After you created the netgroup, you can add members using the **ipa netgroup-add-member** command:

```
# ipa netgroup-add-member --users=user_name --groups=group_name --
hosts=host_name \
    --hostgroups=host_group_name --netgroups=netgroup_name group_name
```

To set more than one member, use a comma-separated list inside a set of curly braces. For example:

```
[root@server ~]# ipa netgroup-add-member --users={user1;user2,user3} \
--groups={group1,group2} example-group
```

21.4. EXPOSING AUTOMOUNT MAPS TO NIS CLIENTS

If any automount maps are already defined, you must manually add them to the NIS configuration in IdM. This ensures the maps are exposed to NIS clients.

The NIS server is managed by a special plug-in entry in the IdM LDAP directory. Each NIS domain and map used by the NIS server is added as a sub-entry in this container. The NIS domain entry contains:

- the name of the NIS domain
- the name of the NIS map
- information on how to find the directory entries to use as the NIS map's contents
- information on which attributes to use as the NIS map's key and value

Most of these settings are the same for every map.

21.4.1. Adding an Automount Map

IdM stores the automount maps, grouped by the automount location, in the **cn=automount** branch of the IdM directory tree. You can add the NIS domain and maps using the LDAP protocol.

For example, to add an automount map named **auto.example** in the **default** location for the **example.com** domain:

```
[root@server ~]# ldapadd -h server.example.com -x -D "cn=Directory
Manager" -W

dn: nis-domain=example.com+nis-map=auto.example,cn=NIS
Server,cn=plugins,cn=config
objectClass: extensibleObject
nis-domain: example.com
nis-map: auto.example
nis-filter: (objectclass=automount)
nis-key-format: %{automountKey}
nis-value-format: %{automountInformation}
nis-base:
automountmapname=auto.example,cn=default,cn=automount,dc=example,dc=com
```

**NOTE**

Set the ***nis-domain*** attribute to the name of your NIS domain.

The value set in the ***nis-base*** attribute must correspond:

- To an existing automount map set using the **ipa automountmap-*** commands.
- To an existing automount location set using the **ipa automountlocation-*** commands.

After you set the entry, you can verify the automount map:

```
[root@server ~]# ypcat -k -d example.com -h server.example.com
auto.example
```

21.5. MIGRATING FROM NIS TO IDM

Migrating from an existing NIS server to Identity Management (IdM) requires the following steps:

1. [Enable the NIS Listener in Identity Management](#)
2. [Export and import the existing data from NIS](#)

21.5.1. Preparing Netgroup Entries in IdM

Before migrating, identify what kind of identities are being managed in the current NIS server:

User Entries

Determine what applications are using the user information provided by NIS. While some utilities, such as **sudo**, require NIS netgroups, several others can use regular UNIX groups.

To migrate:

1. Create the corresponding user accounts in IdM. See [Section 21.5.3.1, “Migrating User Entries”](#).
2. If you additionally require netgroups:
 - a. Add the netgroups. See [Section 21.3.1, “Adding a Netgroup”](#).
 - b. Add the users to the netgroups. See [Section 21.5.3.4, “Migrating Netgroup Entries”](#).

Host Entries

When you create a host group in IdM, a corresponding shadow NIS group is automatically created. Do not use the **ipa netgroup-*** commands on these shadow NIS groups. Use the **ipa netgroup-*** commands only to manage *native* netgroups created via the **netgroup-add** command.

For a Direct Conversion

If every user and host entry must use the same name, you can create the entries using the same names in IdM:

1. Create an entry for every user referenced in a netgroup.
2. Create an entry for every host referenced in a netgroup.
3. Create a netgroup with the same name as the original netgroup.
4. Add the users and hosts as direct members of the netgroup. If the users and hosts are members of groups or host groups, you can alternatively add these groups to the netgroup.

21.5.2. Enabling the NIS Listener in Identity Management

See [Section 21.2, “Enabling NIS in Identity Management”](#).

21.5.3. Exporting and Importing the Existing NIS Data

A NIS server can contain information about users, groups, hosts, netgroups, and automount maps. You can migrate these entry types to IdM.

In the following sections, we export the data from the current NIS server using the **ypcat** command, and use the output to import the entries to IdM using the corresponding **ipa *-add** commands.

21.5.3.1. Migrating User Entries

The NIS **passwd** map contains information about users, such as names, UIDs, primary group, GECOS, shell, and home directory. Use this data to migrate NIS user accounts to IdM:

1. *Optional:* If you require weak password support, see [Section 21.5.4, “Enabling Weak Password Hashing for NIS User Authentication”](#).
2. Create the **/root/nis-users.sh** script with the following content:

```
#!/bin/sh
# $1 is the NIS domain, $2 is the NIS master server
ypcat -d $1 -h $2 passwd > /dev/shm/nis-map.passwd 2>&1

IFS=$'\n'
for line in $(cat /dev/shm/nis-map.passwd) ; do
    IFS=' '
    username=$(echo $line | cut -f1 -d:)
    # Not collecting encrypted password because we need cleartext
    password
    # to create kerberos key
    uid=$(echo $line | cut -f3 -d:)
    gid=$(echo $line | cut -f4 -d:)
    gecost=$(echo $line | cut -f5 -d:)
    homedir=$(echo $line | cut -f6 -d:)
    shell=$(echo $line | cut -f7 -d:)
```

```
# Now create this entry
echo passw0rd1 | ipa user-add $username --first=NIS --last=USER \
--password --gidnumber=$gid --uid=$uid --gecos=$gecos --
homedir=$homedir \
--shell=$shell
ipa user-show $username
done
```

3. Authenticate as the IdM **admin** user:

```
[root@nis-server ~]# kinit admin
```

4. Run the script. For example:

```
[root@nis-server ~]# sh /root/nis-users.sh nisdomain nis-
master.example.com
```



NOTE

This script uses hard-coded values for first name, last name, and sets the password to **passw0rd1**. The user must change the temporary password at the next log in.

21.5.3.2. Migrating Group Entries

The NIS **group** map contains information about groups, such as group names, GIDs, or group members. Use this data to migrate NIS groups to IdM:

1. Create the **/root/nis-groups.sh** script with the following content:

```
#!/bin/sh
# $1 is the NIS domain, $2 is the NIS master server
ypcat -d $1 -h $2 group > /dev/shm/nis-map.group 2>&1

IFS=$'\n'
for line in $(cat /dev/shm/nis-map.group); do
    IFS=' '
    groupname=$(echo $line | cut -f1 -d:)
    # Not collecting encrypted password because we need cleartext
    # password
    # to create kerberos key
    gid=$(echo $line | cut -f3 -d:)
    members=$(echo $line | cut -f4 -d:)

    # Now create this entry
    ipa group-add $groupname --desc=NIS_GROUP_$groupname --gid=$gid
    if [ -n "$members" ]; then
        ipa group-add-member $groupname --users={$members}
    fi
    ipa group-show $groupname
done
```

2. Authenticate as the IdM **admin** user:

```
[root@nis-server ~]# kinit admin
```

3. Run the script. For example:

```
[root@nis-server ~]# sh /root/nis-groups.sh nisdomain nis-  
master.example.com
```

21.5.3.3. Migrating Host Entries

The NIS **hosts** map contains information about hosts, such as host names and IP addresses. Use this data to migrate NIS host entries to IdM:

1. Create the **/root/nis-hosts.sh** script with the following content:

```
#!/bin/sh
# $1 is the NIS domain, $2 is the NIS master server
ypcat -d $1 -h $2 hosts | egrep -v "localhost|127.0.0.1" >
/dev/shm/nis-map.hosts 2>&1

IFS=$'\n'
for line in $(cat /dev/shm/nis-map.hosts); do
    IFS=' '
    ipaddress=$(echo $line | awk '{print $1}')
    hostname=$(echo $line | awk '{print $2}')
    master=$(ipa env xmlrpc_uri | tr -d '[:space:]' | cut -f3 -d: | cut
-f3 -d/)
    domain=$(ipa env domain | tr -d '[:space:]' | cut -f2 -d:)
    if [ $(echo $hostname | grep "\." | wc -l) -eq 0 ] ; then
        hostname=$(echo $hostname.$domain)
    fi
    zone=$(echo $hostname | cut -f2- -d.)
    if [ $(ipa dnszone-show $zone 2>/dev/null | wc -l) -eq 0 ] ; then
        ipa dnszone-add --name-server=$master --admin-email=root.$master
    fi
    ptrzone=$(echo $ipaddress | awk -F. '{print $3 "." $2 "." $1 ".in-  
addr.arpa."}')
    if [ $(ipa dnszone-show $ptrzone 2>/dev/null | wc -l) -eq 0 ] ;
then
        ipa dnszone-add $ptrzone --name-server=$master --admin-  
email=root.$master
    fi
    # Now create this entry
    ipa host-add $hostname --ip-address=$ipaddress
    ipa host-show $hostname
done
```

2. Authenticate as the IdM **admin** user:

```
[root@nis-server ~]# kinit admin
```

3. Run the script. For example:


```
[root@nis-server ~]# sh /root/nis-hosts.sh nisdomain nis-
master.example.com
```



NOTE

This script does not migrate special host configurations, such as aliases.

21.5.3.4. Migrating Netgroup Entries

The NIS **netgroup** map contains information about netgroups. Use this data to migrate NIS netgroups to IdM:

1. Create the **/root/nis-netgroups.sh** script with the following content:

```
#!/bin/sh
# $1 is the NIS domain, $2 is the NIS master server
ypcat -k -d $1 -h $2 netgroup > /dev/shm/nis-map.netgroup 2>&1

IFS=$'\n'
for line in $(cat /dev/shm/nis-map.netgroup); do
    IFS=' '
    netgroupname=$(echo $line | awk '{print $1}')
    triples=$(echo $line | sed "s/^\$netgroupname //")
    echo "ipa netgroup-add $netgroupname --desc=NIS_NG_$netgroupname"
    if [ $(echo $line | grep "(," | wc -l) -gt 0 ]; then
        echo "ipa netgroup-mod $netgroupname --hostcat=all"
    fi
    if [ $(echo $line | grep ",," | wc -l) -gt 0 ]; then
        echo "ipa netgroup-mod $netgroupname --usercat=all"
    fi

    for triple in $triples; do
        triple=$(echo $triple | sed -e 's/-//g' -e 's/(///' -e 's/)///')
        if [ $(echo $triple | grep ",.*," | wc -l) -gt 0 ]; then
            hostname=$(echo $triple | cut -f1 -d,)
            username=$(echo $triple | cut -f2 -d,)
            domain=$(echo $triple | cut -f3 -d,)
            hosts=""; users=""; doms="";
            [ -n "$hostname" ] && hosts="--hosts=$hostname"
            [ -n "$username" ] && users="--users=$username"
            [ -n "$domain" ] && doms="--nisdomain=$domain"
            echo "ipa netgroup-add-member $hosts $users $doms"
        else
            netgroup=$triple
            echo "ipa netgroup-add $netgroup --desc=NIS_NG_$netgroup"
        fi
    done
done
```

2. Authenticate as the IdM **admin** user:

```
[root@nis-server ~]# kinit admin
```

3. Run the script. For example:

```
[root@nis-server ~]# sh /root/nis-netgroups.sh nisdomain nis-master.example.com
```

21.5.3.5. Migrating Automount Maps

Automount maps are a series of nested and interrelated entries that define the location (the parent entry), the associated keys, and maps. To migrate NIS automount maps to IdM:

1. Create the `/root/nis-automounts.sh` script with the following content:

```
#!/bin/sh
# $1 is for the automount entry in ipa

ipa automountlocation-add $1

# $2 is the NIS domain, $3 is the NIS master server, $4 is the map
name
ypcat -k -d $2 -h $3 $4 > /dev/shm/nis-map.$4 2>&1

ipa automountmap-add $1 $4

basedn=$(ipa env basedn | tr -d '[:space:]' | cut -f2 -d:)
cat > /tmp/amap.ldif <<EOF
dn: nis-domain=$2+nis-map=$4,cn=NIS Server,cn=plugins,cn=config
objectClass: extensibleObject
nis-domain: $2
nis-map: $4
nis-base: automountmapname=$4,cn=$1,cn=automount,$basedn
nis-filter: (objectclass=*)
nis-key-format: %{automountKey}
nis-value-format: %{automountInformation}
EOF
ldapadd -x -h $3 -D "cn=Directory Manager" -W -f /tmp/amap.ldif

IFS=$'\n'
for line in $(cat /dev/shm/nis-map.$4); do
    IFS=" "
    key=$(echo "$line" | awk '{print $1}')
    info=$(echo "$line" | sed -e "s#^$key[ \t]*##")
    ipa automountkey-add nis $4 --key="$key" --info="$info"
done
```

The script exports the NIS automount information, generates an LDAP Data Interchange Format (LDIF) for the automount location and associated map, and imports the LDIF file into the IdM Directory Server. For further details, see [Section 21.4, “Exposing Automount Maps to NIS Clients”](#).

2. Authenticate as the IdM **admin** user:

```
[root@nis-server ~]# kinit admin
```

3. Run the script. For example:

```
[root@nis-server ~]# sh /root/nis-automounts.sh location nisdomain \
    nis-master.example.com map_name
```

21.5.4. Enabling Weak Password Hashing for NIS User Authentication

Using the Directory Server component's default setting, passwords stored in the ***userPassword*** attribute are hashed using the salted secure hash algorithm (SSHA). If your NIS clients require a weak hashing algorithm for passwords, update the password storage scheme setting.

Enabling a weak password hashing scheme affects only passwords stored in ***userPassword*** attribute. Note that Kerberos does not use this attribute and therefore Kerberos encryption is not affected by this setting.

For example, to enable **CRYPT** hashed passwords:

```
[root@server ~]# ldapmodify -D "cn=Directory Manager" -W -p 389 -h
ipaserver.example.com -x
dn: cn=config
changetype: modify
replace: passwordStorageScheme
passwordStorageScheme: crypt
```



NOTE

Because password hashes cannot be decrypted, Directory Server does not convert existing password hashes. The server applies the new password storage only to passwords set after you changed the storage scheme.

PART V. ADMINISTRATION: MANAGING AUTHENTICATION

CHAPTER 22. USER AUTHENTICATION

This chapter describes managing user authentication mechanisms, including information on how to manage users' passwords, SSH keys, and certificates, or how to configure one-time password (OTP) and smart-card authentication.



NOTE

For documentation on how to log in to Identity Management (IdM) using Kerberos, see [Chapter 5, *The Basics of Managing the IdM Server and Services*](#)

22.1. USER PASSWORDS

22.1.1. Changing and Resetting User Passwords

Regular users without the permission to change other users' passwords can change only their own personal password. Personal passwords changed in this way:

- Must meet the IdM password policies. For details on configuring password policies, see [Chapter 28, *Defining Password Policies*](#).

Administrators and users with password change rights can set initial passwords for new users and reset passwords for existing users. Passwords changed in this way:

- Do not have to meet the IdM password policies
- Expire after the first successful login. When this happens, IdM prompts the user to change the expired password immediately. To disable this behavior, see [Section 22.1.2, “Enabling Password Reset Without Prompting for a Password Change at the Next Login”](#).



NOTE

The LDAP Directory Manager (DM) user can change user passwords using LDAP tools. The new password can override any IdM password policies. Passwords set by DM do not expire after the first login.

22.1.1.1. Web UI: Changing Your Own Personal Password

1. In the top right corner, click **User name** → **Change password**.




Figure 22.1. Resetting Password

2. Enter the new password.

22.1.1.2. Web UI: Resetting Another User's Password

1. Select **Identity** → **Users**.
2. Click the name of the user to edit.
3. Click **Actions** → **Reset password**.




Figure 22.2. Resetting Password

4. Enter the new password, and click **Reset Password**.

Figure 22.3. Confirming New Password

22.1.1.3. Command Line: Changing or Resetting Another User's Password

To change your own personal password or to change or reset another user's password, add the **--password** option to the **ipa user-mod** command. The command will prompt you for the new password.

```
$ ipa user-mod user --password
Password:
Enter Password again to verify:
-----
Modified user "user"
-----
...
```

22.1.2. Enabling Password Reset Without Prompting for a Password Change at the Next Login

By default, when an administrator resets another user's password, the password expires after the first successful login. See [Section 22.1.1, “Changing and Resetting User Passwords”](#) for details.

To ensure that passwords set by administrators do not expire when used for the first time, make these changes on every Identity Management server in the domain:

- Edit the password synchronization entry:
cn=ipa_pwd_extop,cn=plugins,cn=config.
- Specify the administrative user accounts in the **passSyncManagersDNs** attribute. The attribute is multi-valued.

For example, to specify the **admin** user by using the **ldapmodify** utility:

```
$ ldapmodify -x -D "cn=Directory Manager" -W -h ldap.example.com -p 389

dn: cn=ipa_pwd_extop,cn=plugins,cn=config
changetype: modify
add: passSyncManagersDNs
passSyncManagersDNs: uid=admin,cn=users,cn=accounts,dc=example,dc=com
```

**WARNING**

Specify only the users who require these additional privileges. All users listed under *passSyncManagerDNs* can:

- Perform password change operations without requiring a subsequent password reset
- Bypass the password policy so that no strength or history enforcement is applied

22.1.3. Unlocking User Accounts After Password Failures

If a user attempts to log in using an incorrect password a certain number of times, IdM will lock the user account, which prevents the user from logging in. Note that IdM does not display any warning message that the user account has been locked.

**NOTE**

For information on setting the exact number of allowed failed attempts and the duration of the lockout, see [Chapter 28, Defining Password Policies](#).

IdM automatically unlocks the user account after a specified amount of time has passed. Alternatively, the administrator can unlock the user account manually.

Unlocking a User Account Manually

To unlock a user account, use the **ipa user-unlock** command.

```
$ ipa user-unlock user
-----
Unlocked account "user"
-----
```

After this, the user is able to log in again.

22.1.3.1. Checking the Status of a User Account

To display the number of failed login attempts for a user, use the **ipa user-status** command. If the displayed number exceeds the number of allowed failed login attempts, the user account is locked.

```
$ ipa user-status user
-----
Account disabled: False
-----
Server: example.com
Failed logins: 8
Last successful authentication: 20160229080309Z
Last failed authentication: 20160229080317Z
```


Time now: 2016-02-29T08:04:46Z

Number of entries returned 1

22.2. ONE-TIME PASSWORDS



IMPORTANT

The IdM solution for OTP authentication is only supported for clients running Red Hat Enterprise Linux 7.1 or later.

One-time password (OTP) is a password valid for only one authentication session and becomes invalid after use. Unlike a traditional static password, OTP generated by an authentication token keeps changing. OTPs are used as part of two-factor authentication:

1. The user authenticates with a traditional password.
2. The user provides an OTP code generated by a recognized OTP token.

Two-factor authentication is considered safer than authentication using a traditional password alone. Even if a potential intruder intercepts the OTP during login, the intercepted OTP will already be invalid by that point because it can only be used for successful authentication once.



WARNING

The following security and other limitations currently relate to the OTP support in IdM:

- The most important security limitation is the potential vulnerability to replay attacks across the system. Replication is asynchronous, and an OTP code can therefore be reused during the replication period. A user might be able to log on to two servers at the same time. However, this vulnerability is usually difficult to exploit due to comprehensive encryption.
- It is not possible to obtain a ticket-granting ticket (TGT) using a client that does not support OTP authentication. This might affect certain use cases, such as authentication using the **mod_auth_kerb** module or the Generic Security Services API (GSSAPI).

22.2.1. How OTP Authentication Works in IdM

22.2.1.1. OTP Tokens Supported in IdM

Software and Hardware Tokens

IdM supports both software and hardware tokens.

User-managed and Administrator-managed Tokens

Users can manage their own tokens, or the administrator can manage their tokens for them:

User-managed tokens

Users have full control over user-managed tokens in Identity Management: they are allowed to create, edit, or delete their tokens.

Administrator-managed tokens

The administrator adds administrator-managed tokens to the users' accounts. Users themselves have read-only access for such tokens: they do not have the permission to manage or modify the tokens and they are not required to configure them in any way.

Note that users cannot delete or deactivate a token if it is their only active token at the moment. As an administrator, you cannot delete or deactivate your last active token, but you can delete or deactivate the last active token of another user.

Supported OTP Algorithms

Identity Management supports the following two standard OTP mechanisms:

- The HMAC-Based One-Time Password (HOTP) algorithm is based on a counter. HMAC stands for Hashed Message Authentication Code.
- The Time-Based One-Time Password (TOTP) algorithm is an extension of HOTP to support time-based moving factor.

22.2.1.2. Available OTP Authentication Methods

When enabling OTP authentication, you can choose from the following authentication methods:

Two-factor authentication (password + OTP)

With this method, the user is always required to enter both a standard password and an OTP code.

Password

With this method, the user still has the option to authenticate using a standard password only.

RADIUS proxy server authentication

For information on configuring a RADIUS server for OTP validation, see [Section 22.2.6, “Migrating from a Proprietary OTP Solution”](#).

Global and User-specific Authentication Methods

You can configure these authentication methods either globally or for individual users:

- By default, user-specific authentication method settings take precedence over global settings. If no authentication method is set for a user, the globally-defined methods apply.

- You can disable per-user authentication method settings for any user. This ensures IdM ignores the per-user settings and always applies the global settings for the user.


Combining Multiple Authentication Methods

If you set multiple methods at once, either one of them will be sufficient for successful authentication. For example:

- If you configure both two-factor and password authentication, the user must provide the password (first factor), but providing the OTP (second factor) is optional when using the command line:

```
First Factor:  
Second Factor (optional):
```

- In the web UI, the user must still provide both factors.



NOTE

Individual hosts or services might be configured to require a certain authentication method, for example OTP. If you attempt to authenticate to such hosts or services using the first factor only, you will be denied access. See [Section 22.3, “Restricting Access to Services and Hosts Based on How Users Authenticate”](#).

However, a minor exception exists when RADIUS and another authentication method are configured:

- Kerberos will always use RADIUS, but LDAP will not. LDAP only recognizes the password and two-factor authentication methods.
- If you use an external two-factor authentication provider, use Kerberos from your applications. If you want to let users authenticate with a password only, use LDAP. It is recommended that the applications leverage Apache modules and SSSD, which allows to configure either Kerberos or LDAP.

22.2.1.3. GNOME Keyring Service Support

IdM integrates OTP authentication with the GNOME Keyring service. Note that GNOME Keyring integration requires the user to enter the first and second factors separately:

```
First factor: static_password  
Second factor: one-time_password
```

22.2.1.4. Offline Authentication with OTP

IdM supports offline OTP authentication. However, to be able to log in offline, the user must first authenticate when the system is online by entering the static password and OTP separately:

```
First factor: static_password  
Second factor: one-time_password
```

If both passwords are entered separately like this when logging in online, the user will

subsequently be able to authenticate even if the central authentication server is unavailable. Note that IdM only prompts for the first-factor traditional static password when the user authenticates offline.

IdM also supports entering both the static password and OTP together in one string in the **First factor** prompt. However, note that this is not compatible with offline OTP authentication. If the user enters both factors in a single prompt, IdM will always have to contact the central authentication server when authenticating, which requires the system to be online.



IMPORTANT

If you use OTP authentication on devices that also operate offline, such as laptops, Red Hat recommends to enter the static password and OTP separately to make sure offline authentication will be available. Otherwise, IdM will not allow you to log in after the system goes offline.

If you want to benefit from OTP offline authentication, apart from entering the static and OTP passwords separately, also make sure to meet the following conditions:

- The ***cache_credentials*** option in the ***/etc/sss/sss.conf*** file is set to **True**, which enables caching the first factor password.
- The first-factor static password meets the password length requirement defined in the ***cache_credentials_minimal_first_factor_length*** option set in ***/etc/sss/sss.conf***. The default minimal length is 8 characters. For more information about the option, see the ***sss.conf(5)*** man page.

Note that even if the ***krb5_store_password_if_offline*** option is set to **true** in ***/etc/sss/sss.conf***, SSSD does not attempt to refresh the Kerberos ticket-granting ticket (TGT) when the system goes online again because the OTP might already be invalid at that point. To obtain a TGT in this situation, the user must authenticate again using both factors.

22.2.2. Enabling OTP Authentication

For details on the available authentication methods related to OTP, see [Section 22.2.1.2, “Available OTP Authentication Methods”](#).


To enable OTP authentication using:

- the web UI, see [the section called “Web UI: Enabling OTP Authentication”](#).
- the command line, see [the section called “Command Line: Enabling OTP Authentication”](#).

Web UI: Enabling OTP Authentication

To set authentication methods globally for all users:

1. Select **IPA Server** → **Configuration**.
2. In the **User Options** area, select the required **Default user authentication types**.



Default user authentication types ⓘ


- ☐ Disable per-user override
- ☐ Password
- ☐ Radius
- ☒ Two factor authentication (password + OTP)

Figure 22.4. User Authentication Methods

To ensure the global settings are not overridden with per-user settings, select **Disable per-user override**. If you do not select **Disable per-user override**, authentication methods configured per user take precedence over the global settings.

To set authentication methods individually on a per-user basis:

1. Select **Identity** → **Users**, and click the name of the user to edit.
2. In the **Account Settings** area, select the required **User authentication types**.



User authentication types ⓘ

- ☐ Password
- ☐ Radius
- ☒ Two factor authentication (password + OTP)

Figure 22.5. User Authentication Methods

Command Line: Enabling OTP Authentication

To set authentication methods globally for all users:

1. Run the **ipa config-mod --user-auth-type** command. For example, to set the global authentication method to two-factor authentication:

```
$ ipa config-mod --user-auth-type=otp
```

For a list of values accepted by **--user-auth-type**, run the **ipa config-mod --help** command.

2. To disable per-user overrides, thus ensuring the global settings are not overridden with per-user settings, add the **--user-auth-type=disabled** option as well. For example, to set the global authentication method to two-factor authentication and disable per-user overrides:

```
$ ipa config-mod --user-auth-type=otp --user-auth-type=disabled
```

If you do not set **--user-auth-type=disabled**, authentication methods configured per user take precedence over the global settings.

To set authentication methods individually for a specified user:

- Run the **ipa user-mod --user-auth-type** command. For example, to set that user will be required to use two-factor authentication:

```
$ ipa user-mod user --user-auth-type=otp
```

To set multiple authentication methods, add **--user-auth-type** multiple times. For example, to configure both password and two-factor authentication globally for all users:

```
$ ipa config-mod --user-auth-type=otp --user-auth-type=password
```

22.2.3. Adding a User-Managed Software Token

1. Log in with your standard password.
2. Make sure the **FreeOTP Authenticator** application is installed on your mobile device. To download **FreeOTP Authenticator**, see [the FreeOTP source page](#).
3. Create the software token in the IdM web UI or from the command line.
 - To create the token in the web UI, click **Add** under the **OTP tokens** tab. If you are logged-in as the administrator, the **OTP Tokens** tab is accessible through the **Authentication** tab.




Figure 22.6. Adding an OTP Token for a User

- To create the token from the command line, run the **ipa otptoken-add** command.

```
$ ipa otptoken-add
-----
Added OTP token ""
-----
Unique ID: 7060091b-4e40-47fd-8354-cb32fec548a
Type: TOTP
...
```

For more information about **ipa otptoken-add**, run the command with the **--help** option added.

4. A QR code is displayed in the web UI or on the command line. Scan the QR code with **FreeOTP Authenticator** to provision the token to the mobile device.

22.2.4. Adding a User-Managed YubiKey Hardware Token

A programmable hardware token, such as a YubiKey token, can only be added from the command line. To add a YubiKey hardware token as the user owning the token:

1. Log in with your standard password.
2. Insert your YubiKey token.
3. Run the **ipa otptoken-add-yubikey** command.
 - If the YubiKey has an empty slot available, the command will select the empty slot automatically.
 - If no empty slot is available, you must select a slot manually using the **--slot** option. For example:

```
$ ipa otptoken-add-yubikey --slot=2
```

Note that this overwrites the selected slot.

22.2.5. Adding a Token for a User as the Administrator

To add a software token as the administrator:

1. Make sure you are logged-in as the administrator.
2. Make sure the **FreeOTP Authenticator** application is installed on the mobile device. To download **FreeOTP Authenticator**, see [the FreeOTP source page](#).
3. Create the software token in the IdM web UI or from the command line.
 - To create the token in the web UI, select **Authentication** → **OTP Tokens** and click **Add** at the top of the list of OTP tokens. In the **Add OTP Token** form, select the owner of the token.

Figure 22.7. Adding an Administrator-Managed Software Token

- To create the token from the command line, run the **ipa otptoken-add** command with the **--owner** option. For example:

```
$ ipa otptoken-add --owner=user
-----
Added OTP token ""
-----
Unique ID: 5303baa8-08f9-464e-a74d-3b38de1c041d
Type: TOTP
...
```

4. A QR code is displayed in the web UI or on the command line. Scan the QR code with **FreeOTP Authenticator** to provision the token to the mobile device.

To add a programmable hardware token, such as a YubiKey token, as the administrator:

1. Make sure you are logged-in as the administrator.
2. Insert the YubiKey token.
3. Run the **ipa otptoken-add-yubikey** command with the **--owner** option. For example:

```
$ ipa otptoken-add-yubikey --owner=user
```

22.2.6. Migrating from a Proprietary OTP Solution

To enable migrating a large deployment from a proprietary OTP solution to the IdM-native OTP solution, IdM offers a way to offload OTP validation to a third-party RADIUS server for a subset of users. The administrator creates a set of RADIUS proxies where each proxy can contain multiple individual RADIUS servers. The administrator then assigns one of these proxy sets to a user. As long as the user has a RADIUS proxy set assigned, IdM bypasses all other authentication mechanisms.



NOTE

IdM does not provide any token management or synchronization support for tokens in the third-party system.

To configure a RADIUS server for OTP validation and to add a user to the proxy server:

1. Make sure that the **radius** user authentication method is enabled. See [Section 22.2.2, “Enabling OTP Authentication”](#).
2. Run the **ipa radiusproxy-add proxy_name** command to add a RADIUS proxy. The command prompts you for the required information.
3. Run the **ipa user-mod radiususer --radius=proxy_name** command to assign a user to the added proxy.
4. If required, configure the user name to be sent to RADIUS by running the **ipa user-mod radiususer --radius-username=radius_user** command.

After this, the user OTP authentication will now be processed through the RADIUS proxy server.

When the user is ready to be migrated to the IdM native OTP system, you can simply remove the RADIUS proxy assignment for the user.

22.2.7. Promoting the Current Credentials to Two-Factor Authentication

If both password and two-factor authentication are configured, but you only authenticated using the password, you might be denied access to certain services or hosts (see [Section 22.3, “Restricting Access to Services and Hosts Based on How Users Authenticate”](#)). In this situation, promote your credentials from one-factor to two-factor authentication by authenticating again:

1. Lock your screen. The default keyboard shortcut to lock the screen is **Super key+L**.

2. Unlock your screen. When asked for credentials, use both password and OTP.

22.2.8. Resynchronizing an OTP Token


See [Section B.4.3, “OTP Token Out of Sync”](#).

22.3. RESTRICTING ACCESS TO SERVICES AND HOSTS BASED ON HOW USERS AUTHENTICATE

The authentication mechanisms supported by IdM vary in their authentication strength. For example, authentication using a one-time password (OTP) in combination with a standard password is considered safer than authentication using a standard password only. This section shows how to limit access to services and hosts based on how the user authenticates.

For example, you can configure:

- services critical to security, such as VPN, to require a strong authentication method
- noncritical services, such as local logins, to allow authentication using a weaker, but more convenient authentication method



RHEL_404973_1016

Figure 22.8. Example of Authenticating Using Different Methods

Authentication Indicators

Access to services and hosts is defined by *authentication indicators*:

- Indicators included in a service or host entry define what authentication methods the user can use to access that service or host.
- Indicators included in the user's ticket-granting ticket (TGT) show what authentication method was used to obtain the ticket.

If the indicator in the principal does not match the indicator in the TGT, the user is denied access.

22.3.1. Configuring a Host or a Service to Require a Specific Authentication Method

To configure a host or a service using:

- the web UI, see [the section called “Web UI: Configuring a Host or a Service to Require a Specific Authentication Method”](#)
- the command line, see [the section called “Command Line: Configuring a Host or a Service to Require a Specific Authentication Method”](#)

Web UI: Configuring a Host or a Service to Require a Specific Authentication Method

1. Select **Identity** → **Hosts** or **Identity** → **Services**.
2. Click the name of the required host or service.
3. Under **Authentication indicators**, select the required authentication method.
 - For example, selecting **OTP** ensures that only users who used a valid OTP code with their password will be allowed to access the host or service.
 - If you select both **OTP** and **RADIUS**, either OTP or RADIUS will be sufficient to allow access.
4. Click **Save** at the top of the page.

Command Line: Configuring a Host or a Service to Require a Specific Authentication Method

1. *Optional.* Use the **ipa host-find** or **ipa service-find** commands to identify the host or service.
2. Use the **ipa host-mod** or the **ipa service-mod** command with the **--auth-ind** option to add the required authentication indicator. For a list of the values accepted by **--auth-ind**, see the output of the **ipa host-mod --help** or **ipa service-mod --help** commands.

For example, **--auth-ind=otp** ensures that only users who used a valid OTP code with their password will be allowed to access the host or service:

```
$ ipa host-mod server.example.com --auth-ind=otp
-----
Modified host "server.example.com"
-----
Host name: server.example.com
...
Authentication Indicators: otp
...
```

If you add indicators for both OTP and RADIUS, either OTP or RADIUS will be sufficient to allow access.

22.4. MANAGING PUBLIC SSH KEYS FOR USERS

Identity Management allows you to upload a public SSH key to a user entry. The user who has access to the corresponding private SSH key can use **ssh** to log into an IdM machine without using Kerberos credentials. If **pam_krb5** is configured properly or if SSSD is used as the IdM server's identity provider, the user also receives a Kerberos ticket-granting ticket (TGT) after login; see [the section called “Obtaining Kerberos Tickets Automatically”](#) for more details.

Note that users can still authenticate by providing their Kerberos credentials if they are logging in from a machine where their private SSH key file is not available.

Caching and Retrieving SSH Keys Automatically

During an IdM server or client installation, SSSD is automatically configured on the machine to cache and retrieve user and host SSH keys. This allows IdM to serve as a universal and centralized repository of SSH keys.

If the server or client was not configured during installation, you can configure SSSD on the machine manually. For information on how to do this, see the [System-Level Authentication Guide](#). Note that caching SSH keys by SSSD requires administrative privileges on the local machines.

SSH Key Format Requirements

IdM accepts the following two SSH key formats:

OpenSSH-style key

See [RFC 4716](#) for more details about this format.

Raw RFC 4253-style key

See [RFC 4253](#) for more details about this format.

Note that IdM automatically converts RFC 4253-style keys into OpenSSH-style keys before saving them into the IdM LDAP server.

A key file, such as **id_rsa.pub**, consists of three parts: the key type, the key itself, and an additional comment or identifier. In the following example, the key type is RSA and the comment associates the key with the **client.example.com** host name:

```
ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDMM4xPu54Kf2dx7C4Ta2F7vnIzuL1i6P21TTKniskjFu
A+r
qW06588e7v14Im4VejwnNk352gp49A62qSV0zp8IKA9xdtyRmHYCTUvmkcyspZvFRI713zfRKQ
VFyJ0qHmW/m
dCmak7QBxYou2ELSPH3pe8MYTQIu1KDSu5Zbsrqedg1VGkSJxf7mDnCSPNWWzAY9AFB9Lmd2m
2xZmNgVAQEQ
nZXNMaIlroLD/51rmMSkJGHGb1068kEq9Z client.example.com
```

When uploading a key to IdM, you can either upload all three key parts, or only the key itself. If you only upload the key itself, IdM automatically identifies the key type, such as RSA or DSA, from the uploaded key.

22.4.1. Generating an SSH Key

You can generate an SSH key using the OpenSSH **ssh-keygen** utility. The utility displays information about the location of the public key. For example:

```

$ ssh-keygen -t rsa -C user@example.com
Generating public/private rsa key pair.
Enter file in which to save the key (/home/user/.ssh/id_rsa):
Created directory '/home/user/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/user/.ssh/id_rsa.
Your public key has been saved in /home/user/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:GAUIDVVEgly7rs1lTWP6oguHz8BKvyZkpqCqVSsmi7c user@example.com
The key's randomart image is:
+--[ RSA 2048]-----+
|      + .            |
|    + = .           |
|   =  +            |
|  . E S..          |
| .   . . .0        |
| . . . . 00.       |
| . 0 . . +.+0      |
| 0  .0..0+0       |
+-----+


```

To upload an SSH key for a user, use the public key string stored in the displayed file.

22.4.2. Uploading User SSH Keys

22.4.2.1. Web UI: Uploading User SSH Keys

1. Select **Identity** → **Users**.
2. Click the name of the user to edit.
3. Under the **Settings** tab in the **Account Settings** area, click **SSH public keys**: **Add**.



The screenshot shows a web form for user settings. It includes fields for 'Login shell' (set to /bin/sh) and 'Home directory' (set to /home/user) with an 'Undo' button. The 'SSH public keys' section is highlighted with a red rectangle and contains an 'Add' button. At the bottom, there is a 'Certificate' section showing a warning icon and the text 'No Valid Certificate'.

Figure 22.9. SSH public keys in the Account Settings

4. Paste in the Base 64-encoded public key string, and click **Set**.



Figure 22.10. Pasting in the Public Key

5. Click **Save** at the top of the page.

22.4.2.2. Command Line: Uploading User SSH Keys

Use the **ipa user-mod** command and pass the Base 64-encoded public key string using the **--sshpubkey** option.

For example, to upload the key type, the key itself, and the host name identifier:

```
$ ipa user-mod user --sshpubkey="ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDMM4xPu54Kf2dx7C4+hkZZZ8/E8JmjY5f8omlpSFTHG0wXP1h1BQIXvpHpWFNibadkzyT6PPutNcSTdS16owwhwLJUUsbDTIHawca3pGKzTa2F7vnIzuL1i6P21TTKniSkjFuA+W06588e7v14Im4VejwnNk352gp49A62qSV0zp8IKA9xdtyRmHYCTUvmkcyspZvFRI713zfRKQVFyJ0qHmWmdCmak7QBxYou2ELSPH3pe8MYTQIulKDSu5Zbsrqedg1VGkSJxf7mDnCSPNWWzAY9AFB9Lmd2m2xZmNgVAQEQnZXNMaiLroLD/51rmMSkJGHGb1068kEq9Z user@example.com"
```

To upload multiple keys, use **--sshpubkey** multiple times. For example, to upload two SSH keys:

```
--sshpubkey="AAAAB3NzaC1yc2EAAAADAQABAAQDMM4xPu54Kf2dx7C4+hkZZZ8/E8JmjY5f8omlpSFTHG0wXP1h1BQIXvpHpWFNibadkzyT6PPutNcSTdS16owwhwLJUUsbDTIHawca3pGKzTa2F7vnIzuL1i6P21TTKniSkjFuA+W06588e7v14Im4VejwnNk352gp49A62qSV0zp8IKA9xdtyRmHYCTUvmkcyspZvFRI713zfRKQVFyJ0qHmWmdCmak7QBxYou2ELSPH3pe8MYTQIulKDSu5Zbsrqedg1VGkSJxf7mDnCSPNWWzAY9AFB9Lmd2m2xZmNgVAQEQnZXNMaiLroLD/51rmMSkJGHGb1068kEq9Z user@example.com" --sshpubkey="Rj1zYQo...ZEt0TAo="
```



NOTE

Instead of pasting the key string manually into the command line, you can use command redirection and point to the file containing the key. For example:

```
$ ipa user-mod user --sshpubkey="$(cat ~/.ssh/id_rsa.pub)" --sshpubkey="$(cat ~/.ssh/id_rsa2.pub)"
```

22.4.3. Deleting User Keys

To delete an SSH key:

- using the web UI, see [Section 22.4.3.1, “Web UI: Deleting User SSH Keys”](#)
- using the command line, see [Section 22.4.3.2, “Command Line: Deleting User SSH Keys”](#)

22.4.3.1. Web UI: Deleting User SSH Keys

1. Select **Identity** → **Users**.
2. Click the name of the user to edit.
3. Under the **Settings** tab in the **Account Settings** area, click **Delete** next to the key you want to remove.



Figure 22.11. Deleting User SSH Public Key

4. Click **Save** at the top of the page.

22.4.3.2. Command Line: Deleting User SSH Keys

To delete all SSH keys assigned to a user account, add the `--sshpubkey` option to the `ipa user-mod` command without specifying any key:

```
$ ipa user-mod user --sshpubkey=
```

If you only want to delete a specific SSH key or keys, use the `--sshpubkey` option to specify the key or keys you want to keep.

22.5. CONFIGURING SSSD TO PROVIDE A CACHE FOR THE OPENSSH SERVICES

The System Security Services Daemon (SSSD) provides interfaces towards several system services, including OpenSSH. For details, see [the documentation for SSSD](#) in the *System-Level Authentication Guide*.


This section describes how you can configure SSSD to cache SSH keys for machines and users.

22.5.1. How SSSD Works with OpenSSH

OpenSSH is an SSH protocol implementation. OpenSSH creates secure, encrypted connections between two systems based on *public-private key pairs* that identify the authenticating entity. For details, see [OpenSSH](#) in the *System Administrator's Guide*.

SSSD can serve as a credentials cache for SSH public keys for machines and users. In this setup:

1. OpenSSH is configured to reference SSSD to check for cached keys.
2. SSSD uses an Identity Management (IdM) domain, and IdM stores the public keys and host information.

**NOTE**

Only Linux machines in the IdM domain can use SSSD as a key cache for OpenSSH. Other machines, including Windows machines, cannot.

How SSSD Manages Host Keys

To manage host keys, SSSD performs the following:

1. Retrieves the public host key from the host system.
2. Stores the host key in the `/var/lib/sss/pubconf/known_hosts` file.
3. Establishes a connection with the host machine.

See [Section 22.5.2, “Configuring OpenSSH to Use SSSD for Host Keys”](#) for details on the required configuration steps.

How SSSD Manages User Keys

To manage user keys, SSSD performs the following:

1. Retrieves the user's public key from the user entries in the IdM domain.
2. Stores the user key in the `.ssh/sss_authorized_keys` file in the standard authorized keys format.

See [Section 22.5.3, “Configuring OpenSSH to Use SSSD for User Keys”](#) for details on the required configuration steps.

22.5.2. Configuring OpenSSH to Use SSSD for Host Keys

You can change the configuration on a per-user basis or for the whole system.

1. Open the required configuration file.
 - a. To change user-specific configuration, open the `~/.ssh/config` file.
 - b. To change system-wide configuration, open the `/etc/ssh/sshd_config` file.
2. Use the ***ProxyCommand*** option to specify what command will be used to connect to the SSH client (the `sss_ssh_knownhostsproxy` utility with the required arguments and host name).

For details on `sss_ssh_knownhostsproxy`, see the `sss_ssh_knownhostsproxy(1)` man page.

3. Use the ***GlobalKnownHostsFile*** option to specify the location of the SSSD hosts file: `/var/lib/sss/pubconf/known_hosts`. This file will be used instead of the default OpenSSH `known_hosts` file.

The following example configures SSH to look for public keys in the SSSD domain and connect over the supplied port and host:

```
ProxyCommand /usr/bin/sss_ssh_knownhostsproxy -p %p %h
GlobalKnownHostsFile /var/lib/sss/pubconf/known_hosts
```

For details on configuring SSH and on the configuration files, see the `ssh_config(5)` man page.

22.5.3. Configuring OpenSSH to Use SSSD for User Keys

You can change the configuration on a per-user basis or for the whole system.

1. Open the required configuration file.
 - a. To change user-specific configuration, open the `~/.ssh/config` file.
 - b. To change system-wide configuration, open the `/etc/ssh/sshd_config` file.
2. Use the ***AuthorizedKeysCommand*** option to specify the command that will be executed to retrieve user keys.
3. Use the ***AuthorizedKeysCommandUser*** option to specify the user under whose account the command will be run.

The following example configures SSH to run the `sss_ssh_authorizedkeys` utility under the account of **user**.

```
AuthorizedKeysCommand /usr/bin/sss_ssh_authorizedkeys
AuthorizedKeysCommandUser user
```

For details on the `sss_ssh_authorizedkeys`, see the `sss_ssh_authorizedkeys(1)` man page.

For details on configuring SSH and on the configuration files, see the `ssh_config(5)` man page.

22.6. SMART-CARD AUTHENTICATION IN IDENTITY MANAGEMENT

For information on smart-card authentication in Identity Management, see [Chapter 23, *Smart-card Authentication in Identity Management*](#).

22.7. USER CERTIFICATES

For information on user certificates, see [Chapter 24, *Managing Certificates for Users, Hosts, and Services*](#).

CHAPTER 23. SMART-CARD AUTHENTICATION IN IDENTITY MANAGEMENT

Authentication based on smart cards is an alternative to passwords. User credentials are stored on the smart card, and special software and hardware is used to access them. The user places the smart card into a reader and supplies the PIN code for the smart card.

This chapter describes how an administrator can configure smart card-based authentication in Identity Management and how users can use smart cards to authenticate to Identity Management.

23.1. MANAGING SMART CARD LINKS IN THE IDENTITY MANAGEMENT SERVER

Before a user can use a smart card for authentication in the Identity Management domain, the administrator must link the certificate from the user's smart card with the corresponding user account in Identity Management. Linking a smart card certificate to a user account enables the user to authenticate with the smart card as the required role. This section describes how to manage links between a user smart card and one or more user accounts in the Identity Management server.

Before you can link the smart card to a user account, the smart card certificate must be available:

- If you need to extract the certificate from the smart card, see [Section 23.1.1, “Exporting a Certificate From a Smart Card”](#).

For details on creating the links between a certificate and a user account, see:

- [Section 23.1.2, “Linking User Accounts to Smart Card Certificates”](#)

If you need to find a user account that corresponds to a certain smart card certificate, see:

- [Section 23.1.3, “Finding Users That Match a Specified Certificate”](#)

23.1.1. Exporting a Certificate From a Smart Card

To export the certificate:

1. Place the smart card into the reader.
2. Use the following command to list the certificates on the smart card. In the output, locate the certificate to use for authentication, and note its nickname:

```
$ certutil -L -d /etc/pki/nssdb/ -h all
Certificate Nickname           Trust Attributes
                               SSL,S/MIME,JAR/XPI

my_certificate                 CT,C,C
```

3. Extract the certificate to a file using the certificate nickname. For example, to extract the certificate in the Base64 format to a file named **user.crt**:

```
$ certutil -L -d /etc/pki/nssdb/ -n 'my_certificate' -r | base64 -w
0 > user.crt
```

The **base64** utility is part of the **coreutils** package.

23.1.2. Linking User Accounts to Smart Card Certificates

As a security officer, you can manage links between a user's smart card and one or more user role accounts in the Identity Management Server. This enables the user to authenticate as a required role.

Create the link using one of the following options:

- Using the full certificate blob:
 - For Identity Management users, see [Section 23.1.2.1, “Creating a Link Between a Certificate and a User Account”](#). You can also remove such link using [Section 23.1.2.2, “Removing a Link Between a Certificate and a User Account”](#).
 - For Active Directory users, see [Section 23.1.2.3, “Linking an Active Directory User Account and a Smart Card”](#).
- Using certificate mapping: [Section 23.1.2.4, “Configuring Identity Mapping”](#)

23.1.2.1. Creating a Link Between a Certificate and a User Account

To link an Identity Management user account and a certificate, store the certificate in the user account.

Perform these steps on any Identity Management system.

Command Line: Creating a Link Between a Certificate and a User Account

1. Log in as the Identity Management administrator:

```
$ kinit admin
```

2. Add the smart card certificate to the user account using the **ipa user-add-cert** command. For example:

```
$ cat cert.pem | tail -n +2 | head -n -1 | tr -d '\r\n' | ipa user-
add-cert idm_user
```

Web UI: Creating a Link Between a Certificate and a User Account

1. Select **Identity** → **Users**, and click on the required user account.
2. Click **Add** next to the **Certificates** entry, and enter the certificate.
3. Click **Save** at the top of the user account page.

Additional Resources

- For details on adding and removing certificates issued by an external certificate authority (CA), see [Section 24.2, “Managing Certificates Issued by External CAs”](#).

23.1.2.2. Removing a Link Between a Certificate and a User Account

To remove a link between an Identity Management user account and a certificate, remove the certificate from the user account.

Perform these steps on any Identity Management system.

Command Line: Creating a Link Between a Certificate and a User Account

1. Log in as the Identity Management administrator:

```
$ kinit admin
```

2. Find the required user account:

```
$ ipa user-show idm_user
User login: idm_user
First name: first_name
Last name: last_name
...
Certificate: MIIC3...
```

3. Remove the certificate from the account:

```
$ ipa user-remove-cert idm_user --certificate MIIC3...
```

Web UI: Removing a Link Between a Certificate and a User Account

1. Select **Identity** → **Users**, and click on the required user account.
2. Click **Actions** next to the certificate to delete, and select **Delete**.

Additional Resources

For details on adding and removing certificates issued by an external certificate authority (CA), see [Section 24.2, “Managing Certificates Issued by External CAs”](#).

23.1.2.3. Linking an Active Directory User Account and a Smart Card

If it is possible to modify the user entry in Active Directory, store the user certificate in the user entry in Active Directory. See Active Directory documentation for details. This ensures that Identity Management can read the smart card certificate from the user object in Active Directory.

If it is not possible to modify the user entry in Active Directory, store the user certificate in an ID view. You can use the Default Trust View in Identity Management to store the certificate, or you can create a new ID view. The information in the ID view overrides the information in the user object in Active Directory, which also enables you to set up different smart card links for systems enrolled in Active Directory and systems enrolled in Identity Management.

Perform these steps on an Identity Management server.

Command Line: Linking an Active Directory User Account and a Smart Card

1. Log in as the Identity Management administrator:

```
$ kinit admin
```

2. Create an environment variable (**CERT**) for the user certificate:

```
$ CERT=`cat cert.pem | tail -n +2 | head -n -1 | tr -d '\r\n'`
```

3. Add the user certificate to the ID view by creating a new ID override. In this procedure, we are using the Default Trust View:

```
$ ipa idoverrideuser-add 'Default Trust View' ad_user@ad.example.com  
--certificate $CERT
```

Web UI: Linking an Active Directory User Account and a Smart Card

1. Select **Identity** → **ID Views**, and click on the required ID view.
2. Add the user certificate to the ID view by creating a new ID override. Click **Add**, and fill out the required information in the **Add User ID override** form.

Additional Resources

- For details on managing ID views, see [Chapter 18, ID Views](#).
- For details on the Default Trust View, see [Using ID Views in Active Directory Environments](#).

23.1.2.4. Configuring Identity Mapping

As a security officer, you can manage links between a user's smart card and one or more user role accounts using identity matching and mapping rules. This enables you to allow users to access to Identity Management even if it is not possible to store the smart card certificate in the user entry or in an ID override, for example if you do not have physical access to the smart card.

For a basic overview of identity mapping, see:

- [Section 23.1.2.4.1, “Identity Mapping in Identity Management”](#)

For the procedures to configure identity mapping, see:

- [Section 23.1.2.4.2, “Creating a Certificate Identity Mapping Rule”](#)
- [Section 23.1.2.4.3, “Linking a User Account and a Smart Card Certificate”](#)

For various examples, see:

- [Section 23.1.2.4.4, “Examples of Identity Mapping Rules”](#)
- [Section 23.1.2.4.5, “Examples of Translating the Issuer from a Certificate to a Matching Rule”](#)

23.1.2.4.1. Identity Mapping in Identity Management

Identity mapping is configured by creating *identity mapping rules*. Identity Management supports the following components in identity mapping rules. All components are optional:

Mapping rule

A mapping rule associates (or *maps*) a certificate with one or more user accounts. The rule defines an LDAP search filter that associates a certificate with the intended user account.

Certificates issued by different certificate authorities (CAs) might have different properties and might be used in different domains. Therefore, Identity Management does not apply mapping rules unconditionally, but only to the appropriate certificates. The appropriate certificates are defined using *matching rules*.

Matching rule

A matching rule selects a certificate or CA to which you want to apply the mapping rule.

Domain list

The domain list specifies the DNS domain names in which you want Identity Management to search the users when processing identity mapping rules.



NOTE

If you do not specify any domains, Identity Management searches the users only in the local domain to which the client belongs.

Priority

When multiple rules are applicable to a certificate, the rule with the highest priority takes precedence. All other rules are ignored.

- The lower the numerical value, the higher the priority of the identity mapping rule. For example, a rule with a priority 1 has higher priority than a rule with a priority 2.
- If a rule has no priority value defined, it has the lowest priority.

23.1.2.4.2. Creating a Certificate Identity Mapping Rule

Creating a certificate identity mapping rule ensures that Identity Management can correctly map a smart card certificate to a user account.

Command Line: Creating a Certificate Identity Mapping Rule

Perform these steps on the server:

1. Log in as the administrator:

```
$ kinit admin
```

2. Create the rule by using the **ipa certmaprule-add** command. To specify the components for the identity mapping rule, use these options:
 - **--maprule** defines the mapping rule
 - **--matchrule** defines the matching rule
 - **--domain** defines the domain in which you want to search the user entry

- **--priority** defines the priority of the identity mapping rule

For example, to create a simple identity mapping rule that consists only of a mapping rule and a matching rule:

```
$ ipa certmaprule-add rule_name --matchrule '<ISSUER>CN=Smart Card
CA,0=EXAMPLE.ORG' --maprule '(ipacertmapdata=X509:<I>
{issuer_dn!nss_x500}<S>{subject_dn!nss_x500})'
-----
Added Certificate Identity Mapping Rule "rule_name"
-----
Rule name: rule_name
Mapping rule: (ipacertmapdata=X509:<I>{issuer_dn!nss_x500}<S>
{subject_dn!nss_x500})
Matching rule: <ISSUER>CN=Smart Card CA,0=EXAMPLE.ORG
Enabled: TRUE
```

This rule now links the subject and issuer from the smart card certificate to the value of the **ipacertmapdata** attribute in the user account.

Web UI: Creating a Certificate Identity Mapping Rule

1. Select **Authentication → Certificate Identity Mapping Rules**.
2. Click **Add**.
3. Fill out the components of the rule, and click **Add**.

Additional Resources

- For details on the syntax of the certificate mapping and matching rules, see the `sss-certmap(5)` man page.
- For details on using the **ipa certmaprule-add** command, execute it with the **--help** option.
- For additional commands for managing identity mapping, use the **ipa help certmap** command.

23.1.2.4.3. Linking a User Account and a Smart Card Certificate

To link a specific user account with a smart card certificate, save the subject and issuer from the certificate in the **ipacertmapdata** attribute.

Command Line: Linking a User Account and a Smart Card Certificate

- If you have access to the certificate, use the full certificate blob:

```
$ CERT=`cat cert.pem | tail -n +2 | head -n -1 | tr -d '\r\n'`
$ ipa user-add-certmapdata idm_user --certificate $CERT
-----
Added certificate mappings to user "idm_user"
-----
User login: idm_user
Certificate mapping data: X509:<I>0=EXAMPLE.ORG,CN=Smart Card
CA<S>CN=test,0=EXAMPLE.ORG
```

- If you do not have access to the certificate, but know the subject and issuer, use the **--subject** and **--issuer** options:

```
$ ipa user-add-certmapdata idm_user --subject
"O=EXAMPLE.ORG,CN=test" --issuer "CN=Smart Card CA,O=EXAMPLE.ORG"
-----
Added certificate mappings to user "idm_user"
-----
User login: idm_user
Certificate mapping data: X509:<I>O=EXAMPLE.ORG,CN=Smart Card
CA<S>CN=test,O=EXAMPLE.ORG
```


- If you are comfortable with the mapping format, provide the mapping data directly:

```
$ ipa user-add-certmapdata idm_user 'X509:<I>O=EXAMPLE.ORG,CN=Smart
Card CA<S>CN=test,O=EXAMPLE.ORG'
-----
Added certificate mappings to user "idm_user"
-----
User login: idm_user
Certificate mapping data: X509:<I>O=EXAMPLE.ORG,CN=Smart Card
CA<S>CN=test,O=EXAMPLE.ORG
```

The user can now use the smart card to log in to the Identity Management server.

Web UI: Linking a User Account and a Smart Card Certificate

1. Click **Identity** → **Users**, and click on the required user login.
2. Click **Add** next to the **Certificate mapping data** entry.



The screenshot shows a web interface for managing user accounts. At the top, there is a 'Certificates' section with an 'Add' button. Below this, there is a 'Certificate mapping data' section, which is highlighted with a red rectangular box. This section also has an 'Add' button. Further down, there is a 'User authentication types' section with three options: 'Password', 'RADIUS', and 'Two factor authentication (password + OTP)', each with an unchecked checkbox.

Figure 23.1. Adding certificate mapping data

3. In the **Add Certificate Mapping Data** form, fill out the required information. Specify one of the following:
 - The full certificate blob under **Certificate**
 - The subject and issuer under **Issuer and subject**
 - The mapping data directly under **Certificate mapping data**

Additional Resources

- For details on the **ipa user-add-certmapdata** command, execute it with the **--help** option.

23.1.2.4.4. Examples of Identity Mapping Rules

Example 23.1. Active Directory Certificates for Identity Management Users

```
$ ipa certmaprule-add ad_cert_for_ipa_users \
  --maprule='(ipacertmapdata=X509:<I>{issuer_dn!nss_x500}<S>
{subject_dn!nss_x500})' \
  --matchrule='<ISSUER>CN=AD-ROOT-CA,DC=ad,DC=example,DC=com' \
  --domain=idm.example.com
```

Example 23.2. Active Directory Certificates for Active Directory Users

```
$ ipa certmaprule-add ad_cert_for_ad_users \
  --maprule='(altSecurityIdentities=X509:<I>{issuer_dn!ad_x500}<S>
{subject_dn!ad_x500})' \
  --matchrule='<ISSUER>CN=AD-ROOT-CA,DC=ad,DC=example,DC=com' \
  --domain=ad.example.com
```

Example 23.3. Active Directory Certificates for Both Identity Management and Active Directory Users

```
$ ipa certmaprule-add ad_cert_for_ipa_and_ad_users \
  --maprule='(|(ipacertmapdata=X509:<I>{issuer_dn!nss_x500}<S>
{subject_dn!nss_x500})(altSecurityIdentities=X509:<I>{issuer_dn!ad_x500}
<S>{subject_dn!ad_x500}))' \
  --matchrule='<ISSUER>CN=AD-ROOT-CA,DC=ad,DC=example,DC=com' \
  --domain=ad.example.com
```

In the above example, the filter definition in the **--maprule** option includes these criteria:

- **ipacertmapdata=X509:<I>{issuer_dn!nss_x500}<S>{subject_dn!nss_x500}** is a filter that links the subject and issuer from a smart card certificate to the value of the **ipacertmapdata** attribute in an Identity Management user account
- **altSecurityIdentities=X509:<I>{issuer_dn!ad_x500}<S>{subject_dn!ad_x500}** is a filter that links the subject and issuer from a smart card certificate to the value of the **altSecurityIdentities** attribute in an Active Directory user account

The filter definition in the **--maprule** option accepts the logical operator **|** (or), so that you can specify multiple criteria. In this case, the rule maps all user accounts that meet at least one of the criteria.

Example 23.4. Identity Management Certificates for Identity Management and Active Directory Users

```
$ ipa certmaprule-add ipa_cert_for_ad_users \
  --maprule='(|(userCertificate;binary={cert!bin})(ipacertmapdata=X509:
<I>{issuer_dn!nss_x500}<S>{subject_dn!nss_x500})
(altSecurityIdentities=X509:<I>{issuer_dn!ad_x500}<S>
{subject_dn!ad_x500}))' \
  --matchrule='<ISSUER>CN=Certificate Authority,0=REALM.EXAMPLE.COM' \
  --domain=idm.example.com --domain=ad.example.com
```

In the above example, the filter definition in the **--maprule** option includes these criteria:

- **userCertificate;binary={cert!bin}** is a filter that returns Identity Management or Active Directory user entries that include the whole certificate
- **ipacertmapdata=X509:<I>{issuer_dn!nss_x500}<S>{subject_dn!nss_x500}** is a filter that links the subject and issuer from a smart card certificate to the value of the **ipacertmapdata** attribute in an Identity Management user account
- **altSecurityIdentities=X509:<I>{issuer_dn!ad_x500}<S>{subject_dn!ad_x500}** is a filter that links the subject and issuer from a smart card certificate to the value of the **altSecurityIdentities** attribute in an Active Directory user account

The filter definition in the **--maprule** option accepts the logical operator **|** (or), so that you can specify multiple criteria. In this case, the rule maps all user accounts that meet at least one of the criteria.

23.1.2.4.5. Examples of Translating the Issuer from a Certificate to a Matching Rule

To get the issuer format required by a matching rule, reverse the path, and replace the delimiters (/) with commas:

Example 23.5. Translating the Issuer from a Certificate Issued by Identity Management

Example of a certificate issued by Identity Management:

```
# openssl x509 -in user.crt -noout -issuer
issuer= /O=REALM.EXAMPLE.COM/CN=Certificate Authority
```

The issuer of this certificate expressed in the format required by a matching rule:

```
'<ISSUER>CN=Certificate Authority,0=REALM.EXAMPLE.COM'
```

Example 23.6. Translating the Issuer from a Certificate with an Email Included

Example of a certificate that includes an email address:

```
# openssl x509 -in expired_user.pem -noout -issuer
issuer= /C=US/ST=North Carolina/L=Raleigh/O=Red
Hat/OU=QE/CN=ExampleCA/emailAddress=admin@example.com
```

The issuer of this certificate expressed in the format required by a matching rule:

```
'<ISSUER>emailAddress=admin@example.com,CN=ExampleCA,OU=QE,O=Red
Hat,L=Raleigh,ST=North Carolina,C=US'
```

23.1.2.5. Additional Resources

- To verify the smart-card certificate links, see [Section 23.1.3, “Finding Users That Match a Specified Certificate”](#).
- For more details on identity mapping for certificates, see [Matching and Mapping Certificates](#) in the upstream SSSD documentation.

23.1.3. Finding Users That Match a Specified Certificate

To list all employees whose role account matches the certificate, present the Identity Management server with an employee's smart card certificate.

Perform these steps on any Identity Management system.

Command Line: Finding Users That Match a Specified Certificate

1. Log in as the administrator:

```
$ kinit admin
```

2. To find the user, specify one of the following:

- The name of the certificate file:

```
$ ipa certmap-match cert.pem
-----
1 user matched
-----
Domain: IDM.EXAMPLE.COM
User logins: idm_user
-----
Number of entries returned 1
-----
```

- The contents of the certificate:

```
$ ipa certmap-match --certificate="MII...."
-----
1 user matched
-----
Domain: IDM.EXAMPLE.COM
User logins: idm_user
```

```
-----  
Number of entries returned 1  
-----
```

This command returns also users in a trusted Active Directory domain if their user entries contain the full certificate blob:

```
$ ipa certmap-match --certificate="MII...."  
-----  
2 users matched  
-----  
Domain: ad.domain.com  
User logins: ad_user  
  
Domain: IDM.EXAMPLE.COM  
User logins: idm_user  
-----  
Number of entries returned 2  
-----
```

Web UI: Finding Users That Match a Specified Certificate

- 1. Click **Authentication** → **Certificate Identity Mapping Rules** → **Certificate Mapping Match**.
- 2. Enter the contents of the certificate in the **Certificate** field, and click **Match**. Identity Management displays the users who match the certificate under **Matched Users**.

Matched Users	
User Login	Domain
user	EXAMPLE.COM
1 user matched	

Figure 23.2. List of users who match a certificate

Additional Resources

- For details on the commands for certificate identity mapping, use the **ipa help certmap** command.
- For details on the **ipa certmap-match** command, execute it with the **--help** option.

23.1.4. Additional Resources

- For information on managing personal certificates and keys using the Enterprise Security Client, an application for Red Hat Certificate System, see [Managing Smart Cards with the Enterprise Security Client](#) in the Certificate System documentation.

23.2. AUTHENTICATING TO AN IDENTITY MANAGEMENT CLIENT WITH A SMART CARD

As an Identity Management user with multiple role accounts in the Identity Management server, you can authenticate with your smart card to a desktop client system joined to the Identity Management domain. This enables you to use the client system as the selected role.

For a basic overview of the supported options, see:

- [Section 23.2.1, “Smart Card-based Authentication Options Supported on Identity Management Clients”](#)

For information on configuring the environment to enable the authentication, see:

- [Section 23.2.2, “Preparing the Identity Management Client for Smart-card Authentication”](#)

For information on how to authenticate, see:

- [Section 23.2.3, “Authenticating on an Identity Management Client with a Smart Card Using the Console Login”](#)
- [Section 23.2.4, “Authenticating on an Identity Management Client with a Smart Card Using SSH”](#)

23.2.1. Smart Card-based Authentication Options Supported on Identity Management Clients

Users in Identity Management can use the following options when authenticating using a smart card on Identity Management clients.

Local authentication

Local authentication includes authentication using:

- the text console
- the graphical console, such as the Gnome Display Manager (GDM)
- local authentication services, such as **su** or **sudo**

Remote authentication with ssh

Certificates on a smart card are stored together with the PIN-protected SSH private key.

Smart card-based authentication using other services, such as FTP, is not supported.

23.2.2. Preparing the Identity Management Client for Smart-card Authentication

As the Identity Management administrator, perform these steps:

1. On the server, create a shell script to configure the client.

- a. Use the **ipa-adviser config-client-for-smart-card-auth** command, and save its output to a file:

```
# ipa-adviser config-client-for-smart-card-auth >
client_smart_card_script.sh
```

- b. Open the script file, and review its contents.
- c. Add execute permissions to the file using the **chmod** utility:

```
# chmod +x client_smart_card_script.sh
```

2. Copy the script to the client, and run it. Add the path to the PEM file with the certificate authority (CA) that signed the smart card certificate:

```
# ./client_smart_card_script.sh CA_cert.pem
```

Additionally, if an external certificate authority (CA) signed the certificate on the smart card, add the smart card CA as a trusted CA:

1. On the Identity Management server, install the CA certificate:

```
# ipa-cacert-manage -n "SmartCard CA" -t CT,C,C install ca.pem
# ipa-certupdate
```

Repeat **ipa-certupdate** also on all replicas and clients.

2. Restart the HTTP server:

```
# systemctl restart httpd
```

Repeat **systemctl restart httpd** also on all replicas.



NOTE

SSSD enables administrators to tune the certificate verification process with the **certificate_verification** parameter, for example if the Online Certificate Status Protocol (OCSP) servers defined in the certificate are not reachable from the client. For more information, see the `sssd.conf(5)` man page.

23.2.3. Authenticating on an Identity Management Client with a Smart Card Using the Console Login

To authenticate as an Identity Management user, enter the user name and PIN.

- When logging in from the command line:

```
client login: idm_user
PIN for PIV Card Holder pin (PIV_II) for user
idm_user@idm.example.com:
```

- When logging in using the Gnome Desktop Manager (GDM), GDM prompts you for the smart card PIN after you select the required user:




Figure 23.3. Entering the smart card PIN in the Gnome Desktop Manager

To authenticate as an Active Directory user, enter the user name in a format that uses the NetBIOS domain name: `AD.EXAMPLE.COM\ad_user` or `ad_user@AD.EXAMPLE.COM`.

If the authentication fails, see [Section A.4, “Investigating Smart Card Authentication Failures”](#).

23.2.4. Authenticating on an Identity Management Client with a Smart Card Using SSH

When using the `ssh` utility, specify the path to the smart card reader module. For example:

```
$ ssh -I /usr/lib/libmypkcs11.so -l user@example.com host.example.com
Enter PIN for 'Smart Card':
```

If the authentication fails, see [Section A.4, “Investigating Smart Card Authentication Failures”](#).

23.2.5. Additional Resources

- For details on smart-card authentication with OpenSSH, see [Using Smart Cards to Supply Credentials to OpenSSH](#) in the *Security Guide*.

23.3. AUTHENTICATING TO AN IDENTITY MANAGEMENT SYSTEM REMOTELY WITH A SMART CARD

As an Identity Management user with multiple role accounts in the Identity Management server, you can authenticate with your smart card from a local system (not enrolled into the Identity Management domain) to a remote system (enrolled in the Identity Management domain) by using the `ssh` utility. This enables you to use the remote system as the selected role.

For information on configuring the environment to enable the authentication, see:

- [Section 23.3.1, “Preparing the Local System for Smart-card Authentication”](#)

- [Section 23.3.2, “Preparing the Remote Identity Management System for Smart-card Authentication”](#)
- [Section 23.3.3, “Linking the Smart Card Certificate and the User Entry in Active Directory”](#)

For information on how to authenticate, see:

- [Section 23.3.4, “Authenticating to the Remote System from the Local System”](#)

23.3.1. Preparing the Local System for Smart-card Authentication

As the administrator, perform these steps on the local system:

1. Install the `opensc` package:

```
# yum install opensc
```

2. Make sure the `pcscd` service for the smart-card daemon is started and enabled:

```
# systemctl start pcscd.socket pcscd.service  
# systemctl enable pcscd.socket pcscd.service
```

Additionally, if an external certificate authority (CA) signed the certificate on the smart card, add the smart card CA as a trusted CA:

1. On the Identity Management server, install the CA certificate:

```
# ipa-cacert-manage -n "SmartCard CA" -t CT,C,C install ca.pem  
# ipa-certupdate
```

Repeat **ipa-certupdate** also on all replicas and clients.

2. Restart the HTTP server on the Identity Management server:

```
# systemctl restart httpd
```

Repeat **systemctl restart httpd** also on all replicas.

23.3.2. Preparing the Remote Identity Management System for Smart-card Authentication

As the administrator, perform these steps:

1. Install the smart card certificate authority (CA) certificate in the `/etc/pki/nssdb/` database on the remote system:

```
# certutil -A -d /etc/pki/nssdb/ -n "SmartCard CA" -t CT,C,C -i  
ca.pem
```

2. Make sure the `sssd-dbus` package is installed.

23.3.3. Linking the Smart Card Certificate and the User Entry in Active Directory

If the user entry is stored in Active Directory, the administrator must link the entry with the smart card certificate. See [Section 23.1.2.3, “Linking an Active Directory User Account and a Smart Card”](#).

23.3.4. Authenticating to the Remote System from the Local System

On the local system, perform these steps:

1. Insert the smart card.
2. Launch **ssh**, and specify the PKCS#11 library with the **-I** option:

- As an Identity Management user:

```
$ ssh -I /usr/lib64/opensc-pkcs11.so -l idm_user
server.idm.example.com

Enter PIN for 'PIV_II (PIV Card Holder pin)':
Last login: Thu Apr  6 12:49:32 2017 from 10.36.116.42
```

- As an Active Directory user:

```
$ ssh -I /usr/lib64/opensc-pkcs11.so -l ad_user@ad.example.com
server.idm.example.com

Enter PIN for 'PIV_II (PIV Card Holder pin)':
Last login: Thu Apr  6 12:49:32 2017 from 10.36.116.42
```

3. *Optional.* Use the **id** utility to check that you are logged in as the intended user.

- As an Identity Management user:

```
$ id
uid=1928200001(idm_user) gid=1928200001(idm_user)
groups=1928200001(idm_user)
context=unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023
```

- As an Active Directory user:

```
$ id
uid=1171201116(ad_user@ad.example.com)
gid=1171201116(ad_user@ad.example.com)
groups=1171201116(ad_user@ad.example.com),1171200513(domain
users@ad.example.com)
context=unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023
```

If the authentication fails, see [Section A.4, “Investigating Smart Card Authentication Failures”](#).

23.3.5. Additional Resources

- Authentication using **ssh** with a smart card does not obtain a ticket-granting ticket (TGT) on the remote system. To obtain a TGT on the remote system, the administrator must configure Kerberos on the local system and enable Kerberos delegation. For an example of the required configuration, see [this Kerberos knowledge base entry](#).
- For details on smart-card authentication with OpenSSH, see [Using Smart Cards to Supply Credentials to OpenSSH](#) in the *Security Guide*.

23.4. CONFIGURING A USER NAME HINT POLICY FOR SMART-CARD AUTHENTICATION

As an Identity Management administrator, you can configure a *user name hint* policy for smart cards linked with multiple accounts.

23.4.1. User Name Hints in Identity Management

The user name hint policy configures Identity Management to prompt smart card users for their user name. When a user tries to authenticate with a smart card certificate that matches multiple user accounts in Identity Management, one of the following occurs:

- If the user name hint policy is enabled, the user is prompted for a user name and then can proceed with authentication.
- If the user name hint policy is disabled, the authentication fails without prompting.

Identity Management adds the user name hint to applications that would by default prompt for a smart card PIN without asking for a user name. On Red Hat Enterprise Linux, this is currently only the Gnome Desktop Manager (GDM) login.




Figure 23.4. User name hint in the Gnome Desktop Manager

Identity Management does not add the user name hint to applications that ask for a user name by default, for example:

- The Identity Management web UI authentication, because the GUI always displays the **Username** field
- **ssh** authentication, because **ssh** uses the current user's login name or the name provided with the **-l** option or in the **username@host** format

- Console authentication, where the login name is always supplied

In these situations, authentication with a certificate that matches multiple users is always allowed.

23.4.2. Enabling User Name Hints in Identity Management

The Identity Management administrator sets the user name hint policy centrally. The policy applies to all hosts enrolled into the Identity Management domain.

Perform these steps on any Identity Management system.

Command Line: Enabling User Name Hints in Identity Management

1. Log in as the Identity Management administrator:

```
$ kinit admin
Password for admin@IDM.EXAMPLE.COM:
```

2. Enable user name hints by using the **ipa certmapconfig-mod** command with the **--promptusername=True** option.

```
$ ipa certmapconfig-mod --promptusername=TRUE
Prompt for the username: TRUE
```

To disable user name hints, use the **--promptusername=False** option.

Web UI: Enabling User Name Hints in Identity Management

1. Click **Authentication** → **Certificate Identity Mapping Rules** → **Certificate Identity Mapping Global Configuration**.
2. Select **Prompt for the username**, and click **Save**.




Figure 23.5. Enabling user name hints in the web UI

Additional Resources

- For details on the **ipa certmapconfig-mod** command, execute it with the **--help** option.

23.5. PKINIT SMART-CARD AUTHENTICATION IN IDENTITY MANAGEMENT

Identity Management users can authenticate with a smart card to a desktop client system

joined to Identity Management and get a Kerberos ticket-granting ticket (TGT) automatically. The users can use the ticket for further single sign-on (SSO) authentication from the client.

23.5.1. Preparing the Identity Management Client for PKINIT Authentication

As the Identity Management administrator, perform these steps on the client where you want the users to authenticate:

1. On the server, create a shell script to configure the client.
 - a. Use the **ipa-adviser config-client-for-smart-card-auth** command, and save its output to a file:

```
# ipa-adviser config-client-for-smart-card-auth >  
client_smart_card_script.sh
```

- b. Open the script file, and review its contents.
- c. Add execute permissions to the file using the **chmod** utility:

```
# chmod +x client_smart_card_script.sh
```

2. Copy the script to the client, and run it. Add the path to the PEM file with the certificate authority (CA) that signed the smart card certificate:

```
# ./client_smart_card_script.sh CA_cert.pem
```

3. Make sure the **krb5-pkinit** package is installed.

Additionally, if an external certificate authority (CA) signed the certificate on the smart card, add the smart card CA as a trusted CA:

1. On the Identity Management server, install the CA certificate:

```
# ipa-cacert-manage -n "SmartCard CA" -t CT,C,C install ca.pem  
# ipa-certupdate
```

Repeat **ipa-certupdate** also on all replicas and clients.

2. Restart the HTTP server:

```
# systemctl restart httpd
```

Repeat **systemctl restart httpd** also on all replicas.

**NOTE**

SSSD enables administrators to tune the certificate verification process with the ***certificate_verification*** parameter, for example if the Online Certificate Status Protocol (OCSP) servers defined in the certificate are not reachable from the client. For more information, see the `sssd.conf(5)` man page.

23.5.2. As an Identity Management User: Authenticate Using PKINIT on an Identity Management Client

Authenticate using the **kinit** utility on an Identity Management client:

```
$ kinit -X X509_user_identity='PKCS11:opensc-pkcs11.so' idm_user
```

The **-X** option specifies the **opensc-pkcs11.so** module as the pre-authentication attribute. For details, see the `kinit(1)` man page.

23.5.3. As an Active Directory User: Authenticate Using PKINIT on an Identity Management Client

Prerequisites

As the administrator, configure the environment to support PKINIT authentication for Active Directory users:

- Configure the Active Directory server to trust the certificate authority (CA) that issued the smart card certificate. Import the CA in the NTAuth store (see [Microsoft support](#)), and add the CA as a trusted CA. See Active Directory documentation for details.
- Configure the Kerberos client to trust the CA that issued the smart card certificate:
 1. On the Identity Management client, open the **/etc/krb5.conf** file.
 2. Add the following lines to the file:

```
[libdefaults]
[... file truncated ...]
pkinit_eku_checking = kpServerAuth
pkinit_kdc_hostname = adserver.ad.domain.com
```

- If the user certificates do not contain a certificate revocation list (CRL) distribution point extension, configure Active Directory to ignore revocation errors:
 1. Save the following REG-formatted content in a plain text file, and double-click the file to import it to the Windows Registry:

```
Windows Registry Editor Version 5.00

[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Kdc]
"UseCachedCRLOnlyAndIgnoreRevocationUnknownErrors"=dword:00000001
```

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA\Kerberos  
\Parameters]  
"UseCachedCRLOnlyAndIgnoreRevocationUnknownErrors"=dword:00000001
```

Alternatively, set the values manually using the **regedit.exe** application.

2. Reboot the Windows system to apply the changes.

Procedure

Authenticate using the **kinit** utility on an Identity Management client. Specify the Active Directory user with the user name and domain name:

```
$ kinit -X X509_user_identity='PKCS11:opensc-pkcs11.so'  
ad_user@AD.DOMAIN.COM
```

The **-X** option specifies the **opensc-pkcs11.so** module as the pre-authentication attribute. For details, see the `kinit(1)` man page.

23.6. AUTHENTICATING TO THE IDENTITY MANAGEMENT WEB UI WITH A SMART CARD

As an Identity Management user with multiple role accounts in the Identity Management server, you can authenticate with your smart card to the Identity Management web UI as a selected role. This enables you to use the web UI as the selected role.



NOTE

Only Identity Management users can log in to the web UI with a smart card. Active Directory users can log in with their user name and password. For details, see [Section 5.4.2.3, “Authenticating to the IdM Web UI as an AD User”](#).

For information on configuring the environment to enable the authentication, see:

- [Section 23.6.1, “Preparing the Identity Management Server for Smart-card Authentication in the Web UI”](#)
- [Section 23.6.2, “Preparing the Browser for Smart-card Authentication”](#)

For information on how to authenticate, see:

- [Section 23.6.3, “Authenticating to the Identity Management Web UI with a Smart Card as an Identity Management User”](#)

23.6.1. Preparing the Identity Management Server for Smart-card Authentication in the Web UI

As the Identity Management administrator:

1. On an Identity Management server, create a shell script to configure the server.
 - a. Use the **ipa-adviser config-server-for-smart-card-auth** command, and save its output to a file:

```
# ipa-adviser config-server-for-smart-card-auth >
```

```
server_smart_card_script.sh
```

- b. Open the script file, and review its contents.
- c. Add execute permissions to the file using the **chmod** utility:

```
# chmod +x server_smart_card_script.sh
```

2. Run the script on all servers in the Identity Management domain.
3. Make sure the `sssd-dbus` package is installed.

Additionally, if an external certificate authority (CA) signed the certificate on the smart card:

1. On an Identity Management server, add the CA certificate to the NSS database used by the HTTP server:

```
# ipa-cacert-manage -n "SmartCard CA" -t CT,C,C install ca.pem
# ipa-certupdate
```

Repeat **ipa-certupdate** on all replicas and clients.

2. Restart the HTTP server and the Kerberos server:

```
# systemctl restart httpd
# systemctl restart krb5kdc
```

Repeat the commands on all replicas.

23.6.2. Preparing the Browser for Smart-card Authentication

To configure the browser for smart-card authentication, perform these steps on the client from which the user launches the web browser to access the web UI. The system on which the browser is running does not have to be part of the Identity Management domain. In this procedure, we are using the Firefox browser.

1. Launch Firefox.
2. Configure Firefox to read the certificate from the smart card.
 - a. Select **Edit** → **Preferences** → **Advanced** → **Certificates** → **Security Devices**




Figure 23.6. Configuring security devices in Firefox

- b. Click **Load**. In the **Load PKCS#11 Device** window, fill out the following information:

- **Module Name:** OpenSC
- **Module filename:** /usr/lib64/opensc-pkcs11.so




Figure 23.7. Device Manager in Firefox

- c. Click **OK** to confirm. Then click **OK** to close the Device Manager.

Firefox can now use smart card certificates for authentication.

23.6.3. Authenticating to the Identity Management Web UI with a Smart Card as an Identity Management User

To authenticate:

1. Insert the smart card into the smart card reader.

2. In the browser, navigate to the Identity Management web UI at **`https://ipaserver.example.com/ipa/ui`**.
3. If the smart card certificate is linked to a single user account, do not fill out the **Username** field.

If the smart card certificate is linked to multiple user accounts, fill out the **Username** field to specify the required account.

4. Click **Login Using Certificate**.


The image shows the Red Hat Identity Management login web UI. It has a dark grey background. At the top, it says "RED HAT® IDENTITY MANAGEMENT" in white. Below that, there are two input fields: "Username" with a placeholder "Username" and "Password" with a placeholder "Password or Password+One-Time-Password". At the bottom right, there are three buttons: "Login Using Certificate" (highlighted with a red rectangle), "Sync OTP Token", and "Login".

Figure 23.8. Login Using Certificate in the Identity Management web UI

5. Enter the smart card PIN when prompted.


The image shows a "Password Required" dialog box. It has a title bar with "Password Required" and a close button (X). Inside, there is a key icon and the text "Please enter the master password for the PIV_II (PIV Card Holder pin)." Below the text is a password input field with a blue border and a cursor. At the bottom right, there are "Cancel" and "OK" buttons. A mouse cursor is pointing at the "OK" button.

Figure 23.9. Entering the smart card PIN

6. A new window opens, proposing the certificate to use. Select the smart card certificate.




Figure 23.10. Selecting the smart card certificate

You are now authenticated as the user who corresponds to the smart card certificate.

Additional Resources

- If the authentication fails, see [Section A.4, “Investigating Smart Card Authentication Failures”](#).

23.6.4. Additional Resources

- For details on the Identity Management web UI, see [Section 5.4, “The IdM Web UI”](#).

23.7. INTEGRATING IDENTITY MANAGEMENT SMART-CARD AUTHENTICATION WITH WEB APPLICATIONS

As a developer whose applications use the Identity Management server as an authentication back end through the Identity Management web infrastructure Apache modules, you can configure the applications to enable authentication of users with multiple role accounts linked to their smart card. This enables these users to use the application under allowed role accounts.

23.7.1. Prerequisites for Web Application Authentication with Smart Cards

On the server where the Apache web application is running:

- Enroll the server as a client in the Identity Management domain.
- Install the `sssd-dbus` and `mod_lookup_identity` packages.
- Make sure Apache has a working HTTPS connection configured using the `mod_nss` module.

23.7.2. Configuring Identity Management Smart-card Authentication for a Web Application

1. Enable TLS renegotiation in the **mod_nss** configuration in the **/etc/httpd/conf.d/nss.conf** file:

```
NSSRenegotiation
NSSRequireSafeNegotiation on
```

2. Make sure that the CA issuing the user certificates is trusted for the client certificates in the **mod_nss** certificate database. The default location for the database is **/etc/httpd/alias**.
3. Add the web application. In this procedure, we are using an almost minimal example consisting of a login page and a protected area.
 - The **/login** end point only lets the user provide a user name and sends the user to a protected part of the application.
 - The **/app** end point checks the **REMOTE_USER** environment variable. If the login was successful, the variable contains the ID of the logged-in user. Otherwise, the variable is unset.
4. Create a directory, and set its group to **apache** and the mode to at least **750**. In this procedure, we are using a directory named **/var/www/app/**.
5. Create a file, and set its group to **apache** and the mode to at least **750**. In this procedure, we are using a file named **/var/www/app/login.py**.

Save the following contents to the file:

```
#!/usr/bin/env python

def application(environ, start_response):
    status = '200 OK'
    response_body = """
<!DOCTYPE html>
<html>
  <head>
    <title>Login</title>
  </head>
  <body>
    <form action='/app' method='get'>
      Username: <input type='text' name='username'>
      <input type='submit' value='Login with certificate'>
    </form>
  </body>
</html>
"""
    response_headers = [
        ('Content-Type', 'text/html'),
        ('Content-Length', str(len(response_body)))
    ]
    start_response(status, response_headers)
    return [response_body]
```

6. Create a file, and set its group to **apache** and the mode to at least **750**. In this procedure, we are using a file named **/var/www/app/protected.py**.

Save the following contents in the file:

```
#!/usr/bin/env python

def application(environ, start_response):
    try:
        user = environ['REMOTE_USER']
    except KeyError:
        status = '400 Bad Request'
        response_body = 'Login failed.\n'
    else:
        status = '200 OK'
        response_body = 'Login succeeded. Username:
{}'.format(user)

    response_headers = [
        ('Content-Type', 'text/plain'),
        ('Content-Length', str(len(response_body)))
    ]
    start_response(status, response_headers)
    return [response_body]
```

7. Create a configuration file for your application. In this procedure, we are using a file named **/etc/httpd/conf.d/app.conf** with the following contents:

```
<IfModule !lookup_identity_module>
    LoadModule lookup_identity_module modules/mod_lookup_identity.so
</IfModule>

WSGIScriptAlias /login /var/www/app/login.py
WSGIScriptAlias /app /var/www/app/protected.py

<Location "/app">
    NSSVerifyClient require
    NSSUserName SSL_CLIENT_CERT
    LookupUserByCertificate On
    LookupUserByCertificateParamName "username"
</Location>
```

In this file:

- The first part loads **mod_lookup_identity** if it is not already loaded.
- The next part maps the **/login** and **/app** end points to the respective Web Server Gateway Interface (WSGI) scripts.
- The last part configures **mod_nss** for the **/app** end point so that it requires a client certificate during the TLS handshake and uses it. In addition, it configures an optional request parameter **username** to look up the identity of the user.

CHAPTER 24. MANAGING CERTIFICATES FOR USERS, HOSTS, AND SERVICES

Identity Management (IdM) supports two types of certificate authorities (CAs):

Integrated IdM CA

Integrated CAs can create, revoke, and issue certificates for users, hosts, and services. For more details, see [Section 24.1, “Managing Certificates with the Integrated IdM CAs”](#).

IdM supports creating lightweight sub-CAs. For more details, see [Section 26.1, “Lightweight Sub-CAs”](#)

External CA

An external CA is a CA other than the integrated IdM CA.

Using IdM tools, you add certificates issued by these CAs to users, services, or hosts as well as remove them. For more details, see [Section 24.2, “Managing Certificates Issued by External CAs”](#).

Each user, host, or service can have multiple certificates assigned.



NOTE

For more details on the supported CA configurations of the IdM server, see [Section 2.3.2, “Determining What CA Configuration to Use”](#).

24.1. MANAGING CERTIFICATES WITH THE INTEGRATED IDM CAS

24.1.1. Requesting New Certificates for a User, Host, or Service

To request a certificate using:

- the IdM web UI, see [the section called “Web UI: Requesting New Certificates”](#)
- the command line, see [the section called “Command Line: Requesting New Certificates”](#).

Note that you must generate the certificate request itself with a third-party tool. The following procedures use the **certutil** and **openssl** utilities.



IMPORTANT

Services typically run on dedicated service nodes on which the private keys are stored. Copying a service's private key to the IdM server is considered insecure. Therefore, when requesting a certificate for a service, create the CSR on the service node.

Web UI: Requesting New Certificates

1. Under the **Identity** tab, select the **Users, Hosts, or Services** subtab.

2. Click the name of the user, host, or service to open its configuration page.

Hosts

Search

Refresh

Delete

Add

Actions

☐	Host name	Description	Enrolled
☐	server.example.com		True

Showing 1 to 1 of 1 entries.

Figure 24.1. List of Hosts

3. Click **Actions** → **New Certificate**.
4. Optional: Select the issuing CA and profile ID.
5. Follow the instructions on the screen for using **certutil**.
6. Click **Issue**.

Command Line: Requesting New Certificates

Request a new certificate using **certutil** in standard situations - see [Section 24.1.1.1, “Requesting New Certificates Using certutil”](#). Request a new certificate using **openSSL** to enable a Kerberos alias to use a host or service certificate - see [Section 24.1.1.2, “Requesting New Certificates Using openSSL”](#).

24.1.1.1. Requesting New Certificates Using certutil

1. Create a new temporary certificate database, for instance:

```
# certutil -N -d ~/certdb/
```

2. Create the certificate signing request (CSR) and redirect the output to a file. For example, to create a CSR for a 4096 bit certificate and to set the subject to `CN=server.example.com,O=EXAMPLE.COM`:

```
# certutil -R -d ~/certdb/ -a -g 4096 -s
"CN=server.example.com,O=EXAMPLE.COM" -8 server.example.com >
certificate_request.csr
```

3. Submit the certificate request file to the server. Be sure to specify the Kerberos principal to associate with the newly-issued certificate:

```
# ipa cert-request certificate_request.csr --
principal=host/server.example.com
```

IdM uses the following defaults:

- Certificate profile: **caIPAServiceCert**

To select a custom profile, use the **--profile-id** option with the **ipa cert-request** command.

- Integrated CA: **ipa** (IdM root CA)

To select a sub-CA, use the **--ca** option with the **ipa cert-request** command.

24.1.1.2. Requesting New Certificates Using openssl

1. Create one or more aliases, for example *test1/server.example.com*, *test2/server.example.com*, for your Kerberos principal *test/server.example.com*. See [Section 20.2.1, “Kerberos Principal Alias”](#) for more details.
2. In the CSR, add a subjectAltName for *dnsName* (*server.example.com*) and *otherName* (*test2/server.example.com*). To do this, configure the **openssl.conf** file so that it includes the following line specifying the UPN *otherName* and *subjectAltName*:

```
otherName=1.3.6.1.4.1.311.20.2.3;UTF8:test2/server.example.com@EXAMPLE.COM
DNS.1 = server.example.com
```

3. Create a certificate request using **openssl**:

```
openssl req -new -newkey rsa:2048 -keyout test2service.key -sha256 -nodes -out certificate_request.csr -config openssl.conf
```

24.1.2. Revoking Certificates with the Integrated IdM CAs

If you need to invalidate the certificate before its expiration date, you can revoke it. To revoke a certificate using:

- the IdM web UI, see [the section called “Web UI: Revoking Certificates”](#)
- the command line, see [the section called “Command Line: Revoking Certificates”](#)

A revoked certificate is invalid and cannot be used for authentication. All revocations are permanent, except for reason 6: Certificate Hold.

Table 24.1. Revocation Reasons


ID	Reason	Explanation
0	Unspecified	
1	Key Compromised	The key that issued the certificate is no longer trusted. Possible causes: lost token, improperly accessed file.
2	CA Compromised	The CA that issued the certificate is no longer trusted.

ID	Reason	Explanation
3	Affiliation Changed	Possible causes: • A person has left the company or moved to another department. • A host or service is being retired.
4	Superseded	A newer certificate has replaced the current certificate.
5	Cessation of Operation	The host or service is being decommissioned.
6	Certificate Hold	The certificate is temporarily revoked. You can restore the certificate later.
8	Remove from CRL	The certificate is not included in the certificate revocation list (CRL).
9	Privilege Withdrawn	The user, host, or service is no longer permitted to use the certificate.
10	Attribute Authority (AA) Compromise	The AA certificate is no longer trusted.

Web UI: Revoking Certificates

To revoke a certificate:

1. Open the **Authentication** tab, and select the **Certificates** subtab.
2. Click the serial number of the certificate to open the certificate information page.



Subject ▼		Search 🔍	Refresh ↻	+ Issue
☐	Serial Number	Subject		
☐	1	CN=Certificate Authority,O=EXAMPLE.COM		
☐	2	CN=OCSP Subsystem,O=EXAMPLE.COM		
☐	3	CN=server.example.com,O=EXAMPLE.COM		
☐	4	CN=CA Subsystem O=EXAMPLE.COM		

Figure 24.2. List of Certificates

3. Click **Actions** → **Revoke Certificate**.
4. Select the reason for revoking, and click **Revoke**. See [Table 24.1, “Revocation Reasons”](#) for details.

Command Line: Revoking Certificates

Use the **ipa cert-revoke** command, and specify:

- the certificate serial number
- a number that identifies the reason for the revocation; see [Table 24.1, “Revocation Reasons”](#) for details

For example, to revoke the certificate with serial number **1032** because of reason 1: Key Compromised:

```
$ ipa cert-revoke 1032 --revocation-reason=1
```


24.1.3. Restoring Certificates with the Integrated IdM CAs

If you have revoked a certificate because of reason 6: Certificate Hold, you can restore it again. To restore a certificate using:

- the IdM web UI, see [the section called “Web UI: Restoring Certificates”](#)
- the command line, see [the section called “Command Line: Restoring Certificates”](#)

Web UI: Restoring Certificates

1. Open the **Authentication** tab, and select the **Certificates** subtab.
2. Click the serial number of the certificate to open the certificate information page.



Certificates		
	Subject	
☐	1	CN=Certificate Authority,O=EXAMPLE.COM
☐	2	CN=OCSP Subsystem,O=EXAMPLE.COM
☐	3	CN=server.example.com,O=EXAMPLE.COM
☐	4	CN=CA Subsystem,O=EXAMPLE.COM

Figure 24.3. List of Certificates

3. Click **Actions** → **Restore Certificate**.

Command Line: Restoring Certificates

Use the **ipa cert-remove-hold** command and specify the certificate serial number. For example:

```
$ ipa cert-remove-hold 1032
```

24.2. MANAGING CERTIFICATES ISSUED BY EXTERNAL CAs

24.2.1. Command Line: Adding and Removing Certificates Issued by External CAs

To add a certificate to a user, host, or service:

- **ipa user-add-cert**
- **ipa host-add-cert**
- **ipa service-add-cert**

To remove a certificate from a user, host, or service:

- **ipa user-remove-cert**
- **ipa host-remove-cert**
- **ipa service-remove-cert**

A certificate issued by an external CA is not revoked after you remove it from IdM. This is because the certificate does not exist in the IdM CA database. You can only revoke these certificates manually from the external CA side.

The commands require you to specify the following information:

- the name of the user, host, or service
- the Base64-encoded DER certificate

To run the commands interactively, execute them without adding any options.

To provide the required information directly with the command, use command-line arguments and options:

```
$ ipa user-add-cert user --certificate=MIQTPrajQAwg...
```



NOTE

Instead of copying and pasting the certificate contents into the command line, you can convert the certificate to the DER format and then re-encode it to base64. For example, to add the **user_cert.pem** certificate to **user**:

```
$ ipa user-add-cert user --certificate="$(openssl x509 -outform  
der -in user_cert.pem | base64 -w 0)"
```

24.2.2. Web UI: Adding and Removing Certificates Issued by External CAs

To add a certificate to a user, host, or service:

1. Open the **Identity** tab, and select the **Users**, **Hosts**, or **Services** subtab.
2. Click on the name of the user, host, or service to open its configuration page.
3. Click **Add**, next to the **Certificates** entry.

User: demouser
demouser is a member of:

Settings | User Groups | Netgroups | Roles | HBAC Rules | Sudo Rules

Refresh Revert Save Actions

Identity Settings

Job Title:

First name *:

Last name *:

Full name *:

Display name:

Initials:

GECOS:

Class:

Account Settings

User login: demouser

Password: *****

Password expiration: 2016-07-14 10:14:41Z

UID:

GID:

Principal alias: demouser@IDM.EXAMPLE.COM Delete

Add

Kerberos principal expiration: : UTC

Login shell:

Home directory:

SSH public keys: Add

Certificates: Add

Figure 24.4. Adding a Certificate to a User Account

4. Paste the certificate in Base64 or PEM encoded format into the text field, and click **Add**.
5. Click **Save** to store the changes.

To remove a certificate from a user, host, or service:


1. Open the **Identity** tab, and select the **Users, Hosts, or Services** subtab.
2. Click on the name of the user, host, or service to open its configuration page.
3. Click the **Actions** next to the certificate to delete, and select **Delete**.
4. Click **Save** to store the changes.

24.3. LISTING AND DISPLAYING CERTIFICATES

Listing and Displaying Certificates in the Web UI

To list certificates assigned to a user, host, or service entry:

1. Open the **Identity** tab, and select the **Users, Hosts, or Services** subtab.
2. Click on the name of the user, host, or service to open its configuration page.



☐	Host name	Description	Enrolled
☐	server.example.com		True


Showing 1 to 1 of 1 entries.

Figure 24.5. List of Hosts

3. The configuration page lists all certificates assigned to the entry. Additionally, clicking **Show** displays a particular certificate.

To list all certificates registered on the IdM server:

1. Open the **Authentication** tab, and select the **Certificates** subtab.
2. A list of all certificates is displayed in the **Certificates** section. To display a particular certificate, click on its serial number.



☐	Serial Number	Subject
☐	1	CN=Certificate Authority,O=EXAMPLE.COM
☐	2	CN=OCSP Subsystem,O=EXAMPLE.COM
☐	3	CN=server.example.com,O=EXAMPLE.COM
☐	4	CN=CA Subsystem O=EXAMPLE.COM

Figure 24.6. List of Certificates

Listing Certificates from the Command Line

To list all certificates in the IdM database, run the **ipa cert-find** command.

```
$ ipa cert-find
-----
10 certificates matched
-----
  Serial number (hex): 0x1
  Serial number: 1
  Status: VALID
  Subject: CN=Certificate Authority,O=EXAMPLE.COM
  ...
-----
Number of entries returned 10
-----
```

You can filter the search results by specifying certain certificate properties, such as issue

date or validity date. For example, to search by an issue date interval, use the **--issuedon-from** or **--issuedon-to** options to specify the start and end points or a period of time.

```
ipa cert-find --issuedon-from=2018-01-07 --issuedon-to=2018-02-07
```

For a complete list of options used to filter the search for a certificate, run **ipa cert-find** with the **--help** option added.

Displaying Certificates from the Command Line

To display a certificate, use the **ipa cert-show** command and specify the serial number.

```
$ ipa cert-show 132
Serial number: 132
Certificate:
MIIDtzCCAp+gAwIBAgIBATANBgqhkiG9w0BAQsFADBBMR8wHQYDVQKExZMQUIu
...
LxIQjrEFtJmoBGB/TWRLwGEWylayr4iTEf1ayZ+RGNylLalEAtk9RLjEjg==
Subject: CN=Certificate Authority,O=EXAMPLE.COM
Issuer: CN=Certificate Authority,O=EXAMPLE.COM
Not Before: Sun Jun 08 05:51:11 2014 UTC
Not After: Thu Jun 08 05:51:11 2034 UTC
Serial number (hex): 0x132
Serial number: 132
```

To display the certificates assigned to a user, host, or service entry, use **ipa cert-show** and specify the entry. For example, to display the certificate assigned to a user:

```
$ ipa user-show user
User login: user
...
Certificate: MIICfzCCAwwCAQA...
...
```

You can also save a certificate to a file by adding the **--out** option to **ipa cert-show**.

```
$ ipa cert-show certificate_serial_number --out=path_to_file
```

Note that if the user, host, or service has more than one certificate, the **--out** option exports all of them. The certificate or certificates are exported as PEM objects.

24.4. CERTIFICATE PROFILES

A certificate profile defines the content of certificates belonging to the particular profile, as well as constraints for issuing the certificates, enrollment method, and input and output forms for enrollment. A single certificate profile is associated with issuing a particular type of certificate. Different certificate profiles can be defined for users, services, and hosts in IdM.

The CA uses certificate profiles in signing of certificates to determine:

- whether the CA can accept a certificate signing request (CSR)
- what features and extensions should be present on the certificate

IdM includes the following two certificate profiles by default: **caIPAServiceCert** and **IECUserRoles**. In addition, custom profiles can be imported.

Custom certificate profiles allow issuing certificates for specific, unrelated purposes. For example, it is possible to restrict use of a particular profile to only one user or one group, preventing other users and groups from using that profile to issue a certificate for authentication.

For details on supported certificate profile configuration, see [Defaults Reference](#) and [Constraints Reference](#) in the Red Hat Certificate System *Administration Guide*.



NOTE

By combining certificate profiles and CA ACLs, [Section 24.5, “Certificate Authority ACL Rules”](#), the administrator can define and control access to custom certificate profiles. For a description of using profiles and CA ACLs to issue user certificates, see [Section 24.6, “Using Certificate Profiles and ACLs to Issue User Certificates with the IdM CAs”](#).

24.4.1. Certificate Profile Management from the Command Line

The **certprofile** plug-in for management of IdM profiles allows privileged users to import, modify, or remove IdM certificate profiles. To display all commands supported by the plug-in, run the **ipa certprofile** command:

```
$ ipa certprofile
Manage Certificate Profiles

...

EXAMPLES:

  Import a profile that will not store issued certificates:
    ipa certprofile-import ShortLivedUserCert \
      --file UserCert.profile --desc "User Certificates" \
      --store=false

  Delete a certificate profile:
    ipa certprofile-del ShortLivedUserCert

...
```

Note that to perform the **certprofile** operations, you must be operating as a user who has the required permissions. IdM includes the following certificate profile-related permissions by default:

System: Read Certificate Profiles

Enables users to read all profile attributes.

System: Import Certificate Profile

Enables users to import a certificate profile into IdM.

System: Delete Certificate Profile

Enables users to delete an existing certificate profile.

System: Modify Certificate Profile

Enables users to modify the profile attributes and to disable or enable the profile.

All these permissions are included in the default **CA Administrator** privilege. For more information on IdM role-based access controls and managing permissions, see [Section 10.4, “Defining Role-Based Access Controls”](#).



NOTE

When requesting a certificate, the **--profile-id** option can be added to the **ipa cert-request** command to specify which profile to use. If no profile ID is specified, the default **caIPAServiceCert** profile is used for the certificate.

This section only describes the most important aspects of using the **ipa certprofile** commands for profile management. For complete information about a command, run it with the **--help** option added, for example:

```
$ ipa certprofile-mod --help
Usage: ipa [global-options] certprofile-mod ID [options]

Modify Certificate Profile configuration.
Options:
  -h, --help          show this help message and exit
  --desc=STR          Brief description of this profile
  --store=B00L        Whether to store certs issued using this profile
  ...
```

Importing Certificate Profiles

To import a new certificate profile to IdM, use the **ipa certprofile-import** command. Running the command without any options starts an interactive session in which the **certprofile-import** script prompts you for the information required to import the certificate.

```
$ ipa certprofile-import

Profile ID: smime
Profile description: S/MIME certificates
Store issued certificates [True]: TRUE
Filename of a raw profile. The XML format is not supported.: smime.cfg
-----
Imported profile "smime"
-----
Profile ID: smime
Profile description: S/MIME certificates
Store issued certificates: TRUE
```

The **ipa certprofile-import** command accepts several command-line options. Most notably:

--file

This option passes the file containing the profile configuration directly to **ipa certprofile-import**. For example:

■

```
$ ipa certprofile-import --file=smime.cfg
```

--store

This option sets the **Store issued certificates** attribute. It accepts two values:

- **True**, which delivers the issued certificates to the client and stores them in the target IdM principal's **userCertificate** attribute.
- **False**, which delivers the issued certificates to the client, but does not store them in IdM. This option is most commonly-used when issuing multiple short-term certificates is required.

Import fails if the profile ID specified with **ipa certprofile-import** is already in use or if the profile content is incorrect. For example, the import fails if a required attribute is missing or if the profile ID value defined in the supplied file does not match the profile ID specified with **ipa certprofile-import**.

To obtain a template for a new profile, you can run the **ipa certprofile-show** command with the **--out** option, which exports a specified existing profile to a file. For example:

```
$ ipa certprofile-show caIPAServiceCert --out=file_name
```

You can then edit the exported file as required and import it as a new profile.

Displaying Certificate Profiles

To display all certificate profiles currently stored in IdM, use the **ipa certprofile-find** command:

```
$ ipa certprofile-find
-----
3 profiles matched
-----
  Profile ID: caIPAServiceCert
  Profile description: Standard profile for network services
  Store issued certificates: TRUE

  Profile ID: IECUserRoles
  ...
```

To display information about a particular profile, use the **ipa certprofile-show** command:

```
$ ipa certprofile-show profile_ID
  Profile ID: profile_ID
  Profile description: S/MIME certificates
  Store issued certificates: TRUE
```

Modifying Certificate Profiles

To modify an existing certificate profile, use the **ipa certprofile-mod** command. Pass the required modifications with the command using the command-line options accepted by **ipa certprofile-mod**. For example, to modify the description of a profile and change whether IdM stores the issued certificates:

```
$ ipa certprofile-mod profile_ID --desc="New description" --store=False
```

```
-----
Modified Certificate Profile "profile_ID"
-----
```

```
Profile ID: profile_ID
Profile description: New description
Store issued certificates: FALSE
```

To update the certificate profile configuration, import the file containing the updated configuration using the **--file** option. For example:

```
$ ipa certprofile-mod profile_ID --file=new_configuration.cfg
```

Deleting Certificate Profiles

To remove an existing certificate profile from IdM, use the **ipa certprofile-del** command:

```
$ ipa certprofile-del profile_ID
-----
Deleted profile "profile_ID"
-----
```

24.4.2. Certificate Profile Management from the Web UI

To manage certificate profiles from the IdM web UI:

1. Open the **Authentication** tab and the **Certificates** subtab.
2. Open the **Certificate Profiles** section.

Identity

Policy

Authentication

Network Services

IPA Server

Certificates

OTP Tokens

RADIUS Servers

Certificates

Certificates

Certificate Profiles >

CA ACLs

Certificate Profiles

Search

Refresh

Delete

☐	Profile ID	Profile description	Store issued certificates
☐	IECUserRoles	User profile that includes IECUserRoles extension from request	TRUE
☐	calPAserviceCert	Standard profile for network services	TRUE

Showing 1 to 2 of 2 entries.


Figure 24.7. Certificate Profile Management in the Web UI

In the **Certificate Profiles** section, you can display information about existing profiles, modify their attributes, or delete selected profiles.

For example, to modify an existing certificate profile:

1. Click on the name of the profile to open the profile configuration page.
2. In the profile configuration page, fill in the required information.

3. Click **Save** to confirm the new configuration.



Certificate Profile: smime

Refresh Revert **Save**

Profile ID smime

Profile description * S/Mime certificates

Undo

Store issued certificates ☒

Figure 24.8. Modifying a Certificate Profile in the Web UI

If you enable the **Store issued certificates** option, the issued certificates are delivered to the client as well as stored in the target IdM principal's **userCertificate** attribute. If the option is disabled, the issued certificates are delivered to the client, but not stored in IdM. Storing certificates is often disabled when issuing multiple short-lived certificates is required.

Note that some certificate profile management operations are currently unavailable in the web UI:

- It is not possible to import a certificate profile in the web UI. To import a certificate, use the **ipa certprofile-import** command.
- It is not possible to set, add, or delete attribute and value pairs. To modify the attribute and value pairs, use the **ipa certprofile-mod** command.
- It is not possible to import updated certificate profile configuration. To import a file containing updated profile configuration, use the **ipa certprofile-mod --file=file_name** command.

For more information about the commands used to manage certificate profiles, see [Section 24.4.1, “Certificate Profile Management from the Command Line”](#).

24.4.3. Upgrading IdM Servers with Certificate Profiles

When upgrading an IdM server, the profiles included in the server are all imported and enabled.

If you upgrade multiple server replicas, the profiles of the first upgraded replica are imported. On the other replicas, IdM detects the presence of other profiles and does not import them or resolve any conflicts between the two sets of profiles. If you have custom

profiles defined on replicas, make sure the profiles on all replicas are consistent before upgrading.

24.5. CERTIFICATE AUTHORITY ACL RULES

Certificate Authority access control list (CA ACL) rules define which profiles can be used to issue certificates to which users, services, or hosts. By associating profiles, principals, and groups, CA ACLs permit principals or groups to request certificates using particular profiles:

- an ACL can permit access to multiple profiles
- an ACL can have multiple users, services, hosts, user groups, and host groups associated with it

For example, using CA ACLs, the administrator can restrict use of a profile intended for employees working from an office located in London only to hosts that are members of the London office-related group.



NOTE

By combining certificate profiles, described in [Section 24.4, “Certificate Profiles”](#), and CA ACLs, the administrator can define and control access to custom certificate profiles. For a description of using profiles and CA ACLs to issue user certificates, see [Section 24.6, “Using Certificate Profiles and ACLs to Issue User Certificates with the IdM CAs”](#).

24.5.1. CA ACL Management from the Command Line

The **caacl** plug-in for management of CA ACL rules allows privileged users to add, display, modify, or delete a specified CA ACL. To display all commands supported by the plug-in, run the **ipa caacl** command:

```
$ ipa caacl
Manage CA ACL rules.
```

```
...
```

EXAMPLES:

Create a CA ACL "test" that grants all users access to the "UserCert" profile:

```
ipa caacl-add test --usercat=all
ipa caacl-add-profile test --certprofiles UserCert
```

Display the properties of a named CA ACL:

```
ipa caacl-show test
```

Create a CA ACL to let user "alice" use the "DNP3" profile on "DNP3-CA":

```
ipa caacl-add alice_dnp3
ipa caacl-add-ca alice_dnp3 --cas DNP3-CA
ipa caacl-add-profile alice_dnp3 --certprofiles DNP3
ipa caacl-add-user alice_dnp3 --user=alice
```

```
...
```

Note that to perform the **caacl** operations, you must be operating as a user who has the

required permissions. IdM includes the following CA ACL-related permissions by default:

System: Read CA ACLs

Enables the user to read all attributes of the CA ACL.

System: Add CA ACL

Enables the user to add a new CA ACL.

System: Delete CA ACL

Enables the user to delete an existing CA ACL.

System: Modify CA ACL

Enables the user to modify an attribute of the CA ACL and to disable or enable the CA ACL.

System: Manage CA ACL membership

Enables the user to manage the CA, profile, user, host, and service membership in the CA ACL.

All these permissions are included in the default **CA Administrator** privilege. For more information on IdM role-based access controls and managing permissions, see [Section 10.4, “Defining Role-Based Access Controls”](#).

This section describes only the most important aspects of using the **ipa caacl** commands for CA ACL management. For complete information about a command, run it with the **--help** option added, for example:

```
$ ipa caacl-mod --help
Usage: ipa [global-options] caacl-mod NAME [options]

Modify a CA ACL.
Options:
  -h, --help                show this help message and exit
  --desc=STR                 Description
  --cacat=['all']           CA category the ACL applies to
  --profilecat=['all']      Profile category the ACL applies to
  ...
```

Creating CA ACLs

To create a new CA ACL, use the **ipa caacl-add** command. Running the command without any options starts an interactive session in which the **ipa caacl-add** script prompts your for the required information about the new CA ACL.

```
$ ipa caacl-add
ACL name: smime_acl
-----
Added CA ACL "smime_acl"
-----
ACL name: smime_acl
Enabled: TRUE
```

New CA ACLs are enabled by default.

The most notable options accepted by **ipa caacl-add** are the options that associate a CA ACL with a CA, certificate profile, user, host, or service category:

- **--cacat**
- **--profilecat**
- **--usercat**
- **--hostcat**
- **--servicecat**

IdM only accepts the **all** value with these options, which associates the CA ACL with all CAs, profiles, users, hosts, or services. For example, to associate the CA ACL with all users and user groups:

```
$ ipa caacl-add ca_acl_name --usercat=all
```

CA, profile, user, host, and service categories are an alternative to adding particular objects or groups of objects to a CA ACL, which is described in [the section called “Adding Entries to CA ACLs and Removing Entries from CA ACLs”](#). Note that it is not possible to use a category and also add objects or groups of the same type; for example, you cannot use the **--usercat=all** option and then add a user to the CA ACL with the **ipa caacl-add-user --users=*user_name*** command.

NOTE

Requesting a certificate for a user or group using a certificate profile fails if the user or group are not added to the corresponding CA ACL. For example:

```
$ ipa cert-request CSR-FILE --principal user --profile-id profile_id
ipa: ERROR Insufficient access: Principal 'user' is not permitted to use CA '.' with profile 'profile_id' for certificate issuance.
```

You must either add the user or group to the CA ACL, as described in [the section called “Adding Entries to CA ACLs and Removing Entries from CA ACLs”](#), or associate the CA ACL with the **all** user category.

Displaying CA ACLs

To display all CA ACLs, use the **ipa caacl-find** command:

```
$ ipa caacl-find
-----
2 CA ACLs matched
-----
ACL name: hosts_services_caIPAserviceCert
Enabled: TRUE
...
```

Note that **ipa caacl-find** accepts the **--cacat**, **--profilecat**, **--usercat**, **--hostcat**, and **--servicecat** options, which can be used to filter the results of the search to CA ACLs

with the corresponding CA, certificate profile, user, host, or service category. Note that IdM only accepts the **all** category with these options. For more information about the options, see [the section called “Creating CA ACLs”](#).

To display information about a particular CA ACL, use the **ipa caacl-show** command:

```
$ ipa caacl-show ca_acl_name
ACL name: ca_acl_name
Enabled: TRUE
Host category: all
...
```

Modifying CA ACLs

To modify an existing CA ACL, use the **ipa caacl-mod** command. Pass the required modifications using the command-line options accepted by **ipa caacl-mod**. For example, to modify the description of a CA ACL and associate the CA ACL with all certificate profiles:

```
$ ipa caacl-mod ca_acl_name --desc="New description" --profilecat=all
-----
Modified CA ACL "ca_acl_name"
-----
ACL name: smime_acl
Description: New description
Enabled: TRUE
Profile category: all
```

The most notable options accepted by **ipa caacl-mod** are the **--cacat**, **--profilecat**, **--usercat**, **--hostcat**, and **--servicecat** options. For a description of these options, see [the section called “Creating CA ACLs”](#).

Disabling and Enabling CA ACLs

To disable a CA ACL, use the **ipa caacl-disable** command:

```
$ ipa caacl-disable ca_acl_name
-----
Disabled CA ACL "ca_acl_name"
-----
```

A disabled CA ACL is not applied and cannot be used to request a certificate. Disabling a CA ACL does not remove it from IdM.

To enable a disabled CA ACL, use the **ipa caacl-enable** command:

```
$ ipa caacl-enable ca_acl_name
-----
Enabled CA ACL "ca_acl_name"
-----
```

Deleting CA ACLs

To remove an existing CA ACL, use the **ipa caacl-del** command:

```
$ ipa caacl-del ca_acl_name
```

Adding Entries to CA ACLs and Removing Entries from CA ACLs

Using the **ipa caacl-add-*** and **ipa caacl-remove-*** commands, you can add new entries to a CA ACL or remove existing entries.

ipa caacl-add-ca and ipa caacl-remove-ca

Adds or removes a CA.

ipa caacl-add-host and ipa caacl-remove-host

Adds or removes a host or host group.

ipa caacl-add-profile and ipa caacl-remove-profile

Adds or removes a profile.

ipa caacl-add-service and ipa caacl-remove-service

Adds or removes a service.

ipa caacl-add-user and ipa caacl-remove-user

Adds or removes a user or group.

For example:

```
$ ipa caacl-add-user ca_acl_name --groups=group_name
```

Note that it is not possible to add an object or a group of objects to a CA ACL and also use a category of the same object, as described in [the section called “Creating CA ACLs”](#); these settings are mutually exclusive. For example, if you attempt to run the **ipa caacl-add-user --users=*user_name*** command on a CA ACL specified with the **--usercat=all** option, the command fails:

```
$ ipa caacl-add-user ca_acl_name --users=user_name
ipa: ERROR: users cannot be added when user category='all'
```

NOTE

Requesting a certificate for a user or group using a certificate profile fails if the user or group are not added to the corresponding CA ACL. For example:

```
$ ipa cert-request CSR-FILE --principal user --profile-id
profile_id
ipa: ERROR Insufficient access: Principal 'user' is not
permitted to use CA '.' with profile 'profile_id' for
certificate issuance.
```

You must either add the user or group to the CA ACL, or associate the CA ACL with the **all** user category, as described in [the section called “Creating CA ACLs”](#).

For detailed information on the required syntax for these commands and the available options, run the commands with the **--help** option added. For example:

```
$ ipa caacl-add-user --help
```

24.5.2. CA ACL Management from the Web UI

To manage CA ACLs from the IdM web UI:

1. Open the **Authentication** tab and the **Certificates** subtab.
2. Open the **CA ACLs** section.




Figure 24.9. CA ACL Rules Management in the Web UI

In the **CA ACLs** section, you can add new CA ACLs, display information about existing CA ACLs, modify their attributes, as well as enable, disable, or delete selected CA ACLs.

For example, to modify an existing CA ACL:

1. Click on the name of the CA ACL to open the CA ACL configuration page.
2. In the CA ACL configuration page, fill in the required information.

The **Profiles** and **Permitted to have certificates issued** sections allow you to associate the CA ACL with certificate profiles, users or user groups, hosts or host groups, or services. You can either add these objects using the **Add** buttons, or select the **Anyone** option to associate the CA ACL with all users, hosts, or services.

3. Click **Save** to confirm the new configuration.




Figure 24.10. Modifying a CA ACL Rule in the Web UI

24.6. USING CERTIFICATE PROFILES AND ACLS TO ISSUE USER CERTIFICATES WITH THE IDM CAS

Users can request certificates for themselves when permitted by the Certificate Authority access control lists (CA ACLs). The following procedures use certificate profiles and CA ACLs, which are described separately in [Section 24.4, “Certificate Profiles”](#) and [Section 24.5, “Certificate Authority ACL Rules”](#). For more details about using certificate profiles and CA ACLs, see these sections.

Issuing Certificates to Users from the Command Line

1. Create or import a new custom certificate profile for handling requests for user certificates. For example:

```
$ ipa certprofile-import certificate_profile --
file=certificate_profile.cfg --store=True
```

2. Add a new Certificate Authority (CA) ACL that will be used to permit requesting certificates for user entries. For example:

```
$ ipa caacl-add users_certificate_profile --usercat=all
```

3. Add the custom certificate profile to the CA ACL.

```
$ ipa caacl-add-profile users_certificate_profile --
certprofiles=certificate_profile
```

4. Generate a certificate request for the user. For example, using OpenSSL:

```
$ openssl req -new -newkey rsa:2048 -days 365 -nodes -keyout
private.key -out cert.csr -subj '/CN=user'
```

5. Run the **ipa cert-request** command to have the IdM CA issue a new certificate for the user.

```
$ ipa cert-request cert.csr --principal=user --profile-
id=certificate_profile
```

Optionally pass the **--ca *sub-CA_name*** option to the command to request the certificate from a sub-CA instead of the root CA **ipa**.

To make sure the newly-issued certificate is assigned to the user, you can use the **ipa user-show** command:

```
$ ipa user-show user
User login: user
...
Certificate: MIICfzCCAWcCAQA...
...
```

Issuing Certificates to Users in the Web UI

1. Create or import a new custom certificate profile for handling requests for user certificates. Importing profiles is only possible from the command line, for example:


```
$ ipa certprofile-import certificate_profile --
file=certificate_profile.txt --store=True
```

For information about certificate profiles, see [Section 24.4, “Certificate Profiles”](#).

2. In the web UI, under the **Authentication** tab, open the **CA ACLs** section.




Figure 24.11. CA ACL Rules Management in the Web UI

Click **Add** at the top of the list of Certificate Authority (CA) ACLs to add a new CA ACL that permits requesting certificates for user entries.

- a. In the **Add CA ACL** window that opens, fill in the required information about the new CA ACL.

The 'Add CA ACL' dialog box has a title bar with a close button (X). It contains the following fields:

- ACL name ***: A text input field containing 'users_certificate_profile'.
- CA category ⓘ**: A radio button labeled 'All'.
- Description**: A text area containing 'CA ACL for certificate_profile'.

At the bottom left, there is a note: '* Required field'. At the bottom right, there are four buttons: 'Add', 'Add and Add Another', 'Add and Edit', and 'Cancel'.

Figure 24.12. Adding a New CA ACL

Then, click **Add and Edit** to go directly to the CA ACL configuration page.

- b. In the CA ACL configuration page, scroll to the **Profiles** section and click **Add** at the top of the profiles list.




Figure 24.13. Adding a Certificate Profile to the CA ACL

- c. Add the custom certificate profile to the CA ACL by selecting the profile and moving it to the **Prospective** column.




Figure 24.14. Selecting a Certificate Profile

Then, click **Add**.

- d. Scroll to the **Permitted to have certificates issued** section to associate the CA ACL with users or user groups.

You can either add users or groups using the **Add** buttons, or select the **Anyone** option to associate the CA ACL with all users.

Permitted to have certificates issued

User category the ACL applies to: ☒ Anyone ☐ Specified Users and Groups

☐	Users	Delete	+ Add
☐	User Groups	Delete	+ Add

Figure 24.15. Adding Users to the CA ACL

- e. In the **Permitted to have certificates issued** section, you can associate the CA ACL with one or more CAs.

You can either add CAs using the **Add** button, or select the **Any CA** option to associate the CA ACL with all CAs.

CA category the ACL applies to: ☐ Any CA ☒ Specified CAs

If no CAs are specified, requests to the default CA are allowed.

☐	CAs	Delete	+ Add

Figure 24.16. Adding CAs to the CA ACL

- f. At the top of the CA ACL configuration page, click **Save** to confirm the changes to the CA ACL.
3. Request a new certificate for the user.
- Under the **Identity** tab and the **Users** subtab, choose the user for whom the certificate will be requested. Click on the user's user name to open the user entry configuration page.
 - Click **Actions** at the top of the user configuration page, and then click **New Certificate**.




Figure 24.17. Requesting a Certificate for a User

- c. Fill in the required information.

Issue New Certificate for User user
✕

CA ★

Profile ID

- Create a certificate database or use an existing one. To create a new database:

```
# certutil -N -d <database path>
```
- Create a CSR with subject `CN=<uid>,O=<realm>`, for example:

```
# certutil -R -d <database path> -a -g <key size> -s 'CN=user,O=IDM.EXAMPLE.COM'
```
- Copy and paste the CSR (from `-----BEGIN NEW CERTIFICATE REQUEST-----` to `-----END NEW CERTIFICATE REQUEST-----`) into the text area below:

```
-----BEGIN CERTIFICATE REQUEST-----
MIICVTCCAT0CAQAwEDEOMAwGA1UEAwwFdHVzZXIwggEIMA0GCSqGSIb3DQEBAQUA
A4IBDwAwggEKAoIBAQDGH9WgTNE4Zbh8MCpsPx+MWvFyfk9tynyxpLTFg5x2r63
...
-----END CERTIFICATE REQUEST-----
```

Issue Cancel

Figure 24.18. Issuing a Certificate for a User

Then, click **Issue**.

After this, the newly issued certificate is visible in the user configuration page.

CHAPTER 25. STORING AUTHENTICATION SECRETS WITH VAULTS

A vault is a secure location for storing, retrieving, sharing, and recovering secrets. A secret is security-sensitive data that should only be accessible by a limited group of people or entities. For example, secrets include:

- passwords
- PINs
- private SSH keys

Users and services can access the secrets stored in a vault from any machine enrolled in the Identity Management (IdM) domain.



NOTE

Vault is only available from the command line, not from the IdM web UI.

Use cases for vaults include:

Storing personal secrets of a user

See [Section 25.4, “Storing a User's Personal Secret”](#) for details.

Storing a secret for a service

See [Section 25.5, “Storing a Service Secret in a Vault”](#) for details.

Storing a common secret used by multiple users

See [Section 25.6, “Storing a Common Secret for Multiple Users”](#) for details.

Note that to use vaults, you must meet the conditions described in [Section 25.2, “Prerequisites for Using Vaults”](#).

25.1. HOW VAULTS WORK

25.1.1. Vault Owners, Members, and Administrators

IdM distinguishes the following vault user types:

Vault owner

A vault owner is a user or service with basic management privileges on the vault. For example, a vault owner can modify the properties of the vault or add new vault members.


Each vault must have at least one owner. A vault can also have multiple owners.

Vault member

A vault member is a user or service who can access a vault created by another user or service.

Vault administrator

Vault administrators have unrestricted access to all vaults and are allowed to perform all vault operations.



NOTE

Symmetric and asymmetric vaults are protected with a password or key and apply special access control rules (see [Section 25.1.2, “Standard, Symmetric, and Asymmetric Vaults”](#)). The administrator must meet these rules to:

- access secrets in symmetric and asymmetric vaults
- change or reset the vault password or key

A vault administrator is any user with the **Vault Administrators** privilege. See [Section 10.4, “Defining Role-Based Access Controls”](#) for information on defining user privileges.

Certain owner and member privileges depend on the type of the vault. See [Section 25.1.2, “Standard, Symmetric, and Asymmetric Vaults”](#) for details.

Vault User

The output of some commands, such as the **ipa vault-show** command, also displays **Vault user** for user vaults:

```
$ ipa vault-show my_vault
Vault name: my_vault
Type: standard
Owner users: user
Vault user: user
```

The vault user represents the user in whose container the vault is located. For details on vault containers and user vaults, see [Section 25.1.4, “Vault Containers”](#) and [Section 25.1.3, “User, Service, and Shared Vaults”](#).

25.1.2. Standard, Symmetric, and Asymmetric Vaults

The following vault types are based on the level of security and access control:

Standard vault

Vault owners and vault members can archive and retrieve the secrets without having to use a password or key.

Symmetric vault

Secrets in the vault are protected with a symmetric key. Vault members and vault owners can archive and retrieve the secrets, but they must provide the vault password.

Asymmetric vault

Secrets in the vault are protected with an asymmetric key. Users archive the secret using a public key and retrieve it using a private key. Vault members can only archive secrets, while vault owners can both archive and retrieve secrets.

25.1.3. User, Service, and Shared Vaults

The following vault types are based on ownership:

User vault: a private vault for a user

Owner: a single user.

Any user can own one or more user vaults.

Service vault: a private vault for a service

Owner: a single service.

Any service can own one or more service vaults.

Shared vault

Owner: the vault administrator who created the vault. Other vault administrators also have full access to the vault.

Shared vaults can be used by multiple users or services.

25.1.4. Vault Containers

A vault container is a collection of vaults.

IdM provides the following default vault containers:

User container: a private container for a user

This container stores: user vaults for a particular user.

Service container: a private container for a service

This container stores: service vaults for a particular service.

Shared container

This container stores: vaults that can be shared by multiple users or services.

IdM creates user and service containers for each user or service automatically when the first private vault for the user or service is created. After the user or service is deleted, IdM removes the container and its contents.

25.2. PREREQUISITES FOR USING VAULTS

To enable vaults, install the Key Recovery Authority (KRA) Certificate System component on one of the servers in your IdM domain:

```
# ipa-kra-install
```

25.3. GETTING HELP FOR VAULT COMMANDS

To display all commands used to manage vaults and vault containers:

–


```
$ ipa help vault
```

To display detailed help for a particular command, add the **--help** option to the command:

```
$ ipa vault-add --help
```

Vault Commands Fail with `vault not found` Error

Some commands require you to specify the owner or the type of the vault using the following options:

- **--user** or **--service** specify the owner of the vault you want to view

```
$ ipa vault-show user_vault --user user
```

- **--shared** specify that the vault you want to view is a shared vault

For example, if you attempt to view another user's vault without adding **--user**, IdM informs you it did not find the vault:

```
[admin@server ~]$ ipa vault-show user_vault  
ipa: ERROR: user_vault: vault not found
```

25.4. STORING A USER'S PERSONAL SECRET

This section shows how a user can create one or more private vaults to securely store personal secrets. The user then retrieves the secrets when required, on any machine in the domain. For example, the user can archive a personal certificate in a vault, thus storing the certificate securely in a centralized location.

This section includes these procedures:

- [Section 25.4.1, “Archiving a User's Personal Secret”](#)
- [Section 25.4.2, “Retrieving a User's Personal Secret”](#)

In the procedures:

- **user** is the user who wants to create the vault
- **my_vault** is the vault used to store the user's certificate
- the vault type is **standard**, so that accessing the archived certificate does not require the user to provide a vault password
- **secret.txt** is the file containing the certificate that the user wants to store in the vault
- **secret_exported.txt** is the file to which the user exports the archived certificate

25.4.1. Archiving a User's Personal Secret

Create a private user vault and store your certificate in it. The vault type is **standard**, which ensures you will not be required to authenticate when accessing the certificate.

1. Log in as **user**:

```
$ kinit user
```

2. Use the **ipa vault-add** command to create a standard vault:

```
$ ipa vault-add my_vault --type standard
-----
Added vault "my_vault"
-----
Vault name: my_vault
Type: standard
Owner users: user
Vault user: user
```



IMPORTANT

Make sure the first user vault for a user is created by the same user. For example, if another user, such as **admin**, creates the first user vault for **user1**, the owner of the user's vault container will also be **admin**, and **user1** will be unable to access the user vault or create new user vaults. See also [Section B.5.1, “Users Cannot Access Their Vault Due To Insufficient 'add' Privilege”](#).

3. Use the **ipa vault-archive --in** command to archive the **secret.txt** file into the vault:

```
$ ipa vault-archive my_vault --in secret.txt
-----
Archived data into vault "my_vault"
-----
```



NOTE

One vault can only store one secret.

25.4.2. Retrieving a User's Personal Secret

Export the certificate from your private standard vault.

1. Log in as **user**:

```
$ kinit user
```

2. Use the **ipa vault-retrieve --out** command to retrieve the contents of the vault and save them into the **secret_exported.txt** file.

```
$ ipa vault-retrieve my_vault --out secret_exported.txt
-----
Retrieved data from vault "my_vault"
-----
```

25.5. STORING A SERVICE SECRET IN A VAULT

This section shows how an administrator can use vaults to securely store a service secret in a centralized location. The service secret is encrypted with the service public key. The service then retrieves the secret using its private key on any machine in the domain. Only the service and the administrator are allowed to access the secret.

This section includes these procedures:

- [Section 25.5.1, “Creating a User Vault to Store a Service Password”](#)
- [Section 25.5.2, “Provisioning a Service Password from a User Vault to Service Instances”](#)
- [Section 25.5.3, “Retrieving a Service Password for a Service Instance”](#)
- [Section 25.5.4, “Changing Service Vault Password”](#)

In the procedures:

- **admin** is the administrator who manages the service password
- **http_password** is the name of the private user vault created by the administrator
- **password.txt** is the file containing the service password
- **password_vault** is the vault created for the service
- **HTTP/server.example.com** is the service whose password is being archived
- **service-public.pem** is the service public key used to encrypt the password stored in **password_vault**

25.5.1. Creating a User Vault to Store a Service Password

Create an administrator-owned user vault, and use it to store the service password. The vault type is standard, which ensures the administrator is not required to authenticate when accessing the contents of the vault.

1. Log in as the administrator:

```
$ kinit admin
```

2. Create a standard user vault:

```
$ ipa vault-add http_password --type standard
-----
Added vault "http_password"
-----
Vault name: http_password
Type: standard
Owner users: admin
Vault user: admin
```

3. Archive the service password into the vault:

—

```
$ ipa vault-archive http_password --in password.txt
-----
Archived data into vault "http_password"
-----
```

**WARNING**

After archiving the password into the vault, delete **password.txt** from your system.

25.5.2. Provisioning a Service Password from a User Vault to Service Instances

Using an asymmetric vault created for the service, provision the service password to a service instance.

1. Log in as the administrator:

```
$ kinit admin
```

2. Obtain the public key of the service instance. For example, using the **openssl** utility:
 - a. Generate the **service-private.pem** private key.

```
$ openssl genrsa -out service-private.pem 2048
Generating RSA private key, 2048 bit long modulus
.+++
.....+++
e is 65537 (0x10001)
```

- b. Generate the **service-public.pem** public key based on the private key.

```
$ openssl rsa -in service-private.pem -out service-public.pem -
pubout
writing RSA key
```

3. Create an asymmetric vault as the service instance vault, and provide the public key:

```
$ ipa vault-add password_vault --service HTTP/server.example.com --
type asymmetric --public-key-file service-public.pem
-----
Added vault "password_vault"
-----
Vault name: password_vault
Type: asymmetric
```

```
Public key: LS0tLS1C...S0tLS0tCg==
Owner users: admin
Vault service: HTTP/server.example.com@EXAMPLE.COM
```

The password archived into the vault will be protected with the key.

4. Retrieve the service password from the administrator's private vault, and then archive it into the new service vault:

```
$ ipa vault-retrieve http_password --out password.txt
-----
Retrieved data from vault "http_password"
-----
```

```
$ ipa vault-archive password_vault --service HTTP/server.example.com
--in password.txt
-----
Archived data into vault "password_vault"
-----
```

This encrypts the password with the service instance public key.



WARNING

After archiving the password into the vault, delete **password.txt** from your system.

Repeat these steps for every service instance that requires the password. Create a new asymmetric vault for each service instance.

25.5.3. Retrieving a Service Password for a Service Instance

A service instance can retrieve the service vault password using the locally-stored service private key.

1. Log in as the administrator:

```
$ kinit admin
```

2. Obtain a Kerberos ticket for the service:

```
# kinit HTTP/server.example.com -k -t /etc/httpd/conf/ipa.keytab
```

3. Retrieve the service vault password:

```
$ ipa vault-retrieve password_vault --service
HTTP/server.example.com --private-key-file service-private.pem --out
password.txt
```

```
-----
Retrieved data from vault "password_vault"
-----
```

25.5.4. Changing Service Vault Password

If a service instance is compromised, isolate it by changing the service vault password and then re-provisioning the new password to non-compromised service instances only.

1. Archive the new password in the administrator's user vault:

```
$ ipa vault-archive http_password --in new_password.txt
-----
Archived data into vault "http_password"
-----
```

This overwrites the current password stored in the vault.

2. Re-provision the new password to each service instance excluding the compromised instance.

- a. Retrieve the new password from the administrator's vault:

```
$ ipa vault-retrieve http_password --out password.txt
-----
Retrieved data from vault "http_password"
-----
```

- b. Archive the new password into the service instance vault:

```
$ ipa vault-archive password_vault --service
HTTP/server.example.com --in password.txt
-----
Archived data into vault "password_vault"
-----
```



WARNING

After archiving the password into the vault, delete **password.txt** from your system.

25.6. STORING A COMMON SECRET FOR MULTIPLE USERS

This section shows how an administrator can create a shared vault and allow other users to access the secret in the vault. The administrator archives a common password into the vault, and the other users are able to retrieve the password on any machine in the domain.

This section includes these procedures:

- [Section 25.6.2, “Retrieving a Secret from a Shared Vault as a Member User”](#)
- [Section 25.6.1, “Creating the Shared Vault with the Common Secret”](#)

In the procedures:

- **shared_vault** is the vault used to store the common password
- **admin** is the administrator who creates the shared vault
- the vault type is **standard**, so that accessing the archived password does not require the user to provide a vault password
- **secret.txt** is the file containing the common secret
- **user1** and **user2** are the users allowed to access the vault

25.6.1. Creating the Shared Vault with the Common Secret

Create a shared vault and use it to store the common secret. Add the users who will be accessing the secret as vault members. The vault type is standard, which ensures any user accessing the secret will not be required to authenticate.

1. Log in as the administrator:

```
$ kinit admin
```

2. Create the shared vault:

```
$ ipa vault-add shared_vault --shared --type standard
-----
Added vault "shared_vault"
-----
Vault name: shared_vault
Type: standard
Owner users: admin
Shared vault: True
```

3. Archive the secret into the vault. Add the **--shared** option to specify that the vault is in the shared container:

```
$ ipa vault-archive shared_vault --shared --in secret.txt
-----
Archived data into vault "shared_vault"
-----
```



NOTE

One vault can only store one secret.

4. Add **user1** and **user2** as vault members:

```
ipa vault-add-member shared_vault --shared --users={user1,user2}
```

```
Vault name: shared_vault
Type: standard
Owner users: admin
Shared vault: True
Member users: user1, user2
-----
Number of members added 2
-----
```

25.6.2. Retrieving a Secret from a Shared Vault as a Member User

Log in as a member user of the vault, and export the file with the secret from the vault.

1. Log in as the **user1** member user:

```
$ kinit user1
```

2. Retrieve the secret from the shared vault:

```
$ ipa vault-retrieve shared_vault --shared --out secret_exported.txt
-----
Retrieved data from vault "shared_vault"
-----
```


CHAPTER 26. MANAGING CERTIFICATES AND CERTIFICATE AUTHORITIES

26.1. LIGHTWEIGHT SUB-CAS

If your IdM installation is configured with the integrated Certificate System (CS) certificate authority (CA), you are able to create lightweight sub-CAs. They enable you to configure services, like virtual private network (VPN) gateways, to accept only certificates issued by one sub-CA. At the same time, you can configure other services to accept only certificates issued by a different sub-CA or the root CA.

If you revoke the intermediate certificate of a sub-CA, all certificates issued by this sub-CA are automatically invalid.

If you set up IdM using the integrated CA, the automatically created **ipa** CA is the root CA of the certificate system. All sub-CAs you create, are subordinated to this root CA.

26.1.1. Creating a Lightweight Sub-CA


For details on creating a sub-CA, see

- [the section called “Creating a Sub-CA from the Web UI”](#)
- [the section called “Creating a Sub-CA from the Command Line”](#)

Creating a Sub-CA from the Web UI

To create a new sub-CA named *vpn-ca*:

1. Open the **Authentication** tab, and select the **Certificates** subtab.
2. Select **Certificate Authorities** and click **Add**.
3. Enter the name and subject DN for the CA.



Add Certificate Authority [X]

Name *

Subject DN *

Description

* Required field

[Add] [Add and Add Another] [Add and Edit] [Cancel]

Figure 26.1. Adding a CA

The subject DN must be unique in the IdM CA infrastructure.

Creating a Sub-CA from the Command Line

To create a new sub-CA named *vpn-ca*, enter:

```
[root@ipaserver ~]# ipa ca-add vpn-ca --subject="CN=VPN,O=IDM.EXAMPLE.COM"
-----
Created CA "vpn-ca"
-----
Name: vpn-ca
Authority ID: ba83f324-5e50-4114-b109-acca05d6f1dc
Subject DN: CN=VPN,O=IDM.EXAMPLE.COM
Issuer DN: CN=Certificate Authority,O=IDM.EXAMPLE.COM
```

Name

Name of the CA.

Authority ID

Automatically created, individual ID for the CA.

Subject DN

Subject distinguished name (DN). The subject DN must be unique in the IdM CA infrastructure.

Issuer DN

Parent CA that issued the sub-CA certificate. All sub-CAs are created as a child of the IdM root CA.


To verify that the new CA signing certificate has been successfully added to the IdM database, run:

```
[root@ipaserver ~]# certutil -d /etc/pki/pki-tomcat/alias/ -L

Certificate Nickname           Trust
Attributes

SSL,S/MIME,JAR/XPI

caSigningCert cert-pki-ca      CTu,Cu,Cu
Server-Cert cert-pki-ca       u,u,u
auditSigningCert cert-pki-ca   u,u,Pu
caSigningCert cert-pki-ca ba83f324-5e50-4114-b109-acca05d6f1dc u,u,u
ocspSigningCert cert-pki-ca    u,u,u
subsystemCert cert-pki-ca     u,u,u
```

**NOTE**

The new CA certificate is automatically transferred to all replicas when they have a certificate system instance installed.

26.1.2. Removing a Lightweight Sub-CA

For details on deleting a sub-CA, see

- [the section called “Removing a Sub-CA from the Web UI”](#)

- [the section called “Removing a Sub-CA from the Command Line”](#)

Removing a Sub-CA from the Web UI

1. Open the **Authentication** tab, and select the **Certificates** subtab.
2. Select **Certificate Authorities**.
3. Select the sub-CA to remove and click **Delete**.
4. Click **Delete** to confirm.

Removing a Sub-CA from the Command Line

To delete a sub-CA, enter:

```
[root@ipaserver ~]# ipa ca-del vpn-ca
-----
Deleted CA "vpn-ca"
-----
```

26.2. RENEWING CERTIFICATES

For details on:

- automatic certificate renewal, see [Section 26.2.1, “Renewing Certificates Automatically”](#)
- manual certificate renewal, see [Section 26.2.2, “Renewing CA Certificates Manually”](#)

26.2.1. Renewing Certificates Automatically

The **certmonger** service automatically renews the following certificates 28 days before their expiration date:

- CA certificate issued by the IdM CA as the root CA
- Subsystem and server certificates issued by the integrated IdM CA that are used by internal IdM services

To automatically renew sub-CA CA certificates, they must be listed on the **certmonger** tracking list. To update the tracking list:

```
[root@ipaserver ~]# ipa-certupdate
trying https://idmserver.idm.example.com/ipa/json
Forwarding 'schema' to json server
'https://idmserver.idm.example.com/ipa/json'
trying https://idmserver.idm.example.com/ipa/json
Forwarding 'ca_is_enabled' to json server
'https://idmserver.idm.example.com/ipa/json'
Forwarding 'ca_find/1' to json server
'https://idmserver.idm.example.com/ipa/json'
Systemwide CA database updated.
Systemwide CA database updated.
The ipa-certupdate command was successful
```

**NOTE**

If you are using an external CA as the root CA, you must renew the certificates manually, as described in [Section 26.2.2, “Renewing CA Certificates Manually”](#). The **certmonger** service cannot automatically renew certificates signed by an external CA.

For more information on how **certmonger** monitors certificate expiration dates, see [Tracking Certificates with certmonger](#) in the *System-Level Authentication Guide*.

To verify that automatic renewal works as expected, examine **certmonger** log messages in the `/var/log/messages` file:

- After a certificate is renewed, **certmonger** records message like the following to indicate that the renewal operation has succeeded or failed:

```
Certificate named "NSS Certificate DB" in token "auditSigningCert
cert-pki-ca" in database "/var/lib/pki-ca/alias" renew success
```

- As the certificate nears its expiration, **certmonger** logs the following message:

```
certmonger: Certificate named "NSS Certificate DB" in token
"auditSigningCert cert-pki-ca" in database "/var/lib/pki-ca/alias"
will not be valid after 20160204065136.
```

26.2.2. Renewing CA Certificates Manually

You can use the **ipa-cacert-manage** utility to manually renew:

- self-signed IdM CA certificate
- externally-signed IdM CA certificate

The certificates renewed with the **ipa-cacert-manage renew** command use the same key pair and subject name as the old certificates. Renewing a certificate does not remove its previous version to enable certificate rollover.

For details, see the `ipa-cacert-manage(1)` man page.

26.2.2.1. Renewing a Self-Signed IdM CA Certificate Manually

1. Run the **ipa-cacert-manage renew** command. The command does not require you to specify the path to the certificate.
2. The renewed certificate is now present in the LDAP certificate store and in the `/etc/pki/pki-tomcat/alias` NSS database.
3. Run the **ipa-certupdate** utility on all servers and clients to update them with the information about the new certificate from LDAP. You must run **ipa-certupdate** on every server and client separately.

**IMPORTANT**

Always run **ipa-certupdate** after manually installing a certificate. If you do not, the certificate will not be distributed to the other machines.

To make sure the renewed certificate is properly installed, use the **certutil** utility to list the certificates in the database. For example:


```
# certutil -L -d /etc/pki/pki-tomcat/alias
```

26.2.2.2. Renewing an Externally-Signed IdM CA Certificate Manually

1. Run the **ipa-cacert-manage renew --external-ca** command.
2. The command creates the **/var/lib/ipa/ca.crt** CSR file. Submit the CSR to the external CA to get the renewed CA certificate issued.
3. Run **ipa-cacert-manage renew** again, and this time specify the renewed CA certificate and the external CA certificate chain files using the **--external-cert-file** option. For example:

```
# ipa-cacert-manage renew --external-cert-  
file=/tmp/servercert20110601.pem --external-cert-  
file=/tmp/cacert.pem
```

4. The renewed CA certificate and the external CA certificate chain are now present in the LDAP certificate store and in the **/etc/pki/pki-tomcat/alias/** NSS database.
5. Run the **ipa-certupdate** utility on all servers and clients to update them with the information about the new certificate from LDAP. You must run **ipa-certupdate** on every server and client separately.

**IMPORTANT**

Always run **ipa-certupdate** after manually installing a certificate. If you do not, the certificate will not be distributed to the other machines.

To make sure the renewed certificate is properly installed, use the **certutil** utility to list the certificates in the database. For example:

```
# certutil -L -d /etc/pki/pki-tomcat/alias/
```

26.3. INSTALLING A CA CERTIFICATE MANUALLY


To install a new certificate to IdM, use the **ipa-cacert-manage install** command. For example, the command allows you to change the current certificate when it is nearing its expiration date.

1. Run the **ipa-cacert-manage install** command, and specify the path to the file containing the certificate. The command accepts PEM-formatted certificate files:

```
[root@server ~]# ipa-cacert-manage install /etc/group/cert.pem
```

The certificate is now present in the LDAP certificate store.

2. Run the **ipa-certupdate** utility on all servers and clients to update them with the information about the new certificate from LDAP. You must run **ipa-certupdate** on every server and client separately.



IMPORTANT

Always run **ipa-certupdate** after manually installing a certificate. If you do not, the certificate will not be distributed to the other machines.

The **ipa-cacert-manage install** command can take the following options:

-n

gives the nickname of the certificate; the default value is the subject name of the certificate

-t

specifies the trust flags for the certificate in the **certutil** format; the default value is **C,,**. For information about the format in which to specify the trust flags, see the **ipa-cacert-manage(1)** man page.

26.4. CHANGING THE CERTIFICATE CHAIN

You can modify the certificate chain by renewing the CA certificate using the **ipa-cacert-manage renew**.

Self-signed CA certificate → externally-signed CA certificate

Add the **--external-ca** option to **ipa-cacert-manage renew**. This renews the self-signed CA certificate as an externally-signed CA certificate.

For details on running the command with this option, see [Section 26.2.2, “Renewing CA Certificates Manually”](#).

Externally-signed CA certificate → self-signed CA certificate

Add the **--self-signed** option to **ipa-cacert-manage renew**. This renews the externally-signed CA certificate as a self-signed CA certificate.

26.5. ALLOWING IDM TO START WITH EXPIRED CERTIFICATES

After the IdM administrative server certificates expire, most IdM services become inaccessible. You can configure the underlying Apache and LDAP services to allow SSL access to the services even if the certificates are expired.

If you allow limited access with expired certificates:

- Apache, Kerberos, DNS, and LDAP services will continue working. With these services active, users will be able to log in to the IdM domain.

- Client services that require SSL for access will still fail. For example, **sudo** will fail because it requires SSSD on IdM clients, and SSSD needs SSL to contact IdM.



IMPORTANT

This procedure is intended only as a temporary workaround. Renew the required certificates as quickly as possible, and then revert the described changes.

1. Configure the **mod_nss** module for the Apache server to not enforce valid certificates.

- a. Open the **/etc/httpd/conf.d/nss.conf** file.
- b. Set the **NSSEnforceValidCerts** parameter to **off**:

```
NSSEnforceValidCerts off
```

2. Restart Apache.

```
# systemctl restart httpd.service
```

3. Make sure that validity checks are disabled for the LDAP directory server. To do this, verify that the **nsslapd-validate-cert** attribute is set to **warn**:

```
# ldapsearch -h server.example.com -p 389 -D "cn=directory manager"
-w secret -LLL -b cn=config -s base "(objectclass=*)" nsslapd-
validate-cert

dn: cn=config
nsslapd-validate-cert: warn
```

If the attribute is not set to **warn**, change it:

```
# ldapmodify -D "cn=directory manager" -w secret -p 389 -h
server.example.com

dn: cn=config
changetype: modify
replace: nsslapd-validate-cert
nsslapd-validate-cert: warn
```

4. Restart the directory server.

```
# systemctl restart dirsrv.target
```

26.6. INSTALLING THIRD-PARTY CERTIFICATES FOR HTTP OR LDAP

Installing a new SSL server certificate for the Apache Web Server, the Directory Server, or both replaces the current SSL certificate with a new one. To do this, you need:

- your private SSL key (**ssl.key** in the procedure below)
- your SSL certificate (**ssl.crt** in the procedure below)

For a list of accepted formats of the key and certificate, see the `ipa-server-certinstall(1)` man page.

Prerequisites

The **ssl.crt** certificate must be signed by a CA known by the service you are loading the certificate into. If this is not the case, install the CA certificate of the CA that signed **ssl.crt** into IdM, as described in [Section 26.3, “Installing a CA Certificate Manually”](#).

This ensures that IdM recognizes the CA, and thus accepts **ssl.crt**.

Installing the Third-Party Certificate

1. Use the **ipa-server-certinstall** utility to install the certificate. Specify where you want to install it:
 - **--http** installs the certificate in the Apache Web Server
 - **--dirsrv** installs the certificate on the Directory Server

For example, to install the SSL certificate into both:

```
# ipa-server-certinstall --http --dirsrv ssl.key ssl.crt
```

2. Restart the server into which you installed the certificate.

- To restart the Apache Web Server:

```
# systemctl restart httpd.service
```

- To restart the Directory Server:

```
# systemctl restart dirsrv@REALM.service
```

3. To verify that the certificate has been correctly installed, make sure it is present in the certificate database.

- To display the Apache certificate database:

```
# certutil -L -d /etc/httpd/alias
```

- To display the Directory Server certificate database:

```
# certutil -L -d /etc/dirsrv/slapd-REALM/
```

26.7. CONFIGURING OCSP RESPONDERS

Every CA integrated with the IdM server uses an internal online certificate status protocol (OCSP) responder. The IdM service allowing to access the OCSP responders is available at **`http://ca-server.example.com/ca/ocsp`**. Clients can connect to this URL to check the validity of a certificate.

**NOTE**

For details on OCSP, see the Red Hat Certificate System documentation. For example, [2.2.4. Revoking Certificates and Checking Status](#) in the *Planning, Installation, and Deployment Guide*.

26.7.1. Changing the CRL Update Interval

The CRL file is automatically generated by the IdM CA every four hours by default. To change this interval:

1. Stop the CA server.

```
# systemctl stop pki-tomcatd@pki-tomcat.service
```

2. Open the `/var/lib/pki/pki-tomcat/conf/ca/CS.cfg` file, and change the `ca.crl.MasterCRL.autoUpdateInterval` value to the new interval setting. For example, to generate the CRL every 60 minutes:

```
ca.crl.MasterCRL.autoUpdateInterval=60
```

3. Start the CA server.

```
# systemctl start pki-tomcatd@pki-tomcat.service
```

26.8. INSTALLING A CA INTO AN EXISTING IDM DOMAIN

If an IdM domain was installed without a Certificate Authority (CA), you can install the CA services subsequently. Depending on your environment, you can install the IdM Certificate Server CA or use an external CA.

**NOTE**

For details on the supported CA configurations, see [Section 2.3.2, “Determining What CA Configuration to Use”](#).

IdM Certificate Server

1. Use the following command to install the IdM Certificate Server CA:

```
[root@ipa-server ~] ipa-ca-install
```

2. Run the **ipa-certupdate** utility on all servers and clients to update them with the information about the new certificate from LDAP. You must run **ipa-certupdate** on every server and client separately.

**IMPORTANT**

Always run **ipa-certupdate** after manually installing a certificate. If you do not, the certificate will not be distributed to the other machines.

External CA

The subsequent installation of an external CA consists of multiple steps:

1. Start the installation:

```
[root@ipa-server ~] ipa-ca-install --external-ca
```

After this step an information is shown that a certificate signing request (CSR) was saved. Submit the CSR to the external CA and copy the issued certificate to the IdM server.

2. Continue the installation with passing the certificates and full path to the external CA files to **ipa-ca-install**:

```
[root@ipa-server ~]# ipa-ca-install --external-cert-file=/root/master.crt --external-cert-file=/root/ca.crt
```

3. Run the **ipa-certupdate** utility on all servers and clients to update them with the information about the new certificate from LDAP. You must run **ipa-certupdate** on every server and client separately.



IMPORTANT

Always run **ipa-certupdate** after manually installing a certificate. If you do not, the certificate will not be distributed to the other machines.

The CA installation does not replace the existing service certificates for the LDAP and web server with ones issued by the new installed CA. For details how to replace the certificates, see [Section 26.9, “Replacing the Web Server's and LDAP Server's Certificate”](#).

26.9. REPLACING THE WEB SERVER'S AND LDAP SERVER'S CERTIFICATE

To replace the service certificates for the web server and LDAP server:

1. Request a new certificate. You can do this using:
 - the integrated CA: see [Section 24.1.1, “Requesting New Certificates for a User, Host, or Service”](#) for details.
 - an external CA: generate a private key and certificate signing request (CSR). For example, using OpenSSL:

```
$ openssl req -new -newkey rsa:2048 -days 365 -nodes -keyout new.key -out new.csr -subj '/CN=idmserver.idm.example.com,O=IDM.EXAMPLE.COM'
```

Submit the CSR to the external CA. The process differs depending on the service to be used as the external CA.

2. Replace the Apache web server's private key and certificate:

```
[root@ipaserver ~]# ipa-server-certinstall -w --pin=password new.key  
new.crt
```

3. Replace the LDAP server's private key and certificate:

```
[root@ipaserver ~]# ipa-server-certinstall -d --pin=password new.key  
new.cert
```

CHAPTER 27. KERBEROS PKINIT AUTHENTICATION IN IDM

Public Key Cryptography for Initial Authentication in Kerberos (PKINIT) is a preauthentication mechanism for Kerberos. As of Red Hat Enterprise Linux 7.4, the Identity Management (IdM) server includes a mechanism for Kerberos PKINIT authentication. The following sections give an overview of the PKINIT implementation in IdM and describe how to configure PKINIT in IdM.

27.1. DEFAULT PKINIT STATUS IN DIFFERENT IDM VERSIONS

The default PKINIT configuration on your IdM servers depends on the version of IdM in Red Hat Enterprise Linux (RHEL) and the certificate authority (CA) configuration. See [Table 27.1, “Default PKINIT configuration in IdM versions”](#).

Table 27.1. Default PKINIT configuration in IdM versions

RHEL version	CA configuration	PKINIT configuration
7.3 and earlier	Without a CA	Local PKINIT: IdM only uses PKINIT for internal purposes on servers.
7.3 and earlier	With an integrated CA	IdM attempts to configure PKINIT by using the certificate signed by the integrated IdM CA. If the attempt fails, IdM configures local PKINIT only.
7.4 and later	Without a CA No external PKINIT certificate provided to IdM	Local PKINIT: IdM only uses PKINIT for internal purposes on servers.
7.4 and later	Without a CA External PKINIT certificate provided to IdM	IdM configures PKINIT by using the external Kerberos key distribution center (KDC) certificate and CA certificate.
7.4 and later	With an integrated CA	IdM configures PKINIT by using the certificate signed by the IdM CA.

At domain level 0, PKINIT is disabled. The default behavior is local PKINIT: IdM only uses PKINIT for internal purposes on servers. See also [Chapter 7, Displaying and Raising the Domain Level](#).

27.2. DISPLAYING THE CURRENT PKINIT CONFIGURATION

IdM provides multiple commands you can use to query the PKINIT configuration in your domain.

To determine the PKINIT status in your domain, use the **ipa pkinit-status** command:

```
$ ipa pkinit-status
Server name: server1.example.com
PKINIT status: enabled
[...output truncated...]
Server name: server2.example.com
PKINIT status: disabled
[...output truncated...]
```

To determine the PKINIT status on the server where you are logged in, use the **ipa-pkinit-manage status** command:

```
# ipa-pkinit-manage status
PKINIT is enabled
The ipa-pkinit-manage command was successful
```

The commands display the PKINIT configuration status as **enabled** or **disabled**:

- **enabled**: PKINIT is configured using a certificate signed by the integrated IdM CA or an external PKINIT certificate. See also [Section 27.1, “Default PKINIT Status in Different IdM Versions”](#).
- **disabled**: IdM only uses PKINIT for internal purposes on IdM servers.

To display the IdM servers with active Kerberos key distribution centers (KDCs) that support PKINIT for IdM clients, use the **ipa config-show** command on any server:

```
$ ipa config-show
Maximum username length: 32
Home directory base: /home
Default shell: /bin/sh
Default users group: ipausers
[...output truncated...]
IPA masters capable of PKINIT: server1.example.com
[...output truncated...]
```

Additional Resources

- For more details on the command-line tools for reporting the PKINIT status, use the **ipa help pkinit** command.

27.3. CONFIGURING PKINIT IN IDM

If your IdM servers are running with PKINIT disabled, use these steps to enable it. For example, a server is running with PKINIT disabled if you passed the **--no-pkinit** option with the **ipa-server-install** or **ipa-replica-install** utilities.

Prerequisites

- Ensure that all IdM servers with a certificate authority (CA) installed are running on the same domain level. See [Chapter 7, *Displaying and Raising the Domain Level*](#) for details.

Procedure

1. If you are using IdM without a CA, use the **ipa-server-certinstall** utility to install an external Kerberos key distribution center (KDC) certificate. The KDC certificate must meet the following conditions:
 - It is issued with the common name **CN=fully_qualified_domain_name,certificate_subject_base**.
 - It includes the Kerberos principal **krbtgt/REALM_NAME@REALM_NAME**.
 - It contains the object identifier (OID) for KDC authentication: **1.3.6.1.5.2.3.5**.

```
# ipa-server-certinstall --kdc kdc.pem
# systemctl restart krb5kdc.service
```

For details, see the `ipa-server-certinstall(1)` man page.

2. Enable PKINIT:

```
$ ipa-pkinit-manage enable
Configuring Kerberos KDC (krb5kdc)
[1/1]: installing X509 Certificate for PKINIT
Done configuring Kerberos KDC (krb5kdc).
The ipa-pkinit-manage command was successful
```

If you are using an IdM CA, the command requests a PKINIT KDC certificate from the CA.

3. To verify the new PKINIT status, see [Section 27.2, “Displaying the Current PKINIT Configuration”](#).

27.4. ADDITIONAL RESOURCES

- For details on Kerberos PKINIT, [PKINIT configuration](#) in the MIT Kerberos Documentation.
- For documentation on configuring PKINIT smart-card authentication in IdM, see [Section 23.5, “PKINIT Smart-card Authentication in Identity Management”](#).

PART VI. ADMINISTRATION: MANAGING POLICIES

CHAPTER 28. DEFINING PASSWORD POLICIES

This chapter describes what password policies in Identity Management (IdM) are and how to manage them.

28.1. WHAT ARE PASSWORD POLICIES AND WHY ARE THEY USEFUL

A *password policy* is a set of rules that passwords must meet.

For example, a password policy can define minimum password length and maximum password lifetime. All users affected by such a policy are required to set a sufficiently long password and change it frequently enough.

Password policies help reduce the risk of someone discovering and misusing a user's password.

28.2. HOW PASSWORD POLICIES WORK IN IDM

All users must have a password that they use to authenticate to the Identity Management (IdM) Kerberos domain. Password policies in IdM define the requirements these user passwords must meet.



NOTE

The IdM password policy is set in the underlying LDAP directory, but is also enforced by the Kerberos Key Distribution Center (KDC).

28.2.1. Supported Password Policy Attributes


Table 28.1, “Password Policy Attributes” lists the attributes that password policies in IdM can define.

Table 28.1. Password Policy Attributes

Attribute	Explanation	Example
Max lifetime	The maximum amount of time (in days) that a password is valid before a user must reset it.	Max lifetime = 90 User passwords are valid only for 90 days. After that, IdM prompts users to change them.
Min lifetime	The minimum amount of time (in hours) that must pass between two password change operations.	Min lifetime = 1 After users change their passwords, they must wait at least 1 hour before changing them again.

Attribute	Explanation	Example
History size	How many previous passwords are stored. A user cannot reuse a password from their password history.	History size = 0 Users can reuse any of their previous passwords.
Character classes	The number of different character classes the user must use in the password. The character classes are: • Uppercase characters • Lowercase characters • Digits • Special characters, such as comma (,), period (.), asterisk (*) • Other UTF-8 characters Using a character three or more times in a row decreases the character class by one. For example: • Secret1 has 3 character classes: uppercase, lowercase, digits • Secret111 has 2 character classes: uppercase, lowercase, digits, and a -1 penalty for using 1 repeatedly	Character classes = 0 The default number of classes required is 0. To configure the number, run the ipa pwpolicy-mod command with the --minclasses option. This command sets the required number of character classes to 1: <pre>\$ ipa pwpolicy-mod --minclasses=1</pre> See also the Important note below this table.
Min length	The minimum number of characters in a password.	Min length = 8 Users cannot use passwords shorter than 8 characters.
Max failures	The maximum number of failed login attempts before IdM locks the user account. See also Section 22.1.3, “Unlocking User Accounts After Password Failures” .	Max failures = 6 IdM locks the user account the user enters a wrong password 7 times in a row.
Failure reset interval	The amount of time (in seconds) after which IdM resets the current number of failed login attempts.	Failure reset interval = 60 If the user waits for more than 1 minute after the number of failed login attempts defined in Max failures , the user can attempt to log in again without risking a user account lock.

Attribute	Explanation	Example
Lockout duration	The amount of time (in seconds) for which the user account is locked after the number of failed login attempts defined in Max failures . See also Section 22.1.3, “Unlocking User Accounts After Password Failures” .	Lockout duration = 600 Users with locked accounts are unable to log in for 10 minutes.



IMPORTANT

Use the English alphabet and common symbols for the character classes requirement if you have a diverse set of hardware that may not have access to international characters and symbols. For more information on character class policies in passwords, see [What characters are valid in a password?](#) in Red Hat Knowledgebase.

28.2.2. Global and Group-specific Password Policies

The default password policy is the *global password policy*. Apart from the global policy, you can create additional *group password policies*.

Global password policy

Installing the initial IdM server automatically creates a global password policy with default settings.

The global policy rules apply to all users without a group password policy.

Group password policies

Group password policies apply to all members of the corresponding user group.

Only one password policy can be in effect at a time for any user. If a user has multiple password policies assigned, one of them takes precedence based on priority. See [Section 28.2.3, “Password Policy Priorities”](#).

28.2.3. Password Policy Priorities

Every group password policy has a *priority* set. The lower the value, the higher the policy's priority. The lowest supported priority value is **0**.

- If multiple password policies are applicable to a user, the policy with the lowest priority value takes precedence. All rules defined in other policies are ignored.
- The password policy with the lowest priority value applies to all password policy attributes, even the attributes that are not defined in the policy.

The global password policy does not have a priority value set. It serves as a fallback policy when no group policy is set for a user. The global policy can never take precedence over a group policy.

[Table 28.2, “Example of Applying Password Policy Attributes Based on Priority”](#) demonstrates how password policy priorities work on an example of a user who belongs to two groups with a policy defined.

Table 28.2. Example of Applying Password Policy Attributes Based on Priority

	Max lifetime	Min length
Policy for group A (priority 0)	60	10
Policy for group B (priority 1)	90	0 (no restriction)
	↓	↓
User (member of group A and group B)	60	10

**NOTE**

The `ipa pwpolicy-show --user=user_name` command shows which policy is currently in effect for a particular user.

28.3. ADDING A NEW PASSWORD POLICY

When adding a new password policy, you must specify:

- a user group to which the policy will apply (see [Section 28.2.2, “Global and Group-specific Password Policies”](#))
- a priority (see [Section 28.2.3, “Password Policy Priorities”](#))

To add a new password policy using:

- the web UI, see [the section called “Web UI: Adding a New Password Policy”](#)
- the command line, see [the section called “Command Line: Adding a New Password Policy”](#)

Web UI: Adding a New Password Policy

1. Select **Policy** → **Password Policies**.
2. Click **Add**.
3. Define the user group and priority.
4. Click **Add** to confirm.

To configure the attributes of the new password policy, see [Section 28.4, “Modifying Password Policy Attributes”](#).

Command Line: Adding a New Password Policy

1. Use the `ipa pwpolicy-add` command. Specify the user group and priority:

```
$ ipa pwpolicy-add
Group: group_name
Priority: priority_level
```

2. *Optional.* Use the **ipa pwpolicy-find** command to verify that the policy has been successfully added:

```
$ ipa pwpolicy-find
```

To configure the attributes of the new password policy, see [Section 28.4, “Modifying Password Policy Attributes”](#).

28.4. MODIFYING PASSWORD POLICY ATTRIBUTES



IMPORTANT

When you modify a password policy, the new rules apply to new passwords only. The changes are not applied retroactively to existing passwords.

For the change to take effect, users must change their existing passwords, or the administrator must reset the passwords of other users. See [Section 22.1.1, “Changing and Resetting User Passwords”](#).



NOTE

For recommendations on secure user passwords, see [Password Security](#) in the *Security Guide*.

To modify a password policy using:

- the web UI, see [the section called “Web UI: Modifying a Password Policy”](#)
- the command line, see [the section called “Command Line: Modifying a Password Policy”](#)

Note that setting a password policy attribute to **0** means no attribute restriction. For example, if you set maximum lifetime to **0**, user passwords never expire.

Web UI: Modifying a Password Policy

1. Select **Policy** → **Password Policies**.
2. Click the policy you want to change.
3. Update the required attributes. For details on the available attributes, see [Section 28.2.1, “Supported Password Policy Attributes”](#).
4. Click **Save** to confirm the changes.

Command Line: Modifying a Password Policy

1. Use the **ipa pwpolicy-mod** command to change the policy's attributes.
 - a. For example, to update the global password policy and set the minimum password length to **10**:

```
$ ipa pwpolicy-mod --minlength=10
```

- b. To update a group policy, add the group name to **ipa pwpolicy-mod**. For example:

```
$ ipa pwpolicy-mod group_name --minlength=10
```

2. *Optional.* Use the **ipa pwpolicy-show** command to display the new policy settings.

- a. To display the global policy:

```
$ ipa pwpolicy-show
```

- b. To display a group policy, add the group name to **ipa pwpolicy-show**:

```
$ ipa pwpolicy-show group_name
```

28.5. CHANGING PASSWORD EXPIRATION DATE WITH IMMEDIATE EFFECT

IdM applies the password policy rules when an existing password changes or when a user enters a new password. See [Section 28.4, “Modifying Password Policy Attributes”](#).

To enforce an immediate change of the expiration date of a user password, reset the **krbPasswordExpiration** attribute value in LDAP. For example, for a single user:

1. Use the **ldapmodify** utility:

```
# ldapmodify -D "cn=Directory Manager" -w secret -h
server.example.com -p 389 -vv

dn: uid=user_name,cn=users,cn=accounts,dc=example,dc=com
changetype: modify
replace: krbPasswordExpiration
krbPasswordExpiration: 20160203203734Z
```

The **krbPasswordExpiration** format follows this template:


- Year (**2016**)
- Month (**02**)
- Day (**03**)
- Current time in hours, minutes, and seconds (**20:37:34**)
- Time zone (**Z**)

2. Press **Ctrl+D** to confirm and send the changes to the server.

To edit multiple entries at once, use the **-f** option with **ldapmodify** to reference an LDIF file.

CHAPTER 29. MANAGING THE KERBEROS DOMAIN

This chapter describes managing the Kerberos Key Distribution Center (KDC) component of the Identity Management server.



IMPORTANT

Do not use the **kadmin** or **kadmin.local** utilities to manage the Identity Management Kerberos policies. Use the native Identity Management command-line tools as described in this guide.

If you attempt to manage the Identity Management policies using the mentioned Kerberos tools, some of the operations will not affect the Identity Management configuration stored in its Directory Server instance.

29.1. MANAGING KERBEROS TICKET POLICIES

Kerberos ticket policies in Identity Management set restrictions on ticket duration and renewal. Using the following procedures, you can configure Kerberos ticket policies for the Kerberos Key Distribution Center (KDC) running on your Identity Management server.

29.1.1. Global and User-specific Kerberos Ticket Policies

You can redefine the global Kerberos ticket policy and define additional policies specifically to individual users.

Global Kerberos ticket policy

The global policy applies to all tickets issued within the Identity Management Kerberos realm.

User-specific Kerberos ticket policies

User-specific policies apply only to the associated user account. For example, a user-specific Kerberos ticket policy can define a longer maximum ticket lifetime for the **admin** user.

User-specific policies take precedence over the global policy.

29.1.2. Configuring the Global Kerberos Ticket Policy

To configure the global Kerberos ticket policy, you can use:

- the Identity Management web UI: see [the section called “Web UI: Configuring the Global Kerberos Ticket Policy”](#)
- the command line: see [the section called “Command Line: Configuring the Global Kerberos Ticket Policy”](#)

Table 29.1. Supported Kerberos Ticket Policy Attributes

Attribute	Explanation	Example
Max renew	The period of time (in seconds) during which the user can renew the Kerberos ticket after its expiry. After the renew period, the user must log in using the kinit utility to get a new ticket. To renew the ticket, use the kinit -R command.	Max renew = 604800 After the ticket expires, the user can renew it within the next 7 days (604,800 seconds).
Max life	The lifetime of a Kerberos ticket (in seconds). The period during which the Kerberos ticket stays active.	Max life = 86400 The ticket expires 24 hours (86,400 seconds) after it was issued.

Web UI: Configuring the Global Kerberos Ticket Policy

1. Select **Policy** → **Kerberos Ticket Policy**.
2. Define the required values:
 - a. In the **Max renew** field, enter the maximum renewal period of Kerberos tickets.
 - b. In the **Max life** field, enter the maximum lifetime of Kerberos tickets.

Figure 29.1. Configuring the Global Kerberos Ticket Policy

3. Click **Save**.

Command Line: Configuring the Global Kerberos Ticket Policy

To modify the global Kerberos ticket policy:

- Use the **ipa krbtpolicy-mod** command, and pass at least one of the following options:
 - **--maxrenew** to define the maximum renewal period of Kerberos tickets

- **--maxlife** to define the maximum lifetime of Kerberos tickets

For example, to change the maximum lifetime:

```
$ ipa krbtpolicy-mod --maxlife=80000
Max life: 80000
Max renew: 604800
```

To reset the global Kerberos ticket policy to the original default values:

1. Use the **ipa krbtpolicy-reset** command.
2. *Optional.* Use the **ipa krbtpolicy-show** command to verify the current settings.

For details on **ipa krbtpolicy-mod** and **ipa krbtpolicy-reset**, pass the **--help** option with them.

29.1.3. Configuring User-specific Kerberos Ticket Policies

To modify the Kerberos ticket policy for a particular user:

1. Use the **ipa krbtpolicy-mod *user_name*** command, and pass at least one of the following options:
 - **--maxrenew** to define the maximum renewal period of Kerberos tickets
 - **--maxlife** to define the maximum lifetime of Kerberos tickets

If you define only one of the attributes, Identity Management will apply the global Kerberos ticket policy value for the other attribute.

For example, to change the maximum lifetime for the **admin** user:

```
$ ipa krbtpolicy-mod admin --maxlife=160000
Max life: 80000
Max renew: 604800
```

2. *Optional.* Use the **ipa krbtpolicy-show *user_name*** command to display the current values for the specified user.

The new policy takes effect immediately on the next Kerberos ticket that the user requests, such as when using the **kinit** utility.

To reset a user-specific Kerberos ticket policy, use the **ipa krbtpolicy-reset *user_name*** command. The command clears the values defined specifically to the user, after which Identity Management applies the global policy values.

For details on **ipa krbtpolicy-mod** and **ipa krbtpolicy-reset**, pass the **--help** option with them.

29.2. REKEYING KERBEROS PRINCIPALS

Rekeying a Kerberos principal adds a new keytab entry with a higher key version number (KVNO) to the principal's keytab. The original entry remains in the keytab, but is no longer used to issue tickets.

1. Find all keytabs issued within the required time period. For example, the following commands use the **ldapsearch** utility to display all host and service principals created between midnight on January 1, 2016, and 11:59 PM on December 31, 2016 in Greenwich Mean Time (GMT):

```
# ldapsearch -x -b "cn=computers,cn=accounts,dc=example,dc=com" "(&
(krblastpwdchange>=20160101000000)
(krblastpwdchange<=20161231235959))" dn krbprincipalname
```

```
# ldapsearch -x -b "cn=services,cn=accounts,dc=example,dc=com" "(&
(krblastpwdchange>=20160101000000)
(krblastpwdchange<=20161231235959))" dn krbprincipalname
```

- The searchbase (**-b**) defines the subtree where **ldapsearch** looks for the principals:
 - Host principals are stored under the **cn=computers,cn=accounts,dc=example,dc=com** subtree.
 - Service principals are stored under the **cn=services,cn=accounts,dc=example,dc=com** subtree.
 - The **krblastpwdchange** parameter filters the search results by the last change date. The parameter accepts the YYYYMMDD format for the date and the HHMMSS format for the time of day in GMT.
 - Specifying the **dn** and **krbprincipalname** attributes limits the search results to the entry name and principal.
2. For every service and host that requires rekeying the principal, use the **ipa-getkeytab** utility to retrieve a new keytab entry. Pass the following options:
 - **--principal (-p)** to specify the principal
 - **--keytab (-k)** to specify the location of the original keytab
 - **--server (-s)** to specify the Identity Management server host name

For example:

- To rekey a host principal with its keytab in the default location of **/etc/krb5.keytab**:

```
# ipa-getkeytab -p host/client.example.com@EXAMPLE.COM -s
server.example.com -k /etc/krb5.keytab
```

- To rekey the keytab for the Apache service in the default location of **/etc/httpd/conf/ipa.keytab**:

```
# ipa-getkeytab -p HTTP/client.example.com@EXAMPLE.COM -s
server.example.com -k /etc/httpd/conf/ipa.keytab
```



IMPORTANT

Some services, such as NFS version 4, support only a limited set of encryption types. Pass the appropriate arguments to the **ipa-getkeytab** command to configure the keytab properly.

3. *Optional.* Verify that you rekeyed the principals successfully. Use the **klist** utility to list all Kerberos tickets. For example, to list all keytab entries in **/etc/krb5.keytab**:

```
# klist -kt /etc/krb5.keytab
Keytab: WRFILE:/etc/krb5.keytab
KVNO Timestamp          Principal
-----
-----
1 06/09/16 05:58:47 host/client.example.com@EXAMPLE.COM(aes256-
cts-hmac-sha1-96)
2 06/09/16 11:23:01 host/client.example.com@EXAMPLE.COM(aes256-
cts-hmac-sha1-96)
1 03/09/16 13:57:16 krbtgt/EXAMPLE.COM@EXAMPLE.COM(aes256-cts-
hmac-sha1-96)
1 03/09/16 13:57:16 HTTP/server.example.com@EXAMPLE.COM(aes256-
cts-hmac-sha1-96)
1 03/09/16 13:57:16 ldap/server.example.com@EXAMPLE.COM(aes256-
cts-hmac-sha1-96)
```

The output shows that the keytab entry for **client.example.com** was rekeyed with a higher KVNO. The original keytab still exists in the database with the previous KVNO.

Tickets issued against the earlier keytab continue to work, while new tickets are issued using the key with the highest KVNO. This avoids any disruption to system operations.

29.3. PROTECTING KEYTABS

To protect Kerberos keytabs from other users with access to the server, restrict access to the keytab to only the keytab owner. It is recommended to protect the keytabs right after they are retrieved.

For example, to protect the Apache keytab at **/etc/httpd/conf/ipa.keytab**:

1. Set the owner of the file to **apache**.

```
# chown apache /etc/httpd/conf/ipa.keytab
```

2. Set the permissions for the file to **0600**. This grants read, write, and execute permissions to the owner.

```
# chmod 0600 /etc/httpd/conf/ipa.keytab
```

29.4. REMOVING KEYTABS

Removing a keytab and creating a new keytab is necessary for example when you unenroll and re-enroll a host or when you experience Kerberos connection errors.

To remove all keytabs on a host, use the **ipa-rmkeytab** utility, and pass these options:

- **--realm (-r)** to specify the Kerberos realm
- **--keytab (-k)** to specify the path to the keytab file

```
# ipa-rmkeytab --realm EXAMPLE.COM --keytab /etc/krb5.keytab
```

To remove a keytab for a specific service, use the **--principal (-p)** option to specify the service principal:

```
# ipa-rmkeytab --principal ldap/client.example.com --keytab  
/etc/krb5.keytab
```

29.5. ADDITIONAL RESOURCES

- For an overview of the Kerberos KDC hosted by Identity Management servers, see [Section 1.2.1.1, “Services Hosted by IdM Servers”](#).
- For Red Hat documentation on Kerberos, see [Using Kerberos](#) in the *System-Level Authentication Guide*.
- For more information on Kerberos concepts, see the [MIT Kerberos documentation](#).

CHAPTER 30. USING SUDO

Identity Management provides a mechanism for predictably and consistently applying **sudo** policies across the IdM domain. Every system in the IdM domain can be configured as a **sudo** client.

30.1. THE SUDO UTILITY IN IDENTITY MANAGEMENT

The **sudo** utility gives administrative access to specified users. When trusted users precede an administrative command with **sudo**, they are prompted for their own password. Then, when they have been authenticated and assuming that the command is permitted, the administrative command is executed as if they were the root user. For more information about **sudo**, see the [System Administrator's Guide](#).

30.1.1. The Identity Management LDAP Schema for sudo

IdM has a specialized LDAP schema for **sudo** entries. The schema supports:

- Host groups as well as netgroups. Note that **sudo** only supports netgroups.
- **sudo** command groups, which contain multiple commands.



NOTE

Because **sudo** does not support host groups or command groups, IdM translates the IdM **sudo** configuration into the native **sudo** configuration when the **sudo** rules are created. For example, IdM creates a corresponding shadow netgroup for every host group, which allows the IdM administrator to create **sudo** rules that reference host groups, while the local **sudo** command uses the corresponding netgroup.

By default, the **sudo** information is not available anonymously over LDAP. Therefore, IdM defines a default **sudo** user at **uid=sudo, cn=sysaccounts, cn=etc, \$SUFFIX**. You can change this user in the LDAP **sudo** configuration file at **/etc/sudo-ldap.conf**.

30.1.2. NIS Domain Name Requirements

The NIS domain name must be set for netgroups and **sudo** to work properly. The **sudo** configuration requires NIS-formatted netgroups and a NIS domain name for netgroups. However, IdM does not require the NIS domain to actually exist. It is also not required to have a NIS server installed.



NOTE

The **ipa-client-install** utility sets a NIS domain name automatically to the IdM domain name by default.

30.2. SUDO RULES IN IDENTITY MANAGEMENT

Using **sudo** rules, you can define *who* can do *what*, *where*, and *as whom*.

- *Who* are the users allowed to use **sudo**.

- *What* are the commands that can be used with **sudo**.
- *Where* are the target hosts on which the users are allowed to use **sudo**.
- *As whom* is the system or other user identity which the users assume to perform tasks.

30.2.1. External Users and Hosts in **sudo** Rules

IdM accepts external entities in **sudo** rules. External entities are entities that are stored outside of the IdM domain, such as users or hosts that are not part of the IdM domain.

For example, you can use **sudo** rules to grant root access to a member of the IT group in IdM, where the root user is not a user defined in the IdM domain. Or, for another example, administrators can block access to certain hosts that are on a network but are not part of the IdM domain.

30.2.2. User Group Support for **sudo** Rules

You can use **sudo** to give access to whole user groups in IdM. IdM supports both Unix and non-POSIX groups. Note that creating non-POSIX groups can cause access problems because any users in a non-POSIX group inherit non-POSIX permissions from the group.

30.2.3. Support for **sudoers** Options

IdM supports **sudoers** options. For a complete list of the available **sudoers** options, see the `sudoers(5)` man page.

Note that IdM does not allow white spaces or line breaks in **sudoers** options. Therefore, instead of supplying multiple options in a comma-separated list, add them separately. For example, to add two **sudoers** options from the command line:

```
$ ipa sudorule-add-option sudo_rule_name  
Sudo Option: first_option  
$ ipa sudorule-add-option sudo_rule_name  
Sudo Option: second_option
```

Similarly, make sure to supply long options on one line. For example, from the command line:

```
$ ipa sudorule-add-option sudo_rule_name  
Sudo Option: env_keep="COLORS DISPLAY EDITOR HOSTNAME HISTSIZE INPUTRC  
KDEDIR LESSSECURE LS_COLORS MAIL PATH PS1 PS2 XAUTHORITY"
```

30.3. CONFIGURING THE LOCATION FOR LOOKING UP **sudo** POLICIES

The centralized IdM database for **sudo** configuration makes the **sudo** policies defined in IdM globally available to all domain hosts. On Red Hat Enterprise Linux 7.1 systems and later, the **ipa-server-install** and **ipa-client-install** utilities automatically configure the system to use the IdM-defined policies by setting SSSD as the data provider for **sudo**.

The location for looking up the **sudo** policies is defined on the **sudoers** line of the **/etc/nsswitch.conf** file. On IdM systems running Red Hat Enterprise Linux 7.1 and later, the default **sudoers** configuration in **nsswitch.conf** is:

```
sudoers: files sss
```

The **files** option specifies that the system uses the **sudo** configuration defined in the **/etc/sudoers** local SSSD configuration file. The **sss** option specifies that the **sudo** configuration defined in IdM is used.

30.3.1. Configuring Hosts to Use IdM **sudo** Policies in Earlier Versions of IdM

To implement the IdM-defined **sudo** policies on IdM systems running Red Hat Enterprise Linux versions earlier than 7.1, configure the local machines manually. You can do this using SSSD or LDAP. Red Hat strongly recommends to use the SSSD-based configuration.

30.3.1.1. Applying the **sudo** Policies to Hosts Using SSSD

Follow these steps on each system that is required to use SSSD for **sudo** rules:

1. Configure **sudo** to look to SSSD for the **sudoers** file.

```
# vim /etc/nsswitch.conf

sudoers: files sss
```

Leaving the **files** option in place allows **sudo** to check its local configuration before checking SSSD for the IdM configuration.

2. Add **sudo** to the list of services managed by the local SSSD client.

```
# vim /etc/sss/sss.conf

[sss]
config_file_version = 2
services = nss, pam, sudo
domains = IPADOMAIN
```

3. Set a name for the NIS domain in the **sudo** configuration. **sudo** uses NIS-style netgroups, so the NIS domain name must be set in the system configuration for **sudo** to be able to find the host groups used in the IdM **sudo** configuration.

1. Enable the **rhel-domainname** service if it is not already enabled to ensure that the NIS domain name will be persistent across reboots.

```
# systemctl enable rhel-domainname.service
```

2. Set the NIS domain name to use with the **sudo** rules.

```
# nisdomainname example.com
```

3. Configure the system authentication settings to persist the NIS domain name. For example:

```
# echo "NISDOMAIN=example.com" >> /etc/sysconfig/network
```


This updates the `/etc/sysconfig/network` and `/etc/yp.conf` files with the NIS domain.

4. Optionally, enable debugging in SSSD to show what LDAP settings it is using.

```
[domain/IPADOMAIN]
debug_level = 6
....
```

The LDAP search base used by SSSD for operations is recorded in the `sssd_DOMAINNAME.log` log.

30.3.1.2. Applying the sudo Policies to Hosts Using LDAP



IMPORTANT

Only use the LDAP-based configuration for clients that do not use SSSD. Red Hat recommends to configure all other clients using the SSSD-based configuration, as described in [Section 30.3.1.1, “Applying the sudo Policies to Hosts Using SSSD”](#).

For information on applying **sudo** policies using LDAP, see the [Identity Management Guide for Red Hat Enterprise Linux 6](#).

The LDAP-based configuration is expected to be used primarily for clients based on Red Hat Enterprise Linux versions earlier than Red Hat Enterprise Linux 7. It is therefore only described in the documentation for Red Hat Enterprise Linux 6.

30.4. ADDING `sudo` COMMANDS, COMMAND GROUPS, AND RULES

30.4.1. Adding `sudo` Commands

Adding `sudo` Commands in the Web UI

1. Under the **Policy** tab, click **Sudo** → **Sudo Commands**.
2. Click **Add** at the top of the list.
3. Fill out the information about the command. Enter the full system path to the command executable.

Figure 30.1. Adding a New sudo Command

4. Click **Add**. Alternatively, click **Add and Add Another** to start adding another entry or **Add and Edit** to start editing the new entry.

Adding sudo Commands from the Command Line


To add a **sudo** command, use the **ipa sudocmd-add** command. Provide the full system path to the command executable. For example, to add the **/usr/bin/less** command and a description:

```
$ ipa sudocmd-add /usr/bin/less --desc="For reading log files"
-----
Added sudo command "/usr/bin/less"
-----
sudo Command: /usr/bin/less
Description: For reading log files
```

30.4.2. Adding sudo Command Groups

Adding sudo Command Groups in the Web UI

1. Under the **Policy** tab, click **Sudo** → **Sudo Command Groups**.
2. Click **Add** at the top of the list.
3. Fill out the information about the command group.



Add Sudo Command Group

Sudo Command *

Group


Description

* Required field

Add Add and Add Another Add and Edit Cancel

Figure 30.2. Adding a New sudo Command Group

- Click **Add and Edit** to start editing the command group.
- Under the **Sudo Commands** tab, click **Add** to add a **sudo** command to the group. Select the required commands and move them to the **Prospective** column using the > button.



Add Sudo Commands into Sudo Command Group files

Filter available Sudo Commands

Available

	Sudo Command	Description
☐		
☒	/usr/bin/less	For reading log files.
☒	/usr/bin/vim	For editing files.

Prospective

	Sudo Command	Description
☐		

Add Cancel

Figure 30.3. Adding Commands to a sudo Command Group

- Click **Add**.

Adding sudo Command Groups from the Command Line

- Create the command group using the **ipa sudocmdgroup-add** command. For example, to create the **files** command group and add its description:

```
$ ipa sudocmdgroup-add files --desc="File editing commands"
-----
```

```
Added sudo command group "files"
-----
sudo Command Group: files
Description: File editing commands
```

2. Include a **sudo** command in the group using the **ipa sudocmdgroup-add-member** command. Note that you can only include commands that have already been added to IdM, as described in [Section 30.4.1, “Adding sudo Commands”](#).

```
$ ipa sudocmdgroup-add-member files --sudocmds "/usr/bin/vim"
sudo Command Group: files
Description: File editing commands
Member sudo commands: /usr/bin/vim
-----
Number of members added 1
-----
```

30.4.3. Adding sudo Rules

Adding sudo Rules in the Web UI

1. Under the **Policy** tab, click **Sudo** → **Sudo Rules**.
2. Click **Add** at the top of the list.
3. Enter the name for the rule.




Figure 30.4. Naming a New sudo Rule

4. Click **Add**. Alternatively, click **Add and Add Another** to start adding another entry or **Add and Edit** to start editing the new entry.

For information on how to edit the new **sudo** rule, see [Section 30.6, “Modifying sudo Rules”](#).

Adding sudo Rules from the Command Line

To add a new **sudo** rule, use the **ipa sudorule-add** command. For example, to add a rule named **files-commands**:

```
$ ipa sudorule-add files-commands
-----
Added Sudo Rule "files-commands"
```

```
-----  
Rule name: files-commands  
Enabled: TRUE
```

For more information on using **ipa sudorule-add** and the options it accepts, run the command with the **--help** option added.

For information on how to edit the new **sudo** rule, see [Section 30.6, “Modifying sudo Rules”](#).

For a complete example of adding a new **sudo** rule and editing it from the command line, see [Example 30.1, “Adding and Modifying a New sudo Rule from the Command Line”](#).

30.5. MODIFYING SUDO COMMANDS AND COMMAND GROUPS

Modifying sudo Commands and Command Groups in the Web UI

1. Under the **Policy** tab, click **Sudo** → **Sudo Commands** or **Sudo** → **Sudo Command Groups**.
2. Click the name of the command or command group to display its configuration page.
3. Change the settings as required. On some configuration pages, the **Save** button is available at the top of the page. On these pages, you must click the button to confirm the changes.

Modifying sudo Commands and Command Groups from the Command Line

To modify a command or command group, use the following commands:

- **ipa sudocmd-mod**
- **ipa sudocmdgroup-mod**

Add command-line options to the above-mentioned commands to update the **sudo** command or command group attributes. For example, to add a new description for the **/usr/bin/less** command:

```
$ ipa sudocmd-mod /usr/bin/less --desc="For reading log files"  
-----  
Modified Sudo Command "/usr/bin/less"  
-----  
Sudo Command: /usr/bin/less  
Description: For reading log files  
Sudo Command Groups: files
```

For more information about these commands and the options they accept, run them with the **--help** option added.

30.6. MODIFYING SUDO RULES

Modifying sudo Rules in the Web UI

1. Under the **Policy** tab, click **Sudo** → **Sudo Rules**.
2. Click the name of the rule to display its configuration page.

3. Change the settings as required. On some configuration pages, the **Save** button is available at the top of the page. On these pages, click the button to confirm the changes.

The **sudo** rule configuration page includes several configuration areas:

The General area

In this area, you can modify the rule's description and **sudo order**. The **sudo order** field accepts integers and defines the order in which IdM evaluates the rules. The rule with the highest **sudo order** value is evaluated first.

The Options area

In this area, you can add **sudoers** options to the rule.

1. Click **Add** above the options list.




Figure 30.5. Adding a sudo Option

2. Enter the **sudoers** option. For example, to specify that **sudo** will not prompt the user to authenticate, add the **!authenticate** option:



Figure 30.6. Entering a sudoers Option

For more information on **sudoers** options, see the `sudoers(5)` man page.

3. Click **Add**.


The Who area

In this area, you can select the users or user groups to which the **sudo** rule will be applied. These users will be entitled to use **sudo** as defined in the rule.

To specify that all system users will be able to use **sudo** as defined in the rule, select **Anyone**.

To apply the rule to specific users or groups only, select **Specified Users and Groups** and then follow these steps:

1. Click **Add** above the users or user groups list.



Who


User category the rule applies to: ☐ Anyone ☒ Specified Users and Groups

☐	Users	External	Delete	+ Add
☐	manager			
☐	employee			
☐	helpdesk			

☐	User Groups	Delete	+ Add
☐	admins		

Figure 30.7. Adding Users to a sudo Rule

2. Select the users or user groups to add to the rule, and click the > arrow button to move them to the **Prospective** column. To add an external user, specify the user in the **External** field, and then click the > arrow button.



Add Users into Sudo Rule files-commands

Filter available Users

Available

☐	Users
☐	xyz

External

>

<

Prospective

☐	Users
☒	employee
☒	helpdesk
☒	manager

Figure 30.8. Selecting Users for a sudo Rule

3. Click **Add**.


The Access This Host area

In this area, you can select the hosts on which the **sudo** rule will be in effect. These are the hosts where the users will be granted **sudo** permissions.

To specify that the rule will be in effect on all hosts, select **Anyone**.

To apply the rule to specific hosts or host groups only, select **Specified Hosts and Groups** and then follow these steps:

1. Click **Add** above the hosts list.




Access this host

Host category the rule applies to: ☐ Any Host ☒ Specified Hosts and Groups

☐	Hosts	External	Delete	+ Add
☐	Host Groups		Delete	+ Add

Figure 30.9. Adding Hosts to a sudo Rule

2. Select the hosts or host groups to include with the rule, and click the > arrow button to move them to the **Prospective** column. To add an external host, specify the host in the **External** field, and then click the > arrow button.



Add Hosts into Sudo Rule files-commands

Filter available Hosts Filter

Available			Prospective	
☐	Hosts	>	☐	Hosts
☐	qe-server.example.com		☒	server.example.com
		<		

External

Add Cancel

Figure 30.10. Selecting Hosts for a sudo Rule

3. Click **Add**.

The Run Commands area

In this area, you can select the commands to be included in the **sudo** rule. You can specify that users will be either allowed or denied to use specific commands.

To specify that users will be allowed to use any command with **sudo**, select **Any Command**.

To associate the rule with specific commands or command groups, select **Specified Commands and Groups** and then follow these steps:

1. Click one of the **Add** buttons to add a command or a command group.

To specify allowed commands or command groups, use the **Allow** area. To specify denied commands or command groups, use the **Deny** area.

The screenshot shows a web interface titled "Run Commands". Below the title, it says "Command category the rule applies to:" followed by two radio buttons: "Any Command" (unselected) and "Specified Commands and Groups" (selected). Below this, there are two main sections: "Allow" and "Deny". Each section contains a table with two rows. The "Allow" section has rows for "Sudo Allow Commands" and "Sudo Allow Command Groups". The "Deny" section has rows for "Sudo Deny Commands" and "Sudo Deny Command Groups". Each row has a checkbox on the left, a text label in the middle, a "Delete" button (trash icon) on the right, and a "+ Add" button on the far right. The "+ Add" buttons in all four rows are highlighted with red rectangular boxes.

Figure 30.11. Adding Commands to a sudo Rule

2. Select the commands or command groups to include with the rule, and click the > arrow button to move them to the **Prospective** column.




Figure 30.12. Selecting Commands for a sudo Rule

3. Click **Add**.

The As Whom area

In this area, you can configure the **sudo** rule to run the given commands as a specific, non-root user.

Note that if you add a group of RunAs users, UIDs of the members of the group will be used to run the command. If you add a RunAs group, the GID of the group will be used to run the command.

To specify that the rule will be run as any user on the system, select **Anyone**. To specify that the rule will be run as any group on the system, select **Any Group**.

1. Click **Add** above the users list.




Figure 30.13. Configuring sudo Rules to Execute Commands as a Specific User

2. Select the required users or groups, and use the > arrow button to move them to the **Prospective** column. To add an external entity, specify it in the **External** field, and then click the > arrow button.

Figure 30.14. Selecting Users for the Command

3. Click **Add**.

Modifying sudo Rules from the Command Line

The IdM command-line utilities allow you to configure several **sudo** rule areas:

General sudo rules management

To change the general configuration for a **sudo** rule, use the **ipa sudorule-mod** command. The most common options accepted by the command are:

- The **--desc** option to change the **sudo** rule description. For example:

```
$ ipa sudorule-mod sudo_rule_name --desc="sudo_rule_description"
```

- The **--order** option to define the order of the specified rule. For example:

```
$ ipa sudorule-mod sudo_rule_name --order=3
```

- Options to specify a category of entities: **--usercat** (user category), **--hostcat** (host category), **--cmdcat** (command category), **--runasusercat** (run-as user category), and **--runasgroupcat** (run-as group category). These options only accept the **all** value that associates the rule with all users, hosts, commands, run-as users, or run-as groups.

For example, to specify that all users will be able to use **sudo** as defined in the **sudo_rule** rule:

```
$ ipa sudorule-mod sudo_rule --usercat=all
```

Note that if the rule is already associated with a specific entity, you must remove it before defining the corresponding **all** category. For example, if **sudo_rule** was previously associated with a specific user using the **ipa sudorule-add-user** command, you must first use the **ipa sudorule-remove-user** command to remove the user.

For more details and a complete list of options accepted by **ipa sudorule-mod**, run the command with the **--help** option added.

Managing sudo options

To add a **sudoers** option, use the **ipa sudorule-add-option** command.

For example, to specify that users using **sudo** based on the **files-commands** rule will not be required to authenticate, add the **!authenticate** option:

```
$ ipa sudorule-add-option files-commands
Sudo Option: !authenticate
-----
Added option "!authenticate" to Sudo Rule "files-commands"
-----
```

For more information on **sudoers** options, see the **sudoers(5)** man page.

To remove a **sudoers** option, use the **ipa sudorule-remove-option** command. For example:

```
$ ipa sudorule-remove-option files-commands
Sudo Option: authenticate
-----
Removed option "authenticate" from Sudo Rule "files-commands"
-----
```

Managing who is granted the permission to use sudo

To specify an individual user, add the **--users** option to the **ipa sudorule-add-user** command. To specify a user group, add the **--groups** option to **ipa sudorule-add-user**.

For example, to add **user** and **user_group** to the **files-commands** rule:

```
$ ipa sudorule-add-user files-commands --users=user --groups=user_group
...
-----
Number of members added 2
-----
```

To remove an individual user or group, use the **ipa sudorule-remove-user**. For example, to remove a user:

```
$ ipa sudorule-remove-user files-commands
[member user]: user
[member group]:
...
```

```
-----  
Number of members removed 1  
-----
```

Managing where the users are granted the sudo permissions

To specify a host, add the **--hosts** option to the **ipa sudorule-add-host** command. To specify a host group, add the **--hostgroups** option to **ipa sudorule-add-host**.

For example, to add **example.com** and **host_group** to the **files-commands** rule:

```
$ ipa sudorule-add-host files-commands --hosts=example.com --  
hostgroups=host_group  
...  
-----  
Number of members added 2  
-----
```

To remove a host or host group, use the **ipa sudorule-remove-host** command. For example:

```
$ ipa sudorule-remove-host files-commands  
[member host]: example.com  
[member host group]:  
...  
-----  
Number of members removed 1  
-----
```

Managing what commands can be used with sudo

You can specify that users will be either allowed or denied to use specific commands.

To specify an allowed command or command group, add the **--sudocmds** or **--sudocmdgroups** option to the **ipa sudorule-add-allow-command**. To specify a denied command or command group, add the **--sudocmds** or **--sudocmdgroups** option to the **ipa sudorule-add-deny-command** command.

For example, to add the **/usr/bin/less** command and the **files** command group as allowed to the **files-commands** rule:

```
$ ipa sudorule-add-allow-command files-commands --sudocmds=/usr/bin/less  
--sudocmdgroups=files  
...  
-----  
Number of members added 2  
-----
```

To remove a command or command group from a rule, use the **ipa sudorule-remove-allow-command** or **ipa sudorule-remove-deny-command** commands. For example:

```
$ ipa sudorule-remove-allow-command files-commands  
[member sudo command]: /usr/bin/less  
[member sudo command group]:  
...
```

```
-----
Number of members removed 1
-----
```

Note that the **--sudocmds** option only accepts commands added to IdM, as described in [Section 30.4.1, “Adding sudo Commands”](#).

Managing as whom the sudo commands are run

To use the UIDs of an individual user or users in a group as the identity under which the commands are run, use the **--users** or **--groups** options with the **ipa sudorule-add-runasuser** command.

To use the GID of a user group as the identity for the commands, use the **ipa sudorule-add-runasgroup --groups** command.

If you specify no user or group, **sudo** commands will be run as root.

For example, to specify that the identity of **user** will be used to execute the commands in the **sudo** rule:

```
$ ipa sudorule-add-runasuser files-commands --users=user
...
RunAs Users: user
...
```

For more information on the **ipa sudorule-*** commands, see the output of the **ipa help sudorule** command or run a particular command with the **--help** option added.

Example 30.1. Adding and Modifying a New sudo Rule from the Command Line

To allow a specific user group to use **sudo** with any command on selected servers:

1. Obtain a Kerberos ticket for the **admin** user or any other user allowed to manage **sudo** rules.

```
$ kinit admin
Password for admin@EXAMPLE.COM:
```

2. Add a new **sudo** rule to IdM.

```
$ ipa sudorule-add new_sudo_rule --desc="Rule for user_group"
-----
Added Sudo Rule "new_sudo_rule"
-----
Rule name: new_sudo_rule
Description: Rule for user_group
Enabled: TRUE
```

3. Define the *who*: specify the group of users who will be entitled to use the **sudo** rule.

```
$ ipa sudorule-add-user new_sudo_rule --groups=user_group
Rule name: new_sudo_rule
```

```
Description: Rule for user_group
Enabled: TRUE
User Groups: user_group
-----
Number of members added 1
-----
```

4. Define the *where*: specify the group of hosts where the users will be granted the **sudo** permissions.

```
$ ipa sudorule-add-host new_sudo_rule --hostgroups=host_group
Rule name: new_sudo_rule
Description: Rule for user_group
Enabled: TRUE
User Groups: user_group
Host Groups: host_group
-----
Number of members added 1
-----
```

5. Define the *what*: to allow the users to run any **sudo** command, add the **all** command category to the rule.

```
$ ipa sudorule-mod new_sudo_rule --cmdcat=all
-----
Modified Sudo Rule "new_sudo_rule"
-----
Rule name: new_sudo_rule
Description: Rule for user_group
Enabled: TRUE
Command category: all
User Groups: user_group
Host Groups: host_group
```

6. To let the **sudo** commands be executed as root, do not specify any run-as users or groups.
7. Add the **!authenticate sudoers** option to specify that the users will not be required to authenticate when using the **sudo** command.

```
$ ipa sudorule-add-option new_sudo_rule
Sudo Option: !authenticate
-----
Added option "!authenticate" to Sudo Rule "new_sudo_rule"
-----
Rule name: new_sudo_rule
Description: Rule for user_group
Enabled: TRUE
Command category: all
User Groups: user_group
Host Groups: host_group
Sudo Option: !authenticate
```

8. Display the new **sudo** rule configuration to verify it is correct.

```
$ ipa sudorule-show new_sudo_rule
Rule name: new_sudo_rule
Description: Rule for user_group
Enabled: TRUE
Command category: all
User Groups: user_group
Host Groups: host_group
Sudo Option: !authenticate
```

30.7. LISTING AND DISPLAYING SUDO COMMANDS, COMMAND GROUPS, AND RULES

Listing and Displaying sudo Commands, Command Groups, and Rules in the Web UI

1. Under the **Policy** tab, click **Sudo** and select **Sudo Rules**, **Sudo Commands**, or **Sudo Command Groups**.
2. Click the name of the rule, command, or command group to display its configuration page.

Listing and Displaying sudo Commands, Command Groups, and Rules from the Command Line

To list all commands, command groups, and rules, use the following commands:

- **ipa sudocmd-find**
- **ipa sudocmdgroup-find**
- **ipa sudorule-find**

To display information about a particular command, command group, or rule, use the following commands:

- **ipa sudocmd-show**
- **ipa sudocmdgroup-show**
- **ipa sudorule-show**

For example, to display information about the **/usr/bin/less** command:

```
$ ipa sudocmd-show /usr/bin/less
Sudo Command: /usr/bin/less
Description: For reading log files.
Sudo Command Groups: files
```

For more information about these commands and the options they accept, run them with the **--help** option added.

30.8. DISABLING AND ENABLING SUDO RULES

Disabling a **sudo** rule temporarily deactivates it. A disabled rule is not removed from IdM and can be enabled again.

Disabling and Enabling sudo Rules from the Web UI

1. Under the **Policy** tab, click **Sudo** → **Sudo Rule**.
2. Select the rule to disable and click **Disable** or **Enable**.




Figure 30.15. Disabling or Enabling a sudo Rule

Disabling and Enabling sudo Rules from the Command Line

To disable a rule, use the **ipa sudo-rule-disable** command.

```
$ ipa sudorule-disable sudo_rule_name
-----
Disabled Sudo Rule "sudo_rule_name"
-----
```

To re-enable a rule, use the **ipa sudorule-enable** command.

```
$ ipa sudorule-enable sudo_rule_name
-----
Enabled Sudo Rule "sudo_rule_name"
-----
```

30.9. REMOVING SUDO COMMANDS, COMMAND GROUPS, AND RULES

Removing sudo Commands, Command Groups, and Rules in the Web UI

1. Under the **Policy** tab, click **Sudo** and select **Sudo Rules**, **Sudo Commands**, or **Sudo Command Groups**.
2. Select the command, command group, or rule to delete, and click **Delete**.




Figure 30.16. Deleting a sudo Command

Removing sudo Commands, Command Groups, and Rules from the Command Line

To delete a command, command group, or rule, use the following commands:

- `ipa sudocmd-del`
- `ipa sudocmdgroup-del`
- `ipa sudorule-del`

For more information about these commands and the options they accept, run them with the `--help` option added.

CHAPTER 31. CONFIGURING HOST-BASED ACCESS CONTROL

This chapter describes *host-based access control* (HBAC) in Identity Management (IdM) and explains how you can use HBAC to manage access control in your IdM domain.

31.1. HOW HOST-BASED ACCESS CONTROL WORKS IN IDM

Host-based access control defines which users (or user groups) can access specified hosts (or host groups) by using specified services (or services in a service group). For example, you can:

- Limit access to a specified system in your domain to members of a specific user group.
- Allow only a specific service to be used to access the systems in your domain.

The administrator configures host-based access control by using a set of allowing rules named *HBAC rules*. By default, IdM is configured with a default HBAC rule named **allow_all**, which allows universal access in the whole IdM domain.

Applying HBAC Rules to Groups

For centralized and simplified access control management, you can apply HBAC rules to whole user, host, or service groups instead of individual users, hosts, or services.




Figure 31.1. Host Groups and Host-Based Access Control


When applying HBAC rules to groups, consider using *automember rules*. See [Section 13.6, “Defining Automatic Group Membership for Users and Hosts”](#).

31.2. CONFIGURING HOST-BASED ACCESS CONTROL IN AN IDM DOMAIN

To configure your domain for host-based access control:

1. [Create HBAC rules](#)
2. [Test the new HBAC rules](#)

3. Disable the default `allow_all` HBAC rule



IMPORTANT

Do not disable the `allow_all` rule before creating custom HBAC rules. If you do this, no users will be able to access any hosts.

31.2.1. Creating HBAC Rules

To create an HBAC rule, you can use:

- the IdM web UI (see [the section called “Web UI: Creating an HBAC Rule”](#))
- the command line (see [the section called “Command Line: Creating HBAC Rules”](#))

For examples, see [the section called “Examples of HBAC Rules”](#).

Web UI: Creating an HBAC Rule

1. Select **Policy** → **Host-Based Access Control** → **HBAC Rules**.
2. Click **Add** to start adding a new rule.
3. Enter a name for the rule, and click **Add and Edit** to go directly to the HBAC rule configuration page.
4. In the **Who** area, specify the target users.
 - To apply the HBAC rule to specified users or groups only, select **Specified Users and Groups**. Then click **Add** to add the users or groups.
 - To apply the HBAC rule to all users, select **Anyone**.

Figure 31.2. Specifying a Target User for an HBAC Rule

5. In the **Accessing** area, specify the target hosts:
 - To apply the HBAC rule to specified hosts or groups only, select **Specified Hosts and Groups**. Then click **Add** to add the hosts or groups.
 - To apply the HBAC rule to all hosts, select **Any Host**.
6. In the **Via Service** area, specify the target HBAC services:
 - To apply the HBAC rule to specified services or groups only, select **Specified**

Services and Groups. Then click **Add** to add the services or groups.

- To apply the HBAC rule to all services, select **Any Service**.



NOTE

Only the most common services and service groups are configured for HBAC rules by default.

- To display the list of services that are currently available, select **Policy → Host-Based Access Control → HBAC Services**.
- To display the list of service groups that are currently available, select **Policy → Host-Based Access Control → HBAC Service Groups**.

To add more services and service groups, see [Section 31.3, “Adding HBAC Service Entries for Custom HBAC Services”](#) and [Section 31.4, “Adding HBAC Service Groups”](#).

7. Changing certain settings on the HBAC rule configuration page highlights the **Save** button at the top of the page. If this happens, click the button to confirm the changes.

Command Line: Creating HBAC Rules

1. Use the **ipa hbacrule-add** command to add the rule.

```
$ ipa hbacrule-add
Rule name: rule_name
-----
Added HBAC rule "rule_name"
-----
Rule name: rule_name
Enabled: TRUE
```

2. Specify the target users.

- To apply the HBAC rule to specified users or groups only, use the **ipa hbacrule-add-user** command.

For example, to add a group:

```
$ ipa hbacrule-add-user
Rule name: rule_name
[member user]:
[member group]: group_name
Rule name: rule_name
Enabled: TRUE
User Groups: group_name
-----
Number of members added 1
-----
```

To add multiple users or groups, use the **--users** and **--groups** options:

```
$ ipa hbacrule-add-user rule_name --users=user1 --users=user2 --
users=user3
  Rule name: rule_name
  Enabled: TRUE
  Users: user1, user2, user3
-----
Number of members added 3
-----
```

- To apply the HBAC rule to all users, use the **ipa hbacrule-mod** command and specify the **all** user category:

```
$ ipa hbacrule-mod rule_name --usercat=all
-----
Modified HBAC rule "rule_name"
-----
  Rule name: rule_name
  User category: all
  Enabled: TRUE
```



NOTE

If the HBAC rule is associated with individual users or groups, **ipa hbacrule-mod --usercat=all** fails. In this situation, remove the users and groups using the **ipa hbacrule-remove-user** command.

For details, run **ipa hbacrule-remove-user** with the **--help** option.

3. Specify the target hosts.

- To apply the HBAC rule to specified hosts or groups only, use the **ipa hbacrule-add-host** command.

For example, to add a single host:

```
$ ipa hbacrule-add-host
Rule name: rule_name
[member host]: host.example.com
[member host group]:
  Rule name: rule_name
  Enabled: TRUE
  Hosts: host.example.com
-----
Number of members added 1
-----
```

To add multiple hosts or groups, use the **--hosts** and **--hostgroups** options:

```
$ ipa hbacrule-add-host rule_name --hosts=host1 --hosts=host2 --
hosts=host3
  Rule name: rule_name
  Enabled: TRUE
  Hosts: host1, host2, host3
```

```
-----
Number of members added 3
-----
```

- To apply the HBAC rule to all hosts, use the **ipa hbacrule-mod** command and specify the **all** host category:

```
$ ipa hbacrule-mod rule_name --hostcat=all
-----
Modified HBAC rule "rule_name"
-----
Rule name: rule_name
Host category: all
Enabled: TRUE
```



NOTE

If the HBAC rule is associated with individual hosts or groups, **ipa hbacrule-mod --hostcat=all** fails. In this situation, remove the hosts and groups using the **ipa hbacrule-remove-host** command.

For details, run **ipa hbacrule-remove-host** with the **--help** option.

4. Specify the target HBAC services.

- To apply the HBAC rule to specified services or groups only, use the **ipa hbacrule-add-service** command.

For example, to add a single service:

```
$ ipa hbacrule-add-service
Rule name: rule_name
[member HBAC service]: ftp
[member HBAC service group]:
Rule name: rule_name
Enabled: TRUE
Services: ftp
-----
Number of members added 1
-----
```

To add multiple services or groups, you can use the **--hbacsvcs** and **--hbacsvcgroups** options:

```
$ ipa hbacrule-add-service rule_name --hbacsvcs=su --
hbacsvcs=sudo
Rule name: rule_name
Enabled: TRUE
Services: su, sudo
-----
Number of members added 2
-----
```

**NOTE**

Only the most common services and service groups are configured for HBAC rules. To add more, see [Section 31.3, “Adding HBAC Service Entries for Custom HBAC Services”](#) and [Section 31.4, “Adding HBAC Service Groups”](#).

- To apply the HBAC rule to all services, use the **ipa hbacrule-mod** command and specify the **all** service category:

```
$ ipa hbacrule-mod rule_name --servicecat=all
-----
Modified HBAC rule "rule_name"
-----
Rule name: rule_name
Service category: all
Enabled: TRUE
```

**NOTE**

If the HBAC rule is associated with individual services or groups, **ipa hbacrule-mod --servicecat=all** fails. In this situation, remove the services and groups using the **ipa hbacrule-remove-service** command.

For details, run **ipa hbacrule-remove-service** with the **--help** option.

5. *Optional.* Verify that the HBAC rule has been added correctly.
 - a. Use the **ipa hbacrule-find** command to verify that the HBAC rule has been added to IdM.
 - b. Use the **ipa hbacrule-show** command to verify the properties of the HBAC rule.

For details, run the commands with the **--help** option.

Examples of HBAC Rules

Example 31.1. Granting a Single User Access to All Hosts Using Any Service

To allow the **admin** user to access all systems in the domain using any service, create a new HBAC rule and set:

- the user to **admin**
- the host to **Any host** (in the web UI), or use **--hostcat=all** with **ipa hbacrule-add** (when adding the rule) or **ipa hbacrule-mod**
- the service to **Any service** (in the web UI), or use **--servicecat=all** with **ipa hbacrule-add** (when adding the rule) or **ipa hbacrule-mod**


Example 31.2. Ensuring That Only Specific Services Can Be Used to Access a Host

To make sure that all users must use **sudo**-related services to access the host named **host.example.com**, create a new HBAC rule and set:

- the user to **Anyone** (in the web UI), or use **--usercat=all** with **ipa hbacrule-add** (when adding the rule) or **ipa hbacrule-mod**
- the host to **host.example.com**
- the HBAC service group to **Sudo**, which is a default group for **sudo** and related services

31.2.2. Testing HBAC Rules

IdM enables you to test your HBAC configuration in various situations using simulated scenarios. By performing these simulated test runs, you can discover misconfiguration problems or security risks before deploying HBAC rules in production.



IMPORTANT

Always test custom HBAC rules before you start using them in production.

Note that IdM does not test the effect of HBAC rules on trusted Active Directory (AD) users. Because AD data is not stored in the IdM LDAP directory, IdM cannot resolve group membership of AD users when simulating HBAC scenarios.

To test an HBAC rule, you can use:

- the IdM web UI (see [the section called “Web UI: Testing an HBAC Rule”](#))
- the command line (see [the section called “Command Line: Testing an HBAC Rule”](#))

Web UI: Testing an HBAC Rule

1. Select **Policy** → **Host-Based Access Control** → **HBAC Test**.
2. On the **Who** screen: Specify the user under whose identity you want to perform the test, and click **Next**.

Who

Who Accessing Via Service Rules Run Test

WHO 

	User login	First name	Last name	Status
☐	admin		Administrator	✓ Enabled
☒	user1	user	user	✓ Enabled
☐	user2	user	user	✓ Enabled
☐	user3	user	user	✓ Enabled

Showing 1 to 4 of 4 entries.


☐ Specify external User:

 Next

Figure 31.3. Specifying the Target User for an HBAC Test

- On the **Accessing** screen: Specify the host that the user will attempt to access, and click **Next**.
- On the **Via Service** screen: Specify the service that the user will attempt to use, and click **Next**.
- On the **Rules** screen: Select the HBAC rules you want to test, and click **Next**. If you do not select any rule, all rules will be tested.

Select **Include Enabled** to run the test on all rules whose status is **Enabled**. Select **Include Disabled** to run the test on all rules whose status is **Disabled**. To view and change the status of HBAC rules, select **Policy → Host-Based Access Control → HBAC Rules**.



IMPORTANT

If the test runs on multiple rules, it will pass successfully if at least one of the selected rules allows access.

- On the **Run Test** screen: Click **Run Test**.




Figure 31.4. Running an HBAC Test

7. Review the test results:

- If you see **ACCESS DENIED**, the user was not granted access in the test.
- If you see **ACCESS GRANTED**, the user was able to access the host successfully.




Figure 31.5. Reviewing HBAC Test Results

By default, IdM lists all the tested HBAC rules when displaying the test results.

- Select **Matched** to display the rules that allowed successful access.
- Select **Unmatched** to display the rules that prevented access.

Command Line: Testing an HBAC Rule

Use the **ipa hbactest** command and specify at least:

- the user under whose identity you want to perform the test
- the host that the user will attempt to access
- the service that the user will attempt to use

For example, when specifying these values interactively:

```
$ ipa hbactest
User name: user1
```

```

Target host: example.com
Service: sudo
-----
Access granted: False
-----
Not matched rules: rule1

```

By default, IdM runs the test on all HBAC rules whose status is **enabled**. To specify different HBAC rules:

- Use the **--rules** option to define one or more HBAC rules.
- Use the **--disabled** option to test all HBAC rules whose status is **disabled**.

To see the current status of HBAC rules, run the **ipa hbacrule-find** command.

Example 31.3. Testing an HBAC Rule from the Command Line

In the following test, an HBAC rule named **rule2** prevented **user1** from accessing **example.com** using the **sudo** service:

```

$ ipa hbactest --user=user1 --host=example.com --service=sudo --
rules=rule1
-----
Access granted: False
-----
Not matched rules: rule1

```

Example 31.4. Testing Multiple HBAC Rules from the Command Line

When testing multiple HBAC rules, the test passes if at least one rule allows the user successful access.

```

$ ipa hbactest --user=user1 --host=example.com --service=sudo --
rules=rule1 --rules=rule2
-----
Access granted: True
-----
Matched rules: rule2
Not matched rules: rule1

```

In the output:

- **Matched rules** list the rules that allowed successful access.
- **Not matched rules** list the rules that prevented access.

31.2.3. Disabling HBAC Rules

Disabling an HBAC rule deactivates the rule, but does not delete it. If you disable an HBAC rule, you can re-enable it later.

**NOTE**

For example, disabling HBAC rules is useful after you configure custom HBAC rules for the first time. To ensure that your new configuration is not overridden by the default **allow_all** HBAC rule, you must disable **allow_all**.

To disable an HBAC rule, you can use:

- the IdM web UI (see [the section called “Web UI: Disabling an HBAC Rule”](#))
- the command line (see [the section called “Command Line: Disabling an HBAC Rule”](#))

Web UI: Disabling an HBAC Rule

1. Select **Policy** → **Host-Based Access Control** → **HBAC Rules**.
2. Select the HBAC rule you want to disable, and click **Disable**.




Figure 31.6. Disabling the allow_all HBAC Rule

Command Line: Disabling an HBAC Rule

Use the **ipa hbacrule-disable** command. For example, to disable the **allow_all** rule:

```
$ ipa hbacrule-disable allow_all
-----
Disabled HBAC rule "allow_all"
-----
```

31.3. ADDING HBAC SERVICE ENTRIES FOR CUSTOM HBAC SERVICES

Only the most common services and service groups are configured for HBAC rules by default. However, you can also configure any other pluggable authentication module (PAM) service as an HBAC service. This enables you to define the custom PAM service in an HBAC rule.

**NOTE**

Adding a service as an HBAC service is not the same as adding a service to the domain. Adding a service to the domain (described in [Section 16.1, “Adding and Editing Service Entries and Keytabs”](#)) makes the service a recognized resource available to other resources in the domain, but it does not enable you to use the service in HBAC rules.

To add an HBAC service entry, you can use:

- the IdM web UI (see [the section called “Web UI: Adding an HBAC Service Entry”](#))
- the command line (see [the section called “Command Line: Adding an HBAC Service Entry”](#))

Web UI: Adding an HBAC Service Entry

1. Select **Policy** → **Host-Based Access Control** → **HBAC Services**.
2. Click **Add** to add an HBAC service entry.
3. Enter a name for the service, and click **Add**.

Command Line: Adding an HBAC Service Entry

Use the `ipa hbacsvc-add` command. For example, to add an entry for the `tftp` service:

```
$ ipa hbacsvc-add tftp
-----
Added HBAC service "tftp"
-----
Service name: tftp
```

31.4. ADDING HBAC SERVICE GROUPS

HBAC service groups can simplify HBAC rules management: instead of adding individual services to an HBAC rule, you can add a whole service group.

To add an HBAC service group, you can use:

- the IdM web UI (see [the section called “Web UI: Adding an HBAC Service Group”](#))
- the command line (see [the section called “Command Line: Adding an HBAC Service Group”](#))

Web UI: Adding an HBAC Service Group

1. Select **Policy** → **Host-Based Access Control** → **HBAC Service Groups**.
2. Click **Add** to add an HBAC service group.
3. Enter a name for the service group, and click **Add and Edit**.
4. On the service group configuration page, click **Add** to add an HBAC service as a member of the group.




Figure 31.7. Adding HBAC Services to an HBAC Service Group

Command Line: Adding an HBAC Service Group

1. Use the **ipa hbacsvgroup-add** command to add an HBAC service group. For example, to add a group named **login**:

```
$ ipa hbacsvgroup-add
Service group name: login
-----
Added HBAC service group "login"
-----
Service group name: login
```

2. Use the **ipa hbacsvgroup-add-member** command to add an HBAC service as a member of the group. For example, to add the **sshd** service to the **login** group:

```
$ ipa hbacsvgroup-add-member
Service group name: login
[member HBAC service]: sshd
Service group name: login
Member HBAC service: sshd
-----
Number of members added 1
-----
```

CHAPTER 32. DEFINING SELINUX USER MAPS

Security-enhanced Linux (SELinux) sets rules over what system users can access processes, files, directories, and system settings. Both the system administrator and system applications can define *security contexts* that restrict or allow access from other applications.

As part of defining centralized security policies in the Identity Management domain, Identity Management provides a way to map IdM users to existing SELinux user contexts and grant or restrict access to clients and services within the IdM domain, per host, based on the defined SELinux policies.

32.1. ABOUT IDENTITY MANAGEMENT, SELINUX, AND MAPPING USERS

Identity Management does not create or modify the SELinux contexts on a system. Rather, it uses strings that might match existing contexts on the target hosts as the basis for mapping IdM users in the domain to SELinux users on a system.

Security-enhanced Linux defines kernel-level, mandatory access controls for how processes can interact with other resources on a system. Based on the expected behavior of processes on the system, and on their security implications, specific rules called policies are set. This is in contrast to higher-level discretionary access controls which are concerned primarily with file ownership and user identity. Every resource on a system is assigned a context. Resources include users, applications, files, and processes.

System users are associated with an SELinux *role*. The role is assigned both a multilayer security context (MLS) and a multi-category security context (MCS). The MLS and MCS contexts confine users so that they can only access certain processes, files, and operations on the system.

To get the full list of available SELinux users:

```
[root@server1 ~]# semanage user -l
```

SELinux User SELinux Roles	Labelling Prefix	MLS/ MCS Level	MLS/ MCS Range
guest_u guest_r	user	s0	s0
root	user	s0	s0-s0:c0.c1023
staff_r sysadm_r system_r	unconfined_r		
staff_u	user	s0	s0-s0:c0.c1023
staff_r sysadm_r system_r	unconfined_r		
sysadm_u	user	s0	s0-s0:c0.c1023
sysadm_r			
system_u	user	s0	s0-s0:c0.c1023
system_r unconfined_r			
unconfined_u	user	s0	s0-s0:c0.c1023
system_r unconfined_r			
user_u	user	s0	s0
user_r			
xguest_u	user	s0	s0
xguest_r			

For more information about SELinux in Red Hat Enterprise Linux, see [Red Hat Enterprise Linux 7 SELinux User's and Administrator's Guide](#).

SELinux users and policies function at the system level, not the network level. This means that SELinux users are configured independently on each system. While this is acceptable in many situations, as SELinux has common defined system users and SELinux-aware services define their own policies, it causes problems when remote users and systems access local resources. Remote users and services can be assigned a default guest context without knowing what their actual SELinux user and role should be.

Identity Management can integrate an identity domain with local SELinux services. Identity Management can map IdM users to configured SELinux roles *per host*, *per host group*, or based on an *HBAC rule*. Mapping SELinux and IdM users improves user administration:

- Remote users can be granted appropriate SELinux user contexts based on their IdM group assignments. This also allows administrators to consistently apply the same policies to the same users without having to create local accounts or reconfigure SELinux.
- The SELinux context associated with a user is centralized.
- SELinux policies can be planned and related to domain-wide security policies through settings like IdM host-based access control rules.
- Administrators gain environment-wide visibility and control over how users and systems are assigned in SELinux.

An SELinux user map defines two separate relationships that exist between three parts: the SELinux user for the system, an IdM user, and an IdM host. First, the SELinux user map defines a relationship between the SELinux user and the IdM host (the local or target system). Second, it defines a relationship between the SELinux user and the IdM user.

This arrangement allows administrators to set different SELinux users for the same IdM users, depending on which host they are accessing.

The core of an SELinux mapping rule is the SELinux system user. Each map is first associated with an SELinux user. The SELinux users which are available for mapping are configured in the IdM server, so there is a central and universal list. In this way, IdM defines a set of SELinux users it knows about and can associate with an IdM user upon login. By default, these are:

- `unconfined_u` (also used as a default for IdM users)
- `guest_u`
- `xguest_u`
- `user_u`
- `staff_u`

However, this default list can be modified and any *native* SELinux user (see [Section 32.1, “About Identity Management, SELinux, and Mapping Users”](#)) can be added or removed from the central IdM SELinux users list.

In the IdM server configuration, each SELinux user is configured with not only its user name but also its MLS and MCS range, *SELinux_user:MLS[:MCS]*. The IPA server uses this format to identify the SELinux user when configuring maps.

The IdM user and host configuration is very flexible. Users and hosts can be explicitly and individually assigned to an SELinux user map, or user groups or host groups can be explicitly assigned to the map.

You can also associate SELinux mapping rules with host-based access control rules to make administration easier, to avoid duplicating the same rule in two places, and to keep the rules synchronized. As long as the host-based access control rule defines a user and a host, you can use it for an SELinux user map. Host-based access control rules (described in [Chapter 31, Configuring Host-Based Access Control](#)) help integrate SELinux user maps with other access controls in IdM and can help limit or allow host-based user access for remote users, as well as define local security contexts.



NOTE

If a host-based access control rule is associated with an SELinux user map, the host-based access control rule cannot be deleted until it is removed from the SELinux user map configuration.

SELinux user maps work with the System Security Services Daemon (SSSD) and the **pam_selinux** module. When a remote user attempts to log into a machine, SSSD checks its IdM identity provider to collect the user information, including any SELinux maps. The PAM module then processes the user and assigns it the appropriate SELinux user context. SSSD caching enables the mapping to work offline.

32.2. CONFIGURING SELINUX USER MAP ORDER AND DEFAULTS

An SELinux user map is the association between an SELinux user on a client and an IdM user.

The available SELinux user map order is part of the IdM server configuration. The SELinux user map order is a list of the SELinux users, in an order from the most to the least confined. The SELinux user entry itself has this format:

```
SELinux_user:MLS[:MCS]
```

The individual user entries are separated with a dollar sign (\$).

Since there is no requirement on user entries to have an SELinux map, many entries might be unmapped. The IdM server configuration sets a default SELinux user, one of the users from the total SELinux map list, to use for unmapped IdM user entries. This way, even unmapped IdM users have a functional SELinux context. The default SELinux user for unmapped IdM user entries is **unconfined_u**, the default SELinux user for system users on Red Hat Enterprise Linux.

This configuration defines the map order of available system SELinux users. This does not define any IdM user SELinux policies. The IdM user - SELinux user map must be defined and then users are added to the map. For details, see [Section 32.3, “Mapping SELinux Users and IdM Users”](#).

32.2.1. In the Web UI

1. In the top menu, click the **IPA Server** main tab and the **Configuration** subtab.
2. Scroll to the bottom of the list of server configuration areas, to **SELINUX OPTIONS**.
3. Edit the SELinux user configuration, the **SELinux user map order**, the **Default SELinux user**, or both.

The screenshot shows the 'Configuration' subtab of the 'IPA Server' main tab. The 'SELINUX OPTIONS' section is highlighted with a red box. It contains the following fields:

- SELinux user map order:** A text input field containing the value `guest_u:s0$guest_u:s0$user_u:s0$staff_u:s0-s0:c0.c1023$unconfined_u:s0-s0:c0.c1023`.
- Default SELinux user:** A text input field containing the value `unconfined_u:s0-s0:c0.c1023`.

Other visible sections include 'Group Options' with fields for 'Group search fields' (containing 'cn,description') and 'Default group objectclasses' (containing 'top', 'groupofnames', 'nestedgroup', 'ipausergroup', 'ipaobject'). The 'Service Options' section shows 'Default PAC types' with checkboxes for 'MS-PAC' (checked), 'PAD' (unchecked), and 'nfs:NONE' (checked).

4. Click the **Update** link at the top of the page to save the changes.

32.2.2. In the CLI

To view the list of SELinux users, set in the IdM server configuration, which are available to be mapped:

```
[user1]@server ~]$ ipa config-show
...
SELinux user map order: guest_u:s0$guest_u:s0$user_u:s0$staff_u:s0-
s0:c0.c1023$unconfined_u:s0-s0:c0.c1023
Default SELinux user: unconfined_u:s0-s0:c0.c1023
```

To edit the SELinux user settings, use the **config-mod** command:

Example 32.1. List of SELinux Users

To edit the list of SELinux users to be available for mapping, use the **--ipaselinuusermaporder** option. The list orders the SELinux users from the most to the least confined, for example:

```
[user1@server ~]$ ipa config-mod --
ipaselinuusermaporder="unconfined_u:s0-
s0:c0.c1023$guest_u:s0$guest_u:s0$user_u:s0-s0:c0.c1023$staff_u:s0-
s0:c0.c1023"
```



NOTE

The default SELinux user, used for unmapped entries, must be included in the user map list or the edit operation fails. Likewise, if the default is edited, it must be changed to a user in the SELinux map list or the map list must be updated first.

Example 32.2. Default SELinux User

IdM users are not required to have a specific SELinux user mapped to their account. However, the local system still checks the IdM entry for an SELinux user to use for the IdM user account.

To modify the default SELinux user, use the `--ipaselinlinuxusermapdefault` option. For example:

```
[user1@server ~]$ ipa config-mod --ipaselinlinuxusermapdefault="guest_u:s0"
```

32.3. MAPPING SELINUX USERS AND IDM USERS

An SELinux map associates an SELinux user context on a local system with an IdM user, or users, within the domain. An SELinux map has three parts: the SELinux user context and an IdM user-host pairing. The IdM user-host pair can be defined in one of two ways: it can be set for explicit users, or user groups, on explicit hosts, or host groups; or it can be defined using a host-based access control rule.


32.3.1. In the Web UI

1. In the top menu, click the **Policy** main tab and the **SELinux User Mappings** subtab.
2. In the list of mappings, click the **Add** button to create a new map.

Rule name	SELinux User	Status	Description
☐ system_unconfined	unconfined_u:s0-s0:c0.c1023	✓ Enabled	
☐ test01	xguest_u:s0	— Disabled	Test_SELinux_User_Map_01

Showing 1 to 2 of 2 entries.

3. Enter the name for the map and the SELinux user. The format of the SELinux user has to be identical with how it appears in the IdM server configuration. SELinux users have the format `SELinux_user:MLS[:MCS]`.



Add SELinux User Map

Rule name *

SELinux User *


* Required field

- Click **Add and Edit** to add the IdM user information.
- To set a host-based access control rule, select the rule from the drop-down menu in the **General** area of the configuration. Using a host-based access control rule also introduces access controls on what hosts a remote user can use to access a target machine. **Only one host-based access control rule can be assigned.**



NOTE

The host-based access control rule must contain users and hosts, not just services.



Identity Policy Authentication Network Services IPA Server

Host Based Access Control Sudo SELinux User Maps Password Policies Kerberos Ticket Policy

SELinux User Maps > example-map

✓ SELinux User Map: example-map

Settings

Refresh Revert Save Actions

General

Rule name example-map

Description

SELinux User * staff_u:s0-s0:c0.c1023

HBAC Rule web_admin Undo

Alternatively, scroll down the **Users** and **Hosts** areas, and click the **Add** link to assign users, user groups, hosts, or host groups to the SELinux map.

SELinux User *

HBAC Rule

User

User category the rule applies to: ☐ Anyone ☒ Specified Users and Groups

Users	Delete	+ Add
☐ jsmith		Add
☐ User Groups		Add

Host

Host category the rule applies to: ☐ Any Host ☒ Specified Hosts and Groups

Hosts	Delete	+ Add
☐ test.example.com		Add
☐ Host Groups		Add

Select the users (or hosts or groups) on the left, click the right arrows button (>>) to move them to the **Prospective** column, and click the **Add** button to add them to the rule.

Add Users into SELinux User Map example-map

Filter available Users [Filter](#)

Available

Users
☐ admin
☐ jdoe
☐ jsmith
☒ pbrown
☒ agreeen

Prospective

Users
☐ Users

[Add](#) [Cancel](#)



NOTE

Only one option can be used: either a host-based access control rule can be given or the users and hosts can be set manually. Both options cannot be used at the same time.

- Click the **Update** link at the top to save the changes to the SELinux user map.

32.3.2. In the CLI

An SELinux map rule has three fundamental parts:

- The SELinux user: **--selinuxuser**
- The user or user groups which are associated with the SELinux user: **--users** or **--groups**
- The host or host groups which are associated with the SELinux user: **--hosts** or **--hostgroups**
- Alternatively, a host-based access control rule which specifies both hosts and users in it: **--hbacrule**

A rule can be created with all information at once using the **selinuxusermap-add** command. Users and hosts can be added to a rule after it is created by using the **selinuxusermap-add-user** and **selinuxusermap-add-host** commands, respectively.

Example 32.3. Creating a New SELinux Map

The **--selinuxuser** value must be the SELinux user name exactly as it appears in the IdM server configuration. SELinux users have the format *SELinux_user:MLS[:MCS]*.

The user, or user group; and the host, or host group must be specified for the SELinux mapping to be valid. The user, host, and group options can be used multiple times or can be used once with a comma-separated listed inside curly braces, for example **--option={val1,val2,val3}**.

```
[user1@server ~]$ ipa selinuxusermap-add --selinuxuser="xguest_u:s0"
selinux1
[user1@server ~]$ ipa selinuxusermap-add-user --users=user1 --
users=user2 --users=user3 selinux1
[user1@server ~]$ ipa selinuxusermap-add-host --hosts=server.example.com
--hosts=test.example.com selinux1
```

Example 32.4. Creating an SELinux Map with a Host-Based Access Control Rule

The **--hbacrule** value identifies the host-based access control rule to use for mapping. Using a host-based access control rule introduces access controls on what hosts a remote user can use to access a target machine, along with applying SELinux contexts after the remote user has logged into the target machine.

The access control rule must specify both users and hosts appropriately so that the SELinux map can construct the SELinux user, IdM user, and host triple.

Only one host-based access control rule can be specified.

```
[user1@server ~]$ ipa selinuxusermap-add --hbacrule=webserver --
selinuxuser="xguest_u:s0" selinux1
```

Host-based access control rules are described in [Chapter 31, Configuring Host-Based Access Control](#).

Example 32.5. Adding a User to an SELinux Map

Users and hosts can be added to an already existing map. This is done using a specific command, either **selinuxusermap-add-user** or **selinuxusermap-add-host**.

```
[user1@server ~]$ ipa selinuxusermap-add-user --users=user1 selinux1
```

If the **selinuxusermap-mod** command is used with the **--hbacrule** option to modify an already existing SELinux map, the new SELinux map overwrites the previous SELinux map.

Example 32.6. Removing a User from an SELinux Map

A specific user or host can be removed from an SELinux map by using either the **selinuxusermap-remove-host** or **selinuxusermap-remove-user** command. For example:

```
[user1@server ~]$ ipa selinuxusermap-remove-user --users=user2 selinux1
```

PART VII. ADMINISTRATION: MANAGING NETWORK SERVICES

CHAPTER 33. MANAGING DNS

An Identity Management server can be installed without integrated DNS services so that it uses an external DNS service or with DNS configured. See [Section 2.3, “Installing an IdM Server: Introduction”](#) and [Section 2.3.1, “Determining Whether to Use Integrated DNS”](#) for details.

If the DNS service is configured within the domain, IdM offers the administrator a significant amount of flexibility and control over DNS settings. For example, DNS entries for the domain, such as host entries, locations, or records, can be managed using native IdM tools, and clients can update their own DNS records dynamically.

Most documentation material and tutorials available for BIND version 9.9 are also applicable to IdM DNS, because majority of configuration options work in the same way in BIND and IdM. This chapter mostly focuses on notable differences between BIND and IdM.

33.1. BIND IN IDENTITY MANAGEMENT

IdM integrates BIND DNS server version 9.9 with an LDAP database used for data replication and with Kerberos for DNS update signing using the GSS-TSIG protocol [3]. This enables convenient DNS management using IdM tools and at the same time increases resiliency because IdM-integrated DNS servers support multi-master operations, allowing all IdM-integrated DNS servers to accept DNS updates from clients without having a single point of failure.

The default IdM DNS configuration is suitable for internal networks that are not accessible from the public Internet. If the IdM DNS server is accessible from the public Internet, Red Hat recommends applying the usual hardening applicable to the BIND service, described in the [Red Hat Enterprise Linux Networking Guide](#).




NOTE

It is not possible to run BIND integrated with IdM inside a **chroot** environment.

BIND integrated with IdM communicates with the Directory Server using the **bind-dyndb-ldap** plug-in. IdM creates a **dynamic-db** configuration section in the **/etc/named.conf** file for the BIND service, which configures the **bind-dyndb-ldap** plug-in for the BIND **named-pkcs11** service.

The most notable difference between standard BIND and IdM DNS is that IdM stores all DNS information as LDAP entries. Every domain name is represented as an LDAP entry, and every resource record is stored as an LDAP attribute of the LDAP entry. For example, the following **client1.example.com.** domain name contains three A records and one AAAA record:

```
dn: idnsname=client1,idnsname=example.com.,cn=dns,dc=idm,dc=example,dc=com
objectclass: top
objectclass: idnsrecord
idnsname: client1
Arecord: 192.0.2.1
Arecord: 192.0.2.2
Arecord: 192.0.2.3
AAAArecord: 2001:DB8::ABCD
```


**IMPORTANT**

To edit DNS data or BIND configuration, always use the IdM tools described in this chapter.

33.2. SUPPORTED DNS ZONE TYPES

IdM supports two DNS zone types: *master* and *forward* zones.

**NOTE**

This guide uses the BIND terminology for zone types which is different from the terminology used for Microsoft Windows DNS. Master zones in BIND serve the same purpose as *forward lookup zones* and *reverse lookup zones* in Microsoft Windows DNS. Forward zones in BIND serve the same purpose as *conditional forwarders* in Microsoft Windows DNS.

Master DNS zones

Master DNS zones contain authoritative DNS data and can accept dynamic DNS updates. This behavior is equivalent to the **type master** setting in standard BIND configuration. Master zones are managed using the **ipa dnszone-*** commands.

In compliance with standard DNS rules, every master zone must contain SOA and NS records. IdM generates these records automatically when the DNS zone is created, but the NS records must be manually copied to the parent zone to create proper delegation.

In accordance with standard BIND behavior, forwarding configuration specified for master zones only affects queries for names for which the server is not authoritative.

Example 33.1. Example Scenario for DNS Forwarding

The IdM server contains the **test.example.** master zone. This zone contains an NS delegation record for the **sub.test.example.** name. In addition, the **test.example.** zone is configured with the **192.0.2.254** forwarder IP address.

A client querying the name **nonexistent.test.example.** receives the **NXDomain** answer, and no forwarding occurs because the IdM server is authoritative for this name.

On the other hand, querying for the **sub.test.example.** name is forwarded to the configured forwarder **192.0.2.254** because the IdM server is not authoritative for this name.

Forward DNS zones

Forward DNS zones do not contain any authoritative data. All queries for names belonging to a forward DNS zone are forwarded to a specified forwarder. This behavior is equivalent to the **type forward** setting in standard BIND configuration. Forward zones are managed using the **ipa dnsforwardzone-*** commands.

33.3. DNS CONFIGURATION PRIORITIES

Many DNS configuration options can be configured on three different levels.

Zone-specific configuration

The level of configuration specific for a particular zone defined in IdM has the highest priority. Zone-specific configuration is managed using the **ipa dnszone-*** and **ipa dnsforwardzone-*** commands.

Global DNS configuration

If no zone-specific configuration is defined, IdM uses global DNS configuration stored in LDAP. Global DNS configuration is managed using the **ipa dnsconfig-*** commands. Settings defined in global DNS configuration are applied to all IdM DNS servers.

Configuration in `/etc/named.conf`

Configuration defined in the `/etc/named.conf` file on each IdM DNS server has the lowest priority. It is specific for each server and must be edited manually.

The `/etc/named.conf` file is usually only used to specify DNS forwarding to a local DNS cache; other options are managed using the commands for zone-specific and global DNS configuration mentioned above.

DNS options can be configured on multiple levels at once. In such cases, configuration with the highest priority takes precedence over configuration defined at lower levels.

33.4. MANAGING MASTER DNS ZONES

33.4.1. Adding and Removing Master DNS Zones

Adding Master DNS Zones in the Web UI

1. Open the **Network Services** tab, and select the **DNS** subtab, followed by the **DNS Zones** section.




Figure 33.1. Managing DNS Master Zones

2. To add a new master zone, click **Add** at the top of the list of all zones.




Figure 33.2. Adding a Master DNS Zone

3. Provide the zone name, and click **Add**.

Figure 33.3. Entering a New Master Zone

Adding Master DNS Zones from the Command Line

The **ipa dnszone-add** command adds a new zone to the DNS domain. Adding a new zone requires you to specify the name of the new subdomain. You can pass the subdomain name directly with the command:

```
$ ipa dnszone-add newserver.example.com
```

If you do not pass the name to **ipa dnszone-add**, the script prompts for it automatically.

The **ipa dnszone-add** command also accepts various command-line options. For a complete list of these options, run the **ipa dnszone-add --help** command.

Removing Master DNS Zones

To remove a master DNS zone in the web UI, in the list of all zones, select the check box by the zone name and click **Delete**.




Figure 33.4. Removing a Master DNS Zone

To remove a master DNS zone from the command line, use the **ipa dnszone-del** command. For example:

```
$ ipa dnszone-del server.example.com
```

33.4.2. Adding Additional Configuration for Master DNS Zones

IdM creates a new zone with certain default configuration, such as the refresh periods, transfer settings, or cache settings.

DNS Zone Configuration Attributes

The available zone settings are listed in [Table 33.1, “Zone Attributes”](#). Along with setting the actual information for the zone, the settings define how the DNS server handles the *start of authority* (SOA) record entries and how it updates its records from the DNS name server.

Table 33.1. Zone Attributes

Attribute	Command-Line Option	Description
Authoritative name server	--name-server	Sets the domain name of the master DNS name server, also known as SOA MNAME. By default, each IdM server advertises itself in the SOA MNAME field. Consequently, the value stored in LDAP using --name-server is ignored.
Administrator e-mail address	--admin-email	Sets the email address to use for the zone administrator. This defaults to the root account on the host.
SOA serial	--serial	Sets a serial number in the SOA record. Note that IdM sets the version number automatically and users are not expected to modify it.

Attribute	Command-Line Option	Description
SOA refresh	--refresh	Sets the interval, in seconds, for a secondary DNS server to wait before requesting updates from the primary DNS server.
SOA retry	--retry	Sets the time, in seconds, to wait before retrying a failed refresh operation.
SOA expire	--expire	Sets the time, in seconds, that a secondary DNS server will try to perform a refresh update before ending the operation attempt.
SOA minimum	--minimum	Sets the time to live (TTL) value in seconds for negative caching according to RFC 2308 .
SOA time to live	--ttl	Sets TTL in seconds for records at zone apex. In zone example.com , for instance, all records (A, NS, or SOA) under name example.com are configured, but no other domain names, like test.example.com , are affected.
Default time to live	--default-ttl	Sets the default time to live (TTL) value in seconds for negative caching for all values in a zone that never had an individual TTL value set before. Requires a restart of the named-pkcs11 service on all IdM DNS servers after changes to take effect.
BIND update policy	--update-policy	Sets the permissions allowed to clients in the DNS zone. See Dynamic Update Policies in the BIND 9 Administrator Reference Manual for more information on update policy syntax.
Dynamic update	--dynamic-update=TRUE FALSE	Enables dynamic updates to DNS records for clients. Note that if this is set to false, IdM client machines will not be able to add or update their IP address. See Section 33.5.1, “Enabling Dynamic DNS Updates” for more information.
Allow transfer	--allow-transfer=string	Gives a list of IP addresses or network names which are allowed to transfer the given zone, separated by semicolons (;). Zone transfers are disabled by default. The default --allow-transfer value is none .
Allow query	--allow-query	Gives a list of IP addresses or network names which are allowed to issue DNS queries, separated by semicolons (;).

Attribute	Command-Line Option	Description
Allow PTR sync	--allow-sync-ptr=1 0	Sets whether A or AAAA records (forward records) for the zone will be automatically synchronized with the PTR (reverse) records.
Zone forwarders	--forwarder=IP_address	Specifies a forwarder specifically configured for the DNS zone. This is separate from any global forwarders used in the IdM domain. To specify multiple forwarders, use the option multiple times.
Forward policy	--forward-policy=none only first	Specifies the forward policy. For information about the supported policies, see the section called “Forward Policies”

Editing the Zone Configuration in the Web UI

To manage DNS master zones from the web UI, open the **Network Services** tab, and select the **DNS** subtab, followed by the **DNS Zones** section.




Figure 33.5. DNS Master Zones Management

To edit an existing master zone in the **DNS Zones** section:

1. Click on the zone name in the list of all zones to open the DNS zone page.




Figure 33.6. Editing a Master Zone

- 2. Click **Settings**, and then change the zone configuration as required.




Figure 33.7. The Settings Tab in the Master Zone Edit Page

For information about the available settings, see [Table 33.1, “Zone Attributes”](#).

- 3. Click **Save** to confirm the new configuration.



NOTE

If you are changing the default time to live (TTL) of a zone, restart the **named-pkcs11** service on all IdM DNS servers to make the changes take effect. All other settings are automatically activated immediately.

Editing the Zone Configuration from the Command Line

To modify an existing master DNS zone from the command line, use the **ipa dnszone-mod** command. For information about the available settings, see [Table 33.1, “Zone Attributes”](#).

If an attribute does not exist in the DNS zone entry, the **ipa dnszone-mod** command adds the attribute. If the attribute exists, the command overwrites the current value with the specified value.

For detailed information about **ipa dnszone-mod** and its options, run the **ipa dnszone-mod --help** command.


**NOTE**

If you are changing the default time to live (TTL) of a zone, restart the **named-pkcs11** service on all IdM DNS servers to make the changes take effect. All other settings are automatically activated immediately.

33.4.3. Enabling Zone Transfers

Name servers maintain authoritative data for the zones; changes made to the zones must be sent to and distributed among the name servers for the DNS domain. A *zone transfer* copies all resource records from one name server to another.

IdM supports zone transfers according to the [RFC 5936](#) (AXFR) and [RFC 1995](#) (IXFR) standards.

**IMPORTANT**

The IdM-integrated DNS is multi-master. SOA serial numbers in IdM zones are not synchronized between IdM servers. For this reason, configure DNS slave servers to only use one IdM master server. This prevents zone transfer failures caused by non-synchronized SOA serial numbers.

Enabling Zone Transfers in the UI

Open the DNS zone page, as described in [the section called “Editing the Zone Configuration in the Web UI”](#), and switch to the **Settings** tab.

Under **Allow transfer**, specify the name servers to which the zone records will be transferred.

Allow transfer	192.0.2.1	Undo
	198.51.100.1	Undo
	203.0.113.1	Undo
	Add	Undo All

Figure 33.8. Enabling Zone Transfers

Click **Save** at the top of the DNS zone page to confirm the new configuration.

Enabling Zone Transfers from the Command Line

To enable zone transfers from the command line, add the **--allow-transfer** option to the **ipa dnszone-mod** command. Specify the list of name servers to which the zone records will be transferred using **--allow-transfer**. For example:

```
[user@server ~]$ ipa dnszone-mod --allow-transfer=192.0.2.1;198.51.100.1;203.0.113.1 example.com
```


Once zone transfers are enabled in the **bind** service, IdM DNS zones can be transferred, by name, by clients such as the **dig** utility:

```
[root@server ~]# dig @ipa-server zone_name AXFR
```

33.4.4. Adding Records to DNS Zones

IdM supports many different record types. The following four are used most frequently:

A

This is a basic map for a host name and an ordinary IPv4 address. The record name of an A record is a host name, such as **www**. The **IP Address** value of an A record is a standard IPv4 address, such as **192.0.2.1**.

For more information about A records, see [RFC 1035](#).

AAAA

This is a basic map for a host name and an IPv6 address. The record name of an AAAA record is a host name, such as **www**. The **IP Address** value is a standard hexadecimal IPv6 address, such as **2001:DB8::1111**.

For more information about AAAA records, see [RFC 3596](#).

SRV


Service (SRV) resource records map service names to the DNS name of the server that is providing that particular service. For example, this record type can map a service like an LDAP directory to the server which manages it.

The record name of an SRV record has the format **_service._protocol**, such as **_ldap._tcp**. The configuration options for SRV records include priority, weight, port number, and host name for the target service.

For more information about SRV records, see [RFC 2782](#).

PTR

A pointer record type (PTR) record adds a reverse DNS record, which maps an IP address to a domain name.



NOTE

All reverse DNS lookups for IPv4 addresses use reverse entries that are defined in the **in-addr.arpa** domain. The reverse address, in human-readable form, is the exact reverse of the regular IP address, with the **in-addr.arpa** domain appended to it. For example, for the network address **192.0.2.0/24**, the reverse zone is **2.0.192.in-addr.arpa**.

The record name of a PTR record must be in the standard format specified in [RFC 1035](#), extended in [RFC 2317](#), and [RFC 3596](#). The host name value must be a canonical host name of the host for which you want to create the record. For more information, see [Example 33.8, “PTR Record”](#).



NOTE

Reverse zones can also be configured for IPv6 addresses, with zones in the **.ip6.arpa.** domain. For more information about IPv6 reverse zones, see [RFC 3596](#).

When adding DNS resource records, note that many of the records require different data. For example, a CNAME record requires a host name, while an A record requires an IP address. In the web UI, the fields in the form for adding a new record are updated automatically to reflect what data is required for the currently selected type of record.

DNS Wildcard Support

IdM supports the special record ***** in a DNS zone as wildcard.

Example 33.2. Demonstrating DNS Wildcard Results

1. Configure the following in your DNS zone *example.com*:
 - A wildcard A record ***.example.com**.
 - An MX record for **mail.example.com**, but no A record for this host.
 - No record for **demo.example.com**.
2. Query existing and non-existent DNS records and types. You will receive the following results:

```
# host -t MX mail.example.com.
mail.example.com mail is handled by 10 server.example.com.

# host -t MX demo.example.com.
demo.example.com. has no MX record.

# host -t A mail.example.com.
mail.example.com has no A record

# host -t A demo.example.com.
random.example.com has address 192.168.1.1
```

For more details, see [RFC1034](#).

Adding DNS Resource Records from the Web UI

1. Open the DNS zone page, as described in [the section called “Editing the Zone Configuration in the Web UI”](#).
2. In the **DNS Resource Records** section, click **Add** to add a new record.




Figure 33.9. Adding a New DNS Resource Record

3. Select the type of record to create and fill out the other fields as required.

Add DNS Resource Record

Record name *

Record Type

Hostname *

* Required field

Figure 33.10. Defining a New DNS Resource Record

4. Click **Add** to confirm the new record.

Adding DNS Resource Records from the Command Line

To add a DNS resource record of any type from the command line, use the **ipa dnsrecord-add** command. The command follows this syntax:

```
$ ipa dnsrecord-add zone_name record_name --record_type_option=data
```

The *zone_name* is the name of the DNS zone to which the record is being added. The *record_name* is an identifier for the new DNS resource record.

Table 33.2, “Common **ipa dnsrecord-add** Options” lists options for the most common

resource record types: A (IPv4), AAAA (IPv6), SRV, and PTR. Lists of entries can be set by using the option multiple times with the same command invocation or, in Bash, by listing the options in a comma-separated list inside curly braces, such as **--option={val1,val2,val3}**.

For more detailed information on how to use **ipa dnsrecord-add** and which DNS record types are supported by IdM, run the **ipa dnsrecord-add --help** command.

Table 33.2. Common ipa dnsrecord-add Options

General Record Options	
Option	Description
--ttl=number	Sets the time to live for the record.
--structured	Parses the raw DNS records and returns them in a structured format.

"A" Record Options	
Option	Description
--a-rec=ARECORD	Passes a list of A records.
--a-ip-address=string	Gives the IP address for the record.

"AAAA" Record Options	
Option	Description
--aaaa-rec=AAAAARECORD	Passes a list of AAAA (IPv6) records.
--aaaa-ip-address=string	Gives the IPv6 address for the record.

"PTR" Record Options	
Option	Description
--ptr-rec=PTRRECORD	Passes a list of PTR records.
--ptr-hostname=string	Gives the host name for the record.

"SRV" Record Options	
Option	Description
--srv-rec=SRVRECORD	Passes a list of SRV records.
--srv-priority=number	Sets the priority of the record. There can be multiple SRV records for a service type. The priority (0 - 65535) sets the rank of the record; the lower the number, the higher the priority. A service has to use the record with the highest priority first.
--srv-weight=number	Sets the weight of the record. This helps determine the order of SRV records with the same priority. The set weights should add up to 100, representing the probability (in percentages) that a particular record is used.
--srv-port=number	Gives the port for the service on the target host.
--srv-target=string	Gives the domain name of the target host. This can be a single period (.) if the service is not available in the domain.

33.4.5. Examples of Adding or Modifying DNS Resource Records from the Command Line

Example 33.3. Adding a IPv4 Record

The following example creates the record **www.example.com** with the IP address **192.0.2.123**.

```
$ ipa dnsrecord-add example.com www --a-rec 192.0.2.123
```

Example 33.4. Adding a IPv4 Wildcard Record

The following example creates a wildcard A record with the IP address **192.0.2.123**:

```
$ ipa dnsrecord-add example.com "*" --a-rec 192.0.2.123
```

Example 33.5. Modifying a IPv4 Record

When creating a record, the option to specify the A record value is **--a-record**. However, when modifying an A record, the **--a-record** option is used to specify the current value for the A record. The new value is set with the **--a-ip-address** option.

```
$ ipa dnsrecord-mod example.com www --a-rec 192.0.2.123 --a-ip-address 192.0.2.1
```

Example 33.6. Adding an IPv6 Record

The following example creates the record **www.example.com** with the IP address **2001:db8::1231:5675**.

```
$ ipa dnsrecord-add example.com www --aaaa-rec 2001:db8::1231:5675
```

Example 33.7. Adding an SRV Record

In the following example, **_ldap._tcp** defines the service type and the connection protocol for the SRV record. The **--srv-rec** option defines the priority, weight, port, and target values.

For example:

```
[root@server ~]# ipa dnsrecord-add server.example.com _ldap._tcp --srv-rec="0 51 389 server1.example.com."
[root@server ~]# ipa dnsrecord-add server.example.com _ldap._tcp --srv-rec="1 49 389 server2.example.com."
```

The weight values (**51** and **49** in this example) add up to 100 and represent the probability (in percentages) that a particular record is used.

Example 33.8. PTR Record

When adding the reverse DNS record, the zone name used with the **ipa dnsrecord-add** command is *reverse*, compared to the usage for adding other DNS records:

```
$ ipa dnsrecord-add reverseNetworkIpAddress hostIpAddress --ptr-rec FQDN
```

Typically, *hostIpAddress* is the last octet of the IP address in a given network.

For example, this adds a PTR record for **server4.example.com** with IPv4 address 192.0.2.4:

```
$ ipa dnsrecord-add 2.0.192.in-addr.arpa 4 --ptr-rec
server4.example.com.
```

The next example adds a reverse DNS entry to the **0.0.0.0.0.0.0.0.8.b.d.0.1.0.0.2.ip6.arpa**. IPv6 reverse zone for the host **server2.example.com** with the IP address **2001:DB8::1111**:

```
$ ipa dnsrecord-add 0.0.0.0.0.0.0.0.8.b.d.0.1.0.0.2.ip6.arpa.
1.1.1.0.0.0.0.0.0.0.0.0.0.0.0 --ptr-rec server2.example.com.
```

33.4.6. Deleting Records from DNS Zones

Deleting Records in the Web UI

To delete only a specific record type from the resource record:

1. Open the DNS zone page, as described in [the section called “Editing the Zone Configuration in the Web UI”](#).
2. In the **DNS Resource Records** section, click the name of the resource record.




Figure 33.11. Selecting a DNS Resource Record

3. Select the check box by the name of the record type to delete.




Figure 33.12. Deleting a DNS Resource Record

After this, only the selected record type is deleted; the other configuration is left intact.

To delete all records for the resource in the zone:

1. Open the DNS zone page, as described in [the section called “Editing the Zone Configuration in the Web UI”](#).

2. In the **DNS Resource Records** section, select the check box by the name of the resource record to delete, and then click **Delete** at the top of the list of zone records.




Figure 33.13. Deleting an Entire Resource Record

After this, the entire resource record is deleted.

Deleting Records from the Command Line

To remove records from a zone, use the **ipa dnsrecord-del** command and add the **-recordType-rec** option together with the record value.

For example, to remove the A type record:

```
$ ipa dnsrecord-del example.com www --a-rec 192.0.2.1
```

If you run **ipa dnsrecord-del** without any options, the command prompts for information about the record to delete. Note that passing the **--del-all** option with the command removes all associated records for the zone.

For detailed information on how to use **ipa dnsrecord-del** and a complete list of options accepted by the command, run the **ipa dnsrecord-del --help** command.

33.4.7. Disabling and Enabling Zones

IdM allows the administrator to disable and enable DNS zones. While deleting a DNS zone, described in [the section called “Removing Master DNS Zones”](#), completely removes the zone entry and all the associated configuration, disabling the zone removes it from activity without permanently removing the zone from IdM. A disabled zone can also be enabled again.

Disabling and Enabling Zones in the Web UI

To manage DNS zones from the Web UI, open the **Network Services** tab, and select the **DNS** subtab, followed by the **DNS Zones** section.




Figure 33.14. Managing DNS Zones

To disable a zone, select the check box next to the zone name and click **Disable**.



Figure 33.15. Disabling a DNS Zone

Similarly, to enable a disabled zone, select the check box next to the zone name and click **Enable**.

Disabling and Enabling DNS Zones from the Command Line

To disable a DNS zone from the command line, use the **ipa dnszone-disable** command. For example:

```
[user@server ~]$ ipa dnszone-disable zone.example.com
-----
Disabled DNS zone "example.com"
-----
```

To re-enable a disabled zone, use the **ipa dnszone-enable** command.

33.5. MANAGING DYNAMIC DNS UPDATES

33.5.1. Enabling Dynamic DNS Updates

Dynamic DNS updates are disabled by default for new DNS zones in IdM. With dynamic updates disabled, the **ipa-client-install** script cannot add a DNS record pointing to the new client.

**NOTE**

Enabling dynamic updates can potentially pose a security risk. However, if enabling dynamic updates is acceptable in your environment, you can do it to make client installations easier.

Enabling dynamic updates requires the following:

- The DNS zone must be configured to allow dynamic updates
- The local clients must be configured to send dynamic updates

33.5.1.1. Configuring the DNS Zone to Allow Dynamic Updates

Enabling Dynamic DNS Updates in the Web UI

1. Open the **Network Services** tab, and select the **DNS** subtab, followed by the **DNS Zones** section.




Figure 33.16. DNS Zone Management

2. Click on the zone name in the list of all zones to open the DNS zone page.




Figure 33.17. Editing a Master Zone

3. Click **Settings** to switch to the DNS zone settings tab.




Figure 33.18. The Settings Tab in the Master Zone Edit Page

4. Scroll down to the **Dynamic update** field, and set the value to **True**.




Figure 33.19. Enabling Dynamic DNS Updates

5. Click **Save** at the top of the page to confirm the new configuration.

Enabling Dynamic DNS Updates from the Command Line

To allow dynamic updates to the DNS zones from the command line, use the **ipa dnszone-mod** command with the **--dynamic-update=TRUE** option. For example:

```
[user@server ~]$ ipa dnszone-mod server.example.com --dynamic-update=TRUE
```

33.5.1.2. Configuring the Clients to Send Dynamic Updates

Clients are automatically set up to send DNS updates when they are enrolled in the domain, by using the **--enable-dns-updates** option with the **ipa-client-install** script.

```
[root@client ~]# ipa-client-install --enable-dns-updates
```

The DNS zone has a time to live (TTL) value set for records within its SOA configuration. However, the TTL for the dynamic updates is managed on the local system by the System Security Service Daemon (SSSD). To change the TTL value for the dynamic updates, edit

the SSSD file to set a value; the default is 1200 seconds.

1. Open the SSSD configuration file.

```
[root@server ~]# vim /etc/sss/sss.conf
```

2. Find the domain section for the IdM domain.

```
[domain/ipa.example.com]
```

3. If dynamic updates have not been enabled for the client, then set the **dyndns_update** value to true.

```
dyndns_updates = true
```

4. Add or edit the **dyndns_ttl** parameter to set the value, in seconds.

```
dyndns_ttl = 2400
```

33.5.2. Synchronizing A/AAAA and PTR Records


A and AAAA records are configured separately from PTR records in reverse zones. Because these records are configured independently, it is possible for A/AAAA records to exist without corresponding PTR records, and vice versa.

There are some DNS setting requirements for PTR synchronization to work:

- Both forward and reverse zones must be managed by the IdM server.
- Both zones must have dynamic updates enabled.

Enabling dynamic updates is covered in [Section 33.5.1, “Enabling Dynamic DNS Updates”](#).

- PTR synchronization must be enabled for the master forward zone, not for the reverse zone.
- The PTR record will be updated only if the name of the requesting client matches the name in the PTR record.



IMPORTANT

Changes made through the IdM web UI, through the IdM command-line tools, or by editing the LDAP entry directly **do not** update the PTR record. Only changes made by the DNS service itself trigger PTR record synchronization.



WARNING

A client system can update its own IP address. This means that a compromised client can be used to overwrite PTR records by changing its IP address.

Configuring PTR Record Synchronization in the Web UI

Note that PTR record synchronization must be configured on the zone where A or AAAA records are stored, not on the reverse DNS zone where PTR records are located.

- 1. Open the **Network Services** tab, and select the **DNS** subtab, followed by the **DNS Zones** section.




Figure 33.20. DNS Zone Management

- 2. Click on the zone name in the list of all zones to open the DNS zone page.




Figure 33.21. Editing a DNS Zone

- 3. Click **Settings** to switch to the DNS zone settings tab.




Figure 33.22. The Settings Tab in the Master Zone Edit Page

4. Select the **Allow PTR sync** check box.




Figure 33.23. Enabling PTR Synchronization

5. Click **Save** at the top of the page to confirm the new configuration.

Configuring PTR Record Synchronization from the Command Line

Note that PTR record synchronization must be configured on the zone where A or AAAA records are stored, not on the reverse DNS zone where PTR records are located.

To configure a DNS zone to allow its forward and reverse entries to be synchronized automatically, set the `--allow-sync-ptr` option to `1` when the zone is created or when it is edited. For example, using the `ipa dnszone-mod` command when editing an existing zone:

```
[user@server ~]$ ipa dnszone-mod --allow-sync-ptr=1 zone.example.com
```

The default `--allow-sync-ptr` value is `0`, which disables synchronization.

33.5.3. Updating DNS Dynamic Update Policies

DNS domains maintained by IdM servers can accept a DNS dynamic update according to RFC 3007^[4].

The rules that determine which records can be modified by a specific client follow the same syntax as the `update-policy` statement in the `/etc/named.conf` file. For more information on dynamic update policies, see the [BIND 9 documentation](#).

Note that if dynamic DNS updates are disabled for the DNS zone, all DNS updates are declined without reflecting the dynamic update policy statement. For information on enabling dynamic DNS updates, see [Section 33.5.1, “Enabling Dynamic DNS Updates”](#).

Updating DNS Update Policies in the Web UI

- 1. Open the **Network Services** tab, and select the **DNS** subtab, followed by the **DNS Zones** section.




Figure 33.24. DNS Zone Management

- 2. Click on the zone name in the list of all zones to open the DNS zone page.




Figure 33.25. Editing a DNS Zone

- 3. Click **Settings** to switch to the DNS zone settings tab.




Figure 33.26. The Settings Tab in the Master Zone Edit Page

4. Set the required update policies in a semi-colon separated list in the **BIND update policy** text box.




Figure 33.27. DNS Update Policy Settings

5. Click **Save** at the top of the DNS zone page to confirm the new configuration.

Updating DNS Update Policies from the Command Line

To set the DNS update policy from the command line, use the **--update-policy** option and add the access control rule in a statement after the option. For example:

```
$ ipa dnszone-mod zone.example.com --update-policy "grant EXAMPLE.COM
krb5-self * A; grant EXAMPLE.COM krb5-self * AAAA; grant EXAMPLE.COM
krb5-self * SSHFP;"
```

33.6. MANAGING DNS FORWARDING

DNS forwarding affects how DNS queries are answered. By default, the BIND service integrated with IdM is configured to act as both an authoritative and recursive DNS server.

When a DNS client queries a name belonging to a DNS zone for which the IdM server is authoritative, BIND replies with data contained in the configured zone. Authoritative data always takes precedence over any other data.

When a DNS client queries a name for which the IdM server is not authoritative, BIND attempts to resolve the query using other DNS servers. If no forwarders are defined, BIND asks the root servers on the Internet and uses recursive resolution algorithm to answer the DNS query.

In some cases, it is not desirable to let BIND contact other DNS servers directly and perform the recursion based on data available on the Internet. These cases include:

- *Split DNS* configuration, also known as *DNS views* configuration, where DNS servers return different answers to different clients. Split DNS configuration is typical for environments where some DNS names are available inside the company network, but not from the outside.
- Configurations where a firewall restricts access to DNS on the Internet.
- Configurations with centralized filtering or logging on the DNS level.
- Configurations with forwarding to a local DNS cache, which helps optimize network traffic.

In such configurations, BIND does not use full recursion on the public Internet. Instead, it uses another DNS server, a so-called *forwarder*, to resolve the query. When BIND is configured to use a forwarder, queries and answers are forwarded back and forth between the IdM server and the forwarder, and the IdM server acts as the DNS cache for non-authoritative data.

Forward Policies

IdM supports the *first* and *only* standard BIND forward policies, as well as the *none* IdM-specific forward policy.

Forward first (default)

DNS queries are forwarded to the configured forwarder. If a query fails because of a server error or timeout, BIND falls back to the recursive resolution using servers on the Internet. The forward first policy is the default policy. It is suitable for traffic optimization.

Forward only

DNS queries are forwarded to the configured forwarder. If a query fails because of a server error or timeout, BIND returns an error to the client. The forward only policy is recommended for environments with split DNS configuration.

None: Forwarding disabled

DNS queries are not forwarded. Disabling forwarding is only useful as a zone-specific override for global forwarding configuration. This option is the IdM equivalent of specifying an empty list of forwarders in BIND configuration.

Forwarding Does Not Combine Data from IdM and Other DNS Servers

Forwarding cannot be used to combine data in IdM with data from other DNS servers. You can only forward queries for specific subzones of the master zone in IdM DNS: see [the section called “Zone Delegation in IdM DNS Master Zone”](#).

By default, the BIND service does not forward queries to another server if the queried DNS name belongs to a zone for which the IdM server is authoritative. In such a situation, if the queried DNS name cannot be found in the IdM database, the **NXDOMAIN** answer is returned. Forwarding is not used.

Example 33.9. Example Scenario

The IdM server is authoritative for the **test.example.** DNS zone. BIND is configured to forward queries to the DNS server with the **192.0.2.254** IP address.

When a client sends a query for the **nonexistent.test.example.** DNS name, BIND detects that the IdM server is authoritative for the **test.example.** zone and does not forward the query to the **192.0.2.254.** server. As a result, the DNS client receives the **NXDomain** answer, informing the user that the queried domain does not exist.

Zone Delegation in IdM DNS Master Zone

It is possible to forward queries for specific subzones of a master zone in IdM DNS. For example, if the IdM DNS handles the zone **idm.example.com**, you can delegate the authority for the **sub_zone1.idm.example.com** subzone to a different DNS server. To achieve this behavior, you need to use forwarding, as described above, along with a nameserver record which delegates the subzone to a different DNS server. In the following example, **sub_zone1** is the subzone, and **192.0.2.1** is the IP address of the DNS server the subzone is delegated to:

```
$ ipa dnsrecord-add idm.example.com. sub_zone1 --ns-rec=192.0.2.1
```

Adding the forward zone then looks like this:

```
$ ipa dnsforwardzone-add sub_zone1.idm.example.com. --forwarder 192.0.2.1
```

33.6.1. Configuring Global Forwarders

Global forwarders are DNS servers used for resolving all DNS queries for which an IdM server is not authoritative, as described in [Section 33.6, “Managing DNS Forwarding”](#).

The administrator can configure IP addresses and forward policies for global forwarding in the following two ways:

Using the `ipa dnsconfig-mod` command or the IdM web UI

Configuration set using these native IdM tools is immediately applied to all IdM DNS servers. As explained in [Section 33.3, “DNS Configuration Priorities”](#), global DNS configuration has higher priority than local configuration defined in the **/etc/named.conf** files.


By editing the **/etc/named.conf** file

Manually editing the **/etc/named.conf** on every IdM DNS server allows using a different global forwarder and policy on each of the servers. Note that the BIND service must be restarted after changing **/etc/named.conf**.

Configuring Forwarders in the Web UI

To define the DNS global configuration in the IdM web UI:

1. Click the **Network Services** tab, and select the **DNS** subtab, followed by the **DNS Global Configuration** section.
2. To add a new global forwarder, click **Add** and enter the IP address. To define a new forward policy, select it from the list of available policies.



The screenshot shows the 'DNS Global Configuration' web interface. At the top, there are tabs for 'Identity', 'Policy', 'Authentication', 'Network Services' (selected), and 'IPA Server'. Below the tabs, there's a sub-tab for 'DNS'. The main heading is 'DNS Global Configuration'. There are three buttons: 'Refresh', 'Revert', and 'Save'. Under the 'Options' section, there's a checkbox for 'Allow PTR sync'. Below that, there's a list of 'Global forwarders' with two entries: '192.0.2.253' and '192.0.2.254', each with an 'Undo' button. There are also 'Add' and 'Undo All' buttons. At the bottom, there's a 'Forward policy' section with three radio buttons: 'Forward first' (selected), 'Forward only', and 'Forwarding disabled'.

Figure 33.28. Editing Global DNS Configuration in the Web UI

3. Click **Save** to confirm the new configuration.

Configuring Forwarders from the Command Line

To set a global list of forwarders from the command line, use the **ipa dnsconfig-mod** command. It edits the DNS global configuration by editing the LDAP data. The **ipa dnsconfig-mod** command and its options affect all IdM DNS servers at once and override any local configuration.

For example, to edit the list of global forwarders using **ipa dnsconfig-mod**:

```
[user@server ~]$ ipa dnsconfig-mod --forwarder=192.0.2.254
Global forwarders: 192.0.2.254
```

33.6.2. Configuring Forward Zones

Forward zones do not contain any authoritative data and instruct the name server to only forward queries for names belonging into a particular zone to a configured forwarder.

IMPORTANT

Do not use forward zones unless absolutely required. Limit their use to overriding global forwarding configuration. In most cases, **it is sufficient to only configure global forwarding**, described in [Section 33.6.1, “Configuring Global Forwarders”](#), and forward zones are not necessary.

Forward zones are a non-standard solution, and using them can lead to unexpected and problematic behavior. When creating a new DNS zone, Red Hat recommends to always use standard DNS delegation using NS records and to avoid forward zones.

For information on the supported forward policies, see [the section called “Forward Policies”](#).

For further information about the BIND service, see the [Red Hat Enterprise Linux Networking Guide](#), the BIND 9 Administrator Reference Manual included in the `/usr/share/doc/bind-version_number/` directory, or external sources^[5].

Configuring Forward Zones in the Web UI

To manage forward zones in the web UI, click the **Network Services** tab, and select the **DNS** subtab, followed by the **DNS Forward Zones** section.




Figure 33.29. Managing DNS Forward Zones

In the **DNS Forward Zones** section, the administrator can handle all required operations regarding forward zones: show current list of forward zones, add a new forward zone, delete a forward zone, display a forward zone, allow to modify forwarders and forward policy per a forward zone, and disable or enable a forward zone.

Configuring Forward Zones from the Command Line

To manage forward zones from the command line, use the **ipa dnsforwardzone-*** commands described below.

NOTE

The **ipa dnsforwardzone-*** commands behave consistently with the **ipa dnszone-*** commands used to manage master zones.

The **ipa dnsforwardzone-*** commands accept several options; notably, the **--forwarder**, **--forward-policy**, and **--name-from-ip** options. For detailed information about the available options, see [Table 33.1, “Zone Attributes”](#) or run the commands with the **--help** option added, for example:

```
ipa dnsforwardzone-add --help
```

Adding Forward Zones

Use the **dnsforwardzone-add** command to add a new forward zone. It is required to specify at least one forwarder if the forward policy is not set to **none**.

```
[user@server ~]$ ipa dnsforwardzone-add zone.test. --
forwarder=172.16.0.1 --forwarder=172.16.0.2 --forward-policy=first

Zone name: zone.test.
Zone forwarders: 172.16.0.1, 172.16.0.2
Forward policy: first
```

Modifying Forward Zones

Use the **dnsforwardzone-mod** command to modify a forward zone. It is required to specify at least one forwarder if the forward policy is not **none**. Modifications can be performed in several ways.

```
[user@server ~]$ ipa dnsforwardzone-mod zone.test. --
forwarder=172.16.0.3

Zone name: zone.test.
Zone forwarders: 172.16.0.3
Forward policy: first
```

```
[user@server ~]$ ipa dnsforwardzone-mod zone.test. --forward-policy=only

Zone name: zone.test.
Zone forwarders: 172.16.0.3
Forward policy: only
```

Showing Forward Zones

Use the **dnsforwardzone-show** command to display information about a specified forward zone.

```
[user@server ~]$ ipa dnsforwardzone-show zone.test.

Zone name: zone.test.
Zone forwarders: 172.16.0.5
Forward policy: first
```

Finding Forward Zones

Use the **dnsforwardzone-find** command to locate a specified forward zone.

```
[user@server ~]$ ipa dnsforwardzone-find zone.test.

Zone name: zone.test.
Zone forwarders: 172.16.0.3
Forward policy: first
```

```
-----
Number of entries returned 1
-----
```

Deleting Forward Zones

Use the **dnsforwardzone-del** command to delete specified forward zones.

```
[user@server ~]$ ipa dnsforwardzone-del zone.test.
```

```
-----
Deleted forward DNS zone "zone.test."
-----
```

Enabling and Disabling Forward Zones

Use **dnsforwardzone-enable** and **dnsforwardzone-disable** commands to enable and disable forward zones. Note that forward zones are enabled by default.

```
[user@server ~]$ ipa dnsforwardzone-enable zone.test.
```

```
-----
Enabled forward DNS zone "zone.test."
-----
```

```
[user@server ~]$ ipa dnsforwardzone-disable zone.test.
```

```
-----
Disabled forward DNS zone "zone.test."
-----
```

Adding and Removing Permissions

Use **dnsforwardzone-add-permission** and **dnsforwardzone-remove-permission** commands to add or remove system permissions.

```
[user@server ~]$ ipa dnsforwardzone-add-permission zone.test.
```

```
-----
Added system permission "Manage DNS zone zone.test."
-----
```

```
Manage DNS zone zone.test.
```

```
[user@server ~]$ ipa dnsforwardzone-remove-permission zone.test.
```

```
-----
Removed system permission "Manage DNS zone zone.test."
-----
```

```
Manage DNS zone zone.test.
```

33.7. MANAGING REVERSE DNS ZONES

A reverse DNS zone can be identified in the following two ways:

- By the zone name, in the format ***reverse_ipv4_address.in-addr.arpa*** or ***reverse_ipv6_address.ip6.arpa***.

The reverse IP address is created by reversing the order of the components of the IP address. For example, if the IPv4 network is **192.0.2.0/24**, the reverse zone name is **2.0.192.in-addr.arpa.** (with the trailing period).

- By the network address, in the format ***network_ip_address/subnet_mask_bit_count***

To create the reverse zone by its IP network, set the network information to the (forward-style) IP address, with the subnet mask bit count. The bit count must be a multiple of eight for IPv4 addresses or a multiple of four for IPv6 addresses.

Adding a Reverse DNS Zone in the Web UI

1. Open the **Network Services** tab, and select the **DNS** subtab, followed by the **DNS Zones** section.




Figure 33.30. DNS Zone Management

2. Click **Add** at the top of the list of all zones.





Figure 33.31. Adding a Reverse DNS Zone

3. Fill in the zone name or the reverse zone IP network.
 - a. For example, to add a reverse DNS zone by the zone name:



Add DNS Zone [X]

☒ **Zone name ***

☐ **Reverse zone**


IP network

* Required field

[Add] [Add and Add Another] [Add and Edit] [Cancel]

Figure 33.32. Creating a Reverse Zone by Name

- b. Alternatively, to add a reverse DNS zone by the reverse zone IP network:



Add DNS Zone [X]

☐ **Zone name**

☒ **Reverse zone ***

IP network

* Required field

[Add] [Add and Add Another] [Add and Edit] [Cancel]

Figure 33.33. Creating a Reverse Zone by IP Network

The validator for the **Reverse zone IP network** field warns you about an invalid network address during typing. The warning will disappear once you enter the full network address.

4. Click **Add** to confirm the new reverse zone.

Adding a Reverse DNS Zone from the Command Line

To create a reverse DNS zone from the command line, use the **ipa dnszone-add** command.

For example, to create the reverse zone by the zone name:

```
[user@server]$ ipa dnszone-add 2.0.192.in-addr.arpa.
```

Alternatively, to create the reverse zone by the IP network:


```
[user@server ~]$ ipa dnszone-add --name-from-ip=192.0.2.0/24
```

Other Management Operations for Reverse DNS Zones

[Section 33.4, “Managing Master DNS Zones”](#) describes other zone management operations, some of which are also applicable to reverse DNS zone management, such as editing or disabling and enabling DNS zones.

33.8. DEFINING DNS QUERY POLICY

To resolve host names within the DNS domain, a DNS client issues a query to the DNS name server. For some security contexts or for performance, it might be advisable to restrict what clients can query DNS records in the zone.

DNS queries can be configured when the zone is created or when it is modified by using the **--allow-query** option with the **ipa dnszone-mod** command to set a list of clients which are allowed to issue queries.

For example:

```
[user@server ~]$ ipa dnszone-mod --allow-  
query=192.0.2.0/24;2001:DB8::/32;203.0.113.1 example.com
```

The default **--allow-query** value is **any**, which allows the zone to be queried by any client.

33.9. DNS LOCATIONS

33.9.1. DNS-based Service Discovery

DNS-based Service discovery is a process in which a client uses the DNS protocol to locate servers in a network that offer a specific service such as **LDAP** or **Kerberos**. One typical type of operation is to allow clients to locate authentication servers within the closest network infrastructure, because they provide a higher throughput and lower network latency, lowering overall costs.

The major advantages of service discovery are:

- No need for clients to be explicitly configured with names of nearby servers.
- DNS servers are used as central providers of policy. Clients using the same DNS server have access to the same policy about service providers and their preferred order.

In an IdM domain, DNS service records (SRV records) exists for LDAP, Kerberos, and other services. For example, the following command queries the DNS server for hosts providing a TCP-based Kerberos service in an IdM DNS domain:

Example 33.10. DNS Location Independent Results

```
$ dig -t SRV +short _kerberos._tcp.idm.example.com  
0 100 88 idmsvr-01.idm.example.com.  
0 100 88 idmsvr-02.idm.example.com.
```

The output contains the following information:

- **0** (priority): Priority of the target host. A lower value is preferred.
- **100** (weight). Specifies a relative weight for entries with the same priority. For further information, see [RFC 2782, section 3](#).
- **88** (port number): Port number of the service.
- Canonical name of the host providing the service.

In the previous example, the two host names returned have the same priority and weight. In this case, the client uses a random entry from the result list.

When the client instead queries a DNS server configured in a DNS location, the output differs. For IdM servers that are assigned to a location, tailored values are returned. In the example below, the client queries a DNS server in the location **germany**:

Example 33.11. DNS Location-based Results

```
$ dig -t SRV +short _kerberos._tcp.idm.example.com
_kerberos._tcp.germany._locations.idm.example.com.
0 100 88 idmsvr-01.idm.example.com.
50 100 88 idmsvr-02.idm.example.com.
```

The IdM DNS server automatically returns a DNS alias (CNAME) pointing to a DNS location specific SRV record which prefers local servers. This CNAME record is shown in the first line of the output. In the previous example, the host **idmsvr-01.idm.example.com** has the lowest priority value and is therefore preferred. The **idmsvr-02.idm.example.com** has a higher priority and thus is used only as backup for cases when the preferred host is unavailable.

33.9.2. Deployment Considerations for DNS Locations

For IdM DNS servers that are authoritative to the primary IdM DNS domain, IdM can generate location-specific SRV records. Because each IdM DNS server generates location-specific SRV records, you have to install at least one IdM DNS server in each DNS location.

The client's affinity to a DNS location is only defined by the DNS records received by the client. For this reason, you can combine IdM DNS servers with non-IdM DNS slave servers and recursors if the clients doing DNS service discovery resolve location-specific records from IdM DNS servers.

In the majority of deployments with mixed IdM and non-IdM DNS services, DNS recursors select the closest IdM DNS server automatically using round-trip time metrics. Typically, this ensures that clients using non-IdM DNS servers are getting records for the nearest DNS location and thus use the optimal set of IdM servers.

33.9.2.1. DNS Time to Live (TTL)

Clients can cache DNS resource records for an amount of time that is set in the zone's configuration. Because of this caching, a client might not be able to receive the changes until the time to live (TTL) value is expired. The default TTL value in IdM is **1 day**.

If your client computers roam between sites, you should adapt the TTL value for your IdM

DNS zone. Set the value to a lower value than the time clients need to roam between sites. This ensures that cached DNS entries on the client expire before they reconnect to another site and thus query the DNS server to refresh location-specific SRV records.

For further information how to modify the default TTL of a DNS zone, see [Section 33.4.2, “Adding Additional Configuration for Master DNS Zones”](#).

33.9.3. Creating DNS Locations

Creating DNS Locations from the Web UI

1. Open the **IPA Server** tab, and select **Topology** subtab.
2. Click **IPA Locations** in the navigation bar.
3. Click **Add** at the top of the locations list.
4. Fill in the location name.
5. Click the **Add** button to save the location.

Repeat the steps for further locations to add.

Creating DNS Locations from the Command Line

For example, to create a new location **germany**, enter:

```
[root@server ~]# ipa location-add germany
-----
Added IPA location "germany"
-----
Location name: germany
```

Repeat the step for all locations to add.

33.9.4. Assigning an IdM Server to a DNS Location

Assigning an IdM Server to a DNS Location from the Web UI

1. Open the **IPA Server** tab, and select **Topology** subtab.
2. Click **IPA Servers** in the navigation.
3. Click on the IdM server name.
4. Select a DNS location, and optionally set a service weight:

IPA Server: idmserver-01.idm.example.com

Refresh Revert Save

Server name	idmserver-01.idm.example.com.
Min domain level	0
Max domain level	1
Managed suffixes	domain ca
Location	germany
Service weight	100

Figure 33.34. Assigning a Server to a DNS Location

- Click **Save**.
- Restart the **named-pkcs11** service on the host you assigned in the previous steps the DNS location to:

```
[root@idmserver-01 ~]# systemctl restart named-pkcs11
```

Repeat the steps for further IdM servers you want to assign a DNS location to.

Assigning an IdM Server to a DNS Location from the Command Line

- Optional: List all configured DNS locations:

```
[root@server ~]# ipa location-find
-----
2 IPA locations matched
-----
Location name: australia
Location name: germany
-----
Number of entries returned: 2
-----
```

- Assign the server to the DNS location. For example, to assign the location **germany** to the server *idmserver-01.idm.example.com*, run:

```
[root@server ~]# ipa server-mod idmserver-01.idm.example.com --
location=germany
ipa: WARNING: Service named-pkcs11.service requires restart on IPA
server
idmserver-01.idm.example.com to apply configuration changes.
-----
Modified IPA server "idmserver-01.idm.example.com"
-----
```

```
Servername: idmserver-01.idm.example.com
Min domain level: 0
Max domain level: 1
Location: germany
Enabled server roles: DNS server, NTP server
```

3. Restart the **named-pkcs11** service on the host you assigned in the previous steps the DNS location to:

```
[root@idmserver-01 ~]# systemctl restart named-pkcs11
```

Repeat the steps for further IdM servers you want to assign a DNS location to.

33.10. UPDATING DNS RECORDS SYSTEMATICALLY WHEN USING EXTERNAL DNS

When using external DNS, Identity Management does not update the DNS records automatically after a change in the topology. The following procedures explain how you can update the DNS records managed by an external DNS service systematically, which reduces the need for manual DNS updates.

For a basic overview, see [Section 33.10.1, “Updating External DNS in Identity Management”](#).

For procedures and examples, see:

- [Section 33.10.2, “GUI: Updating External DNS Records”](#) if you use a GUI to manage the external DNS records
- [Section 33.10.3, “Command Line: Updating External DNS Records Using **nsupdate**”](#) if you use the **nsupdate** utility to manage the external DNS records

33.10.1. Updating External DNS in Identity Management

Updating DNS records removes old or invalid DNS records and adds new records.

You must update DNS records after a change in your topology, for example:

- After installing or uninstalling a replica
- After installing a CA, DNS, KRA, or Active Directory trust on an Identity Management server

33.10.2. GUI: Updating External DNS Records

1. Display the records that you must update. Use the **ipa dns-update-system-records --dry-run** command.

```
$ ipa dns-update-system-records --dry-run
IPA DNS records:
    _kerberos-master._tcp.example.com. 86400 IN SRV 0 100 88
ipa.example.com.
```

```
_kerberos-master._udp.example.com. 86400 IN SRV 0 100 88
ipa.example.com.
[... output truncated ...]
```

2. Use the external DNS GUI to update the records.

33.10.3. Command Line: Updating External DNS Records Using `nsupdate`

Generating a File with the DNS Records for `nsupdate`

1. Use the `ipa dns-update-system-records --dry-run --out dns_records_file.nsupdate` command with the `--out` option. The option specifies the path of the file to generate:

```
$ ipa dns-update-system-records --dry-run --out
dns_records_file.nsupdate
IPA DNS records:
_kerberos-master._tcp.example.com. 86400 IN SRV 0 100 88
ipa.example.com.
_kerberos-master._udp.example.com. 86400 IN SRV 0 100 88
ipa.example.com.
[... output truncated ...]
```

The generated file contains the required DNS records in the format accepted by the **nsupdate** utility.

2. The generated records rely on:
 - Automatic detection of the zone in which the records are to be updated
 - Automatic detection of the zone's authoritative server

If you are using an atypical DNS setup or if zone delegations are missing, **nsupdate** might not be able to find the right zone and server. In this case, add the following options to the beginning of the generated file:

- **server** specifies the server name or port of the authoritative DNS server to which **nsupdate** sends the records
- **zone** specifies the zone name of the zone where **nsupdate** places the records

Example:

```
$ cat dns_records_file.nsupdate
zone example.com.
server 192.0.2.1
; IPA DNS records
update delete _kerberos-master._tcp.example.com. SRV
update add _kerberos-master._tcp.example.com. 86400 IN SRV 0 100 88
ipa.example.com.
[... output truncated ...]
```

Submitting the Dynamic DNS Update Request to the Name Server

When sending a request using **nsupdate**, make sure you properly secure it. You can secure the request using these mechanisms:

Transaction Signature (TSIG) protocol

TSIG enables you to use **nsupdate** with a shared key. See [Procedure 33.1, “Sending an nsupdate Request Secured Using TSIG”](#).

GSS algorithm for TSIG (GSS-TSIG)

GSS-TSIG uses the GSS-API interface to obtain the secret TSIG key. GSS-TSIG is an extension to the TSIG protocol. See [Procedure 33.2, “Sending an nsupdate Request Secured Using GSS-TSIG”](#)

Procedure 33.1. Sending an nsupdate Request Secured Using TSIG

1. Make sure you meet these prerequisites:
 - Your DNS server must be configured for TSIG. See these server configuration examples: [BIND](#), [PowerDNS](#)
 - Both the DNS server and its client must have the shared key.
2. Run **nsupdate**, and provide the shared secret using one of these options:

- **-k** to provide the TSIG authentication key:


```
$ nsupdate -k tsig_key.file dns_records_file.nsupdate
```

- **-y** to generate a signature from the name of the key and from the Base64-encoded shared secret:

```
$ nsupdate -y algorithm:keyname:secret dns_records_file.nsupdate
```

Procedure 33.2. Sending an nsupdate Request Secured Using GSS-TSIG

1. Make sure you meet these prerequisites:
 - Your DNS server must be configured for GSS-TSIG. See these server configuration examples: [BIND](#), [PowerDNS](#), [Windows DNS](#).



NOTE

This procedure assumes that Kerberos V5 protocol is used as the technology for GSS-API.

2. To submit the DNS update request, authenticate with a principal allowed to update the records, and run **nsupdate** with the **-g** option to enable the GSS-TSIG mode:

```
$ kinit principal_allowed_to_update_records@REALM
$ nsupdate -g dns_records_file.nsupdate
```

Additional Resources

- the `nsupdate(8)` man page
- [RFC 2845](#) describes the TSIG protocol

- [RFC 3645](#) describes the GSS-TSIG algorithm

33.11. INSTALLING DNS SERVICES INTO AN EXISTING SERVER

It is possible to install DNS services into an IdM server that was originally installed without them. To do this, make sure the `ipa-server-dns` package is installed, and then use the `ipa-dns-install` utility.

Configuring DNS services using `ipa-dns-install` follows the same principles as installing DNS with the `ipa-server-install` utility, as described in [Section 2.3.3, “Installing a Server with Integrated DNS”](#).

For more information about `ipa-dns-install`, see the `ipa-dns-install(1)` man page.

33.11.1. Setting up Additional Name Servers

IdM adds the newly-configured IdM DNS server to the list of DNS servers in the `/etc/resolv.conf` file. It is recommended to manually add other DNS servers as backup servers in case the IdM server becomes unavailable. For example:

```
search example.com

; the IdM server
nameserver 192.0.2.1

; backup DNS servers
nameserver 198.51.100.1
nameserver 198.51.100.2
```

For more details about configuring `/etc/resolv.conf`, see the `resolv.conf(5)` man page.

[3] For more information about GSS-TSIG, see [RFC 3545](#).

[4] For the full text of RFC 3007, see <http://tools.ietf.org/html/rfc3007>

[5] For more information, see the [BIND 9 Configuration Reference](#).

CHAPTER 34. USING AUTOMOUNT

Automount is a way to manage, organize, and access directories across multiple systems. Automount automatically mounts a directory whenever access to it is requested. This works exceptionally well within an IdM domain since it allows directories on clients within the domain to be shared easily. This is especially important with user home directories, see [Section 11.1, “Setting up User Home Directories”](#).

In IdM, automount works with the internal LDAP directory and also with DNS services if configured.

34.1. ABOUT AUTOMOUNT AND IDM

Automount provides a coherent structure to the way that directories are organized. Every directory is called a *mount point* or a *key*. Multiple keys that are grouped together create a *map*, and maps are associated according to their physical or conceptual *location*.

The base configuration file for automount is the **auto.master** file in the **/etc** directory. If necessary, there can be multiple **auto.master** configuration files in separate server locations.

When the **autofs** utility is configured on a server and the server is a client in an IdM domain, then all configuration information for automount is stored in the IdM directory. Rather than in separate text files, the **autofs** configuration containing maps, locations, and keys are stored as LDAP entries. For example, the default map file, **auto.master**, is stored as:

```
dn: automountmapname=auto.master,cn=default,cn=automount,dc=example,dc=com
objectClass: automountMap
objectClass: top
automountMapName: auto.master
```



IMPORTANT

Identity Management works with an existing **autofs** deployment but does not set up or configure **autofs** itself.

Each new location is added as a container entry under **cn=automount,dc=example,dc=com**, and each map and each key are stored beneath that location.

As with other IdM domain services, automount works with IdM natively. The automount configuration can be managed by IdM tools:

- The **ipa automountlocation*** commands for *Locations*,
- The **ipa automountmap*** commands for direct and indirect *maps*,
- The **ipa automountkey*** commands for *keys*.

For automount to work within the IdM domain, the NFS server must be configured as an IdM client. Configuring NFS itself is covered in the [Red Hat Enterprise Linux Storage Administration Guide](#).

34.2. CONFIGURING AUTOMOUNT

in Identity Management, configuring automount entries like locations and maps requires an existing autofs/NFS server. Creating automount entries does not create the underlying **autofs** configuration. **Autofs** can be configured manually using LDAP or SSSD as a data store, or it can be configured automatically.



NOTE

Before changing the automount configuration, test that for at least one user, their **/home** directory can be mounted from the command line successfully. Making sure that NFS is working properly makes it easier to troubleshoot any potential IdM automount configuration errors later.

34.2.1. Configuring NFS Automatically

After a system is configured as an IdM client, which includes IdM servers and replicas that are configured as domain clients as part of their configuration, **autofs** can be configured to use the IdM domain as its NFS domain and have **autofs** services enabled.

By default, the **ipa-client-automount** utility automatically configures the NFS configuration files, **/etc/sysconfig/nfs** and **/etc/idmapd.conf**. It also configures SSSD to manage the credentials for NFS. If the **ipa-client-automount** command is run without any options, it runs a DNS discovery scan to identify an available IdM server and creates a default location called **default**.

```
[root@ipa-server ~]# ipa-client-automount
Searching for IPA server...
IPA server: DNS discovery
Location: default
Continue to configure the system with these values? [no]: yes
Configured /etc/nsswitch.conf
Configured /etc/sysconfig/nfs
Configured /etc/idmapd.conf
Started rpcidmapd
Started rpcgssd
Restarting sssd, waiting for it to become available.
Started autofs
```

It is possible to specify an IdM server to use and to create an automount location other than default:

```
[root@server ~]# ipa-client-automount --server=ipaserver.example.com --
location=boston
```

Along with setting up NFS, the **ipa-client-automount** utility configures SSSD to cache automount maps, in case the external IdM store is ever inaccessible. Configuring SSSD does two things:

- It adds service configuration information to the SSSD configuration. The IdM domain entry is given settings for the autofs provider and the mount location.

```
autofs_provider = ipa
ipa_automount_location = default
```

And NFS is added to the list of supported services (**services = nss,pam,autofs...**) and given a blank configuration entry (**autofs**).

- The Name Service Switch (NSS) service information is updated to check SSSD first for automount information, and then the local files.

```
automount: sss files
```

There may be some instances, such as highly secure environments, where it is not appropriate for a client to cache automount maps. In that case, the **ipa-client-automount** command can be run with the **--no-sssd** option, which changes all of the required NFS configuration files, but does not change the SSSD configuration.

```
[root@server ~]# ipa-client-automount --no-sssd
```

If **--no-sssd** is used, the list of configuration files updated by **ipa-client-automount** is different:

- The command updates **/etc/sysconfig/autofs** instead of **/etc/sysconfig/nfs**.
- The command configures **/etc/autofs_ldap_auth.conf** with the IdM LDAP configuration.
- The command configures **/etc/nsswitch.conf** to use the LDAP services for automount maps.



NOTE

The **ipa-client-automount** command can only be run once. If there is an error in the configuration, then the configuration files need to be edited manually.

34.2.2. Configuring autofs Manually to Use SSSD and Identity Management

1. Edit the **/etc/sysconfig/autofs** file to specify the schema attributes that autofs searches for:

```
#
# Other common LDAP naming
#
MAP_OBJECT_CLASS="automountMap"
ENTRY_OBJECT_CLASS="automount"
MAP_ATTRIBUTE="automountMapName"
ENTRY_ATTRIBUTE="automountKey"
VALUE_ATTRIBUTE="automountInformation"
```

2. Specify the LDAP configuration. There are two ways to do this. The simplest is to let the automount service discover the LDAP server and locations on its own:

```
LDAP_URI="ldap:///dc=example,dc=com"
```

Alternatively, explicitly set which LDAP server to use and the base DN for LDAP searches:

```
LDAP_URI="ldap://ipa.example.com"
SEARCH_BASE="cn=location,cn=automount,dc=example,dc=com"
```



NOTE

The default value for *location* is **default**. If additional locations are added ([Section 34.4, “Configuring Locations”](#)), then the client can be pointed to use those locations, instead.

3. Edit the `/etc/autofs_ldap_auth.conf` file so that autofs allows client authentication with the IdM LDAP server.
 - Change ***authrequired*** to `yes`.
 - Set the principal to the Kerberos host principal for the NFS client server, `host/fqdn@REALM`. The principal name is used to connect to the IdM directory as part of GSS client authentication.

```
<autofs_ldap_sasl_conf
    usetls="no"
    tlsrequired="no"
    authrequired="yes"
    authtype="GSSAPI"
    clientprinc="host/server.example.com@EXAMPLE.COM"
/>
```

If necessary, run **`klist -k`** to get the exact host principal information.

4. Configure autofs as one of the services which SSSD manages.

1. Open the SSSD configuration file.

```
[root@server ~]# vim /etc/sss/sss.conf
```

2. Add the autofs service to the list of services handled by SSSD.

```
[sss]
services = nss,pam,autofs
```

3. Create a new **[autofs]** section. This can be left blank; the default settings for an autofs service work with most infrastructures.

```
[nss]

[pam]

[sudo]

[autofs]
```

```
[ssh]
```

```
[pac]
```

4. Optionally, set a search base for the autofs entries. By default, this is the LDAP search base, but a subtree can be specified in the **ldap_autofs_search_base** parameter.

```
[domain/EXAMPLE]
```

```
...
```

```
ldap_search_base = "dc=example,dc=com"
```

```
ldap_autofs_search_base = "ou=automount,dc=example,dc=com"
```

5. Restart SSSD:

```
[root@server ~]# systemctl restart sssd.service
```

6. Check the **/etc/nsswitch.conf** file, so that SSSD is listed as a source for automount configuration:

```
automount: sss files
```

7. Restart autofs:

```
[root@server ~]# systemctl restart autofs.service
```

8. Test the configuration by listing a user's **/home** directory:

```
[root@server ~]# ls /home/userName
```

If this does not mount the remote file system, check the **/var/log/messages** file for errors. If necessary, increase the debug level in the **/etc/sysconfig/autofs** file by setting the **LOGGING** parameter to **debug**.



NOTE

If there are problems with automount, then cross-reference the automount attempts with the 389 Directory Server access logs for the IdM instance, which will show the attempted access, user, and search base.

It is also simple to run automount in the foreground with debug logging on.

```
automount -f -d
```

This prints the debug log information directly, without having to cross-check the LDAP access log with automount's log.

34.2.3. Configuring Automount on Solaris

**NOTE**

Solaris uses a different schema for autofs configuration than the schema used by Identity Management. Identity Management uses the 2307bis-style automount schema which is defined for 389 Directory Server (and used in IdM's internal Directory Server instance).

1. If the NFS server is running on Red Hat Enterprise Linux, specify on the Solaris machine that NFSv3 is the maximum supported version. Edit the `/etc/default/nfs` file and set the following parameter:

```
NFS_CLIENT_VERSMAX=3
```

2. Use the **ldapclient** command to configure the host to use LDAP:

```
ldapclient -v manual -a authenticationMethod=none
-a defaultSearchBase=dc=example,dc=com
-a defaultServerList=ipa.example.com
-a
serviceSearchDescriptor=passwd:cn=users,cn=accounts,dc=example,dc=com
-a
serviceSearchDescriptor=group:cn=groups,cn=compat,dc=example,dc=com
-a
serviceSearchDescriptor=auto_master:automountMapName=auto.master,cn=location,cn=automount,dc=example,dc=com?one
-a
serviceSearchDescriptor=auto_home:automountMapName=auto_home,cn=location,cn=automount,dc=example,dc=com?one
-a objectClassMap=shadow:shadowAccount=posixAccount
-a searchTimeLimit=15
-a bindTimeLimit=5
```

3. Enable **automount**:

```
# svcadm enable svc:/system/filesystem/autofs
```

4. Test the configuration.

1. Check the LDAP configuration:

```
# ldapclient -l auto_master

dn:
automountkey=/home,automountmapname=auto.master,cn=location,cn=automount,dc=example,dc=com
objectClass: automount
objectClass: top
automountKey: /home
automountInformation: auto.home
```

2. List a user's **/home** directory:

```
# ls /home/userName
```

34.3. SETTING UP A KERBEROS-AWARE NFS SERVER

Identity Management can be used to set up a Kerberos-aware NFS server.



NOTE

The NFS server does not need to be running on Red Hat Enterprise Linux.

34.3.1. Setting up a Kerberos-aware NFS Server

1. Obtain a Kerberos ticket before running IdM tools.

```
[jsmith@server ~]$ kinit admin
```

2. If the NFS host machine has not been added as a client to the IdM domain, then create the host entry. See [Section 12.3, “Adding Host Entries”](#).
3. Create the NFS service entry in the IdM domain. For example:

```
[jsmith@server ~]$ ipa service-add nfs/nfs-server.example.com
```

For more information, see [Section 16.1, “Adding and Editing Service Entries and Keytabs”](#).

4. Generate an NFS service keytab for the NFS server using the **ipa-getkeytab** command, and save the keys directly to the host keytab. For example:

```
[jsmith@server ~]$ ipa-getkeytab -s ipaserver.example.com -p  
nfs/nfs-server.example.com -k /etc/krb5.keytab
```



NOTE

Verify that the NFS service has been properly configured in IdM, with its keytab, by checking the service entry:

```
[jsmith@server ~]$ ipa service-show nfs/nfs-  
server.example.com  
Principal: NFS/nfs-server.example.com@EXAMPLE.COM  
Keytab: True
```



NOTE

This procedure assumes that the NFS server is running on a Red Hat Enterprise Linux system or a UNIX system which can run **ipa-getkeytab**.

If the NFS server is running on a system which cannot run **ipa-getkeytab**, then create the keytab using system tools. Two things must be done:

- The key must be created in the **/root** (or equivalent) directory.
- The **ktutil** command can merge the keys into the system **/etc/krb5.keytab** file. The [ktutil man page](#) describes how to use the tool.

5. Install the NFS packages. For example:

```
[root@nfs-server ~]# yum install nfs-utils
```

6. Configure weak crypto support. This is required for every NFS client if *any* client (such as a Red Hat Enterprise Linux 5 client) in the domain will use older encryption options like DES.

1. Edit the **krb5.conf** file to allow weak crypto.

```
[root@nfs-server ~]# vim /etc/krb5.conf

allow_weak_crypto = true
```

2. Update the IdM server Kerberos configuration to support the DES encryption type.

```
[jsmith@ipaserver ~]$ ldapmodify -x -D "cn=directory manager" -w
password -h ipaserver.example.com -p 389

dn: cn=EXAMPLEREALM,cn=kerberos,dc=example,dc=com
changetype: modify
add: krbSupportedEncSaltTypes
krbSupportedEncSaltTypes: des-cbc-crc:normal
-
add: krbSupportedEncSaltTypes
krbSupportedEncSaltTypes: des-cbc-crc:special
-
add: krbDefaultEncSaltTypes
krbDefaultEncSaltTypes: des-cbc-crc:special
```

7. Run the **ipa-client-automount** command to configure the NFS settings.

By default, this enables secure NFS in the **/etc/sysconfig/nfs** file and sets the IdM DNS domain in the **Domain** parameter in the **/etc/idmapd.conf** file.

8. Edit the **/etc/exports** file and add the Kerberos information:


```
/export *(rw,sec=krb5:krb5i:krb5p)
```

- Restart the NFS server and related services.

```
[root@nfs-server ~]# systemctl restart nfs.service
[root@nfs-server ~]# systemctl restart nfs-server.service
[root@nfs-server ~]# systemctl restart nfs-secure.service
[root@nfs-server ~]# systemctl restart nfs-secure-server.service
```

- Configure the NFS server as an NFS client, following the directions in [Section 34.3.2, “Setting up a Kerberos-aware NFS Client”](#).

34.3.2. Setting up a Kerberos-aware NFS Client

- Obtain a Kerberos ticket before running IdM tools.

```
[jsmith@server ~]$ kinit admin
```

- If the NFS client is not enrolled as a client in the IdM domain, then set up the required host entries, as described in [Section 12.3, “Adding Host Entries”](#).
- Run the **ipa-client-automount** command to configure the NFS settings.

By default, this enables secure NFS in the **/etc/sysconfig/nfs** file and sets the IdM DNS domain in the **Domain** parameter in the **/etc/idmapd.conf** file.

- Start the GSS daemon.

```
[root@nfs-client-server ~]# systemctl start rpc-gssd.service
[root@nfs-client-server ~]# systemctl start rpcbind.service
[root@nfs-client-server ~]# systemctl start nfs-idmapd.service
```

- Mount the directory.

```
[root@nfs-client-server ~]# echo "$NFSSERVER:/this /mnt/this nfs4
sec=krb5i,rw,proto=tcp,port=2049" >>/etc/fstab
[root@nfs-client-server ~]# mount -av
```

- Configure SSSD on the client system to manage home directories and renew Kerberos tickets.

- Enable SSSD with the **--enablemkhomedir** option.

```
[root@nfs-client-server ~]# authconfig --update --enablesssd --
enablesssdauth --enablemkhomedir
```

- Restart the OpenSSH client.

```
[root@nfs-client-server ~]# systemctl restart sshd.service
```

- Edit the IdM domain section in the SSSD configuration file to set the keytab renewal options.

```
[root@nfs-client-server ~]# vim /etc/sss/sss.conf

[domain/EXAMPLE.COM]
cache_credentials = True
krb5_store_password_if_offline = True
ipa_domain = example.com
id_provider = ipa
auth_provider = ipa
...
krb5_renewable_lifetime = 50d
krb5_renew_interval = 3600
```

4. Restart SSSD.

```
[root@nfs-client-server ~]# systemctl restart sssd.service
```

34.4. CONFIGURING LOCATIONS

A location is a set of maps, which are all stored in **auto.master**, and a location can store multiple maps. The location entry only works as a container for map entries; it is not an automount configuration in and of itself.




IMPORTANT


Identity Management does not set up or configure autofs. That must be done separately. Identity Management works with an existing autofs deployment.

34.4.1. Configuring Locations through the Web UI

1. Click the **Policy** tab.
2. Click the **Automount** subtab.
3. Click the **Add** link at the top of the list of automount locations.



4. Enter the name for the new location.



5. Click the **Add and Edit** button to go to the map configuration for the new location. Create maps, as described in [Section 34.5.1.1, “Configuring Direct Maps from the Web UI”](#) and [Section 34.5.2.1, “Configuring Indirect Maps from the Web UI”](#).

34.4.2. Configuring Locations through the Command Line

To create a map, using the **automountlocation-add** and give the location name.

```
$ ipa automountlocation-add location
```

For example:

```
$ ipa automountlocation-add raleigh
-----
```

```
Added automount location "raleigh"
```

```
-----
```

```
Location: raleigh
```

When a new location is created, two maps are automatically created for it, **auto.master** and **auto.direct**. **auto.master** is the root map for all automount maps for the location. **auto.direct** is the default map for direct mounts and is mounted on `/-`.

To view all of the maps configured for a location as if they were deployed on a filesystem, use the **automountlocation-tofiles** command:

```
$ ipa automountlocation-tofiles raleigh
/etc/auto.master:
/-      /etc/auto.direct
-----
/etc/auto.direct:
```

34.5. CONFIGURING MAPS

Configuring maps not only creates the maps, it associates mount points through the keys and it assigns mount options that should be used when the directory is accessed. IdM supports both direct and indirect maps.



NOTE

Different clients can use different map sets. Map sets use a tree structure, so maps *cannot* be shared between locations.



IMPORTANT

Identity Management does not set up or configure autofs. That must be done separately. Identity Management works with an existing autofs deployment.


34.5.1. Configuring Direct Maps

Direct maps define exact locations, meaning absolute paths, to the file mount point. In the location entry, a direct map is identified by the preceding forward slash:


```
-----
/etc/auto.direct:
/shared/man server.example.com:/shared/man
```

34.5.1.1. Configuring Direct Maps from the Web UI


1. Click the **Policy** tab.
2. Click the **Automount** subtab.
3. Click name of the automount location to which to add the map.



4. In the **Automount Maps** tab, click the **+ Add** link to create a new map.



5. In pop-up window, select the **Direct** radio button and enter the name of the new map.



Add Automount Map


Map Type: ☒ Direct ☐ Indirect

Map:

Description:

Buttons: Add, Add and Add Another, Add and Edit, Cancel

6. In the **Automount Keys** tab, click the + **Add** link to create a new key for the map.



Host Based Access Control Sudo **Automount** Password Policies Kerberos Ticket Policy

Automount Locations » boston » auto.man


AUTOMOUNT MAP: auto.man

Automount Keys Settings

Refresh Delete + Add

Key	Mount information
☒ /manpages	-ro,soft, ipaserver.example.com:/home/manpages

7. Enter the mount point. The key defines the actual mount point in the key name. The **Info** field sets the network location of the directory, as well as **anymount** options to use.



Add Automount Key

Key:

Mount information:

Buttons: Add, Add and Add Another, Add and Edit, Cancel

8. Click the **Add** button to save the new key.

34.5.1.2. Configuring Direct Maps from the Command Line

The key defines the actual mount point (in the key name) and any options. A map is a direct or indirect map based on the format of its key.

Each location is created with an **auto.direct** item. The simplest configuration is to define a direct mapping by adding an automount key to the existing direct map entry. It is also possible to create different direct map entries.

Add the key for the direct map to the location's **auto.direct** file. The **--key** option identifies the mount point, and **--info** gives the network location of the directory, as well as any **mount** options to use. For example:

```
$ ipa automountkey-add raleigh auto.direct --key=/share --  
info="ro,soft,ipaserver.example.com:/home/share"  
Key: /share  
Mount information: ro,soft,ipaserver.example.com:/home/share
```

Mount options are described in the mount manpage, <http://linux.die.net/man/8/mount>.

On Solaris, add the direct map and key using the **ldapclient** command to add the LDAP entry directly:


```
ldapclient -a  
serviceSearchDescriptor=auto_direct:automountMapName=auto.direct,cn=locati  
on,cn=automount,dc=example,dc=com?one
```

34.5.2. Configuring Indirect Maps


An indirect map essentially specifies a relative path for maps. A parent entry sets the base directory for all of the indirect maps. The indirect map key sets a sub directory; whenever the indirect map location is loaded, the key is appended to that base directory. For example, if the base directory is **/docs** and the key is **man**, then the map is **/docs/man**.

34.5.2.1. Configuring Indirect Maps from the Web UI

1. Click the **Policy** tab.
2. Click the **Automount** subtab.
3. Click name of the automount location to which to add the map.



4. In the **Automount Maps** tab, click the + **Add** link to create a new map.



5. In pop-up window, select the **Indirect** radio button and enter the required information for the indirect map:

- The name of the new map
- The mount point. The **Mount** field sets the base directory to use for all the indirect map keys.
- Optionally, a parent map. The default parent is **auto.master**, but if another map exists which should be used, that can be specified in the **Parent Map** field.

6. Click the **Add** button to save the new key.

34.5.2.2. Configuring Indirect Maps from the Command Line

The primary difference between a direct map and an indirect map is that there is no forward slash in front of an indirect key.

```
-----
/etc/auto.share:
man      ipa.example.com:/docs/man
-----
```

1. Create an indirect map to set the base entry using the **automountmap-add-indirect** command. The **--mount** option sets the base directory to use for all the indirect map keys. The default parent entry is **auto.master**, but if another map exists which should be used, that can be specified using the **--parentmap** option.

```
$ ipa automountmap-add-indirect location mapName --mount=directory
[--parentmap=mapName]
```

For example:

```
$ ipa automountmap-add-indirect raleigh auto.share --mount=/share
-----
Added automount map "auto.share"
-----
```

2. Add the indirect key for the mount location:

```
$ ipa automountkey-add raleigh auto.share --key=docs --
info="ipa.example.com:/export/docs"
-----
Added automount key "docs"
-----
Key: docs
Mount information: ipa.example.com:/export/docs
```

3. To verify the configuration, check the location file list using **automountlocation-tofiles**:

```
$ ipa automountlocation-tofiles raleigh
/etc/auto.master:
/-      /etc/auto.direct
/share  /etc/auto.share
-----
/etc/auto.direct:
-----
/etc/auto.share:
man      ipa.example.com:/export/docs
```

On Solaris, add the indirect map using the **ldapclient** command to add the LDAP entry directly:

```
ldapclient -a
serviceSearchDescriptor=auto_share:automountMapName=auto.share,cn=location
,cn=automount,dc=example,dc=com?one
```

34.5.3. Importing Automount Maps

If there are existing automount maps, these can be imported into the IdM automount configuration.

```
ipa automountlocation-import location map_file [--continuous]
```

The only required information is the IdM automount location and the full path and name of the map file. The **--continuous** option tells the **automountlocation-import** command to continue through the map file, even if the command encounters errors.

For example:

```
$ ipa automountlocation-import raleigh /etc/custom.map
```

PART VIII. SECURITY HARDENING

CHAPTER 35. CONFIGURING TLS FOR IDENTITY MANAGEMENT

This document describes how to configure an Identity Management server to require the TLS protocol version 1.2 in Red Hat Enterprise Linux 7.3 and later.

TLS 1.2 is considered more secure than previous versions of TLS. If your IdM server is deployed in an environment with high security requirements, you can configure it to forbid communication using protocols that are less secure than TLS 1.2.



IMPORTANT

Repeat these steps on every IdM server where you want to use TLS 1.2.

35.1. CONFIGURING THE HTTPD DAEMON

1. Open the `/etc/httpd/conf.d/nss.conf` file, and set the following values for the `NSSProtocol` and `NSSCipherSuite` entries:

```
NSSProtocol TLSv1.2
NSSCipherSuite
+ecdhc_ecdsa_aes_128_sha,+ecdhc_ecdsa_aes_256_sha,+ecdhc_rsa_aes_128
_sha,+ecdhc_rsa_aes_256_sha,+rsa_aes_128_sha,+rsa_aes_256_sha
```

Alternatively, use the following commands to set the values for you:

```
# sed -i 's/^NSSProtocol .*/NSSProtocol TLSv1.2/'
/etc/httpd/conf.d/nss.conf
# sed -i 's/^NSSCipherSuite .*/NSSCipherSuite
+ecdhc_ecdsa_aes_128_sha,+ecdhc_ecdsa_aes_256_sha,+ecdhc_rsa_aes_128
_sha,+ecdhc_rsa_aes_256_sha,+rsa_aes_128_sha,+rsa_aes_256_sha/'
/etc/httpd/conf.d/nss.conf
```

2. Restart the `httpd` daemon:

```
# systemctl restart httpd
```

35.2. CONFIGURING THE DIRECTORY SERVER COMPONENT

To configure Directory Server (DS) manually:

1. Stop DS:

```
# systemctl stop dirsrv@EXAMPLE-COM.service
```

2. Open the `/etc/dirsrv/slapd-EXAMPLE-COM/dse.ldif` file, and modify the `cn=encryption,cn=config` entry to set the following:

```
sslVersionMin: TLS1.2
```

3. Start DS:

```
# systemctl start dirsrv@EXAMPLE-COM.service
```

Alternatively, to configure DS automatically using the **ldapmodify** utility:

1. Use **ldapmodify** to make the configuration changes for you:

```
ldapmodify -h localhost -p 389 -D 'cn=directory manager' -W << EOF
dn: cn=encryption,cn=config
changeType: modify
replace: sslVersionMin
sslVersionMin: TLS1.2
EOF
```

2. Restart DS to load the new configuration:

```
# systemctl restart dirsrv@EXAMPLE-COM.service
```

35.3. CONFIGURING THE CERTIFICATE SERVER COMPONENT

1. To configure Certificate Server (CS) manually, open the **/etc/pki/pki-tomcat/server.xml** file. Set all occurrences of the **sslVersionRangeStream** and **sslVersionRangeDatagram** parameters to the following values:

```
sslVersionRangeStream="tls1_2:tls1_2"
sslVersionRangeDatagram="tls1_2:tls1_2"
```

Alternatively, use the following command to replace the values for you:

```
# sed -i 's/tls1_[01]:tls1_2/tls1_2:tls1_2/g' /etc/pki/pki-
tomcat/server.xml
```

2. Restart CS:

```
# systemctl restart pki-tomcatd@pki-tomcat.service
```

35.4. RESULT

The Identity Management server is configured to require TLS 1.2. Identity Management clients that only support previous TLS versions are no longer able to communicate with the Identity Management server.

CHAPTER 36. DISABLING ANONYMOUS BINDS


Accessing domain resources and running client tools always require Kerberos authentication. However, the back end LDAP directory used by the IdM server allows anonymous binds by default. This potentially opens up all of the domain configuration to unauthorized users, including information about users, machines, groups, services, netgroups, and DNS configuration.

It is possible to disable anonymous binds on the 389 Directory Server instance by using LDAP tools to reset the ***nsslapd-allow-anonymous-access*** attribute.

1. Change the ***nsslapd-allow-anonymous-access*** attribute to **rootdse**.

```
$ ldapmodify -x -D "cn=Directory Manager" -W -h server.example.com -
p 389 -ZZ
Enter LDAP Password:
dn: cn=config
changetype: modify
replace: nsslapd-allow-anonymous-access
nsslapd-allow-anonymous-access: rootdse

modifying entry "cn=config"
```



IMPORTANT

Anonymous access can be completely allowed (on) or completely blocked (off). However, completely blocking anonymous access also blocks external clients from checking the server configuration. LDAP and web clients are not necessarily domain clients, so they connect anonymously to read the root DSE file to get connection information.

The **rootdse** allows access to the root DSE and server configuration *without* any access to the directory data.

2. Restart the 389 Directory Server instance to load the new setting.

```
# systemctl restart dirsrv.target
```

PART IX. PERFORMANCE TUNING

CHAPTER 37. PERFORMANCE TUNING FOR BULK PROVISIONING OF ENTRIES

Adding a large number of entries using the usual workflow, such as [Chapter 11, Managing User Accounts](#) for adding users, can be very slow. This chapter describes how to tune the process to ensure the provisioning is completed as quickly as possible.

As part of the procedure:

- Identity Management (IdM) reads entries to be provisioned from an LDIF file and then imports them to the target IdM LDAP instance.
- The administrator sets custom values for certain attributes, such as cache sizes, and disables the MemberOf and Schema Compatibility plug-ins. The procedure includes running the **fixup-memberof.pl** plug-in on the provisioned entries to compensate for disabling MemberOf.

This procedure has been designed and tested to provision the following entry types: user, user group, host, host group, sudo rules, and host-based access control (HBAC) rules.

Recommendations and Prerequisites for Bulk Provisioning

Recommendations:

- When provisioning a large number of entries (10,000 or more), do not allow any LDAP client to access the server on which the entries are provisioned or to rely on the information from the server. For example, you can achieve this by disabling ports 389 and 636 on the server and using LDAPAPI to work over Unix sockets.

Reason: The MemberOf plug-in is disabled on the server, which means that membership information on the server is not valid.

- Stop applications that are not required to be running during the provisioning.

Reason: This helps free as much memory on the machine as possible. The free memory will be used by the file system cache, thus improving the performance of the provisioning.

Note that the procedure below already includes steps to stop the IdM services, and restart only the Directory Server (DS) instance. IdM services, especially **tomcat**, consume a lot of memory, but are not used during the provisioning.

- Run the procedure on a fresh IdM deployment with only one server. Create replicas only after the provisioning has been completed.

Reason: The provisioning throughput is much faster than replication. In a deployment with more than one server, information on the replicas would become significantly outdated.

Prerequisites:

- Generate an LDIF file containing the entries you want to provision. For example, if you are migrating an existing IdM deployment, create the LDIF file by exporting all the entries using the **ldapsearch** utility.

For details on the LDIF format, see [About the LDIF File Format](#) in the Red Hat Directory Server Administration Guide.

Backing up the Current DS Tuning Parameter Values

1. Retrieve the current values for the DS tuning parameters:

- the database cache size and database locks:

```
# ldapsearch -D "cn=directory manager" -w secret -b
"cn=config,cn=ldb database,cn=plugins,cn=config" nsslapd-
dbcachesize nsslapd-db-locks

...
nsslapd-dbcachesize: 10000000
nsslapd-db-locks: 50000
...
```

- the entry cache size and DN cache size:

```
# ldapsearch -D "cn=directory manager" -w secret -b
"cn=userRoot,cn=ldb database,cn=plugins,cn=config" nsslapd-
cachememsize nsslapd-dncachememsize

...
nsslapd-cachememsize: 10485760
nsslapd-dncachememsize: 10485760
...
```

2. Make note of the obtained values. You will reset the parameters back to these values after you finish the provisioning.

Adjusting the Database, Domain Entry, and DN Cache Size

For the database cache size:

1. Determine the required value.

The recommended value is typically between 200 MB and 500 MB. The value appropriate for your use case depends on the memory available on your system:

- More than 8 GB of memory → 500 MB
- 8 GB - 4 GB of memory → 200 MB
- Less than 4 GB of memory → 100 MB

2. Set the determined value by using this template:

```
dn: cn=config,cn=ldb database,cn=plugins,cn=config
changetype: modify
replace: nsslapd-dbcachesize
nsslapd-dbcachesize: db_cache_size_in_bytes
```

For an example of modifying LDAP attributes using the **ldapmodify** utility, see [Example 37.1, “Using **ldapmodify** to Change an LDAP Attribute”](#).

Example 37.1. Using **ldapmodify** to Change an LDAP Attribute

1. Run the **ldapmodify** command, and then add the statements to modify the attribute value. For example:

```
# ldapmodify -D "cn=directory manager" -w secret -x
dn: cn=config,cn=ldb database,cn=plugins,cn=config
changetype: modify
replace: nsslapd-dbcachesize
nsslapd-dbcachesize: 200000000
```

2. Press **Ctrl+D** to confirm and send the changes to the server. If the operation finishes successfully, the following message is displayed:

```
modifying entry "cn=config,cn=ldb database,cn=plugins,cn=config"
```

For the domain entry cache size:

1. Determine the required value.

The recommended value is between 100 MB and 400 MB. The appropriate value depends on the memory available on your system:

- More than 4 GB of memory → 400 MB
- 2 GB - 4 GB of memory → 200 MB
- Less than 2 GB of memory → 100 MB

If you are provisioning a large static group, it is recommended that the entry cache is large enough to fit all entries: groups and members.

2. Set the determined value by using this template:

```
dn: cn=userRoot,cn=ldb database,cn=plugins,cn=config
changetype: modify
replace: nsslapd-cachememsize
nsslapd-cachememsize: entry_cache_size_in_bytes
```

For the domain name (DN) cache size:

1. For the best possible performance, it is recommended that the DN cache fits all the DNs of the provisioned entries. To estimate the value appropriate for your use case:

- a. Determine the number of all DN entries in the file. The DN entries are on lines starting with **dn:** . For example, using **# grep**, **sed**, and **wc**:

```
# grep '^dn: ' ldif_file | sed 's/^dn: //' | wc -l
92200
```

- b. Determine the size of all DN entry strings in the LDIF file.

```
# grep '^dn: ' ldif_file | sed 's/^dn: //' | wc -c
9802460
```

- c. Get the average DN size: divide the size of all DN entry strings by the number of all the DN entries in the file.

For example: $9,802,460 / 92,200 \approx 106$

- d. Get the average memory size: multiple the average DN size by 2, and then add 32 to the result.

For example: $(106 * 2) + 32 = 244$

- e. Get the appropriate DN cache size: multiply the average memory size by the total number of DN entries in the LDIF file.

For example: $244 * 92,200 = 22,496,800$

2. Set the determined value by using this template:

```
dn: cn=userRoot,cn=ldbm database,cn=plugins,cn=config
changetype: modify
Replace: nsslapd-dncachememsize
Nsslapd-dncachememsize: dn_cache_size
```

Disabling Unnecessary Services and Adjusting Database Locks

1. Disable the MemberOf and Schema Compatibility plug-ins:

```
dn: cn=MemberOf Plugin,cn=plugins,cn=config
changetype: modify
replace: nsslapd-pluginEnabled
nsslapd-pluginEnabled: off
```

```
dn: cn=Schema Compatibility,cn=plugins,cn=config
changetype: modify
replace: nsslapd-pluginEnabled
nsslapd-pluginEnabled: off
```

Disabling MemberOf significantly speeds up the provisioning. Disabling Schema Compatibility also helps reduce the duration of the operation.

For an example of modifying LDAP attributes using the **ldapmodify** utility, see [Example 37.1, “Using **ldapmodify** to Change an LDAP Attribute”](#).

2. If no replicas are installed in your topology (as recommended in [the section called “Recommendations and Prerequisites for Bulk Provisioning”](#)), disable the Content Synchronization and Retro Changelog plug-ins:

```
dn: cn=Content Synchronization,cn=plugins,cn=config
changetype: modify
replace: nsslapd-pluginEnabled
nsslapd-pluginEnabled: off
```

```
dn: cn=Retro Changelog Plugin,cn=plugins,cn=config
changetype: modify
replace: nsslapd-pluginEnabled
nsslapd-pluginEnabled: off
```

■

Disabling these additional plug-ins helps improve the performance of the provisioning.

3. Stop the IdM server. This also stops the DS instance.

```
# ipactl stop
```

Stopping DS is required to set the number of database locks in the next step. You will restart it again later.

4. Adjust the number of database locks. The appropriate value equals half the number of provisioned entries.
 - the minimum value is 10,000
 - the maximum value is 200,000

Because DS is stopped, you must set the value by modifying the `/etc/dirsrv/slapd-EXAMPLE-COM/dse.ldif` file:

```
dn: cn=config,cn=ldbm database,cn=plugins,cn=config
...
nsslapd-db-locks: db_lock_number
```

IdM accesses a large number of database pages when computing membership. The more pages it accesses, the more locks are required for the provisioning.

5. Start DS:

```
# systemctl start dirsrv.target
```

Importing the Entries

To import the new entries from the LDIF file to the IdM LDAP instance. For example, using the **ldapadd** utility:

```
# ldapadd -D "binddn" -y password_file -f ldif_file
```

For details on using **ldapadd**, see the `ldapadd(1)` man page.

Re-enabling the Disabled Services and Restoring the Original Attribute Values

1. Enable MemberOf:

```
dn: cn=MemberOf Plugin,cn=plugins,cn=config
changetype: modify
replace: nsslapd-pluginEnabled
nsslapd-pluginEnabled: on
```

For an example of modifying LDAP attributes using the **ldapmodify** utility, see [Example 37.1, “Using **ldapmodify** to Change an LDAP Attribute”](#).

2. Restart DS:

■

```
# systemctl restart dirsrv.target
```

Restarting DS at this point is required because you enabled MemberOf in the previous step.

3. Run the **fixup-memberof.pl** script with the **(objectClass=*)** filter to regenerate and update the **memberOf** attribute on all provisioned entries. For example:

```
# fixup-memberof.pl -D "cn=directory manager" -j password_file -Z  
server_id -b "suffix" -f "(objectClass=*)" -P LDAP
```

Running **fixup-memberof.pl** is necessary because the MemberOf plug-in was disabled when you imported the entries. To be able to continue with the provisioning, the script must complete successfully.

For details on **fixup-memberof.pl**, see the `fixup-memberof.pl(8)` man page.

4. Enable the Schema Compatibility plug-in:

```
dn: cn=Schema Compatibility,cn=plugins,cn=config  
changetype: modify  
replace: nsslapd-pluginEnabled  
nsslapd-pluginEnabled: on
```

5. If you disabled the Content Synchronization and Retro Changelog plug-ins in [the section called “Disabling Unnecessary Services and Adjusting Database Locks”](#), re-enable them:

```
dn: cn=Content Synchronization,cn=plugins,cn=config  
changetype: modify  
replace: nsslapd-pluginEnabled  
nsslapd-pluginEnabled: on
```

```
dn: cn=Retro Changelog Plugin,cn=plugins,cn=config  
changetype: modify  
replace: nsslapd-pluginEnabled  
nsslapd-pluginEnabled: on
```

6. Restore the original values for the database cache, entry cache, and DN cache size that you backed up in [the section called “Backing up the Current DS Tuning Parameter Values”](#):

```
dn: cn=config,cn=ldbm database,cn=plugins,cn=config  
changetype: modify  
replace: nsslapd-dbcachesize  
nsslapd-dbcachesize: backup_db_cache_size  
  
dn: cn=userRoot,cn=ldbm database,cn=plugins,cn=config  
changetype: modify  
Replace: nsslapd-dncachememsize  
Nsslapd-dncachememsize: backup_dn_cache_size  
-  
replace: nsslapd-cachememsize  
nsslapd-cachememsize: backup_entry_cache_size
```

7. Stop DS:

```
# systemctl stop dirsrv.target
```

8. Restore the original value for database locks that you backed up in [the section called “Backing up the Current DS Tuning Parameter Values”](#). Because DS is stopped, you must set the value by modifying the `/etc/dirsrv/slapd-EXAMPLE-COM/dse.ldif` file:

```
dn: cn=config,cn=ldbm database,cn=plugins,cn=config
...
nsslapd-db-locks: backup_db_lock_number
```

9. Start the IdM server:

```
# ipactl start
```

This starts all IdM services, including DS.

PART X. MIGRATION

CHAPTER 38. MIGRATING FROM AN LDAP DIRECTORY TO IDM

As an administrator, you previously deployed an LDAP server for authentication and identity lookups and now you want to migrate the back end to Identity Management. You want to use the IdM migration tool to transfer user accounts, including passwords, and group, without losing data. Additionally you want to avoid expensive configuration updates on the clients.

The migration process described here, assumes a simple deployment scenario with one name space in LDAP and one in IdM. For more complex environments, such as multiple name spaces or custom schema, contact the Red Hat support services.

38.1. AN OVERVIEW OF AN LDAP TO IDM MIGRATION

The actual migration part of moving from an LDAP server to Identity Management — the process of moving the data from one server to the other — is fairly straightforward. The process is simple: move data, move passwords, and move clients.

The most expensive part of the migration is deciding how clients are going to be configured to use Identity Management. For each client in the infrastructure, you need to decide what services (such as Kerberos and SSSD) are being used and what services can be used in the final IdM deployment.

A secondary, but significant, consideration is planning how to migrate passwords. Identity Management requires Kerberos hashes for every user account in addition to passwords. Some of the considerations and migration paths for passwords are covered in [Section 38.1.2, “Planning Password Migration”](#).

38.1.1. Planning the Client Configuration

Identity Management can support a number of different client configurations, with varying degrees of functionality, flexibility, and security. Decide which configuration is best *for each individual client* based on its operating system, functional area (such as development machines, production servers, or user laptops), and your IT maintenance priorities.



IMPORTANT

The different client configurations *are not mutually exclusive*. Most environments will have a mix of different ways that clients use to connect to the IdM domain. Administrators must decide which scenario is best for each individual client.

38.1.1.1. Initial Client Configuration (Pre-Migration)

Before deciding where you want to go with the client configuration in Identity Management, first establish where you are before the migration.

The initial state for almost all LDAP deployments that will be migrated is that there is an LDAP service providing identity and authentication services.




Figure 38.1. Basic LDAP Directory and Client Configuration

Linux and Unix clients use PAM_LDAP and NSS_LDAP libraries to connect directly to the LDAP services. These libraries allow clients to retrieve user information from the LDAP directory *as if* the data were stored in `/etc/passwd` or `/etc/shadow`. (In real life, the infrastructure may be more complex if a client uses LDAP for identity lookups and Kerberos for authentication or other configurations.)

There are structural differences between an LDAP directory and an IdM server, particularly in schema support and the structure of the directory tree. (For more background on those differences, see [Section 1.1.2, “Contrasting Identity Management with a Standard LDAP Directory”](#).) While those differences may impact data (especially with the directory tree, which affects entry names), they have little impact on the *client configuration*, so it really has little impact on migrating clients to Identity Management.

38.1.1.2. Recommended Configuration for Red Hat Enterprise Linux Clients

Red Hat Enterprise Linux has a service called the *System Security Services Daemon* (SSSD). SSSD uses special PAM and NSS libraries (`pam_sss` and `nss_sss`, respectively) which allow SSSD to be integrated very closely with Identity Management and leverage the full authentication and identity features in Identity Management. SSSD has a number of useful features, like caching identity information so that users can log in even if the connection is lost to the central server; these are described in the *System-Level Authentication Guide*.

Unlike generic LDAP directory services (using `pam_ldap` and `nss_ldap`), SSSD establishes relationships between identity and authentication information by defining *domains*. A domain in SSSD defines four back end functions: authentication, identity lookups, access, and password changes. The SSSD domain is then configured to use a *provider* to supply the information for any one (or all) of those four functions. An identity provider is always required in the domain configuration. The other three providers are optional; if an authentication, access, or password provider is not defined, then the identity provider is used for that function.

SSSD can use Identity Management for all of its back end functions. This is the ideal configuration because it provides the full range of Identity Management functionality, unlike generic LDAP identity providers or Kerberos authentication. For example, during daily operation, SSSD enforces host-based access control rules and security features in Identity Management.



NOTE

During the migration process from an LDAP directory to Identity Management, SSSD can seamlessly migrate user passwords without additional user interaction.




Figure 38.2. Clients and SSSD with an IdM Back End

The **ipa-client-install** script automatically configured SSSD to use IdM for all four of its back end services, so Red Hat Enterprise Linux clients are set up with the recommended configuration by default.



NOTE

This client configuration is only supported for Red Hat Enterprise Linux 6.1 and later and Red Hat Enterprise Linux 5.7 later, which support the latest versions of SSSD and **ipa-client**. Older versions of Red Hat Enterprise Linux can be configured as described in [Section 38.1.1.3, “Alternative Supported Configuration”](#).

38.1.1.3. Alternative Supported Configuration

Unix and Linux systems such as Mac, Solaris, HP-UX, AIX, and Scientific Linux support all of the services that IdM manages but do not use SSSD. Likewise, older Red Hat Enterprise Linux versions (6.1 and 5.6) support SSSD but have an older version, which does not support IdM as an identity provider.

When it is not possible to use a modern version of SSSD on a system, then clients can be configured to connect to the IdM server as if it were an LDAP directory service for identity lookups (using **nss_ldap**) and to IdM as if it were a regular Kerberos KDC (using **pam_krb5**).




Figure 38.3. Clients and IdM with LDAP and Kerberos

If a Red Hat Enterprise Linux client is using an older version of SSSD, SSSD can still be configured to use the IdM server as its identity provider and its Kerberos authentication domain; this is described in the SSSD configuration section of the *System-Level Authentication Guide*.

Any IdM domain client can be configured to use **nss_ldap** and **pam_krb5** to connect to the

IdM server. For some maintenance situations and IT structures, a scenario that fits the lowest common denominator may be required, using LDAP for both identity and authentication (**nss_ldap** and **pam_ldap**). However, it is generally best practice to use the most secure configuration possible for a client. This means SSSD or LDAP for identities and Kerberos for authentication.

38.1.2. Planning Password Migration

Probably the most visible issue that can impact LDAP-to-Identity Management migration is migrating user passwords.

Identity Management (by default) uses Kerberos for authentication and requires that each user has Kerberos hashes stored in the Identity Management Directory Server in addition to the standard user passwords. To generate these hashes, the user password needs to be available to the IdM server in clear text. When you create a user, the password is available in clear text before it is hashed and stored in Identity Management. However, when the user is migrated from an LDAP directory, the associated user password is already hashed, so the corresponding Kerberos key cannot be generated.



IMPORTANT

Users cannot authenticate to the IdM domain or access IdM resources until they have Kerberos hashes.

If a user does not have a Kerberos hash^[6], that user cannot log into the IdM domain even if he has a user account. There are three options for migrating passwords: forcing a password change, using a web page, and using SSSD.

Migrating users from an existing system provides a smoother transition but also requires parallel management of LDAP directory and IdM during the migration and transition process. If you do not preserve passwords, the migration can be performed more quickly but it requires more manual work by administrators and users.

38.1.2.1. Method 1: Using Temporary Passwords and Requiring a Change

When passwords are changed in Identity Management, they will be created with the appropriate Kerberos hashes. So one alternative for administrators is to force users to change their passwords by resetting all user passwords when user accounts are migrated. The new users are assigned a temporary password which they change at the first login. No passwords are migrated.

For details, see [Section 22.1.1, “Changing and Resetting User Passwords”](#).

38.1.2.2. Method 2: Using the Migration Web Page

When it is running in migration mode, Identity Management has a special web page in its web UI that will capture a cleartext password and create the appropriate Kerberos hash.

```
https://ipaserver.example.com/ipa/migration
```

Administrators could tell users to authenticate once to this web page, which would properly update their user accounts with their password and corresponding Kerberos hash, without requiring password changes.

38.1.2.3. Method 3: Using SSSD (Recommended)

SSSD can work with IdM to mitigate the user impact on migrating by generating the required user keys. For deployments with a lot of users or where users should not be burdened with password changes, this is the best scenario.

1. A user tries to log into a machine with SSSD.
2. SSSD attempts to perform Kerberos authentication against the IdM server.
3. Even though the user exists in the system, the authentication will fail with the error *key type is not supported* because the Kerberos hashes do not yet exist.
4. SSSD then performs a plain text LDAP bind over a secure connection.
5. IdM intercepts this bind request. If the user has a Kerberos principal but no Kerberos hashes, then the IdM identity provider generates the hashes and stores them in the user entry.
6. If authentication is successful, SSSD disconnects from IdM and tries Kerberos authentication again. This time, the request succeeds because the hash exists in the entry.

That entire process is entirely transparent to the user; as far as users know, they simply log into a client service and it works as normal.

38.1.2.4. Migrating Cleartext LDAP Passwords

Although in most deployments LDAP passwords are stored encrypted, there may be some users or some environments that use cleartext passwords for user entries.

When users are migrated from the LDAP server to the IdM server, their cleartext passwords are not migrated over. Identity Management does not allow cleartext passwords. Instead, a Kerberos principal is created for the user, the keytab is set to true, and the password is set as expired. This means that Identity Management requires the user to reset the password at the next login.



NOTE

If passwords are hashed, the password is successfully migrated through SSSD and the migration web page, as in [Section 38.1.2.2, “Method 2: Using the Migration Web Page”](#) and [Section 38.1.2.3, “Method 3: Using SSSD \(Recommended\)”](#).

38.1.2.5. Automatically Resetting Passwords That Do Not Meet Requirements

If user passwords in the original directory do not meet the password policies defined in Identity Management, then the passwords must be reset after migration.

Password resets are done automatically the first time the users attempts to **kinit** into the IdM domain.

```
[jsmith@server ~]$ kinit
Password for jsmith@EXAMPLE.COM:
Password expired. You must change it now.
```

```
Enter new password:
Enter it again:
```

38.1.3. Migration Considerations and Requirements

As you are planning a migration from an LDAP server to Identity Management, make sure that your LDAP environment is able to work with the Identity Management migration script.

38.1.3.1. LDAP Servers Supported for Migration

The migration process from an LDAP server to Identity Management uses a special script, **ipa migrate-ds**, to perform the migration. This script has certain expectations about the structure of the LDAP directory and LDAP entries in order to work. Migration is supported only for LDAPv3-compliant directory services, which include several common directories:

- Sun ONE Directory Server
- Apache Directory Server
- OpenLDAP

Migration from an LDAP server to Identity Management has been tested with Red Hat Directory Server and OpenLDAP.



NOTE

Migration using the migration script is *not* supported for Microsoft Active Directory because it is not an LDAPv3-compliant directory. For assistance with migrating from Active Directory, contact Red Hat Professional Services.

38.1.3.2. Migration Environment Requirements

There are many different possible configuration scenarios for both Red Hat Directory Server and Identity Management, and any of those scenarios may affect the migration process. For the example migration procedures in this chapter, these are the assumptions about the environment:

- A single LDAP directory domain is being migrated to one IdM realm. No consolidation is involved.
- User passwords are stored as a hash in the LDAP directory. For a list of supported hashes, see the **passwordStorageScheme** attribute in the [Password Policy Attributes table in the Red Hat Directory Server 10 Administration Guide](#).
- The LDAP directory instance is both the identity store and the authentication method. Client machines are configured to use **pam_ldap** or **nss_ldap** to connect to the LDAP server.
- Entries use only the standard LDAP schema. Entries that contain custom object classes or attributes are not migrated to Identity Management.

38.1.3.3. Migration — IdM System Requirements

With a moderately-sized directory (around 10,000 users and 10 groups), it is necessary to have a powerful enough target system (the IdM system) to allow the migration to proceed. The minimum requirements for a migration are:

- 4 cores
- 4GB of RAM
- 30GB of disk space
- A SASL buffer size of 2MB (default for an IdM server)

In case of migration errors, increase the buffer size:

```
[root@ipaserver ~]# ldapmodify -x -D 'cn=directory manager' -w
password -h ipaserver.example.com -p 389

dn: cn=config
changetype: modify
replace: nsslapd-sasl-max-buffer-size
nsslapd-sasl-max-buffer-size: 4194304

modifying entry "cn=config"
```

Set the ***nsslapd-sasl-max-buffer-size*** value in bytes.

38.1.3.4. Migration Tools

Identity Management uses a specific command, ***ipa migrate-ds***, to drive the migration process so that LDAP directory data are properly formatted and imported cleanly into the IdM server. When using ***ipa migrate-ds***, the remote system user, specified by the ***--bind-dn*** option, needs to have read access to the ***userPassword*** attribute, otherwise passwords will not be migrated.

The Identity Management server must be configured to run in migration mode, and then the migration script can be used. For details, see [Section 38.3, “Migrating an LDAP Server to Identity Management”](#).

38.1.3.5. Improving Migration Performance

An LDAP migration is essentially a specialized import operation for the 389 Directory Server instance within the IdM server. Tuning the 389 Directory Server instance for better import operation performance can help improve the overall migration performance.

There are two parameters that directly affect import performance:

- The ***nsslapd-cachememsize*** attribute, which defines the size allowed for the entry cache. This is a buffer, that is automatically set to 80% of the total cache memory size. For large import operations, this parameter (and possibly the memory cache itself) can be increased to more efficiently handle a large number of entries or entries with larger attributes.

For details how to modify the attribute using the ***ldapmodify***, see [corresponding section in the Red Hat Directory Server Performance Tuning Guide](#).

- The system **ulimit** configuration option sets the maximum number of allowed processes for a system user. Processing a large database can exceed the limit. If this happens, increase the value:

```
[root@server ~]# ulimit -u 4096
```

For further information, see Red Hat Directory Server *Performance Tuning Guide* at https://access.redhat.com/documentation/en-US/Red_Hat_Directory_Server/10/html-single/Performance_Tuning_Guide/index.html.

38.1.3.6. Migration Sequence

There are four major steps when migrating to Identity Management, but the order varies slightly depending on whether you want to migrate the server first or the clients first.

With a client-based migration, SSSD is used to change the client configuration while an IdM server is configured:

1. Deploy SSSD.
2. Reconfigure clients to connect to the current LDAP server and then fail over to IdM.
3. Install the IdM server.
4. Migrate the user data using the IdM **ipa migrate-ds** script. This exports the data from the LDAP directory, formats for the IdM schema, and then imports it into IdM.
5. Take the LDAP server offline and allow clients to fail over to Identity Management transparently.

With a server migration, the LDAP to Identity Management migration comes first:

1. Install the IdM server.
2. Migrate the user data using the IdM **ipa migrate-ds** script. This exports the data from the LDAP directory, formats it for the IdM schema, and then imports it into IdM.
3. *Optional.* Deploy SSSD.
4. Reconfigure clients to connect to IdM. It is not possible to simply replace the LDAP server. The IdM directory tree — and therefore user entry DN — is different than the previous directory tree.

While it is required that clients be reconfigured, clients do not need to be reconfigured immediately. Updated clients can point to the IdM server while other clients point to the old LDAP directory, allowing a reasonable testing and transition phase after the data are migrated.



NOTE

Do not run both an LDAP directory service and the IdM server for very long in parallel. This introduces the risk of user data being inconsistent between the two services.

Both processes provide a general migration procedure, but it may not work in every environment. Set up a test LDAP environment and test the migration process before attempting to migrate the real LDAP environment.

38.2. EXAMPLES FOR USING `IPA MIGRATE-DS`

The data migration is performed using the `ipa migrate-ds` command. At its simplest, the command takes the LDAP URL of the directory to migrate and exports the data based on common default settings.

```
ipa migrate-ds ldap://ldap.example.com:389
```

Migrated entries

The `migrate-ds` command only migrates accounts containing a `gidNumber` attribute, that is required by the `posixAccount` object class, and a `sn` attribute, that is required by the `person` object class.

Customizing the process

The `ipa migrate-ds` command enables you to customize how data is identified and exported. This is useful if the original directory tree has a unique structure or if some entries or attributes within entries should be excluded. For further details, pass the `--help` to the command.

Bind DN

By default, the DN "`cn=Directory Manager`" is used to bind to the remote LDAP directory. Pass the `--bind-dn` option to the command to specify a custom bind DN. For further information, see [Section 38.1.3.4, "Migration Tools"](#).

Naming context changes

If the Directory Server naming context differs from the one used in Identity Management, the base DN for objects is transformed. For example: `uid=user,ou=people,dc=ldap,dc=example,dc=com` is migrated to `uid=user,ou=people,dc=idm,dc=example,dc=com`. Pass the `--base-dn` to the `ipa migrate-ds` command to set the base DN used on the remote LDAP server for the migration.

38.2.1. Migrating Specific Subtrees

The default directory structure places person entries in the `ou=People` subtree and group entries in the `ou=Groups` subtree. These subtrees are container entries for those different types of directory data. If no options are passed with the `migrate-ds` command, then the utility assumes that the given LDAP directory uses the `ou=People` and `ou=Groups` structure.

Many deployments may have an entirely different directory structure (or may only want to export certain parts of the directory tree). There are two options which allow administrators to specify the RDN of a different user or group subtree on the source LDAP server:

- `--user-container`
- `--group-container`

**NOTE**

In both cases, the subtree must be the RDN only and must be relative to the base DN. For example, the **>ou=Employees,dc=example,dc=com** directory tree can be migrated using **--user-container=ou=Employees**.

For example:

```
[root@ipaserver ~]# ipa migrate-ds --user-container=ou=employees \  
--group-container="ou=employee groups" \  
ldap://ldap.example.com:389
```

Pass the **--scope** option to the **ipa migrate-ds** command, to set a scope:

- **onelevel**: Default. Only entries in the specified container are migrated.
- **subtree**: Entries in the specified container and all subcontainers are migrated.
- **base**: Only the specified object itself is migrated.

38.2.2. Specifically Including or Excluding Entries

By default, the **ipa migrate-ds** script imports every user entry with the **person** object class and every group entry with the **groupOfUniqueNames** or **groupOfNames** object class..

In some migration paths, only specific types of users and groups may need to be exported, or, conversely, specific users and groups may need to be excluded.

One option is to set positively which *types* of users and groups to include. This is done by setting which object classes to search for when looking for user or group entries.

This is a really useful option when there are custom object classes used in an environment for different user types. For example, this migrates only users with the custom **fullTimeEmployee** object class:

```
[root@ipaserver ~]# ipa migrate-ds --user-objectclass=fullTimeEmployee  
ldap://ldap.example.com:389
```

Because of the different types of groups, this is also very useful for migrating only certain types of groups (such as user groups) while excluding other types of groups, like certificate groups. For example:

```
[root@ipaserver ~]# ipa migrate-ds --group-objectclass=groupOfNames --  
group-objectclass=groupOfUniqueNames ldap://ldap.example.com:389
```

Positively specifying user and groups to migrate based on object class implicitly excludes all other users and groups from migration.

Alternatively, it can be useful to migrate all user and group entries except for just a small handful of entries. Specific user or group accounts can be excluded while all others of that type are migrated. For example, this excludes a hobbies group and two users:

```
[root@ipaserver ~]# ipa migrate-ds --exclude-groups="Golfers Group" --  
exclude-users=jsmith --exclude-users=bjensen ldap://ldap.example.com:389
```

Exclude statements are applied to users matching the pattern in the **uid** and to groups matching it in the **cn** attribute.

Specifying an object class to migrate can be used together with excluding specific entries. For example, this specifically includes users with the **fullTimeEmployee** object class, yet excludes three managers:

```
[root@ipaserver ~]# ipa migrate-ds --user-objectclass=fullTimeEmployee --
exclude-users=jsmith --exclude-users=bjensen --exclude-users=mreynolds
ldap://ldap.example.com:389
```

38.2.3. Excluding Entry Attributes


By default, every attribute and object class for a user or group entry is migrated. There are some cases where that may not be realistic, either because of bandwidth and network constraints or because the attribute data are no longer relevant. For example, if users are going to be assigned new user certificates as they join the IdM domain, then there is no reason to migrate the **userCertificate** attribute.

Specific object classes and attributes can be ignored by the **migrate-ds** by using any of several different options:

- **--user-ignore-objectclass**
- **--user-ignore-attribute**
- **--group-ignore-objectclass**
- **--group-ignore-attribute**

For example, to exclude the **userCertificate** attribute and **strongAuthenticationUser** object class for users and the **groupOfCertificates** object class for groups:

```
[root@ipaserver ~]# ipa migrate-ds --user-ignore-attribute=userCertificate
--user-ignore-objectclass=strongAuthenticationUser --group-ignore-
objectclass=groupOfCertificates ldap://ldap.example.com:389
```



NOTE

Make sure not to ignore any required attributes. Also, when excluding object classes, make sure to exclude any attributes which are only supported by that object class.

38.2.4. Setting the Schema to Use

Identity Management uses the RFC2307bis schema to define user, host, host group, and other network identities. However, if the LDAP server used as source for a migration uses the RFC2307 schema instead, pass the **--schema** option to the **ipa migrate-ds** command:

```
[root@ipaserver ~]# ipa migrate-ds --schema=RFC2307
ldap://ldap.example.com:389
```

38.3. MIGRATING AN LDAP SERVER TO IDENTITY MANAGEMENT



IMPORTANT

This is a general migration procedure, but it may not work in every environment.

It is strongly recommended that you set up a test LDAP environment and test the migration process before attempting to migrate the real LDAP environment.

1. Install the IdM server, including any custom LDAP directory schema, on a different machine from the existing LDAP directory.



NOTE

Custom user or group schemas have limited support in IdM. They can cause problems during the migration because of incompatible object definitions.

2. Disable the compat plug-in.

```
[root@server ~]# ipa-compat-manage disable
```

This step is not necessary if the data provided by the compat tree is required during the migration.

3. Restart the IdM Directory Server instance.

```
[root@server ~]# systemctl restart dirsrv.target
```

4. Configure the IdM server to allow migration:

```
[root@server ~]# ipa config-mod --enable-migration=TRUE
```

5. Run the IdM migration script, **ipa migrate-ds**. At its most basic, this requires only the LDAP URL of the LDAP directory instance to migrate:

```
[root@server ~]# ipa migrate-ds ldap://ldap.example.com:389
```

Simply passing the LDAP URL migrates all of the directory data using common default settings. The user and group data can be selectively migrated by specifying other options, as covered in [Section 38.2, “Examples for Using **ipa migrate-ds**”](#).

If the compat plug-in was not disabled in the previous step, pass the **--with-compat** option to **ipa migrate-ds**.

Once the information is exported, the script adds all required IdM object classes and attributes and converts DN's in attributes to match the IdM directory tree, if the naming context differs. For example:

`uid=user,ou=people,dc=ldap,dc=example,dc=com` is migrated to `uid=user,ou=people,dc=idm,dc=example,dc=com`.

6. Re-enable the compat plug-in, if it was disabled before the migration.

```
[root@server ~]# ipa-compat-manage enable
```

7. Restart the IdM Directory Server instance.

```
[root@server ~]# systemctl restart dirsrv.target
```

8. Disable the migration mode:

```
[root@server ~]# ipa config-mod --enable-migration=FALSE
```

9. *Optional.* Reconfigure non-SSSD clients to use Kerberos authentication (**pam_krb5**) instead of LDAP authentication (**pam_ldap**). Use PAM_LDAP modules until all of the users have been migrated; then it is possible to use PAM_KRB5. For further information, see [the corresponding section in the System-level Authentication Guide](#).
10. There are two ways for users to generate their hashed Kerberos password. Both migrate the users password without additional user interaction, as described in [Section 38.1.2, “Planning Password Migration”](#).

1. Using SSSD:

1. Move clients that have SSSD installed from the LDAP back end to the IdM back end, and enroll them as clients with IdM. This downloads the required keys and certificates.

On Red Hat Enterprise Linux clients, this can be done using the **ipa-client-install** command. For example:

```
[root@server ~]# ipa-client-install --enable-dns-update
```

2. Using the IdM migration web page:

1. Instruct users to log into IdM using the migration web page:

```
https://ipaserver.example.com/ipa/migration
```

11. To monitor the user migration process, query the existing LDAP directory to see which user accounts have a password but do not yet have a Kerberos principal key.

```
[user@server ~]$ ldapsearch -LL -x -D 'cn=Directory Manager' -w secret -b 'cn=users,cn=accounts,dc=example,dc=com' '(&(! (krbprincipalkey=*)) (userpassword=*))' uid
```



NOTE

Include the single quotes around the filter so that it is not interpreted by the shell.

12. When the migration of all clients and users is complete, decommission the LDAP directory.

38.4. MIGRATING OVER SSL

To encrypt the data transmission between LDAP and IdM during a migration:

1. Store the certificate of the CA, that issued the remote LDAP server's certificate, in a file on the IdM server. For example: **/etc/ipa/remote.crt**.
2. Follow the steps described in [Section 38.3, “Migrating an LDAP Server to Identity Management”](#). However for an encrypted LDAP connection during the migration, use the **ldaps** protocol in the URL and pass the **--ca-cert-file** option to the command. For example:

```
[root@ipaserver ~]# ipa migrate-ds --ca-cert-  
file=/etc/ipa/remote.crt ldaps://ldap.example.com:636
```

[6] It is possible to use LDAP authentication in Identity Management instead of Kerberos authentication, which means that Kerberos hashes are not required for users. However, this limits the capabilities of Identity Management and is not recommended.

APPENDIX A. TROUBLESHOOTING: GENERAL GUIDELINES

This appendix describes general steps for determining the root cause of a problem, for example by querying logs and service statuses.



NOTE

For lists of specific problems and their solutions, see [Appendix B, *Troubleshooting: Solutions to Specific Problems*](#).

What were you doing when you encountered the problem?

- [Executing a command using the **ipa** utility](#)
- [Authenticating Using **kinit**](#)
- [Authenticating to the IdM web UI](#)
- [Authenticating with a Smart Card](#)
- [Starting a Service](#)

If you know which specific area of IdM is causing the problem, follow these links:

- [DNS](#)
- [Replication](#)

If this guide does not help you find and fix the problem and you proceed to file a customer case, include any notable error output that you determined using these troubleshooting procedures in the case report. See also [Contacting Red Hat Technical Support](#).

A.1. INVESTIGATING FAILURES WHEN EXECUTING THE **IPA** UTILITY

Basic Troubleshooting

1. Add the **--verbose** (**-v**) option to the command. This displays debug information.
2. Add the **-vv** option to the command. This displays the JSON response and request.

Advanced Troubleshooting

[Figure A.1, “The architecture of executing the **ipa cert-show** command”](#) shows which components interact when the user uses the IdM command-line utility. Querying these components can help you investigate where the problem occurred and what caused it.

1. Use the following utilities:
 - **host** to check the DNS resolution of the IdM server or client
 - **ping** to check if the IdM server is available
 - **iptables** to check the current firewall configuration on the IdM server

- **date** to check the current time
- **nc** to try to connect to the required ports, as listed in [Section 2.1.4, “Port Requirements”](#)

For details on using these utilities, see their man pages.

2. Set the **KRB5_TRACE** environment variable to the **/dev/stdout** file to send trace-logging output to **/dev/stdout**:

```
$ KRB5_TRACE=/dev/stdout ipa cert-find
```

Review the Kerberos key distribution center (KDC) log: **/var/log/krb5kdc.log**.

3. Review the Apache error log:
 - a. Enable debug level on the server: Open the **/etc/ipa/server.conf** file, and add the **debug=True** option to the **[global]** section.
 - b. Restart the **httpd** service:

```
# systemctl restart httpd.service
```

- c. Run the command that failed again.
- d. Review the **httpd** error log on the server: **/var/log/httpd/error_log**.

Run the command with the **-vvv** option to display the HTTP request and response.

4. Review the Apache access log: **/var/log/httpd/access_log**.

Review the logs for the Certificate System component:

- **/var/log/pki/pki-ca-spawn.time_of_installation.log**
- **/var/log/pki/pki-tomcat/ca/debug**
- **/var/log/pki/pki-tomcat/ca/system**
- **/var/log/pki/pki-tomcat/ca/selftests.log**
- Use the **# journalctl -u pki-tomcatd@pki-tomcat.service** command to review the **journal** log.

5. Review the Directory Server access log: **/var/log/dirsrv/slapd-IPA-EXAMPLE-COM/access**.




Figure A.1. The architecture of executing the `ipa cert-show` command

Related Information

- See [Section C.2, “Identity Management Log Files and Directories”](#) for descriptions of various Identity Management log files.

A.2. INVESTIGATING `KINIT` AUTHENTICATION FAILURES

General Troubleshooting

1. On the IdM client, display the debug messages from the `kinit` process:

```
$ KRB5_TRACE=/dev/stdout kinit admin
```

2. Verify that:

- The client forward record is correct both on the server and on the affected client:

```
# host client_fully_qualified_domain_name
```

- The server forward record is correct both on the server and on the affected client:

```
# host server_fully_qualified_domain_name
```

```
# host server_IP_address
```

The `host server_IP_address` command must return a fully qualified host name with a trailing dot at the end, such as:

```
server.example.com.
```

3. Review the `/etc/hosts` file on the client, and make sure that:

- All server entries in the file are correct

- In all server entries, the first name is a fully qualified domain name

See also [the section called “The `/etc/hosts` File”](#).

4. Make sure you meet the other conditions in [Section 2.1.3, “Host Name and DNS Configuration”](#).
5. On the IdM server, make sure that the **krb5kdc** and **dirsrv** services are running:

```
# systemctl status krb5kdc
# systemctl status dirsrv.target
```

6. Review the Kerberos key distribution center (KDC) log: **`/var/log/krb5kdc.log`**.
7. If the KDCs are hard-coded in the **`/etc/krb5.conf`** file (the file explicitly sets KDC directives and uses the **`dns_lookup_kdc = false`** setting), use the **`ipactl status`** command on each master server. Check the status of the IdM services on each server listed as KDC by the command:

```
# ipactl status
Directory Service: RUNNING
krb5kdc Service: RUNNING
kadmin Service: RUNNING
named Service: RUNNING
httpd Service: RUNNING
ipa-custodia Service: RUNNING
ntpd Service: RUNNING
pki-tomcatd Service: RUNNING
ipa-otpd Service: RUNNING
ipa-dnskeysyncd Service: RUNNING
ipa: INFO: The ipactl command was successful
```

Troubleshooting Errors Cannot find KDC for realm

If **`kinit`** authentication fails with an error that says **Cannot find KDC for realm "EXAMPLE.COM" while getting initial credentials**, it indicates that KDC is not running on the server or that the client has misconfigured DNS. In this situation, try these steps:

1. If the DNS discovery is enabled in the **`/etc/krb5.conf`** file (the **`dns_lookup_kdc = true`** setting), use the **`dig`** utility to check whether the following records are resolvable:

```
$ dig -t TXT _kerberos.ipa.example.com
$ dig -t SRV _kerberos._udp.ipa.example.com
$ dig -t SRV _kerberos._tcp.ipa.example.com
```

In the following example, one of the **`dig`** commands above failed with this output:

```
; <<>> DiG 9.11.0-P2-RedHat-9.11.0-6.P2.fc25 <<>> -t SRV
_kerberos._tcp.ipa.server.example
;; global options: +cmd
;; connection timed out; no servers could be reached
```

The output indicated that the **`named`** service was not running on the master server.

2. If DNS lookup fails, continue with the steps in [Section A.6, “Troubleshooting DNS”](#).

Related Information

- See [Section C.2, “Identity Management Log Files and Directories”](#) for descriptions of various Identity Management log files.

A.3. INVESTIGATING IDM WEB UI AUTHENTICATION FAILURES

1. Make sure the user can authenticate from the command line using the **kinit** utility. If the authentication fails, see also [Section A.2, “Investigating kinit Authentication Failures”](#).
2. Make sure that the **httpd** and **dirsrv** services on the affected server are running:

```
# systemctl status httpd.service
# systemctl status dirsrv@IPA-EXAMPLE-COM.service
```

3. Make sure no related SELinux Access Vector Cache (AVC) messages are logged in the **/var/log/audit/audit.log** and **/var/log/messages** files.

See [Basic SELinux Troubleshooting in CLI](#) in the Red Hat Knowledgebase for details on resolving AVC messages.

4. Make sure that cookies are enabled on the browser from which you are authenticating.
5. Make sure that the time difference between the IdM server and the system on which you are authenticating is 5 minutes at the most.
6. Review the Apache error log: **/var/log/httpd/error_log**.
7. Enable verbose logging for the authentication process to help diagnose the problem. See [Firefox Configuration for Kerberos Troubleshooting](#) in the *System-Level Authentication Guide* for advice on how to enable verbose logging in Firefox.

If you are having problems when logging in using certificates:

1. In the **/etc/httpd/conf.d/nss.conf** file, change the **LogLevel** attribute to **info**.
2. Restart the Apache server:

```
# systemctl restart httpd
```

3. Try logging in with the certificate again.
4. Review the Apache error log: **/var/log/httpd/error_log**.

The log shows messages recorded by the **mod_lookup_identity** module, including information about whether the module successfully matched the user during the login attempt or not.

Related Information

- See [Section C.2, “Identity Management Log Files and Directories”](#) for descriptions of

various Identity Management log files.

A.4. INVESTIGATING SMART CARD AUTHENTICATION FAILURES

1. Open the `/etc/sss/sss.conf` file, and set the `debug_level` option to 2.
2. Review the `sss_pam.log` and `sss_EXAMPLE.COM.log` files. If you see timeout error message in the files, see [Section B.4.4, “Smart Card Authentication Fails with Timeout Error Messages”](#).

A.5. INVESTIGATING WHY A SERVICE FAILS TO START

1. Review the log for the service that fails to start. See [Section C.2, “Identity Management Log Files and Directories”](#).

For example, the log for Directory Server is at `/var/log/dirsrv/slapd-IPA-EXAMPLE-COM/errors`.

2. Make sure that the server on which the service is running has a fully qualified domain name (FQDN). See [the section called “Verifying the Server Host Name”](#).
3. If the `/etc/hosts` file contains an entry for the server on which the service is running, make sure the fully qualified domain name is listed first. See also [the section called “The /etc/hosts File”](#).
4. Make sure you meet the other conditions in [Section 2.1.3, “Host Name and DNS Configuration”](#).
5. Determine what keys are included in the keytab that is used for authentication of the service. For example, for the `dirsrv` service ticket:

```
# klist -kt /etc/dirsrv/ds.keytab
Keytab name: FILE:/etc/dirsrv/ds.keytab
KVNO Timestamp                Principal
-----
2 01/10/2017 14:54:39 ldap/server.example.com@EXAMPLE.COM
2 01/10/2017 14:54:39 ldap/server.example.com@EXAMPLE.COM
[... output truncated ...]
```

- a. Make sure that the displayed principals match the system's FQDN.
- b. Make sure that the displayed version of the keys (KVNO) in the above-mentioned service keytab match the KVNO in the server keytab. To display the server keytab:

```
$ kinit admin
$ kvno ldap/server.example.com@EXAMPLE.COM
```

- c. Verify that the forward (A, AAAA, or both) and reverse records on the client match the displayed system name and service principal.

6. Verify that the forward (A, AAAA, or both) and reverse records on the client are correct.
7. Make sure that the system time difference on the client and the server is 5 minutes at the most.
8. Services can fail to start after the IdM administrative server certificates expire. To check if this is the cause in your case:
 - a. Use the **getcert list** command to list all certificates tracked by the **certmonger** utility.
 - b. In the output, find the IdM administrative certificates: the **ldap** and **httpd** server certificates.
 - c. Examine the fields labeled **status** and **expires**.

```
# getcert list
Number of certificates and requests being tracked: 8.
[... output truncated ...]
Request ID '20170421124617':
  status: MONITORING
  stuck: no
  key pair storage: type=NSSDB,location='/etc/dirsrv/slapd-IPA-EXAMPLE-COM',nickname='Server-Cert',token='NSS Certificate DB',pinfile='/etc/dirsrv/slapd-IPA-EXAMPLE-COM/pwdfile.txt'
  certificate: type=NSSDB,location='/etc/dirsrv/slapd-IPA-EXAMPLE-COM',nickname='Server-Cert',token='NSS Certificate DB'
  CA: IPA
  issuer: CN=Certificate Authority,O=IPA.EXAMPLE.COM
  subject: CN=ipa.example.com,O=IPA.EXAMPLE.COM
  expires: 2019-04-22 12:46:17 UTC
[... output truncated ...]
Request ID '20170421130535':
  status: MONITORING
  stuck: no
  key pair storage:
  type=NSSDB,location='/etc/httpd/alias',nickname='Server-Cert',token='NSS Certificate DB',pinfile='/etc/httpd/alias/pwdfile.txt'
  certificate:
  type=NSSDB,location='/etc/httpd/alias',nickname='Server-Cert',token='NSS Certificate DB'
  CA: IPA
  issuer: CN=Certificate Authority,O=IPA.EXAMPLE.COM
  subject: CN=ipa.example.com,O=IPA.EXAMPLE.COM
  expires: 2019-04-22 13:05:35 UTC
[... output truncated ...]
```

If you need to start the service even though the certificates are expired, see [Section 26.5, “Allowing IdM to Start with Expired Certificates”](#).

A.6. TROUBLESHOOTING DNS

1. Many DNS problems are caused by misconfiguration. Therefore, make sure you meet the conditions in [Section 2.1.3, “Host Name and DNS Configuration”](#).

2. Use the **dig** utility to check the response from the DNS server:

```
# dig _ldap._tcp.ipa.example.com. SRV

; <<>> DiG 9.9.4-RedHat-9.9.4-48.el7 <<>>
_ldap._tcp.ipa.example.com. SRV
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 17851
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 1,
ADDITIONAL: 5

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
_ldap._tcp.ipa.example.com. IN SRV

;; ANSWER SECTION:
_ldap._tcp.ipa.example.com. 86400 IN SRV      0 100 389
ipaserver.ipa.example.com.

;; AUTHORITY SECTION:
ipa.example.com.          86400 IN NS
ipaserver.ipa.example.com.

;; ADDITIONAL SECTION:
ipaserver.ipa.example.com. 86400 IN A 192.0.21
ipaserver.ipa.example.com 86400 IN AAAA 2001:db8::1
```

3. Use the **host** utility to perform a DNS name lookup:

```
$ host server.ipa.example.com
server.ipa.example.com. 86400 IN A 192.0.21
server.ipa.example.com 86400 IN AAAA 2001:db8::1
```

4. Review the DNS records in LDAP using the **ipa dnszone-show** command:

```
$ ipa dnszone-show zone_name
$ ipa dnsrecord-show zone_name record_name_in_the_zone
```

For details on using the IdM tools to manage DNS, see [Chapter 33, Managing DNS](#).

5. Restart BIND to force resynchronization with LDAP:

```
$ systemctl restart named-pkcs11
```

6. Get a list of the required DNS records:

```
$ ipa dns-update-system-records --dry-run
```

Use the **dig** utility to check if the displayed records are present in DNS. If you use the Identity Management DNS, use the **ipa dns-update-system-records** command to update any missing records.

A.7. TROUBLESHOOTING REPLICATION

Test replication on at least two servers (see [Section 4.6, “Testing the New Replica”](#)). If changes made on one IdM server are not replicated to the other server:

1. Make sure you meet the conditions in [Section 2.1.3, “Host Name and DNS Configuration”](#).
2. Make sure that both servers can resolve each other's forward and reverse DNS records:

```
[root@server1 ~]# dig +short server2.example.com A
[root@server1 ~]# dig +short server2.example.com AAAA
[root@server1 ~]# dig +short -x server2_IPv4_or_IPv6_address
```

```
[root@server2 ~]# dig +short server1.example.com A
[root@server2 ~]# dig +short server1.example.com AAAA
[root@server2 ~]# dig +short -x server1_IPv4_or_IPv6_address
```

3. Make sure that the time difference on both servers is 5 minutes at the most.
4. Review the Directory Server error log on both servers:
/var/log/dirsrv/slapd-SERVER-EXAMPLE-COM/errors.
5. If you see errors related to Kerberos, make sure that the Directory Server keytab is correct and that you can use it to query the other server (**server2** in this example):

```
[root@server1 ~]# kinit -kt /etc/dirsrv/ds.keytab
ldap/server1.example.com
[root@server1 ~]# klist
[root@server1 ~]# ldapsearch -Y GSSAPI -h server1.example.com -b ""
-s base
[root@server1 ~]# ldapsearch -Y GSSAPI -h server2_FQDN. -b "" -s
base
```

Related Information

- See [Section C.2, “Identity Management Log Files and Directories”](#) for descriptions of various Identity Management log files.

APPENDIX B. TROUBLESHOOTING: SOLUTIONS TO SPECIFIC PROBLEMS

For troubleshooting advice for:

- Servers, see [Section B.1, “Identity Management Servers”](#)
- Replicas, see [Section B.2, “Identity Management Replicas”](#)
- Clients, see [Section B.3, “Identity Management Clients”](#)
- Authentication, see [Section B.4, “Logging In and Authentication Problems”](#)
- Vaults, see [Section B.5, “Vaults”](#)

B.1. IDENTITY MANAGEMENT SERVERS

B.1.1. External CA Installation Fails

The `ipa-server-install --external-ca` command fails with the following error:

```
ipa          : CRITICAL failed to configure ca instance Command
'/usr/sbin/pkispawn -s CA -f /tmp/configuration_file' returned non-zero
exit status 1
Configuration of CA failed
```

The `env|grep proxy` command displays variables such as the following:

```
env|grep proxy
http_proxy=http://example.com:8080
ftp_proxy=http://example.com:8080
https_proxy=http://example.com:8080
```

What this means:

The `*_proxy` environmental variables are preventing the server from being installed.

To fix the problem:

1. Use the following shell script to unset the `*_proxy` environmental variables:

```
# for i in ftp http https; do unset ${i}_proxy; done
```

2. Run the `pkidestroy` utility to remove the unsuccessful CA subsystem installation:

```
# pkidestroy -s CA -i pki-tomcat; rm -rf /var/log/pki/pki-tomcat
/etc/sysconfig/pki-tomcat /etc/sysconfig/pki/tomcat/pki-tomcat
/var/lib/pki/pki-tomcat /etc/pki/pki-tomcat /root/ipa.csr
```

3. Remove the failed IdM server installation:

```
# ipa-server-install --uninstall
```

4. Retry running `ipa-server-install --external-ca`.

B.1.2. named Daemon Fails to Start

After installing an IdM server with integrated DNS, the **named-pkcs11** fails to start. The `/var/log/messages` file includes an error message related to the **named-pkcs11** service and the `ldap.so` library:

```
ipaserver named[6886]: failed to dynamically load driver 'ldap.so':
libldap-2.4.so.2: cannot open shared object file: No such file or
directory
```

What this means:

The `bind-chroot` package is installed and is preventing the **named-pkcs11** service from starting.

To fix the problem:

1. Uninstall the `bind-chroot` package.

```
# yum remove bind-chroot
```

2. Restart the IdM server.

```
# ipactl restart
```

B.1.3. Installing a Server Fails on a System with IPv6 Disabled

When attempting to install an IdM server on a system with IPv6 disabled, the following error occurs during the installation process:

```
CRITICAL Failed to restart the directory server
Command '/bin/systemctl restart dirsrv@EXAMPLE.service' returned non-zero
exit status 1
```

What this means:

Installing and running a server requires IPv6 to be enabled on the network. See [Section 2.1.2, “System Requirements”](#).

To fix the problem:

Enable IPv6 on your system. For details, see [How do I disable or enable the IPv6 protocol in Red Hat Enterprise Linux?](#) in Red Hat Knowledgebase.

Note that IPv6 is enabled by default on Red Hat Enterprise Linux 7 systems.

B.2. IDENTITY MANAGEMENT REPLICAS

This guide describes common replication problems for Identity Management in Red Hat Enterprise Linux.

Additional resources:

- For troubleshooting advice on replication in Red Hat Directory Server, see [Solving Naming Conflicts](#) in the *Directory Server Administration Guide*.
- For advice on how to test that replication is working, see [Section 4.6, “Testing the New Replica”](#).
- The Directory Server **repl-monitor** script shows in-progress status of replication, which can help you troubleshoot replication problems. For documentation on the script, see [Monitoring Replication from the Command-Line](#) in the *Directory Server Administration Guide*.

B.2.1. Authenticating AD Users Against a New Replica Fails

After installing a new replica in an Identity Management - Active Directory trust setup, attempts to authenticate Active Directory (AD) users against the IdM replica fail.

What this means:

The replica is neither a trust controller nor trust agent. Because of this, it cannot serve information from the AD trust.

To fix the problem:

Configure the replica as a trust agent. See [Trust Controllers and Trust Agents](#) in the *Windows Integration Guide*.

B.2.2. Replica Starts with SASL, GSS-API, and Kerberos Errors in the Directory Server Logs

When the replica starts, a series of SASL bind errors are recorded in the Directory Server (DS) logs. The errors state the GSS-API connection failed because it could not find a credentials cache:

```
slapd_ldap_sasl_interactive_bind - Error: could not perform interactive
bind for id [] mech [GSSAPI]: error -2 (Local error) (SASL(-1): generic
failure: GSSAPI Error: Unspecified GSS failure. Minor code may provide
more information (Credentials cache file '/tmp/krb5cc_496' not found)) ...
```

Additionally, other messages can occur stating that the server could not obtain Kerberos credentials for the host principal:

```
set_krb5_creds - Could not get initial credentials for principal [ldap/
replica1.example.com] in keytab [WRFILE:/etc/dirsrv/ds.keytab]: -
1765328324 (Generic error)
```

What this means:

IdM uses GSS-API for Kerberos connections. The DS instance keeps the Kerberos credentials cache in memory. When the DS process ends, such as when the IdM replica stops, the credentials cache is destroyed.

When the replica restarts, DS starts before the KDC server starts. Because of this start order, the Kerberos credentials are not yet saved in the credentials cache when DS starts, which is what causes the errors.

After the initial failure, DS re-attempts to establish the GSS-API connection after the KDC starts. This second attempt is successful and ensures that the replica works as expected.

You can ignore the described startup errors as long as the GSS-API connection is successfully established and the replica works as expected. The following message shows that the connection was successful:

```
Replication bind with GSSAPI auth resumed
```

B.2.3. The DNS Forward Record Does Not Match the Reverse Address

When configuring a new replica, installation fails with a series of certificate errors, followed by a DNS error stating the DNS forward record does not match the reverse address.

```
ipa: DEBUG: approved_usage = SSLServer intended_usage = SSLServer
ipa: DEBUG: cert valid True for "CN=replica.example.com,O=EXAMPLE.COM"
ipa: DEBUG: handshake complete, peer = 192.0.2.2:9444
Certificate operation cannot be completed: Unable to communicate with CMS
(Not Found)

...

ipa: DEBUG: Created connection context.ldap2_21534032
ipa: DEBUG: Destroyed connection context.ldap2_21534032
The DNS forward record replica.example.com. does not match the reverse
address replica.example.org
```

What this means:

Multiple host names are used for a single PTR record. The DNS standard allows such configuration, but it causes an IdM replica installation to fail.

To fix the problem:

Verify the DNS configuration, as described in [the section called “Verifying the Forward and Reverse DNS Configuration”](#).

B.2.4. Serial Numbers Not Found Errors



NOTE

This solution is applicable at domain level 0. See [Chapter 7, Displaying and Raising the Domain Level](#) for details.

An error stating that a certificate serial number was not found appears on a replicated server:

```
Certificate operation cannot be completed: EXCEPTION (Certificate serial
number 0x2d not found)
```

What this means:

A certificate replication agreement between two replicas has been removed but a data replication agreement is still in place. Both replicas are still issuing certificates, but information about the certificates is no longer replicated.

Example situation:

1. Replica A issues a certificate to a host.

2. The certificate is not replicated to replica B, because the replicas have no certificate replication agreement established.
3. A user attempts to use replica B to manage the host.
4. Replica B returns an error that it cannot verify the host's certificate serial number. This is because replica B has information about the host in its data directory, but it does not have the host certificate in its certificate directory.

To fix the problem:

1. Enable certificate server replication between the two replicas using the **ipa-csreplica-manage connect** command. See [Section D.3.3, “Creating and Removing Replication Agreements”](#).
2. Re-initialize one of the replicas from the other to synchronize them. See [Section D.3.5, “Re-initializing a Replica”](#).

**WARNING**

Re-initializing overwrites data on the re-initialized replica with the data from the other replica. Some information might be lost.

B.2.5. Cleaning Replica Update Vector (RUV) Errors**NOTE**

This solution is applicable at domain level **0**. See [Chapter 7, *Displaying and Raising the Domain Level*](#) for details.

After a replica has been removed from the IdM topology, obsolete RUV records are now present on one or more remaining replicas.

Possible causes:

- The replica has been removed without properly removing its replication agreements first, as described in [the section called “Removing Replication Agreements”](#).
- The replica has been removed when another replica was offline.

What this means:

The other replicas still expect to receive updates from the removed replica.

**NOTE**

The correct procedure for removing a replica is described in [Section D.3.6, “Removing a Replica”](#).

To fix the problem:

Clean the RUV records on the replica that expects to receive the updates.

1. List the details about the obsolete RUVs using the **ipa-replica-manage list-ruv** command. The command displays the replica IDs:

```
# ipa-replica-manage list-ruv
server1.example.com:389: 6
server2.example.com:389: 5
server3.example.com:389: 4
server4.example.com:389: 12
```

2. Clear the corrupt RUVs using the **ipa-replica-manage clean-ruv *replica_ID*** command. The command removes any RUVs associated with the specified replica.

Repeat the command for every replica with obsolete RUVs. For example:

```
# ipa-replica-manage clean-ruv 6
# ipa-replica-manage clean-ruv 5
# ipa-replica-manage clean-ruv 4
# ipa-replica-manage clean-ruv 12
```



WARNING

Proceed with extreme caution when using **ipa-replica-manage clean-ruv**. Running the command against a valid replica ID will corrupt all the data associated with that replica in the replication database.

If this happens, re-initialize the replica from another replica as described in [Section D.3.5, “Re-initializing a Replica”](#).

3. Run **ipa-replica-manage list-ruv** again.

- If the command no longer displays any corrupt RUVs, the records have been successfully cleaned.
- If the command still displays corrupt RUVs, clear them manually using this task:

```
dn: cn=clean replica_ID, cn=cleanallruv, cn=tasks, cn=config
objectclass: extensibleObject
replica-base-dn: dc=example,dc=com
replica-id: replica_ID
replica-force-cleaning: no
cn: clean replica_ID
```

If you are not sure on which replica to clean the RUVs:

1. Search all your servers for active replica IDs. Make a list of uncorrupted and reliable replica IDs.

To find the IDs of valid replicas, run this LDAP query for all the nodes in your topology:

```
# ldapsearch -p 389 -h IdM_node -D "cn=directory manager" -W -b  
"cn=config" "(objectclass=nsds5replica)" nsDS5ReplicaId
```

2. Run **ipa-replica-manage list-ruv** on every server. Note any replica IDs that are not on the list of uncorrupted replica IDs.
3. Run **ipa-replica-manage clean-ruv *replica_ID*** for every corrupted replica ID.

B.2.6. Recovering a Lost CA Server



NOTE

This solution is applicable at domain level **0**. See [Chapter 7, *Displaying and Raising the Domain Level*](#) for details.

You only had one server with CA installed. This server failed and is now lost.

What this means:

The CA configuration for your IdM domain is no longer available.

To fix the problem:

If you have a backup of the original CA server available, you can restore the server and install the CA on a replica.

1. Recover the CA server from backup. See [Section 9.2, “Restoring a Backup”](#) for details.

This makes the CA server available to the replica.

2. Delete the replication agreements between the initial server and the replica to avoid replication conflicts. See [Section D.3.3, “Creating and Removing Replication Agreements”](#).
3. Install the CA on the replica. See [Section 6.5.2, “Promoting a Replica to a Master CA Server”](#).
4. Decommission the original CA server. See [Section D.3.6, “Removing a Replica”](#).

If you do not have a backup of the original CA server, the CA configuration was lost when the server failed and cannot be recovered.

B.3. IDENTITY MANAGEMENT CLIENTS

This section describes common client problems for IdM in Red Hat Enterprise Linux.

Additional resources:

- To validate your **/etc/sss.conf** file, see [SSSD Configuration Validation](#) in the *System-Level Authentication Guide*.

B.3.1. The Client Is Unable to Resolve Reverse Lookups when Using an External DNS

An external DNS server returns a wrong host name for the IdM server. The following errors related to the IdM server appear in the Kerberos database:

```
Jun 30 11:11:48 server1 krb5kdc[1279](info): AS_REQ (4 etypes {18 17 16 23}) 192.0.2.1: NEEDED_PREAUTH: admin EXAMPLE COM for krbtgt/EXAMPLE COM
EXAMPLE COM, Additional pre-authentication required
Jun 30 11:11:48 server1 krb5kdc[1279](info): AS_REQ (4 etypes {18 17 16 23}) 192.0.2.1: ISSUE: authtime 1309425108, etypes {rep=18 tkt=18 ses=18},
admin EXAMPLE COM for krbtgt/EXAMPLE COM EXAMPLE COM
Jun 30 11:11:49 server1 krb5kdc[1279](info): TGS_REQ (4 etypes {18 17 16 23}) 192.0.2.1: UNKNOWN_SERVER: authtime 0, admin EXAMPLE COM for
HTTP/server1.wrong.example.com@EXAMPLE.COM, Server not found in Kerberos
database
```

What this means:

The external DNS name server returns the wrong host name for the IdM server or returns no answer at all.

To fix the problem:

1. Verify your DNS configuration, and make sure the DNS domains used by IdM are properly delegated. See [Section 2.1.3, “Host Name and DNS Configuration”](#) for details.
2. Verify your reverse (PTR) DNS records settings. See [Chapter 33, Managing DNS](#) for details.

B.3.2. The Client Is Not Added to the DNS Zone

When running the `ipa-client-install` utility, the `nsupdate` utility fails to add the client to the DNS zone.

What this means:

The DNS configuration is incorrect.

To fix the problem:

1. Verify your configuration for DNS delegation from the parent zone to IdM. See [Section 2.1.3, “Host Name and DNS Configuration”](#) for details.
2. Make sure that dynamic updates are allowed in the IdM zone. See [Section 33.5.1, “Enabling Dynamic DNS Updates”](#) for details.

For details on managing DNS in IdM, see [Section 33.7, “Managing Reverse DNS Zones”](#). For details on managing DNS in Red Hat Enterprise Linux, see [Section 11.2.3. “Editing Zone Files”](#) in the *Networking Guide*.

B.3.3. Client Connection Problems

Users cannot log in to a machine. Attempts to access user and group information, such as with the `getent passwd admin` command, fail.

What this means:

Client authentication problems often indicate problems with the System Security Services Daemon (SSSD) service.

To fix the problem:

Examine the SSSD logs in the `/var/log/sss/` directory. The directory includes a log file for the DNS domain, such as `sss_example.com.log`.

If the logs do not include enough information, increase the log level:

1. In the `/etc/sss/sss.conf` file, look up the `[domain/example.com]` section. Adjust the `debug_level` option to record more information in the logs.

```
debug_level = 9
```

2. Restart the `sss` service.

```
# systemctl start sss
```

3. Examine `sss_example.com.log` again. The file now includes more error messages.

B.4. LOGGING IN AND AUTHENTICATION PROBLEMS

B.4.1. Kerberos GSS Failures When Running `ipa` Commands

Immediately after installing a server, Kerberos errors occur when attempting to run an `ipa` command. For example:

```
ipa: ERROR: Kerberos error: ('Unspecified GSS failure. Minor code may provide more information', 851968)/('Decrypt integrity check failed', -1765328353)
```

What this means:

DNS is not properly configured.

To fix the problem:

Verify your DNS configuration.

- See [Section 2.1.3, “Host Name and DNS Configuration”](#) for DNS requirements for IdM servers.
- See [DNS and Realm Settings](#) in the *Windows Integration Guide* for DNS requirements for Active Directory trust.

B.4.2. SSH Connection Fails when Using GSS-API

Users are unable to log in to IdM machines using SSH.

What this means:

When SSH attempts to connect to an IdM resource using GSS-API as the security method, GSS-API first verifies the DNS records. SSH failures are often caused by incorrect reverse DNS entries. The incorrect records prevent SSH from locating the IdM resource.

To fix the problem:

Verify your DNS configuration as described in [Section 2.1.3, “Host Name and DNS Configuration”](#).

As a temporary workaround, you can also disable reverse DNS lookups in the SSH configuration. To do this, set the **GSSAPITrustDNS** to **no** in the **/etc/ssh/ssh_config** file. Instead of using reverse DNS records, SSH will pass the given user name directly to GSS-API.

B.4.3. OTP Token Out of Sync

Authentication using OTP fails because the token is desynchronized.

To fix the problem:

Resynchronize the token. Any user can resynchronize their tokens regardless of the token type and whether or not the user has permission to modify the token settings.

1. In the IdM web UI: Click **Sync OTP Token** on the login page.





Figure B.1. Sync OTP Token

From the command line: Run the **ipa otptoken-sync** command.

2. Provide the information required to resynchronize the token. For example, IdM will ask you to provide your standard password and two subsequent token codes generated by the token.



NOTE

Resynchronization works even if the standard password is expired. After the token is resynchronized using an expired password, log in to IdM to let the system prompt you to change the password.

B.4.4. Smart Card Authentication Fails with Timeout Error Messages

The **sssd_pam.log** and **sssd_EXAMPLE.COM.log** files contain timeout error messages, such as these:

```
Wed Jun 14 18:24:03 2017) [sssd[pam]] [child_handler_setup] (0x2000):
Setting up signal handler up for pid [12370]
(Wed Jun 14 18:24:03 2017) [sssd[pam]] [child_handler_setup] (0x2000):
Signal
```



```
handler set up for pid [12370]
(Wed Jun 14 18:24:08 2017) [sssd[pam]] [pam_initgr_cache_remove] (0x2000):
[idmeng] removed from PAM initgroup cache
(Wed Jun 14 18:24:13 2017) [sssd[pam]] [p11_child_timeout] (0x0020):
Timeout
reached for p11_child.
(Wed Jun 14 18:24:13 2017) [sssd[pam]] [pam_forwarder_cert_cb] (0x0040):
get_cert request failed.
(Wed Jun 14 18:24:13 2017) [sssd[pam]] [pam_reply] (0x0200): pam_reply
called
with result [4]: System error.
```

What this means:

When using forwarded smart card readers or the Online Certificate Status Protocol (OCSP), you might need to adjust certain default values for users to be able to authenticate with smart cards.

To fix the problem:

On the server and on the client from which you want users to authenticate, make these changes in the `/etc/sss/sss.conf` file:

1. In the `[pam]` section, increase the `p11_child_timeout` value to 60 seconds.
2. In the `[domain/EXAMPLE.COM]` section, increase the `krb5_auth_timeout` value to 60 seconds.
3. If you are using OCSP in the certificate, make sure the OCSP server is reachable. If the OCSP server is not directly reachable, configure a proxy OCSP server by adding the following options to `/etc/sss/sss.conf`:

```
certificate_verification =
ocsp_default_responder=http://ocsp.proxy.url,
ocsp_default_responder_signing_cert=nickname
```

Replace *nickname* with the nickname of the OCSP signing certificate in the `/etc/pki/nssdb/` directory.

For details on these options, see the `sss.conf(5)` man page.

B.5. VAULTS

B.5.1. Users Cannot Access Their Vault Due To Insufficient 'add' Privilege

A user is unable to access his or her own user vault or add new user vaults. The following error message appears:

```
ipa: ERROR: Insufficient access: Insufficient 'add' privilege to add the
entry 'cn=testvault,cn=user,cn=users,cn=vaults,cn=kra,dc=example,dc=com'.
```

What this means:

The user's vault container is owned by another user. Typically, this situation occurs after another user, such as **admin**, creates the first user vault for the first user. The first user then cannot access any vaults in his or her own vault container.

To fix the problem:

Add the intended user as the owner of the vault container:

1. Log in as **admin**.

```
$ kinit admin
```

2. Add *user* as the container owner:

```
$ ipa vaultcontainer-add-owner --user=user --users=user
Owner users: admin, user
Vault user: user
-----
Number of owners added 1
-----
```

Both **admin** and *user* now have access to the user's vault container because they are both the owners of the container.

3. *Optional*. Verify that the user can now create a new user vault:

```
$ kinit user
$ ipa vault-add testvault2
-----
Added vault "testvault2"
-----
```

Additional Resources

- [Section 25.4, “Storing a User's Personal Secret”](#)

APPENDIX C. A REFERENCE OF IDENTITY MANAGEMENT FILES AND LOGS

C.1. IDENTITY MANAGEMENT CONFIGURATION FILES AND DIRECTORIES

Table C.1. IdM Server and Client Configuration Files and Directories

Directory or File	Description
<code>/etc/ipa/</code>	The main IdM configuration directory.
<code>/etc/ipa/default.conf</code>	Primary configuration file for IdM. Referenced when servers and clients start and when the user uses the ipa utility.
<code>/etc/ipa/server.conf</code>	An optional configuration file, does not exist by default. Referenced when the IdM server starts. If the file exists, it takes precedence over /etc/ipa/default.conf .
<code>/etc/ipa/cli.conf</code>	An optional configuration file, does not exist by default. Referenced when the user uses the ipa utility. If the file exists, it takes precedence over /etc/ipa/default.conf .
<code>/etc/ipa/ca.crt</code>	The CA certificate issued by the IdM server's CA.
<code>~/.ipa/</code>	The user-specific IdM directory created on the local system the first time the user runs an IdM command. Users can set individual configuration overrides by creating user-specific default.conf , server.conf , or cli.conf files in ~/.ipa/ .
<code>/etc/sss/sss.conf</code>	Configuration for the IdM domain and for IdM services used by SSSD.
<code>/usr/share/sss/sss.api.d/sss-ipa.conf</code>	A schema of IdM-related SSSD options and their values.
<code>/etc/gssproxy/</code>	The directory for the configuration of the GSS-Proxy protocol. The directory contains files for each GSS-API service and a general /etc/gssproxy/gssproxy.conf file.

Table C.2. System Service Files and Directories

Directory or File	Description
/etc/sysconfig/	systemd -specific files

Table C.3. Web UI Files and Directories

Directory or File	Description
/etc/ipa/html/	A symbolic link for the HTML files used by the IdM web UI.
/etc/httpd/conf.d/ipa.conf	Configuration files used by the Apache host for the web UI application.
/etc/httpd/conf.d/ipa-rewrite.conf	
/etc/httpd/conf/ipa.keytab	The keytab file used by the web server.
/usr/share/ipa/	The directory for all HTML files, scripts, and stylesheets used by the web UI.
/usr/share/ipa/ipa.conf	
/usr/share/ipa/updates/	Contains LDAP data, configuration, and schema updates for IdM.
/usr/share/ipa/html/	Contains the HTML files, JavaScript files, and stylesheets used by the web UI.
/usr/share/ipa/migration/	Contains HTML pages, stylesheets, and Python scripts used for running the IdM server in migration mode.
/usr/share/ipa/ui/	Contains the scripts used by the UI to perform IdM operations.
/etc/httpd/conf.d/ipa-pki-proxy.conf	The configuration file for web-server-to-Certificate-System bridging.

Table C.4. Kerberos Files and Directories

Directory or File	Description
/etc/krb5.conf	The Kerberos service configuration file.
/var/lib/sss/pubconf/krb5.include.d/	Includes IdM-specific overrides for Kerberos client configuration.

Table C.5. Directory Server Files and Directories

Directory or File	Description
<code>/var/lib/dirsrv/slapd-<i>REALM_NAME</i>/</code>	The database associated with the Directory Server instance used by the IdM server.
<code>/etc/sysconfig/dirsrv</code>	IdM-specific configuration of the dirsrv systemd service.
<code>/etc/dirsrv/slapd-<i>REALM_NAME</i>/</code>	The configuration and schema files associated with the Directory Server instance used by the IdM server.

Table C.6. Certificate System Files and Directories

Directory or File	Description
<code>/etc/pki/pki-tomcat/ca/</code>	The main directory for the IdM CA instance.
<code>/var/lib/pki/pki-tomcat/conf/ca/CS.cfg</code>	The main configuration file for the IdM CA instance.

Table C.7. Cache Files and Directories

Directory or File	Description
<code>~/.cache/ipa/</code>	Contains a per-server API schema for the IdM client. IdM caches the API schema on the client for one hour.

Table C.8. System Backup Files and Directories

Directory or File	Description
<code>/var/lib/ipa/sysrestore/</code>	Contains backups of the system files and scripts that were reconfigured when the IdM server was installed. Includes the original .conf files for NSS, Kerberos (both krb5.conf and kdc.conf), and NTP.
<code>/var/lib/ipa-client/sysrestore/</code>	Contains backups of the system files and scripts that were reconfigured when the IdM client was installed. Commonly, this is the sssd.conf file for SSSD authentication services.

C.2. IDENTITY MANAGEMENT LOG FILES AND DIRECTORIES

Table C.9. IdM Server and Client Log Files and Directories

Directory or File	Description
<code>/var/log/ipaserver-install.log</code>	The installation log for the IdM server.
<code>/var/log/ipareplica-install.log</code>	The installation log for the IdM replica.
<code>/var/log/ipaclient-install.log</code>	The installation log for the IdM client.
<code>/var/log/sss/</code>	Log files for SSSD.
<code>~/.ipa/log/cli.log</code>	The log file for errors returned by XML-RPC calls and responses by the ipa utility. Created in the home directory for the <i>system user</i> who runs the tools, who might have a different user name than the IdM user.
<code>/etc/logrotate.d/</code>	The log rotation policies for DNS, SSSD, Apache, Tomcat, and Kerberos.

Table C.10. Apache Server Log Files

Directory or File	Description
<code>/var/log/httpd/</code>	Log files for the Apache web server.
<code>/var/log/httpd/access_log</code>	Standard access and error logs for Apache servers. Messages specific to IdM are recorded along with the Apache messages because the IdM web UI and the XML-RPC command-line interface use Apache.
<code>/var/log/httpd/error_log</code>	
For details, see Log Files in the Apache documentation.	

Table C.11. Certificate System Log Files

Directory or File	Description
<code>/var/log/pki/pki-ca-spawn.time_of_installation.log</code>	The installation log for the IdM CA.
<code>/var/log/pki/pki-kra-spawn.time_of_installation.log</code>	The installation log for the IdM KRA.
<code>/var/log/pki/pki-tomcat/</code>	The top level directory for PKI operation logs. Contains CA and KRA logs.

Directory or File	Description
<code>/var/log/pki/pki-tomcat/ca/</code>	Directory with logs related to certificate operations. In IdM, these logs are used for service principals, hosts, and other entities which use certificates.
<code>/var/log/pki/pki-tomcat/kra</code>	Directory with logs related to KRA.
<code>/var/log/messages</code>	Includes certificate error messages among other system messages.
For details, see Configuring Subsystem Logs in the Red Hat Certificate System <i>Administration Guide</i> .	

Table C.12. Directory Server Log Files

Directory or File	Description
<code>/var/log/dirsrv/slapd-<i>REALM_NAME</i>/</code>	Log files associated with the Directory Server instance used by the IdM server. Most operational data recorded here are related to server-replica interactions.
<code>/var/log/dirsrv/slapd-<i>REALM_NAME</i>/access</code>	Contain detailed information about attempted access and operations for the domain Directory Server instance.
<code>/var/log/dirsrv/slapd-<i>REALM_NAME</i>/errors</code>	
<code>/var/log/dirsrv/slapd-<i>REALM_NAME</i>/audit</code>	Contains audit trails of all Directory Server operations when auditing is enabled in the Directory Server configuration.
For details, see Monitoring Server and Database Activity and Log File Reference in the Red Hat Directory Server documentation.	

Table C.13. Kerberos Log Files

Directory or File	Description
<code>/var/log/krb5kdc.log</code>	The primary log file for the Kerberos KDC server.
<code>/var/log/kadmind.log</code>	The primary log file for the Kerberos administration server.
Locations for these files is configured in the krb5.conf file. They can be different on some systems.	

Table C.14. DNS Log Files

Directory or File	Description
<code>/var/log/messages</code>	Includes DNS error messages among other system messages. DNS logging in this file is not enabled by default. To enable it, run the # /usr/sbin/rndc querylog command. To disable logging, run the command again.

Additional Resources

- See [Using the Journal](#) in the *System Administrator's Guide* for information on how to use the **journalctl** utility. You can use **journalctl** to view the logging output of **systemd** unit files.

C.3. IDM DOMAIN SERVICES AND LOG ROTATION

Several IdM domain services use the system **logrotate** service to handle log rotation and compression:

- **named** (DNS)
- **httpd** (Apache)
- **tomcat**
- **sssd**
- **krb5kdc** (the Kerberos domain controller)

The **logrotate** configuration files are stored in the **/etc/logrotate.d/** directory.

Example C.1. Default httpd Log Rotation File at **/etc/logrotate.d/httpd**

```
/var/log/httpd/*log {
    missingok
    notifempty
    sharedscripts
    delaycompress
    postrotate
        /sbin/service httpd reload > /dev/null 2>/dev/null || true
    endscript
}
```


**WARNING**

The **logrotate** policy files for most of the services create a new log file with the same name, default owner, and default permissions as the previous log. However, with the files for **named** and **tomcat**, a special **create** rule sets this behavior with explicit permissions as well as user and group ownership.

Do not change the permissions or the user and group which own the named and tomcat log files. This is required for both IdM operations and SELinux settings. Changing the ownership of the log rotation policy or of the files can cause the IdM domains services to fail.

Additional Resources

- The 389 Directory Server instances used by IdM as a back end and by the Dogtag Certificate System have their own internal log rotation policies. See the [Configuring Log Files](#) in the Red Hat Directory Server *Administration Guide*.
- For details about other potential log rotation settings, such as compression settings or the size of the log files, see the [Log Rotation](#) in the *System Administrator's Guide* or the `logrotate(8)` man page.

APPENDIX D. MANAGING REPLICAS AT DOMAIN LEVEL 0

This appendix describes managing replicas at domain level **0** (see [Chapter 7, *Displaying and Raising the Domain Level*](#)). For documentation on managing replicas at domain level **1**, see:

- [Section 4.5, “Creating the Replica: Introduction”](#)
- [Chapter 6, *Managing Replication Topology*](#)

D.1. REPLICA INFORMATION FILE

During the replica creation process, the **ipa-replica-prepare** utility creates a *replica information file* named after the replica server in the `/var/lib/ipa/` directory. The replica information file is a GPG-encrypted file containing realm and configuration information for the master server.

The **ipa-replica-install** replica setup script configures a Directory Server instance based on the information contained in the replica information file and initiates the *replica initialization* process, during which the script copies over data from the master server to the replica. A replica information file can only be used to install a replica on the specific machine for which it was created. It cannot be used to create multiple replicas on multiple machines.

D.2. CREATING REPLICAS

The following sections describe the most notable replica installation scenarios.

- The procedures and examples are not mutually exclusive; it is possible to use the CA, DNS, and other command-line options simultaneously. Examples in the following sections are called out separately simply to make it more clear what each configuration area requires.
- The **ipa-replica-install** utility accepts a number of other options as well. For a complete list, the `ipa-replica-install(1)` man page.

D.2.1. Installing a Replica without DNS

1. *On the master IdM server*, run the **ipa-replica-prepare** utility and add the fully qualified domain name (FQDN) of the *replica* machine. Note that the **ipa-replica-prepare** script does not validate the IP address or verify if the IP address of the replica is reachable by other servers.



IMPORTANT

The fully qualified domain name must be a valid DNS name, which means only numbers, alphabetic characters, and hyphens (-) are allowed. Other characters, like underscores, in the host name cause DNS failures. Additionally, the host name must be all lower-case; no capital letters are allowed.

For other recommended naming practices, see the [Red Hat Enterprise Linux Security Guide](#).

If the master server is configured with integrated DNS, specify the IP address of the replica machine using the **--ip-address** option. The installation script then asks if you want to configure the reverse zone for the replica. Only pass **--ip-address** if the IdM server was configured with integrated DNS. Otherwise, there is no DNS record to update, and the attempt to create the replica fails when the DNS record operation fails.

Enter the initial master server's Directory Manager (DM) password when prompted. The output of **ipa-replica-prepare** displays the location of the replica information file. For example:

```
[root@server ~]# ipa-replica-prepare replica.example.com --ip-address 192.0.2.2
Directory Manager (existing master) password:

Do you want to configure the reverse zone? [yes]: no
Preparing replica for replica.example.com from server.example.com
Creating SSL certificate for the Directory Server
Creating SSL certificate for the dogtag Directory Server
Saving dogtag Directory Server port
Creating SSL certificate for the Web Server
Exporting RA certificate
Copying additional files
Finalizing configuration
Packaging replica information into /var/lib/ipa/replica-info-replica.example.com.gpg
Adding DNS records for replica.example.com
Waiting for replica.example.com. A or AAAA record to be resolvable
This can be safely interrupted (Ctrl+C)
The ipa-replica-prepare command was successful
```



WARNING

Replica information files contain sensitive information. Take appropriate steps to ensure that they are properly protected.

For other options that can be added to **ipa-replica-prepare**, see the `ipa-replica-prepare(1)` man page.

2. On the replica machine, install the `ipa-server` package.

```
[root@replica ~]# yum install ipa-server
```

3. Copy the replica information file from the initial server to the replica machine:

```
[root@server ~]# scp /var/lib/ipa/replica-info-replica.example.com.gpg root@replica:/var/lib/ipa/
```

4. On the replica machine, run the **ipa-replica-install** utility and add the location

of the replication information file to start the replica initialization process. Enter the original master server's Directory Manager and admin passwords when prompted, and wait for the replica installation script to complete.

```
[root@replica ~]# ipa-replica-install /var/lib/ipa/replica-info-
replica.example.com.gpg
Directory Manager (existing master) password:

Run connection check to master
Check connection from replica to remote master 'server.example.com':

...

Connection from replica to master is OK.
Start listening on required ports for remote master check
Get credentials to log in to remote master
admin@MASTER.EXAMPLE.COM password:

Check SSH connection to remote master

...

Connection from master to replica is OK.

...

Configuring NTP daemon (ntpd)
[1/4]: stopping ntpd
[2/4]: writing configuration

...

Restarting Directory server to apply updates
[1/2]: stopping directory server
[2/2]: starting directory server
Done.
Restarting the directory server
Restarting the KDC
Restarting the web server
```



NOTE

If the replica file being installed does not match the current host name, the replica installation script displays a warning message and asks for confirmation. In some cases, such as on multi-homed machines, you can confirm to continue with the mismatched host names.

For command-line options that can be added to **ipa-replica-install**, see the **ipa-replica-prepare(1)** man page. Note that one of the options **ipa-replica-install** accepts is the **--ip-address** option. When added to **ipa-replica-install**, **--ip-address** only accepts IP addresses associated with the local interface.

D.2.2. Installing a Replica with DNS

To install a replica with integrated DNS, follow the procedure for installing without DNS described in [Section D.2.1, “Installing a Replica without DNS”](#), but add these options to **ipa-replica-install**:

- **--setup-dns**
- **--forwarder**

See [Section 4.5.3, “Installing a Replica with DNS”](#) for details.

For example:

```
[root@replica ~]# ipa-replica-install /var/lib/ipa/replica-info-  
replica.example.com.gpg --setup-dns --forwarder 198.51.100.0
```

After running **ipa-replica-install**, make sure proper DNS entries were created, and optionally add other DNS servers as backup servers. See [Section 4.5.3, “Installing a Replica with DNS”](#) for details.

D.2.3. Installing a Replica with Various CA Configurations



WARNING

Red Hat strongly recommends to keep the CA services installed on more than one server. For information on installing a replica of the initial server including the CA services, see [Section 4.5.4, “Installing a Replica with a CA”](#).

If you install the CA on only one server, you risk losing the CA configuration without a chance of recovery if the CA server fails. See [Section B.2.6, “Recovering a Lost CA Server”](#) for details.

Installing a Replica from a Server with a Certificate System CA Installed

To set up a CA on the replica when the initial server was configured with an integrated Red Hat Certificate System instance (regardless of whether it was a root CA or whether it was subordinate to an external CA), follow the basic installation procedure described in [Section D.2.1, “Installing a Replica without DNS”](#), but add the **--setup-ca** option to the **ipa-replica-install** utility. The **--setup-ca** option copies the CA configuration from the initial server's configuration.

```
[root@replica ~]# ipa-replica-install /var/lib/ipa/replica-info-  
replica.example.com.gpg --setup-ca
```

Installing a Replica from a Server without a Certificate System CA Installed

For a CA-less replica installation, follow the basic procedure described in [Section D.2.1, “Installing a Replica without DNS”](#), but add the following options when running the **ipa-replica-prepare** utility on the initial server:

- **--dirsrv-cert-file**

- **--dirsrv-pin**
- **--http-cert-file**
- **--http-pin**

See [Section 4.5.5, “Installing a Replica from a Server without a CA”](#) for details.

For example:

```
[root@server ~]# ipa-replica-prepare replica.example.com --dirsrv-cert-
file /tmp/server.key --dirsrv-pin secret --http-cert-file /tmp/server.crt
--http-cert-file /tmp/server.key --http-pin secret --dirsrv-cert-file
/tmp/server.crt
```

D.2.4. Adding Additional Replication Agreements

Installing a replica using **ipa-replica-install** creates an initial replication agreement between the master server and the replica. To connect the replica to other servers or replicas, add additional agreements using the **ipa-replica-manage** utility.

If the master server and the new replica have a CA installed, a replication agreement for CA is also created. To add additional CA replication agreements to other servers or replicas, use the **ipa-csreplica-manage** utility.

For more information on adding additional replication agreements, see [Section D.3, “Managing Replicas and Replication Agreements”](#).

D.3. MANAGING REPLICAS AND REPLICATION AGREEMENTS

This chapter provides details on replication agreements and describes how to manage them.



NOTE

For guidelines on setting up additional replication agreements, see [Section 4.2.2, “Replica Topology Recommendations”](#).

D.3.1. Explaining Replication Agreements

Replicas are joined in a *replication agreement* that copies data between them. Replication agreements are bilateral: the data is replicated from the first replica to the other one as well as from the other replica to the first one.



NOTE

An initial replication agreement is set up between two replicas by the **ipa-replica-install** script. See [Chapter 4, *Installing and Uninstalling Identity Management Replicas*](#) for details on installing the initial replica.

Types of Replication Agreements

Identity Management supports the following three types of replication agreements:

- Replication agreements to replicate directory data, such as users, groups, and policies. You can manage these agreements using the **ipa-replica-manage** utility.
- Replication agreements to replicate certificate server data. You can manage these agreements using the **ipa-csreplica-manage** utility.
- Synchronization agreements to replicate user information with an Active Directory server. These agreements are not described in this guide. For documentation on synchronizing IdM and Active Directory, see the [Windows Integration Guide](#).

The **ipa-replica-manage** and **ipa-csreplica-manage** utilities use the same format and arguments. The following sections of this chapter describe the most notable replication management operations performed using these utilities. For detailed information about the utilities, see the `ipa-replica-manage(1)` and `ipa-csreplica-manage(1)` man pages.

D.3.2. Listing Replication Agreements

To list the directory data replication agreements currently configured for a replica, use the **ipa-replica-manage list** command:

1. Run **ipa-replica-manage list** without any arguments to list all replicas in the replication topology. In the output, locate the required replica:

```
$ ipa-replica-manage list
server1.example.com: master
server2.example.com: master
server3.example.com: master
server4.example.com: master
```

2. Add the replica's host name to **ipa-replica-manage list** to list the replication agreements.

```
$ ipa-replica-manage list server1.example.com
server2.example.com: replica
server3.example.com: replica
```

The output displays the replicas to which **server1.example.com** sends updates.

To list certificate server replication agreements, use the **ipa-csreplica-manage list** command.

D.3.3. Creating and Removing Replication Agreements

Creating Replication Agreements

To create a new replication agreement, use the **ipa-replica-manage connect** command:

```
$ ipa-replica-manage connect server1.example.com server2.example.com
```

The command creates a new bilateral replication agreement going from *server1.example.com* to *server2.example.com* and from *server2.example.com* to *server1.example.com*.

If you only specify one server with **ipa-replica-manage connect**, IdM creates a replication agreement between the local host and the specified server.

To create a new certificate server replication agreement, use the **ipa-csreplica-manage connect** command.

Removing Replication Agreements

To remove a replication agreement, use the **ipa-replica-manage disconnect** command:

```
$ ipa-replica-manage disconnect server1.example.com server4.example.com
```

This command disables replication from *server1.example.com* to *server4.example.com* and from *server4.example.com* to *server1.example.com*.

The **ipa-replica-manage disconnect** command only removes the replication agreement. It leaves both servers in the Identity Management replication topology. To remove all replication agreements and data related to a replica, use the **ipa-replica-manage del** command, which removes the replica entirely from the Identity Management domain.

```
$ ipa-replica-manage del server2.example.com
```

To remove a certificate server replication agreement, use the **ipa-csreplica-manage disconnect** command. Similarly, to remove all certificate replication agreements and data between two servers, use the **ipa-csreplica-manage del** command.

D.3.4. Initiating a Manual Replication Update

Data changes between replicas with direct replication agreements between each other are replicated almost instantaneously. However, replicas that are not joined in a direct replication agreement do not receive updates as quickly.

In some situations, it might be necessary to manually initiate an unplanned replication update. For example, before taking a replica offline for maintenance, all the queued changes waiting for the planned update must be sent to one or more other replicas. In this situation, you can initiate a manual replication update before taking the replica offline.

To manually initiate a replication update, use the **ipa-replica-manage force-sync** command. The local host on which you run the command is the replica that receives the update. To specify the replica that sends the update, use the **--from** option.

```
$ ipa-replica-manage force-sync --from server1.example.com
```

To initiate a replication update for certificate server data, use the **ipa-csreplica-manage force-sync** command.

D.3.5. Re-initializing a Replica

If a replica has been offline for a long period of time or its database has been corrupted, you can *re-initialize* it. Re-initialization is analogous to initialization, which is described in [Section 4.5, “Creating the Replica: Introduction”](#). Re-initialization refreshes the replica with an updated set of data.

**NOTE**

Waiting for a regular replication update or initiating a manual replication update will not help in this situation. During these replication updates, replicas only send changed entries to each other. Unlike re-initialization, replication updates do not refresh the whole database.

To re-initialize a data replication agreement on a replica, use the **ipa-replica-manage re-initialize** command. The local host on which you run the command is the re-initialized replica. To specify the replica from which the data is obtained, use the **--from** option:

```
$ ipa-replica-manage re-initialize --from server1.example.com
```

To re-initialize a certificate server replication agreement, use the **ipa-csreplica-manage re-initialize** command.

D.3.6. Removing a Replica

Deleting or *demoting* a replica removes the IdM replica from the topology so that it no longer processes IdM requests. It also removes the host machine itself from the IdM domain.

To delete a replica, perform these steps on the replica:

1. List all replication agreements for the IdM domain. In the output, note the host name of the replica.

```
$ ipa-replica-manage list
server1.example.com: master
server2.example.com: master
server3.example.com: master
server4.example.com: master
```

2. Use the **ipa-replica-manage del** command to remove all agreements configured for the replica as well as all data about the replica.

```
$ ipa-replica-manage del server3.example.com
```

3. If the replica was configured with its own CA, then also use the **ipa-csreplica-manage del** command to remove all certificate server replication agreements.

```
$ ipa-csreplica-manage del server3.example.com
```

**NOTE**

This step is only required if the replica itself was configured with an IdM CA. It is not required if only the master server or other replicas were configured with a CA.

4. Uninstall the IdM server package.

```
$ ipa-server-install --uninstall -U
```

D.4. PROMOTING A REPLICA TO A MASTER CA SERVER

In a topology including multiple replicas, one server acts as the master CA: it manages the renewal of CA subsystem certificates and generates certificate revocation lists (CRLs). By default, the master CA is the initial server from which replicas were created.

If you plan to take the master CA server offline or decommission it, *promote* a replica to take its place as the master CA:

- Make sure the replica is configured to handle CA subsystem certificate renewal. See [Section D.4.1, “Changing Which Server Handles Certificate Renewal”](#).
- Configure the replica to generate CRLs. See [Section 6.5.2.2, “Changing Which Server Generates CRLs”](#).

D.4.1. Changing Which Server Handles Certificate Renewal

To determine which server is the current renewal master:

- On Red Hat Enterprise Linux 7.3 and later, use the **ipa config-show | grep "CA renewal master"** command:

```
$ ipa config-show | grep "CA renewal master"
IPA CA renewal master: server.example.com
```

- On Red Hat Enterprise Linux 7.2 and earlier, use the **ldapsearch** utility. In the following example, the renewal master is **server.example.com**:

```
$ ldapsearch -H ldap://$HOSTNAME -D 'cn=Directory Manager' -W -b
'cn=masters,cn=ipa,cn=etc,dc=example,dc=com' '(&(cn=CA)
(ipaConfigString=caRenewalMaster))' dn
Enter LDAP Password:
# extended LDIF
#
# LDAPv3
# base <cn=masters,cn=ipa,cn=etc,dc=example,dc=com> with scope
subtree
# filter: (&(cn=CA)(ipaConfigString=caRenewalMaster))
# requesting: dn
#
# CA, server.example.com, masters, ipa, etc, example.com
dn:
cn=CA,cn=server.example.com,cn=masters,cn=ipa,cn=etc,dc=example,dc=c
om

# search result
search: 2
result: 0 Success

# numResponses: 2
# numEntries: 1
```

To configure another server to handle certificate renewal, use the **ipa-csreplica-manage** utility:

```
# ipa-csreplica-manage set-renewal-master
```

The command also automatically reconfigures the previous CA from renewal master to clone.

APPENDIX E. REVISION HISTORY

Note that revision numbers relate to the edition of this manual, not to version numbers of Red Hat Enterprise Linux.

Revision 7.0-39 Minor updates.	Wed Mar 14 2018	Filip Hanzelka
Revision 7.0-38 Minor updates.	Wed Feb 28 2018	Lucie Maňásková
Revision 7.0-37 Added <i>Users Cannot Access Their Vault Due To Insufficient 'add' Privilege</i> . Other minor fixes.	Mon Feb 12 2018	Aneta Šteflová Petrová
Revision 7.0-36 Updated <i>Defining SELinux User Maps</i> . Other minor fixes.	Mon Jan 29 2018	Aneta Šteflová Petrová
Revision 7.0-35 Updated <i>Managing Hosts</i> . Other minor fixes.	Fri Dec 15 2017	Aneta Šteflová Petrová
Revision 7.0-34 Added <i>Kerberos PKINIT Authentication in IdM</i> . Updated <i>Defining Access Control for IdM Users</i> . Other minor fixes.	Mon Dec 4 2017	Aneta Šteflová Petrová
Revision 7.0-33 Updated chapters <i>User and Group Schema</i> and <i>Defining Password Policies</i> .	Mon Nov 20 2017	Aneta Šteflová Petrová
Revision 7.0-32 Minor fixes.	Mon Oct 9 2017	Aneta Šteflová Petrová
Revision 7.0-31 Updated a few web UI screenshots and procedures. Minor updates to <i>Smart-card Authentication in Identity Management</i> .	Tue Sep 12 2017	Aneta Šteflová Petrová
Revision 7.0-30 Updated <i>Smart-card Authentication in Identity Management</i> , <i>Identity Management Configuration Files and Directories</i> , and a few screenshots. Other minor updates.	Mon Aug 28 2017	Aneta Šteflová Petrová
Revision 7.0-29 Document version for 7.4 GA publication.	Tue Jul 18 2017	Aneta Šteflová Petrová
Revision 7.0-28 Updated and merged managing user groups, host groups, and automember. Other minor updates.	Mon Apr 24 2017	Aneta Šteflová Petrová
Revision 7.0-27 Added <i>Configuring TLS for Identity Management</i> . Various minor fixes and updates.	Mon Apr 10 2017	Aneta Šteflová Petrová
Revision 7.0-26 Added <i>Post-installation Considerations for Clients and Enabling Password Reset</i> . Fixed broken links. Other minor updates.	Mon Mar 27 2017	Aneta Šteflová Petrová
Revision 7.0-25 Updated chapters on managing the Kerberos domain, upgrading, and HBAC. Other updates in various chapters.	Mon Feb 27 2017	Aneta Šteflová Petrová
Revision 7.0-24 Updated automember and password policies chapters. Added description for NIS support plug-ins. Other minor updates.	Wed Dec 7 2016	Aneta Šteflová Petrová

Revision 7.0-23 Version for 7.3 GA publication.	Tue Oct 18 2016	Aneta Šteflová Petrová
Revision 7.0-22 Added a chapter on using vaults.	Fri Jul 29 2016	Aneta Petrová
Revision 7.0-21 Updated introduction, other minor fixes.	Thu Jul 28 2016	Marc Muehlfeld
Revision 7.0-19 Updated diagrams. Added a section on benefits of using IdM to the intro chapter. Other minor fixes and tweaks.	Tue Jun 28 2016	Aneta Petrová
Revision 7.0-18 Updated introduction, server installation, and troubleshooting chapters. Added a chapter for user, host, and service certificates. Merged changing domain DNS config chapter into other chapters. Other minor fixes.	Fri Jun 10 2016	Aneta Petrová
Revision 7.0-17 Added a diagram for user lifecycle.	Fri May 27 2016	Aneta Petrová
Revision 7.0-16 Added user lifecycle. Updated the User Accounts, User Authentication, and Managing Replicas chapters.	Thu Mar 24 2016	Aneta Petrová
Revision 7.0-15 Updated several DNS sections. Moved restricting domains for PAM services to the System-Level Authentication Guide.	Thu Mar 03 2016	Aneta Petrová
Revision 7.0-14 Added smart cards, ID views, and OTP. Updated some web UI screenshots and the Basics of Management and Restricting Domains chapters. Moved uninstallation procedures into installation chapters. Commented out index. Other minor updates.	Tue Feb 09 2016	Aneta Petrová
Revision 7.0-13 Minor updates to certificate profile management and promoting a replica to master.	Thu Nov 19 2015	Aneta Petrová
Revision 7.0-12 Version for 7.2 GA release with updates to DNS and other sections.	Fri Nov 13 2015	Aneta Petrová
Revision 7.0-11 Version for 7.2 GA release.	Thu Nov 12 2015	Aneta Petrová
Revision 7.0-10 Async update with last-minute edits for 7.1.	Fri Mar 13 2015	Tomáš Čapek
Revision 7.0-8 Version for 7.1 GA release.	Wed Feb 25 2015	Tomáš Čapek
Revision 7.0-6 Rebuild to update the sort order on the splash page.	Fri Dec 05 2014	Tomáš Čapek
Revision 7.0-4 Initial release.	Wed Jun 11 2014	Ella Deon Ballard