

# **freeIPA 1.2.1**

## **Installation and Deployment Guide**

IPA Solutions from the IPA Experts

# freeIPA

# freeIPA 1.2.1 Installation and Deployment Guide

## IPA Solutions from the IPA Experts

### Edition 1.0

Copyright © 2008 Red Hat. This material may only be distributed subject to the terms and conditions set forth in the Open Publication License, V1.0 or later. The latest version of the OPL is presently available at <http://www.opencontent.org/openpub/>.

Red Hat and the Red Hat "Shadow Man" logo are registered trademarks of Red Hat, Inc. in the United States and other countries.

All other trademarks referenced herein are the property of their respective owners.

The GPG fingerprint of the security@redhat.com key is:

CA 20 86 86 2B D6 9D FC 65 F6 EC C4 21 91 80 CD DB 42 A6 0E

1801 Varsity Drive  
Raleigh, NC 27606-2072  
USA  
Phone: +1 919 754 3700  
Phone: 888 733 4281  
Fax: +1 919 754 3701  
PO Box 13588  
Research Triangle Park, NC 27709  
USA

This guide covers the basic considerations that should be addressed before deploying IPA. It also covers the installation and configuration of each of the supported server platforms.

---

<b>Preface</b>	<b>v</b>
1. Audience .....	v
2. Document Conventions .....	v
2.1. Typographic Conventions .....	v
2.2. Pull-quote Conventions .....	vii
2.3. Notes and Warnings .....	vii
3. We Need Feedback! .....	viii
<b>1. Introduction</b>	<b>1</b>
1.1. What is IPA? .....	1
1.2. Components of IPA .....	1
<b>2. Preparing for an IPA Installation</b>	<b>3</b>
2.1. Assumptions .....	3
2.2. Required Ports .....	3
2.3. File Systems .....	4
2.4. DNS .....	4
2.5. Configuring Networking .....	5
2.5.1. Configuring Networking Services .....	5
2.5.2. Configuring the /etc/hosts File .....	5
2.6. Hardware Requirements .....	6
2.7. Software Requirements .....	6
<b>3. Setting up the IPA Server</b>	<b>7</b>
3.1. Installing the IPA Server .....	7
3.2. Configuring the IPA Server .....	7
3.2.1. Testing the Configuration .....	8
3.3. Configuring Your Browser .....	9
3.3.1. Troubleshooting .....	10
3.4. Using a Browser on Another System .....	11
<b>4. Setting up Synchronization Between IPA and Active Directory</b>	<b>13</b>
4.1. Introduction .....	13
4.2. Prerequisites .....	13
4.2.1. Domain Name Considerations .....	13
4.2.2. Setting up Active Directory .....	13
4.3. Setting up Windows Sync on the IPA Server .....	15
4.4. Creating Synchronization Agreements .....	15
4.5. Modifying Synchronization Agreements .....	16
4.5.1. Changing the Default Synchronization Subtree .....	16
4.6. Deleting Synchronization Agreements .....	17
<b>5. Setting up Multi-Master Replication</b>	<b>19</b>
5.1. Preparing the Replica Servers .....	19
5.2. Installing the Server Packages .....	19
5.3. Creating the Replica Information File .....	19
5.4. Configuring an IPA Replica .....	20
5.4.1. Updating DNS for IPA Replicas .....	21
5.5. Managing Multi-Master Replication .....	21
5.6. Troubleshooting Multi-Master Replication .....	22
<b>6. Setting up IPA to run as an Apache Virtual Host</b>	<b>23</b>
<b>A. Revision History</b>	<b>25</b>



---

# Preface

Welcome to the IPA Installation and Deployment Guide. This guide covers the basic considerations that should be addressed before deploying IPA. The decisions made during this phase can have a significant and lasting affect on the effectiveness, efficiency, and scalability of your installation. You should have a good understanding of your deployment requirements before moving on to the installation phase.

This guide also describes the available methods for obtaining and installing the IPA server software, and how to configure the product to best suit your deployment.

## 1. Audience

The IPA Installation and Deployment Guide is intended for system administrators and those responsible for installing and configuring IPA.

This guide assumes a good understanding of either Red Hat Enterprise Linux or Fedora, and a working knowledge of LDAP and Fedora Directory Server.

## 2. Document Conventions

This manual uses several conventions to highlight certain words and phrases and draw attention to specific pieces of information.

In PDF and paper editions, this manual uses typefaces drawn from the *Liberation Fonts*<sup>1</sup> set. The Liberation Fonts set is also used in HTML editions if the set is installed on your system. If not, alternative but equivalent typefaces are displayed. Note: Red Hat Enterprise Linux 5 and later includes the Liberation Fonts set by default.

### 2.1. Typographic Conventions

Four typographic conventions are used to call attention to specific words and phrases. These conventions, and the circumstances they apply to, are as follows.

#### **Mono-spaced Bold**

Used to highlight system input, including shell commands, file names and paths. Also used to highlight key caps and key-combinations. For example:

To see the contents of the file **my\_novel** in your current working directory, enter the **cat my\_novel** command at the shell prompt and then press **Enter**.

The above example includes a file name, a shell command and a key cap, all presented in Mono-spaced Bold and all distinguishable thanks to context.

Key-combinations can be distinguished from key caps by the hyphen connecting each part of a key-combination. For example:

Press **Enter** to execute the command.

Press **Ctrl-Alt-F1** to switch to the first virtual terminal. Press **Ctrl-Alt-F7** to return to your X-Windows session.

---

<sup>1</sup> <https://fedorahosted.org/liberation-fonts/>

The first sentence highlights the particular key cap to press. The second highlights two sets of three key caps, each set pressed simultaneously.

If source code is discussed, class names, methods, functions, variable names and returned values mentioned within a paragraph will be presented as above, in **Mono-spaced Bold**. For example:

File-related classes include **filesystem** for file systems, **file** for files, and **dir** for directories. Each class has its own associated set of permissions.

### Proportional Bold

This denotes words or phrases encountered on a system, including application names; dialogue box text; labelled buttons; check-box and radio button labels; menu titles and sub-menu titles. For example:

Choose **System > Preferences > Mouse** from the main menu bar to launch **Mouse Preferences**. In the **Buttons** tab, click the **Left-handed mouse** check box and click **Close** to switch the primary mouse button from the left to the right (making the mouse suitable for use in the left hand).

The above text includes application names; system-wide menu names and items; application-specific menu names; and buttons and text found within a GUI interface, all presented in Proportional Bold and all distinguishable by context.

Note the **>** shorthand used to indicate traversal through a menu and its sub-menus. This avoids the difficult-to-follow 'Select **Mouse** from the **Preferences** sub-menu in the **System** menu of the main menu bar' approach.

### ***Mono-spaced Bold Italic*** or ***Proportional Bold Italic***

Whether Mono-spaced Bold or Proportional Bold, the addition of Italics indicates replaceable or variable text. Italics denotes text you do not input literally or displayed text that changes depending on circumstance. For example:

To connect to a remote machine using ssh, type **ssh *username@domain.name*** at a shell prompt. If the remote machine is **example.com** and your username on that machine is john, type **ssh *john@example.com***.

To see the version of a currently installed package, use the **rpm -q *package*** command. It will return a result as follows: ***package-version-release***.

Note the words in bold italics above — *username*, *domain.name*, *package*, *version* and *release*. Each word is a placeholder, either for text you enter when issuing a command or for text displayed by the system.

Aside from standard usage for presenting the title of a work, italics denotes the first use of a new or important term. For example:

When the Apache HTTP Server accepts requests, it dispatches child processes or threads to handle them. This group of child processes or threads is known as a *server-pool*. Under Apache HTTP Server 2.0, the responsibility for creating and maintaining these server-pools has been abstracted to a group of modules called *Multi-Processing Modules (MPMs)*. Unlike other modules, only one module from the MPM group can be loaded by the Apache HTTP Server.

## 2.2. Pull-quote Conventions

Two, commonly multi-line, data types are set off visually from the surrounding text.

Output sent to a terminal is set in Mono-spaced Roman and presented thus:

```
books      Desktop  documentation  drafts  mss    photos  stuff  svn
books_tests Desktop1  downloads      images  notes  scripts  svgs
```

Source-code listings are also set in Mono-spaced Roman but are presented and highlighted as follows:

```
package org.jboss.book.jca.ex1;

import javax.naming.InitialContext;

public class ExClient
{
    public static void main(String args[])
        throws Exception
    {
        InitialContext iniCtx = new InitialContext();
        Object          ref    = iniCtx.lookup("EchoBean");
        EchoHome        home   = (EchoHome) ref;
        Echo             echo   = home.create();

        System.out.println("Created Echo");

        System.out.println("Echo.echo('Hello') = " + echo.echo("Hello"));
    }
}
```

## 2.3. Notes and Warnings

Finally, we use three visual styles to draw attention to information that might otherwise be overlooked.



### Note

A Note is a tip or shortcut or alternative approach to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on a trick that makes your life easier.



### Important

Important boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring Important boxes won't cause data loss but may cause irritation and frustration.



### Warning

A Warning should not be ignored. Ignoring warnings will most likely cause data loss.

## 3. We Need Feedback!

If you find a typographical error in this manual, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report in Bugzilla: [https://bugzilla.redhat.com/enter\\_bug.cgi?product=freeIPA](https://bugzilla.redhat.com/enter_bug.cgi?product=freeIPA) against the Documentation component.

When submitting a bug report, be sure to mention the manual's identifier: *Installation\_Guide*

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. If you have found an error, please include the section number and some of the surrounding text so we can find it easily.

# Introduction

## 1.1. What is IPA?

IPA is an integrated security information management solution combining Fedora, Fedora Directory Server, MIT Kerberos, NTP, and DNS. It consists of a web interface and command-line administration tools for user and group management.

IPA stands for Identity, Policy, and Audit. The current version of IPA supports identity management; support for policy and auditing management is scheduled for later versions.

## 1.2. Components of IPA

Many of the components that comprise IPA, for example, Kerberos, `libuser`, etc., are provided by the base operating system. Only the default versions that ship with the base operating system are supported. In some cases, IPA may deliver packages that replace the base operating system packages. In these cases, support is available for the replaced packages.

The current package versions are:

- MIT Kerberos™ version 1.6
- DNS (BIND 9)
- NTP 4.2
- Apache 2.2.3



# Preparing for an IPA Installation

Before you install IPA, ensure that the installation environment is suitably configured. This chapter describes the information that you need to provide, and the impact of integrating IPA into your organization.

You also need to provide certain information during the installation and configuration procedures, including realm names and certain usernames and passwords.

## 2.1. Assumptions

IPA relies on a script to automate many of the installation tasks. This script makes a number of assumptions:

- *System:*

That you are performing the installation on a "clean" system. The script overwrites a number of files without prompting for confirmation. [Section 3.2, "Configuring the IPA Server"](#) lists the services whose scripts and configuration files are modified.

- *Directory Server:*

That there are no existing Directory Server instances. No automatic upgrade facility from an existing Directory Server currently exists.

- *DNS:*

- That the server's machine name is set, and that it resolves to its public IP address (not to localhost).
- That DNS is correctly configured to resolve forward and reverse addresses. The DNS does not need to be on the same machine as the IPA server, but it does need to be fully functional.

## 2.2. Required Ports

IPA makes use of the following ports:

- TCP
  - 80, 443, 8080: HTTP/HTTPS
  - 389, 636: LDAP/LDAPS
  - 88, 464: Kerberos
- UDP
  - 88, 464: Kerberos
  - 123: NTP

Ensure that these ports are available for both the IPA server and for access by other systems; that is, that they are not assigned to another service and that they are not blocked by a firewall. If these ports are not available, IPA will not function correctly.

### 2.3. File Systems

You should be aware of the following with respect to file systems and IPA:

- The default prefix for users' home directories is **/home**
- IPA does not automatically create home directories when users log in.
  - To automatically create home directories, you can use the `pam_mkhome` module. IPA does not force the use of this module because it may try to create home directories even when the shared storage is not available. It is the responsibility of the system administrator to activate this module on the clients if needed.
  - It is possible to use an NFS filer that provides **/home** that can be made available to all client machines.
  - If a suitable directory and mechanism are not available for the creation of home directories, users may not be able to log in.
- IPA does not currently provide automount support.

### 2.4. DNS

It is recommended that you use DNS to facilitate Service Discovery in IPA. Service Discovery refers to the way that IPA clients find (or discover) IPA servers. You can use the basic DNS configuration that is provided with IPA to configure an existing DNS to work with IPA, or pass the `--setup-bind` option to the `ipa-server-install` command to configure a new DNS. The DNS does not need to be on the same machine as the IPA server, but it does need to be correctly configured and fully functional.



#### Note

The `--setup-bind` option is an optional parameter that can be passed to the `ipa-server-install` script. This is provided for convenience only; it is not a supported aspect of IPA. The following article may help you to configure your DNS server: [How to set up a DNS server](#).<sup>1</sup>

To aid in the creation and configuration of a suitable DNS setup, the IPA installation creates a sample zone file. During the installation you will see a message similar to the following:

```
Sample zone file for bind has been created in /tmp/sample.zone.F_uMf4.db
```

You should use this file in your zone file in DNS. Further, you need to ensure that your FQDN does not resolve to your loopback address.

#### IPA, DNS, and NSCD

It is recommended that you avoid or restrict the use of `nscd` (Name Service Caching Daemon) in a IPA deployment. The `nscd` service is extremely useful for reducing the load on the server, and for making clients more responsive, but drawbacks also exist.

`nscd` performs caching operations for all services that perform queries via the `nsswitch` interface, including **getent**. Because `nscd` performs both positive and negative caching, if a request determines that a specific IPA user does not exist, it marks this as a negative cache. Values stored in the cache remain until the cache expires, regardless of any changes that may occur on the server.

The results of such caching is that new users and memberships may not be visible, and users and memberships that have been removed may still be visible.

To alleviate these effects, you can avoid the use of `nscd` altogether, or use a shorter cache time. In particular, consider changing the following values in the `/etc/nscd.conf` file to suit the usage patterns of your deployment:

```
positive-time-to-live  group      3600
negative-time-to-live  group      60
positive-time-to-live  hosts     3600
negative-time-to-live  hosts     20
```

### DNS and Kerberos

Kerberos, too, has very specific DNS requirements. The Kerberos server requires a valid DNS A record, and reverse DNS needs to work correctly. Do not use CNAME or DDNS names, as it can cause major problems later. The IPA installation process includes checks to ensure that the IPA server name is a DNS A record and that its reverse and forward addresses match.

Refer to IPA, Kerberos, and DNS in the IPA Administration Reference for more information on how these technologies work together.

## 2.5. Configuring Networking

### 2.5.1. Configuring Networking Services

The default networking service used by Red Hat Enterprise Linux and Fedora is `NetworkManager`, and due to the way this service works, it can cause problems with IPA and the KDC. Consequently, it is highly recommended that you use the `network` service to manage the networking requirements in an IPA environment, and disable the `NetworkManager` service.

**Procedure 2.1.** To configure networking services for IPA:

1. Boot into single-user mode and run the following commands:

```
# chkconfig NetworkManager off; service NetworkManager stop
```

```
# chkconfig NetworkManagerDispatcher off; service
NetworkManagerDispatcher stop
```

```
# chkconfig network on; service network start
```

2. Ensure that static networking is correctly configured.
3. Restart the system.

### 2.5.2. Configuring the `/etc/hosts` File

You need to ensure that your `/etc/hosts` file is configured correctly. An incorrectly configured file can prevent the IPA command-line tools from functioning correctly, and can prevent the IPA web interface from connecting to the IPA server.

Configure the `/etc/hosts` file to list the FQDN for the IPA server *before* any aliases. Also ensure that the hostname is not part of the `localhost` entry. The following is an example of a valid hosts file:

```
127.0.0.1 localhost.localdomain localhost
::1 localhost6.localdomain6 localhost6
192.168.1.1 ipaserver.example.com ipaserver
```



### Important

Do not omit the IPv4 entry in the `/etc/hosts` file. This entry is required by the IPA web service.

## 2.6. Hardware Requirements

The following table contains guidelines for Fedora Directory Server disk space and memory requirements based on the number of entries that your organization requires. The values shown here assume that the entries in the LDIF file are approximately 100 bytes each and that only the recommended indexes are configurable.

The system requirements for both 32-bit and 64-bit platforms are the same.

Criteria	< 250,000 Entries	250,000 - 1,000,000 Entries	> 1,000,000 Entries
CPU Type (minimum)	P3; 500MHz		
RAM (minimum)	256 MB	512 MB	1 GB
Disk Space (minimum)	2 GB	4 GB	8 GB

Table 2.1. Minimum hardware requirements for IPA.

## 2.7. Software Requirements

freeIPA 1.2.1 Server depends on:

- Fedora 9 or 10
- Fedora Directory Server; installed as an IPA dependency
- MIT Kerberos 1.6; typically installed as part of Fedora

All other freeIPA 1.2.1 requirements are installed as dependencies during the installation process.

# Setting up the IPA Server

## 3.1. Installing the IPA Server

Run the following command to install the IPA server packages:

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

This will install a large number of dependencies, including TurboGears, fedora-ds-base and krb5-server. Approximately 40 dependencies are required, depending on what is already installed.

## 3.2. Configuring the IPA Server

Use the **ipa-server-install** command to install the IPA server, which includes:

- Configuring the Network Time Daemon (ntpd)
- Creating and configuring an instance of Fedora Directory Server
- Creating and configuring a Kerberos Key Distribution Center (krb5kdc)
- Configuring Apache (httpd)
- Configuring TurboGears
- Updating the SELinux targeted policy
- Installing and configuring the Active Directory WinSync plug-in

You can install the server interactively by running the command with no options, or by passing options directly to the **ipa-server-install** command. To view the available command-line options, run the following command: **\$ /usr/sbin/ipa-server-install --help**



### Note

If you are running IPA as a virtualized guest, you should not run the NTP daemon. In this case, you should pass the **-N** (no ntp) option to the **ipa-server-install** command.

**Procedure 3.1.** To install the IPA server interactively:

1. Run the following command:

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

2. Enter the server's host name, realm name and other details when prompted.

The installation script compares the hostname returned by DNS to the hostname found in the `/etc/hosts` file. If the non-fully—qualified domain name appears first, the script aborts.



### Note

The hostname that you enter into the `ipa-server-install` script must be the same as that returned by the `hostname` command, otherwise the Directory Server cannot use its own keytab. This can cause some `ipa-*` commands to fail.

3. Wait until the configuration script completes. It can take several minutes to set up and configure all of the IPA requirements.
4. When the configuration script completes, restart the SSH service so that it re-reads the Name Server Switch (nss) configuration file.

To restart the SSH service, run the following command (existing connections are not terminated):

```
# service sshd restart
```

### 3.2.1. Testing the Configuration

The following examples assume that you are using `EXAMPLE.COM` as your realm.



### Note

The realm is used as the base DN in the directory instance; in this case it will be `dc=example,dc=com`.

When the installation is complete, all of the services should be running.

**Procedure 3.2.** To test your IPA installation:

1. Use the `kinit` command to request a Kerberos ticket:

```
$ kinit admin
```

```
Password for admin@EXAMPLE.COM:
```

2. Use the `klist` command to display the list of Kerberos tickets:

```
$ klist
```

```
Ticket cache: FILE:/tmp/krb5cc_0
Default principal: admin@EXAMPLE.COM
Valid starting Expires Service principal
03/05/08 02:47:53 03/06/08 02:47:50 krbtgt/EXAMPLE.COM@EXAMPLE.COM
Kerberos 4 ticket cache: /tmp/tkt0
klist: You have no tickets cached
```

3. Use the `ipa-finduser` command to search for the `admin` user:

```
$ /usr/sbin/ipa-finduser admin
```

```
cn: Administrator
homedirectory: /home/admin
loginshell: /bin/bash
uid: admin
```

If you receive output similar to the following, ensure that:

- DNS is configured correctly
- If your network includes Active Directory, that it is not in the same domain as the IPA server. Refer to [Section 4.2.1, "Domain Name Considerations"](#) for more information on this topic.

```
Could not initialize GSSAPI: Unspecified GSS failure.
Minor code may provide more information/Server not found in Kerberos
database.
```

### 3.3. Configuring Your Browser

**Firefox** can use your Kerberos credentials for authentication, but you need to specify which domains to communicate with, and using which attributes.

**Procedure 3.3. To configure Firefox for use with IPA:**

1. Open **Firefox**, and type "about:config" in the **Address Bar**.
2. In the **Search** field, type "negotiate".
3. Ensure the following lines reflect your setup. Replace ".example.com" with your own IPA server's domain, including the preceding period (.):

```
network.negotiate-auth.trusted-uris .example.com
network.negotiate-auth.delegation-uris .example.com
network.negotiate-auth.using-native-gsslib true
```

4. • If you are configuring **Firefox** on Microsoft Windows, make the following changes instead:

```
network.negotiate-auth.trusted-uris .example.com
network.auth.use-sspi false
network.negotiate-auth.delegation-uris .example.com
```

5. In **Firefox**, navigate to the IPA server (use the fully-qualified domain name, for example, `http://ipaserver.example.com`). Ensure that there are no Kerberos authentication errors, and that you can see and interact with the Web interface.

### 3.3.1. Troubleshooting

If you have followed the configuration steps and Negotiate authentication is not working, you can turn on verbose logging of the authentication process, and potentially find the cause of the problem.

**Procedure 3.4.** To troubleshoot Negotiate authentication in Firefox or Mozilla:

1. Exit the browser.
2. Open a shell, and run the following commands:

```
export NSPR_LOG_MODULES=negotiateauth:5
export NSPR_LOG_FILE=/tmp/moz.log
```

This will enable verbose logging, and all information will be logged to **/tmp/moz.log**, which may give a clue to the problem. Restart your browser from that shell, and visit the website you were unable to authenticate to earlier.

#### Analyzing the Symptoms

Refer to the following symptoms and possible solutions to help resolve issues with Negotiate authentication.

1. If you receive output similar to the following:

```
-1208550944[90039d0]: entering nsNegotiateAuth::GetNextToken()
-1208550944[90039d0]: gss_init_sec_context() failed: Miscellaneous
failure
No credentials cache found
```

it means that you do not have Kerberos tickets, and need to run **kinit**. Refer to [To test your IPA installation](#): for more information.

2. If you can run **kinit** successfully but you are unable to authenticate, and the log file contains output similar to the following:

```
-1208994096[8d683d8]: entering nsAuthGSSAPI::GetNextToken()
-1208994096[8d683d8]: gss_init_sec_context() failed: Miscellaneous
failure
Server not found in Kerberos database
```

it generally indicates a Kerberos configuration problem. Ensure you have the following in the [domain\_realm] section of the **/etc/krb5.conf** file:

```
.example.com = EXAMPLE.COM
example.com = EXAMPLE.COM
```

3. If nothing appears in the log file it is possible that you are behind a proxy, and that proxy is removing the HTTP headers required for Negotiate authentication. Try to connect to the server using HTTPS instead, which allows the request to pass through unmodified. Then proceed to debug using the log file, as described above.

## 3.4. Using a Browser on Another System

Procedure 3.5. To set up a browser on another system that already has Kerberos set up for a different realm:

1. Copy the `/etc/krb5.conf` file from the IPA server to the client system. Do not overwrite the existing `krb5.conf` file. Run the following command on the IPA server:

```
# scp /etc/krb5.conf root@ipaclient:/etc/krb5_ipa.conf
```

2. On the IPA client, open a shell and run the following commands:

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

```
$ kinit user@EXAMPLE.COM
```

```
$ /usr/bin/firefox
```

3. Configure and test **Firefox** as described in [Section 3.3, “Configuring Your Browser”](#).



# Setting up Synchronization Between IPA and Active Directory

IPA provides bidirectional user identity and password synchronization with Microsoft Active Directory. In order for this synchronization to occur, you need to ensure that certain aspects of both IPA and Active Directory are correctly set up. This is covered in the following sections.

## 4.1. Introduction

To synchronize user identity information between Fedora Directory Server and Windows Active Directory, IPA employs a plug-in that extends the functionality of the Fedora Directory Server Windows Sync utility. This plug-in allows IPA to perform the data manipulation necessary to achieve synchronization between Fedora Directory Server and Windows Active Directory. The IPA Windows Sync plug-in uses the `ipaWinSyncUserAttr` parameter to specify what attributes and values to add to new users that are synchronized from Active Directory.

Refer to the *IPA Administration Reference* for more information on the IPA Windows Sync plug-in.

Refer to the [Fedora Directory Server Administration Guide](#)<sup>1</sup> for more information on the Windows Sync utility.

## 4.2. Prerequisites

### 4.2.1. Domain Name Considerations

IPA clients find, or discover, IPA servers using a process known as *Service Discovery*. This can occur automatically, using DNS, or manually, by entering the IPA server details during the client configuration phase. If your Active Directory installation is in the same domain as the IPA server, it is possible that when you install IPA clients they will not discover the IPA server, but rather the Active Directory DNS. This means that IPA commands run on the client will fail because the client cannot contact the IPA server.

To avoid this situation, use a separate domain for your IPA and Active Directory servers. If this is not possible, use the `--force` parameter when you run the **ipa-client-install** script.

Refer to the [IPA Client Configuration Guide](#)<sup>2</sup> for more information on installing and configuring IPA clients.

### 4.2.2. Setting up Active Directory

The Windows Sync utility requires TLS/SSL to synchronize password changes. Therefore, you need to set up Active Directory as an SSL server. The easiest way to achieve this is to install Microsoft Certificate System in Enterprise Root Mode; Active Directory will then automatically enroll to retrieve its SSL server certificate.

---

<sup>1</sup> [http://www.redhat.com/docs/manuals/dir-server/ag/8.0/Windows\\_Sync.html](http://www.redhat.com/docs/manuals/dir-server/ag/8.0/Windows_Sync.html)

<sup>2</sup> [http://www.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_IPA/1.1/html/Client\\_Configuration\\_Guide/index.html](http://www.redhat.com/docs/en-US/Red_Hat_Enterprise_IPA/1.1/html/Client_Configuration_Guide/index.html)



### Note

You need to install both the `winsync` and `passynch` utilities to synchronize User IDs and attributes as well as passwords.

You need to install the `passynch` utility on all AD domain controllers to enable password synchronization from AD to IPA.

Refer to the [Fedora Project Windows Sync Howto](#)<sup>3</sup> for information on setting up Active Directory as an SSL server.

After you have installed Microsoft Certificate System, you need to save the CA certificate in ASCII (PEM) format. This CA Certificate is required to create the synchronization agreement.

#### Procedure 4.1. To save the CA certificate in ASCII format:

1. Navigate to My Network Places and drill down to the CA distribution point. On Windows 2003 Server this is typically `C:\WINDOWS\system32\certsrv\CertEnroll\`
2. Double-click the security certificate file (`.crt` file) to display the **Certificate** dialog box.
3. On the **Details** tab, click **Copy to File** to start the **Certificate Export Wizard**.
4. Click **Next**, select **Base-64 encoded X.509 (.CER)** and then click **Next**.
5. Specify a suitable directory and file name for the exported file. The file name is not important. Click **Next** to export the certificate, and then click **Finish**. You should receive a message stating that the export was successful.
6. Click **OK** to exit the wizard.

Refer to [Section 4.4, "Creating Synchronization Agreements"](#) for information on how to use the CA Certificate to create the synchronization agreement.

---

<sup>3</sup> <http://directory.fedoraproject.org/wiki/Howto:WindowsSync>

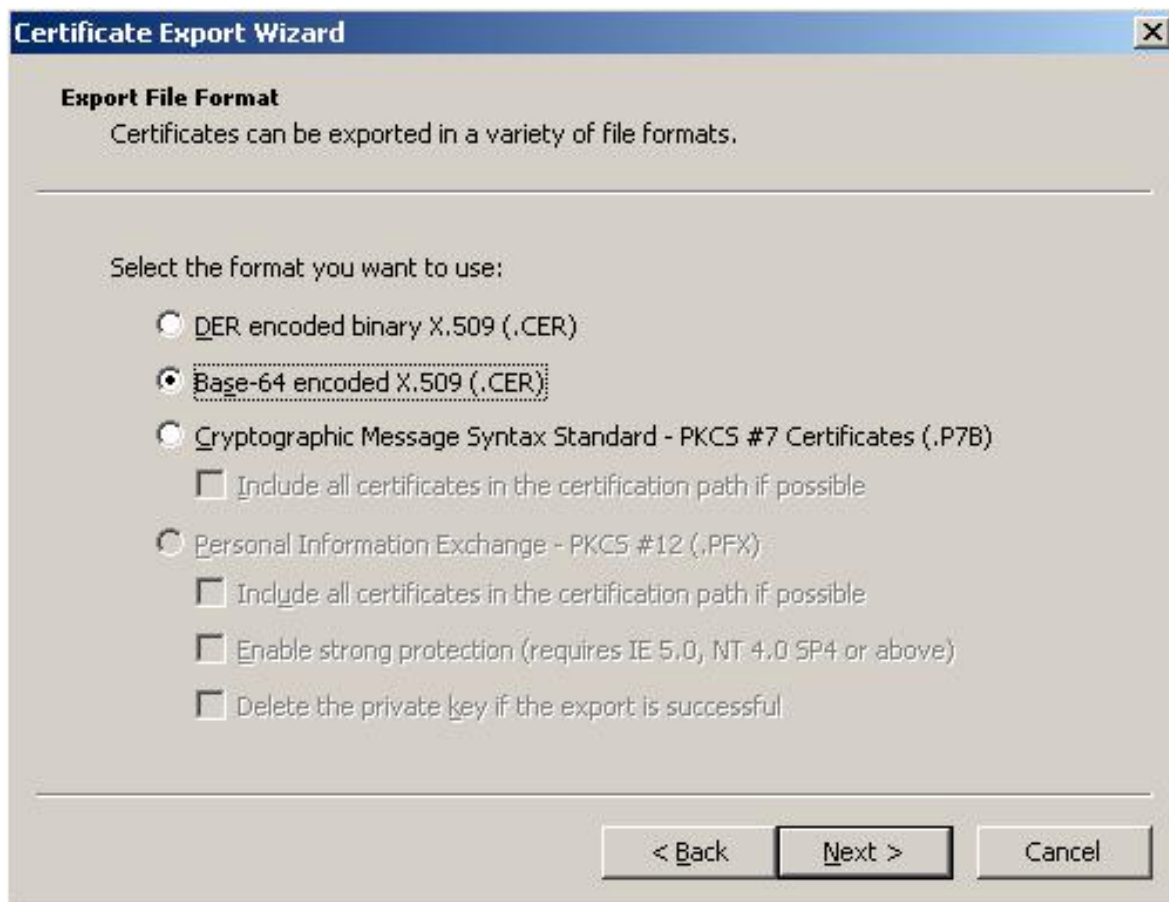


Figure 4.1. Select Base-64 encoded X.509 to export the security certificate as ASCII

### 4.3. Setting up Windows Sync on the IPA Server

The Windows Sync plug-in is installed on the IPA server. The `ipa-server-install` script automatically installs the plug-in configuration entry and enables it by default. The Windows Sync plug-in is only ever called if Windows Sync is used. No other configuration is required.

### 4.4. Creating Synchronization Agreements

Use the `ipa-replica-manage add` command to create the synchronization agreement. The following command-line arguments apply to creating synchronization agreements:

- `--winsync` — specifies that this is a Windows Sync agreement.
- `--binddn` — the full DN of the user to use. The DS will bind to Active Directory as this user to read and write changes. This user requires read, search, and write permissions on the Active Directory subtree, including password changes, as well as permission to use the DirSync control (that is, it must be able to use Replication).
- `--bindpw` — the password for the user specified by the `--binddn` argument.
- `--cacert` — the full path and file name of the ASCII/PEM-encoded Windows Active Directory CA certificate. This certificate will be installed in the Directory Server certificate database as "Imported CA".

- `--win-subtree` — the DN of the Windows subtree containing the users you want to synchronize. The default value is `cn=Users, $SUFFIX` — this is what Windows AD typically uses as the default value.

The following example illustrates adding a new WinSync agreement:

```
ipa-replica-manage add --winsync --binddn
  cn=administrator,cn=users,dc=example,dc=com \
--bindpw password --cacert /path/to/certfile.cer adserver.example.com -v
```

Example 4.1. Adding a WinSync agreement between an IPA server and an AD server.

## 4.5. Modifying Synchronization Agreements

You can change the behavior of the synchronization agreement to suit the changing needs of your organization. You can modify a number of attributes related to the synchronization agreement using default tools provided with IPA.

The following example illustrates changing the synchronization behavior of account lock status. By default, account lock status is synchronized between IPA and AD. This means that accounts that are locked in IPA are also locked (disabled) in AD, and vice versa. You can change this synchronization behavior as follows:

```
$ ldapmodify -x -D "cn=directory manager" -w password
dn: cn=ipa-winsync,cn=plugins,cn=config
changetype: modify
replace: ipaWinSyncAcctDisable
ipaWinSyncAcctDisable: none

modifying entry "cn=ipa-winsync,cn=plugins,cn=config"
```

Example 4.2. Configuring the IPA WinSync agreement to not synchronize account lock status information.

The default value of the `ipaWinSyncAcctDisable` attribute is **both**. If you change this value to **none**, as described in example, account lock status synchronization is completely disabled. Valid values for `ipaWinSyncAcctDisable` are **both**, **to\_ad**, **to\_ds**, and **none**.

### 4.5.1. Changing the Default Synchronization Subtree

When you create synchronization agreements, two default containers are used as the source of the user accounts to synchronize between IPA and Windows Active Directory. IPA uses the `cn=users, cn=accounts, $SUFFIX` subtree as the default container, and Windows uses the `CN=Users, $SUFFIX` subtree. You can use the `--win-subtree` argument to the **ipa-replica-manage add** command to override the default Windows subtree.



#### Note

If you pass such arguments to the bash or other shell, ensure that you quote spaces and other shell metacharacters. For example, the argument `--win-subtree=cn=users,`

```
dc=example, dc=com will fail. The argument --win-subtree="cn=users,  
dc=example, dc=com" will succeed.
```

IPA does not currently support modifying the default synchronization container while you are creating the synchronization agreement. You can, however, change the container after the agreement has been established. To do so, you can either modify the `dse.ldif` file directly (ensure that you stop the directory server before editing this file), or use `ldapmodify` to change `nsds7WindowsReplicaSubtree`.

Refer to the *Changing Synchronization Subtrees* section of the *IPA Administration Reference* for more information on changing the default synchronization subtree.

## 4.6. Deleting Synchronization Agreements

You can use the IPA administration tools to delete existing synchronization agreements. For example, to delete an agreement with the AD server `adserver.example.com`, run the following command:

```
# ipa-replica-manage del adserver.example.com
```

This removes the replication agreement between the IPA and AD servers. To complete the operation, you need to remove the AD certificate from the IPA server. Run the following command to remove the AD certificate:

```
# certutil -D -d /etc/dirsrv/slapd-$REALM/ -n "Imported CA"
```



## Setting up Multi-Master Replication

Replication is the mechanism by which directory data is automatically copied from one Directory Server to another. Updates of any kind, such as adding, modifying, or deleting entries, are automatically mirrored to other Directory Servers using replication.

IPA uses a number of scripts to install, configure, and manage replica servers and replication agreements. These are discussed in the following sections.

### 5.1. Preparing the Replica Servers

Replica servers require the same preparation as master IPA servers. That is, there should be no existing Directory Server installations, the ports required by IPA must be free and available, and the server's machine name must be set and resolve to its public IP address (not to localhost or 127.0.0.1). The replica server must also be able to contact the master LDAP server, which means that DNS or a similar look-up system must be working correctly.

Refer to [Section 3.2, “Configuring the IPA Server”](#) for more information about these and other considerations for installing an IPA server.

### 5.2. Installing the Server Packages

Follow the steps in [Chapter 3, Setting up the IPA Server](#) to install all of the required packages for the replica server.



#### Warning

Do not run the `ipa-server-install` script on the replica servers.

### 5.3. Creating the Replica Information File

You need to create a *replica information file* for each replica that you intend to create. This file contains the realm information required to configure the replica server.



#### Note

Replica information files are version specific. The master and replica servers must have the same version of IPA installed before you can configure the replica.

Before you create the replica information file, ensure that the master IPA server is configured correctly and functioning properly. The master IPA server is the server from which all IPA replica servers are created.

### Procedure 5.1. To create the replica information file:

- Run the following command on the master IPA server, where `ipareplica.example.com` is the FQDN of the server where you are going to create the replica. You need the Directory Server Administrator's password to run this command.

```
# ipa-replica-prepare ipareplica.example.com
```

This will produce output similar to the following:

```
Determining current realm name
Getting domain name from LDAP
Preparing replica for ipareplica.example.com from ipaserver.example.com
Creating SSL certificate for the Directory Server
Creating SSL certificate for the Web Server
Copying additional files
Finalizing configuration
Packaging the replica into replica-info-ipareplica.example.com
```



### Note

Each replica information file is created in the `/var/lib/ipa/` directory as a **GPG**-encrypted file. Each file is named specifically for the replica server for which it is intended. You cannot use the same replica information file for multiple replicas. In the previous example, the resulting file name would be `replica-info-ipareplica.example.com.gpg`



### Warning

Replica information files contain sensitive information. Take appropriate steps to ensure that they are properly protected.

## 5.4. Configuring an IPA Replica

After you have created the replica information file, you need to copy it to the replica server and run the required script to configure the replica.

### Procedure 5.2. To configure an IPA replica:

1. Run the following command on the IPA master server to copy the replica information file to the replica server:

```
# scp /var/lib/ipa/replica-info-ipareplica.example.com.gpg \
root@ipareplica:/var/lib/ipa/
```

2. On the replica server, run the replica installation script, passing it the replica information file you copied from the master:

```
# ipa-replica-install /var/lib/ipa/replica-info-
ipareplica.example.com.gpg
```

The replica installation script runs a test to ensure that the replica file being installed matches the current hostname. If they do not match, the script returns a warning message and asks for confirmation. This could occur on a multi-homed machine, for example, where mismatched hostnames may not be an issue.

3. Enter the Directory Manager (DM) password when prompted.

The script then configures a Directory Server instance based on information in the replica information file, and initiates a replication process. When this has successfully completed, the script continues to set up a complete master replica of the IPA server.



### Note

You can only have a single Directory Server instance on an IPA server, the one used by IPA itself. If the replica installation script detects an existing Directory Server instance, you will be prompted to remove it.

## 5.4.1. Updating DNS for IPA Replicas

After you have configured a new IPA replica, you should update your DNS entries so that IPA clients can discover the new server. For example, for an IPA replica with a server name of \$HOST, you should add the following entries to your zone file:

```
_ldap._tcp          IN SRV 0 100 389 $HOST
_kerberos._tcp      IN SRV 0 100 88 $HOST
_kerberos._udp      IN SRV 0 100 88 $HOST
_kerberos-master._tcp IN SRV 0 100 88 $HOST
_kerberos-master._udp IN SRV 0 100 88 $HOST
_kpasswd._tcp       IN SRV 0 100 464 $HOST
_kpasswd._udp       IN SRV 0 100 464 $HOST
_ntp._udp           IN SRV 0 100 123 $HOST
```

## 5.5. Managing Multi-Master Replication

You can use the **ipa-replica-manage** command to manage certain aspects of replication between IPA servers. This includes listing, adding, and deleting replication agreements, and also manually performing replication initialization and updates.

Initialization is typically only required when you first set up replication, or if a problem arises that causes replication to fail. Initialization erases all data on the target replica (the *consumer*) and re-copies all data from the master. That is, it completely destroys the database on the consumer and rebuilds it with data from the master.

Sending updates is the regular incremental replication protocol. Typically, this is not needed because the server sends changes when required, provided that the replication agreement schedule allows it.

Refer to the `ipa-replica-manage` man page for a full description of the available options.



### Note

There is no web browser interface for managing IPA replicas. You need to use the command line.

Refer to the [Managing Replication](#)<sup>1</sup> section of the *Directory Server Administration Guide* for information about managing replication.

## 5.6. Troubleshooting Multi-Master Replication

Refer to the following sections of the *Directory Server Administration Guide* for information about troubleshooting replication:

- [Solving Common Replication Conflicts](#)<sup>2</sup>
- [Troubleshooting Replication-Related Problems](#)<sup>3</sup>

---

<sup>1</sup> [http://www.redhat.com/docs/manuals/dir-server/ag/8.0/Managing\\_Replication.html](http://www.redhat.com/docs/manuals/dir-server/ag/8.0/Managing_Replication.html)

## Setting up IPA to run as an Apache Virtual Host

If you have a standard Apache instance running on port 80, you can configure IPA to run on a secondary port, for example, on port 8089. You should be aware, however, that in this configuration, IPA does not use SSL; all requests will use standard HTTP.

The following procedure assumes that IPA is configured to run on port 80, and that you want to move it to port 8089.

**Procedure 6.1.** To configure IPA to run on port 8089:

1. Log in as the root user.
2. Edit the `/etc/httpd/conf.d/ipa.conf` file. Add the following three lines to the beginning of the file:

```
Listen 8089
NameVirtualHost *:8089
<VirtualHost *:8089>
```

3. Add the following line to the end of the file:

```
</VirtualHost>
```

This wraps the entire IPA configuration in a virtual host, and ensures that Apache is listening to that port.



### Note

You cannot use port 8080. This port is used by the `ipa_webgui` service.

4. Comment out the following rewrite rules from the `/etc/httpd/conf.d/ipa.conf` file:

```
-----
# Redirect to the fully-qualified hostname. Not redirecting to secure
# port so configuration files can be retrieved without requiring SSL.
RewriteCond %{HTTP_HOST}      !^host.foo.com$ [NC]
RewriteRule ^/(.*)           http://host.foo.com/$1 [L,R=301]

# Redirect to the secure port if not displaying an error or retrieving
# configuration.
RewriteCond %{SERVER_PORT}    !^443$
RewriteCond %{REQUEST_URI}    !^/(errors|config|favicon.ico)
RewriteRule ^/(.*)           https://host.foo.com/$1 [L,R=301,NC]
-----
```

5. Reload the `httpd` service.

```
# service httpd reload
```

This configures IPA to run on port 8089, leaving port 80 free for your normal web site.

---

## Appendix A. Revision History

Revision 1.1    Tue Nov 25 2008                      David O'Brien [davido@redhat.com](mailto:davido@redhat.com)

BZ 465210. Update networking services configuration section.

BZ 469128. Updates from tech review.

BZ 471491. How to manually remove certificate from cert.db after removing a winsync agreement.

BZ 471521. How to manually modify winsync agreement.

BZ 470420. Update /etc/hosts requirements.

Added chapter on Active Directory synchronization.

Updated network configuration section.

Updated hardware requirements section.

Revision 1.0    Wed Apr 30 2008                      David O'Brien [davido@redhat.com](mailto:davido@redhat.com)

Created.

