freeIPA 1.2.1

Administration Guide

IPA Solutions from the IPA Experts

freeIPA 1.2.1 Administration Guide
IPA Solutions from the IPA Experts
Edition 1.0

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 details the tasks and procedures necessary for administering your IPA deployment. It also provides information on how to customize IPA to suit your environment, and information on how to troubleshoot common problems.

Preface

Welcome to the IPA Administration Guide. This guide provides the information necessary to administer your IPA deployment. It includes detailed information on working with user and group accounts, how to set up and manage the password policy, and how to configure various types of access control. It also covers basic troubleshooting techniques to help you resolve any issues that might arise.

1. Audience

The IPA Administration Guide is intended for system administrators and those involved in the ongoing maintenance of IPA.

This guide assumes a good understanding of various operating systems, including `Linux`, `Solaris` and other UNIX systems, `Macintosh` and Microsoft `Windows`. It also assumes a working knowledge of `LDAP` and 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*[1] 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.

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

[1] https://fedorahosted.org/liberation-fonts/

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* against the Documentation component.

When submitting a bug report, be sure to mention the manual's identifier: *Administration_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.

Configuring Users and Groups

1.1. Managing User Accounts

The primary activities associated with managing user accounts, such as creating and deleting accounts, are performed by IPA Administrators. Other activities, such as editing various user account attributes and changing group membership, can be delegated to other accounts.

Refer to *Section 3.1.1.1, "Delegating Administrative Privileges"* for more information.

You can use either the web interface or the command line to manage user accounts. Each interface provides identical functionality, however the web interface displays a greater range of information for each user in an easy to use format.

The web interface displays mandatory fields in a different color. Certain other fields, such as Common Name, Display Name, Initials, Login, and E-mail Address, are populated automatically. You can change these values as required. The UID, GID, and Home Directory are automatically generated by the server.

If you use the command line to add user accounts, you will be prompted for any required information.

Refer to *Section 3.2.3.2, "User Setting Attributes"* for information on the attributes that apply to user accounts, and especially for information regarding users' **`/home`** directories.

1.1.1. Creating User Accounts

You can use the **Add User** page in the web interface, or the **`ipa-adduser`** command on the command line to create user accounts. These procedures are described below.

Note

IPA supports a wide range of username formats, but you need to be aware of any restrictions that may apply to your particular environment. For example, a username that starts with a digit may cause problems for some UNIX systems.

The range of username formats supported by IPA can be described by the following regular expression:

```
[a-zA-Z0-9_.][a-zA-Z0-9_.-]{0,30}[a-zA-Z0-9_.$-]
```

The trailing $ symbol is permitted for Samba 3.x machine support.

1.1.1.1. Using the Web Interface

Procedure 1.1. To create a user account using the web interface:

1. On the IPA homepage, click **Add User** in the **Tasks** list to display the **Add User** page.
2. Enter the required details for the user.

3. If required, add the user account to a group. All users are automatically added to the global group ipausers.

> **Note**
> You can configure the global group to suit your deployment. For example, you may prefer to change its name to include your company name.

4. When you have entered the required account details, click **Add User**.

> **Note**
> It is not essential to provide a password when you create an account. For example, you might create an account for a service (rather than a user), and such an account may not require a password. For a user account, however, you need to provide an initial password so that the user can log in to their account. Users are required to change their initial password the first time they log in.

The following example illustrates using the web interface to add the Identity and Account details for a new user.

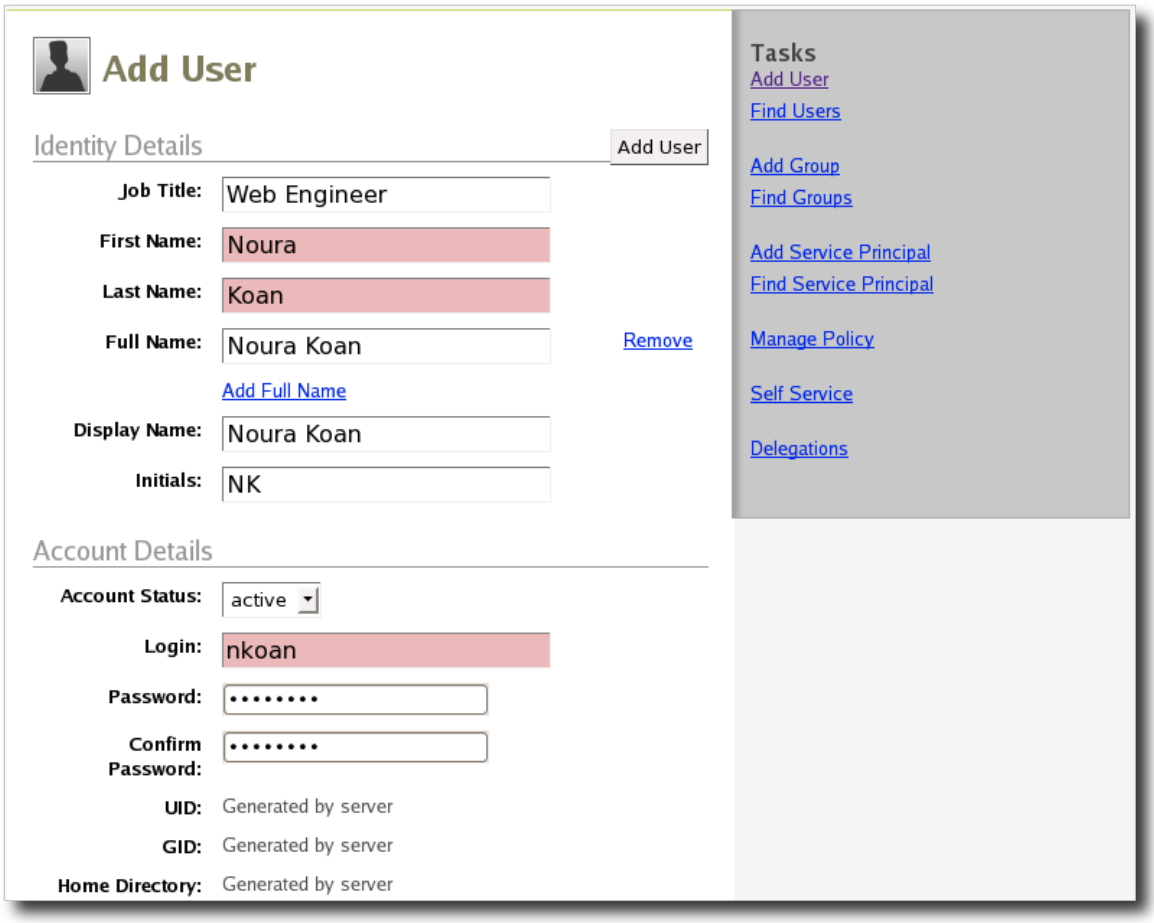

Figure 1.1. Using the web interface to add a new user.

1.1.1.2. Using the Command Line

Use the **ipa-adduser** command to add users to IPA. You can pass attributes directly on the command line, or run the command with no parameters to enter interactive mode. Interactive mode prompts you to enter the basic attributes required to add a new user. You can add further attributes using the **ipa-moduser** command. Use the **ipa-moduser --list** command to view a list of the attributes that you can modify using this command.

Procedure 1.2. To create the user `jlamb` using the command line:

- Open a shell and run the following command:

```
$ /usr/sbin/ipa-adduser -f John -l Lamb -p secret jlamb
```

The following example illustrates using the **ipa-adduser** command in interactive mode to create a user account:

```
$ /usr/sbin/ipa-adduser
First name: Jinny
Last name: Pattanajee
Login name: jpattan
gecos[]: Jinny Pattanajee
home directory [/home/jpattan]:
shell [/bin/sh]:
jpattan successfully added
```

Press **Enter** at each prompt to accept the default values (enclosed in square brackets), or type an alternative.

Refer to the **ipa-adduser** man page for more information.

1.1.2. Editing User Accounts

1.1.2.1. Using the Web Interface

Members of the IPA Administrators group can edit the details of any user account. Other users can edit certain user account details, according to the delegations that have been configured.

Procedure 1.3. To edit a user account using the web interface:

1. Click **Find Users** in the **Tasks** list to display the **Find Users** page.
2. Enter the name or a key word of the user that you want to edit in the search field, and click **Find Users**.
3. In the search results, click the name of the user that you want to edit. The user is displayed on the **View User** page. If the user does not appear in the search results, try using broader search terms.
4. Click **Edit User** to display the **Edit User** page, where you can edit user attributes.
5. Edit the user attributes as required, and click **Update User**. Note that not all fields are immediately editable; select the **Edit Protected Fields** check box to edit the Password, Home Directory, and some other fields.

Warning

It is possible to edit the UID and GID of user accounts, however this is not recommended. Changing these IDs will not cause problems internally for IPA, but it can lead to other issues, such as changes to file ownership and security problems.

1.1.2.2. Using the Command Line

Use the **ipa-moduser** command to modify user account details, such as adding, removing or changing attributes. The following examples illustrate the use of this command:

To update the Zip code, Display Name, and Employee Type for the user jsmith:

```
$ /usr/sbin/ipa-moduser --set postalCode=50211 --set displayName="John Smith" --set employeeType=permanent jsmith
```

To remove the Pager and Home Phone attributes from the same user:

```
$ /usr/sbin/ipa-moduser --del pager --del homePhone jsmith
```

To retrieve a partial list of the default attributes that you can manage with ipa-moduser:

```
$ /usr/sbin/ipa-moduser --list
```

The list of attributes corresponds to those available in the web interface, not including any custom attributes that may have been defined.

1.1.3. Activating and Inactivating User Accounts

IPA user accounts can be set to a status of **Active** or **Inactive**. If you inactivate a user account, that user can no longer log in to IPA, change their password, or perform any other tasks. Any existing connections will remain valid until their Kerberos TGT and other tickets expire, but they will not be able to renew them. The account and all associated information still exists, but is inaccessible by the user.

1.1.3.1. Using the Web Interface

Procedure 1.4. To inactivate a user account using the web interface:

1. Find the user that you want to inactivate as described in *Section 1.1.2, "Editing User Accounts"*.
2. Click **Edit User** to display the **Edit User** page, where you can edit user attributes.
3. In the **Account Details** section, select **inactive** in the **Account Status** drop-down list, and then click **Update User**.

The account remains inactive and inaccessible to the user until reactivated by an IPA Administrator.

1.1.3.2. Using the Command Line

Use the **ipa-lockuser** command to activate or inactivate user accounts.

To lock (inactivate) the `jsmith` user account:

```
$ /usr/sbin/ipa-lockuser jsmith
```

To unlock (activate) the `jsmith` user account:

```
$ /usr/sbin/ipa-lockuser -u jsmith
```

1.1.4. Deleting User Accounts

If you delete an IPA user account, all of the information stored in the entry for that identity is lost. This includes the user's full name, group membership, phone numbers, and passwords. The actual user account and home directory still exist, be they on a server, local machine, or other provider, but they are no longer accessible via IPA.

Unlike inactivation, if you delete a user account, it cannot be retrieved. If you need this user account again, you need to recreate it and add all of the account details manually.

Note

You cannot delete or rename the `admin` account, nor can you remove it from the `admins` group.

1.1.4.1. Using the Web Interface

Procedure 1.5. To delete a user account using the web interface:

1. Find the user that you want to delete as described in *Section 1.1.2, "Editing User Accounts"*.
2. Click **Edit User** to display the **Edit User** page.
3. Click **Delete User**, and then click **OK**.

1.1.4.2. Using the Command Line

Use the **`ipa-deluser`** command to delete user accounts. For example:

To delete the `jsmith` user account:

```
$ /usr/sbin/ipa-deluser jsmith
```

1.2. Managing Groups

IPA uses groups to facilitate the management and administration of both users and permissions. Three groups are created during the installation process: `ipausers`, `admins`, and `editors`. All of these groups are required for IPA operation.

The IPA Administrator is a member of the `admins` group. You cannot delete the IPA Administrator, nor can you remove this user from this group. All other users belong to the global group `ipausers`, and you can create as many additional groups as you require.

Note

Some operating systems limit the number of groups that you can create. For example, `Solaris` and `AIX` allow only 16 groups per user. IPA Administrators need to be aware of this limitation, especially when using nested groups.

The `editors` group is a special group used by the web interface. Members of this group have at least one delegation, which means they can edit records apart from their own.

You can create groups based on the departments within your organization, for example, Devel, Finance, and HR. You can also create groups based on the permissions, or roles, required to manage your departmental or other groups. Refer to *Section 3.1, "Configuring Access Control"* for information on using groups to define roles.

Nested Groups

You can also create nested groups. For example, you can create a group called "Documentation", and then create sub-groups such as "Writers", "Translators", and "Editors". You can add users to each of the sub-groups to suit the needs of your organization.

Note

Any users that you add to a sub-group automatically become members of the parent group.

Warning

Avoid the creation of cyclic groups; that is, groups that contain groups that in turn contain their own ancestors, and avoid creating group names that contain spaces. Either of these conditions can lead to unexpected behavior.

Refer to *Section 3.1, "Configuring Access Control"* for information on using groups to define roles.

1.2.1. Creating Groups

1.2.1.1. Using the Web Interface

Procedure 1.6. To create a group using the web interface:

1. On the IPA homepage, click **Add Group** in the **Tasks** list to display the **Add Group** page.
2. Enter a name and description for the group. The GID (Group ID) is automatically generated by the IPA server.
3. Add any users that you want to include in this group:
   a. Enter the login name or other search term in the **To Add** field, and click **Find**.
   b. Locate the users that you want to include in this group, and click **add**.

4. When you have finished adding members, click **Add Group** to return to the **View Group** page, and display details of the newly-added group.

The following diagram illustrates adding members to a new group.

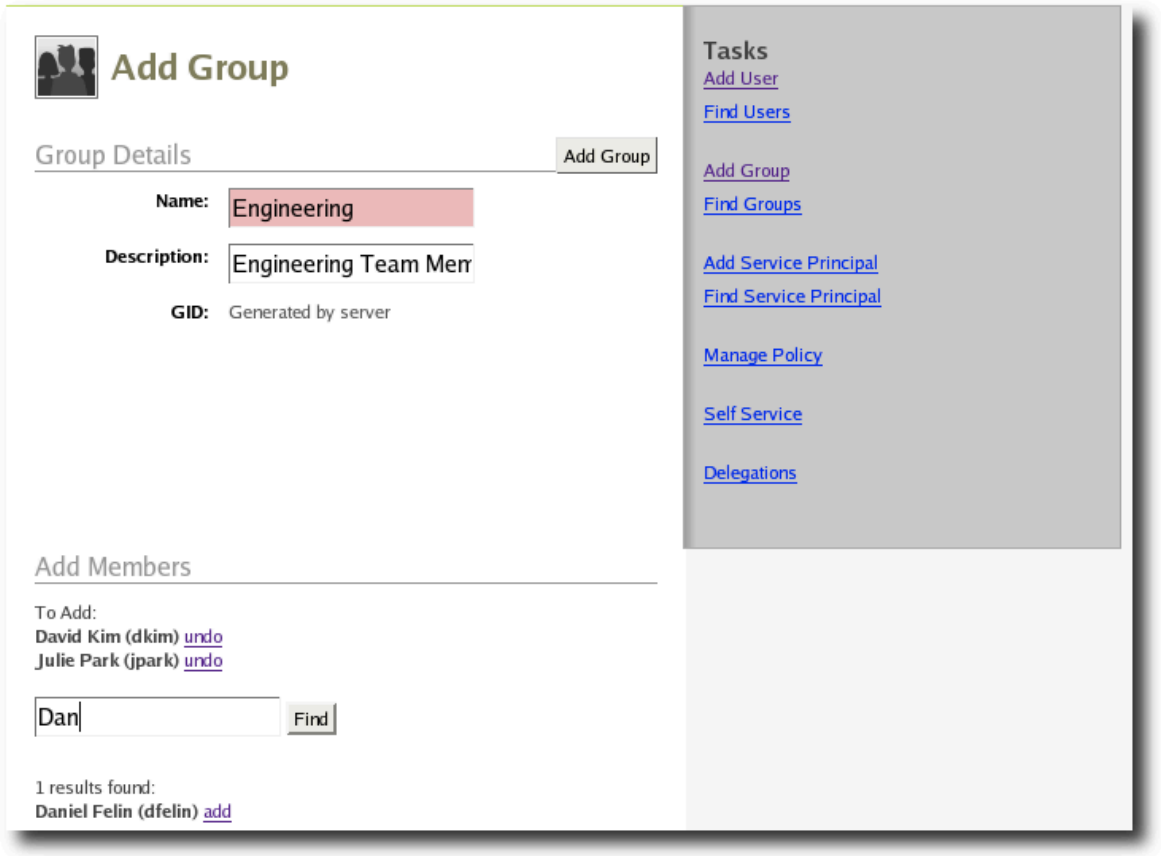

Figure 1.2. Adding members to a new group.

1.2.1.2. Using the Command Line

Use the **ipa-addgroup** command to add groups. You can include attributes on the command-line or use the command interactively. For example,

To create a group called "Engineering" using the command line:

```
$ /usr/sbin/ipa-addgroup
Group name: Engineering
Description: All members of the engineering group
Engineering successfully added
```

Alternatively, include all the required attributes on the command-line:

```
$ /usr/sbin/ipa-addgroup -d "All authors, editors, and translators."
 Documentation
Documentation successfully added
```

The group name and description are mandatory fields. If either of these are not included on the command-line, you will be prompted to include them.

Note

You cannot add users to a newly-created group using the **ipa-addgroup** command. You first need to create the group, and then use the **ipa-modgroup** command to add users. For example:

```
$ /usr/sbin/ipa-modgroup -a user01,user02,user03 Engineering
```

1.2.2. Editing Groups

You can edit any of the attributes that define a group, as well as add or remove members. Some attributes are read-only by default, however you can edit these attributes if required.

1.2.2.1. Using the Web Interface

Procedure 1.7. To edit a group using the web interface:

1. Click **Find Groups** in the **Tasks** list to display the **Find Groups** page.
2. Enter the name or a key word of the group that you want to edit in the search field, and click **Find Groups**.
3. In the search results, click the name of the group that you want to edit. The group is displayed on the **View Group** page. If the group does not appear in the search results, try using broader search terms.
4. Click **Edit Group** to display the **Edit Group** page, where you can edit group attributes.
5. Edit the group attributes as required, and click **Update Group**. Note that if you want to change the Name or GID of the group, you need to select the **Edit Protected Fields** check box.

Warning

Do not change the Group Name or GID unless absolutely necessary, because it can have unexpected effects on permissions, ACIs, and other aspects of IPA functionality.

If you rename a group used in an ACI, the ACI itself is not updated, the result being that the group will fall out of the ACI scope. To avoid this issue, ensure that any changes to group names are reflected in IPA Delegations. IPA does not currently support per-user ACIs, so this issue only affects groups.

1.2.2.2. Using the Command Line

Use the **ipa-modgroup** command to edit groups. The following are some simple examples of using this command. Refer to the **ipa-modgroup** man page for more information.

To add the user `user01` to the `admins` group:

```
$ /usr/sbin/ipa-modgroup -a user01 admins
```

To remove the user `user01` from the `admins` group:

```
$ /usr/sbin/ipa-modgroup -r user01 admins
```

To change the description of the `admins` group to "IPA Administrators":

```
$ /usr/sbin/ipa-modgroup -d "IPA Administrators" admins
```

To add the group `sysadmins` to the `admins` group:

```
$ /usr/sbin/ipa-modgroup -g sysadmins admins
```

To remove the `Editors` group from the `Documentation` group:

```
$ /usr/sbin/ipa-modgroup -e Editors Documentation
```

1.2.3. Activating and Inactivating Groups

IPA groups can be set to a status of **`Active`** or **`Inactive`**. If you inactivate a group, all of the members of that group are also inactivated. This means that they cannot log in to IPA, change passwords, or access resources controlled by IPA. The accounts within an inactivated group still exist, but they are inaccessible.

This also applies to nested groups. If you inactivate a group, then any sub-groups are also inactivated, as are their members. Within these inactive groups, however, you can manually activate individual users or groups if required.

Note

You cannot inactivate the `admins` group.

1.2.3.1. Using the Web Interface

Procedure 1.8. To activate a group using the web interface:

1. Find the group that you want to edit as described in *Section 1.2.2, "Editing Groups"*.
2. Click **Edit Group** to display the **Edit Group** page.
3. Select **inactive** in the **Group Status** drop-down list, and then click **Update Group**.

1.2.3.2. Using the Command Line

Use the **`ipa-modgroup`** command with the `nsaccountlock` option to activate and inactivate groups. For example,

To inactivate the Engineering group:

```
$ /usr/sbin/ipa-modgroup --set nsaccountlock=true Engineering
```

To activate the Finance group:

```
$ /usr/sbin/ipa-modgroup --set nsaccountlock=false Finance
```

1.2.4. Deleting Groups

When you delete a group, only the immediate group is removed; members of the group are not affected. That is, unlike inactivation, there is no cascading effect when you delete a group.

When you delete a group, any delegations that apply to that group are also removed. For example, suppose you added an "EngineeringManager" group specifically to set up delegations for the Engineering Manager. If you delete the EngineeringManager group, then those delegations are also lost. Unlike with inactivation, these cannot be retrieved. If you need this group and delegation again, you need to recreate them.

1.2.4.1. Using the Web Interface

Procedure 1.9. To delete a group using the web interface:

1. Find the group that you want to delete as described in *Section 1.2.2, "Editing Groups"*.
2. Click **Edit Group** to display the **Edit Group** page.
3. Click **Delete Group**, and then click **OK**.

1.2.4.2. Using the Command Line

Use the **ipa-delgroup** command to delete groups. For example,

To delete the Engineering group:

```
$ /usr/sbin/ipa-delgroup Engineering
```

Configuring Authentication

2.1. Managing Certificates and Certificate Authorities

IPA creates a self-signed Certificate Authority (CA) during the installation process. If you have your own or a preferred CA, however, and want to use your own certificates, IPA provides the necessary tools to import certificates for use by Directory Server and the HTTP server. While not a prerequisite for the correct operation of IPA, it is recommended that you save an ASCII copy of your CA certificate as **/usr/share/ipa/html/ca.crt** to ensure that users download the correct certificate.

2.1.1. Installing Your Own Certificate

Use the **ipa-server-certinstall** command to install your own certificate. You can install the certificate for use by Directory Server, HTTP Server, or both.

To install the certificate for use by Directory Server:

```
# /usr/sbin/ipa-server-certinstall -d /path/to/pkcs12.p12
```

2.1.2. Using Your Own Certificate with Firefox

To continue using the **Firefox** auto-configuration feature, you need an object-signing certificate, and you need to regenerate the **/usr/share/ipa/html/configure.jar** file.

Note

The following procedure assumes that the signing certificate is provided as a PKCS#12 file.

Procedure 2.1. To use your own certificate with Firefox:

1. Create a directory to host the certificate database.

   ```
   # mkdir /tmp/signdb
   ```

2. Create the new certificate database.

   ```
   # /usr/bin/certutil -N -d /tmp/signdb
   ```

3. Import the signing certificate.

   ```
   # /usr/bin/pk12util -i /path/to/pkcs12.p12 -d /tmp/signdb
   ```

4. Make a temporary signing directory.

   ```
   # mkdir /tmp/sign
   ```

5. Copy the IPA javascript file to the temporary signing directory.

   ```
   # cp /usr/share/ipa/html/preferences.html /tmp/sign
   ```

6. Use the certificate you created earlier to sign the javascript file and to regenerate the **configure.jar** file.

```
# /usr/bin/signtool -d /tmp/signdb -k Signing_cert_nickname -Z /usr/
share/ipa/html/configure.jar -e .html
```

2.2. Managing Service Principals

Apart from authenticating users, Kerberos can also provide authentication for services that are accessed by users. For example, you can use Kerberos to provide authentication for HTTP, SSH, and other services. In this scenario, mutual authentication must occur between the service and the KDC (rather than between the user and the KDC). That is, each service must have a valid principal (the service principal) on the server, and the service must use a shared secret to authenticate against the KDC. This is true if the service is provided on the same machine as the KDC, or on a separate machine.

2.2.1. Service Principals and Key Tables (keytabs)

Clients use service principals to inform the KDC which service they need a ticket for. The KDC uses the service principal to provide a secret key to the service when the service principal is created. Service principals and their associated keys are stored in a keytab file. Without an appropriate keytab the service cannot authenticate a client, and the KDC cannot provide the client with a ticket.

Service principals are typically released per service, although it is possible for one service principal to be used for more than one service.

The Importance of Service Principals and keytabs

Service principals and their associated keys play a critical role in a Kerberos-aware environment. This is especially true when services are accessed by multiple users. As long as a valid ticket exists for a specific service, users can access that service using their Kerberos credentials.

For example, if a user tries to mount an NFS directory using Kerberos, then both the NFS server and the user require a valid principal, and share a secret key with the KDC. This is established during the IPA NFS configuration on the server. If the secret key is replaced on the server, for example, by getting a new keytab, then you must export the new keytab to any clients that need NFS mount access to the server.

Protecting keytab Files

To protect your keytab files, consider the following general rules with respect to their permissions and ownership:

- Owner: uid of the process that will use the keytab
- Mode: 0600

For example, set the owner of the **Apache** keytab (**/etc/httpd/conf/ipa.keytab**) to **httpd** and the mode to **0600**.

Failure to export an updated keytab can cause problems that are difficult to isolate. For example, existing service connections may continue to function, but no new connections may be possible.

Due to the critical role that keytabs play in authenticating users and services, and the issues that can arise if they are compromised, ensure that all keytab files are appropriately secured, and have suitable file ownership and permissions established.

2.2.2. Creating and Using Service Principals

You can use the web interface to create service principals and also to search for existing service principals. For security and other reasons, however, it is not possible to retrieve a keytab using the web interface. This has to be done either on the command line on the system where the service is accessed, or on the IPA server itself, and the keytab then exported to the client machine.

The following example demonstrates creating a service principal and keytab on a client machine for the SSH service. The client machine is `ipaclient.example.com` and the IPA server is `ipaserver.example.com`:

```
# kinit admin

# ipa-addservice host/ipaclient.example.com@EXAMPLE.COM

# ipa-getkeytab -s ipaserver.example.com -p host/ipaclient.example.com -k /
etc/krb5.keytab
```

Note

The realm name is optional. The IPA server automatically appends the `Kerberos` realm for which it is configured. You cannot specify a different realm.

The hostname 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 **`ipa-getkeytab`** command is part of the *ipa-client* package, which is only available for clients running Red Hat Enterprise Linux 4 or 5, `Fedora` 7, 8, or 9. For other clients, you need to use this procedure on the server and manually copy the keytab to the client.

You can use the `-e` flag to include a comma-separated list of encryption types to include in the keytab. This supersedes any default encryption type. Refer to the **`ipa-getkeytab`** man page for more information.

Warning

The **`ipa-getkeytab`** command resets the secret for the specified principal. This means that all other keytabs for that principal are rendered invalid.

2.2.3. Configuring NFS on the IPA Server

The following procedure describes how to configure `NFS` on the IPA server and to set up an `NFS` service principal.

Procedure 2.2. Configuring NFS on the IPA Server

1. Configure the export directory.

```
# mkdir /export

# chmod 777 /export
```

2. Configure the **/etc/exports** file as follows:

```
/export  *(rw,fsid=0,insecure,no_subtree_check)
/export  gss/krb5(rw,fsid=0,insecure,no_subtree_check)
/export  gss/krb5i(rw,fsid=0,insecure,no_subtree_check)
/export  gss/krb5p(rw,fsid=0,insecure,no_subtree_check)
```

3. To enable secure NFS, add the following line to **/etc/sysconfig/nfs**

```
SECURE_NFS=yes
```

4. Add a service principal and keytab for NFS.

```
# ipa-addservice nfs/ipaserver.example.com

# ipa-getkeytab -s ipaserver.example.com -p nfs/ipaserver.example.com -
k /etc/krb5.keytab
```

Note

The LinuxNFS implementation still has limited encryption type support. You may need to use the `-e des-cbc-crc` option to the **ipa-getkeytab** command for any nfs/<FQDN> service keytab you want to set up, both on the server and on all clients. This instructs the KDC to generate only DES keys.

5. Run the following commands to reload the NFS configuration and restart the required services:

```
# exportfs -a

# restart services

# service nfs restart

# service rpcgssd restart
```

Configuring Authorization

3.1. Configuring Access Control

IPA provides different mechanisms to support different types of access control. *Host-Based Access Control* provides a means of specifying which users can log in on which machines, while *Delegation* provides a means of controlling access to identity data.

3.1.1. Configuring Delegation

IPA supports access control using a process known as delegation. This provides a means of assigning different permissions to the various users and groups that you create, which in turn controls the level of access they have to IPA identity data.

When the IPA server is initially configured, it creates an administrative account named `admin`. Use this initial account to create any further administrative or other accounts that your deployment requires.

Important

The `admin` account is created in the *`cn=users,cn=accounts`* container, the same as all other users. However, the `admin` account cannot be deleted or renamed, nor can it be removed from the `admins` group.

You normally delegate access control to a group, or role. For example, you might create one group called "EngineeringManager" and another called "OfficeManager", as dedicated groups for the Engineering Manager and Office Manager, respectively.

You would then create a delegation to specify the permissions associated with each group. An Engineering Manager might be able to modify the attributes of everyone in the Engineering group, and the Office Manager would be able to change details such as the office fax number, office contacts, and other office-related details.

When you add the Engineering Manager to the EngineeringManager group, that user automatically gains all of the associated permissions. If the person holding that position changes, you simply change the groups that the user belongs to. There is no need to modify any other values.

3.1.1.1. Delegating Administrative Privileges

Procedure 3.1. To delegate administrative privileges:

1. Create the group to which you want to delegate administrative privileges (the source).
2. Create the group over which this group should have administrative privileges (the target).
3. On the IPA homepage, click **Delegations** in the **Tasks** list to open the **Delegations** page.
4. Click **Add New Delegation** to open the **Add Delegation** page.
5. In the **Delegation Name** field, type a descriptive name for the delegation.

6. In the **People in Group** field, enter the CN of the group or type a suitable search term and click **Find**. This is where you select the source, or the group that will receive the delegation.
7. In the search results, click the required group name. If the required group does not appear in the search results, try using broader search terms.
8. In the **Can Modify** list, select the appropriate check boxes for the delegations that you want to apply. For example, you can specify that this group can modify the home directory, login shell, and organizational unit of its subjects.
9. In the **For People in Group** field, enter the CN of the group or type a suitable search term and click **Find**. This is where you select the target, or the group that will be subject to the delegation.
10. Click **Add Delegation** to create the delegation.

Note

You can only create a delegation for one target at a time. If you require that a source have administrative control over several targets, you need to create a separate delegation for each target.

The following diagram illustrates creating a delegation for the Engineering Manager over the Engineering group.

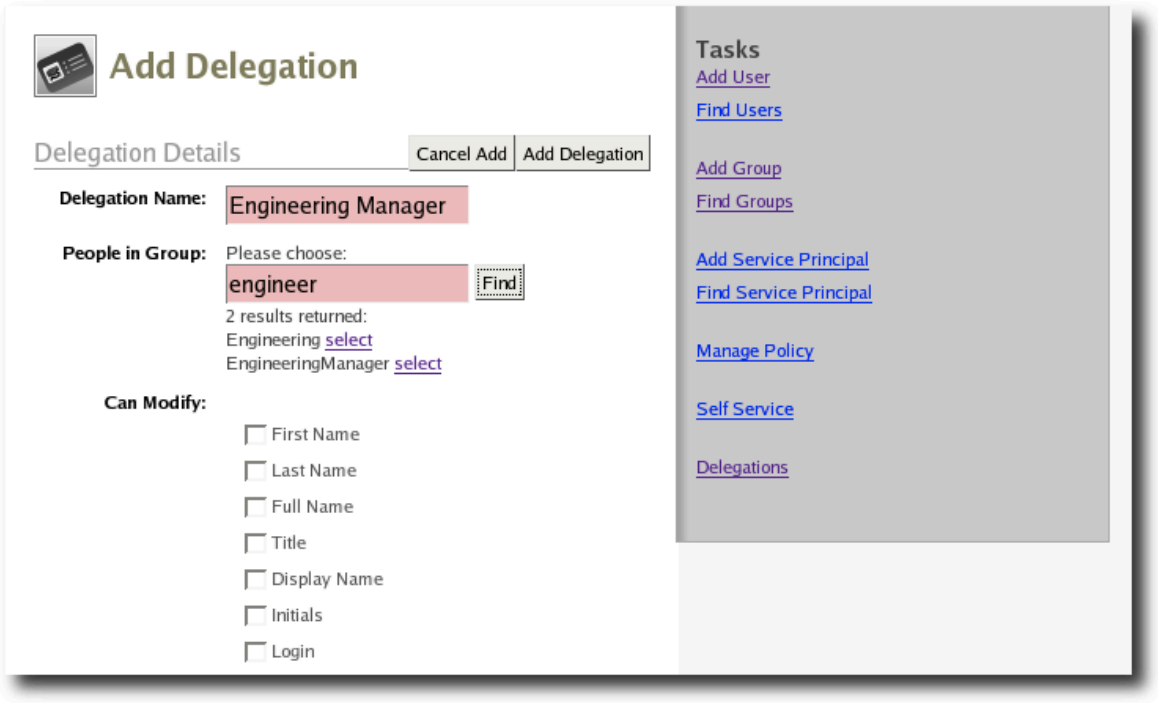

Figure 3.1. Adding a delegation to a group.

3.1.2. Configuring Host-Based Access Control

You can configure Red Hat Enterprise Linux and `Fedora` to allow or deny access to IPA resources and services based on the configuration of the host from which access is attempted. This requires

modification to the **/etc/security/access.conf** and **/etc/pam.d/system-auth** files, as described below.

Procedure 3.2. To configure host-based access control:

1. Modify the **/etc/security/access.conf** file to include the following lines:

```
+ : root : ALL
+ : ipausers : ALL
- : ALL : ALL
```

2. Modify the **/etc/pam.d/system-auth** file to include the following line:

```
account required pam_access.so
```

This configuration specifies that:

- The `root` user can log in.
- All members of the `ipausers` group can log in.
- IPA administrators can not log in (because the `admin` account is not a member of the `ipausers` group).

Note

This example only demonstrates the procedures required to configure host-based access control. It is not intended as a recommended configuration. You need to design your own configuration based on the requirements of your site.

3.2. Managing IPA Policy

The IPA policy specifies various constraints on the way that users can interact with the IPA system as a whole. This affects their user accounts, the details that they can view and edit, minimum password requirements and other details, and also the range of searches that they can perform.

3.2.1. Specifying Search Settings

You can configure various aspects of the IPA search functionality to suit your deployment. For example, you can restrict the number of fields that a user can base a search on, or limit the number of records returned for any particular search.

IPA supports the following search configuration attributes:

- Search Time Limit: The maximum time, in seconds, that a search will run before failing.
- Search Records Limit: The maximum number of records that a search can return. Set this value to zero (0) to specify no limit. The directory server limit (the default value is 2000) still applies.
- User Search Fields: For a user search, specifies the fields to search for the values entered by a user.

- Group Search Fields: For a group search, specifies the fields to search for the values entered by a user.

If you add attributes to the user or group search fields, you should also create a new LDAP index for those attributes to avoid performance degradation. Conversely, the existence of too many indexes can impact write performance, so you need to balance one against the other.

Refer to *Creating Indexes*[1] in the *Directory Server Administration Guide* for information on creating indexes.

3.2.2. Specifying the Password Policy

3.2.2.1. Introduction

IPA supports the specification of various password attributes that help to ensure the security of your system, and also that of individual user accounts. You can specify the password lifetime, length, and the types of characters required in a password, all as part of the IPA Password Policy.

Note

In freeIPA 1.2.1, the password policy is enforced by the KDC. Only a limited number of attributes are currently supported, however this will be extended in later versions.

Because the password policy is enforced by the KDC, any further policy specifications that you implement as part of the Directory Server password policy will not be visible in IPA, and neither will they be enforced.

3.2.2.2. Exceptions to the Password Policy

Different rules apply to changing passwords, depending on your login credentials.

3.2.2.2.1. Changing Passwords as the Directory Manager

If you reset a password using `cn=Directory Manager` credentials (only possible if you manually perform an LDAP password change operation) then you override any checks and the password is set to whatever you specify. The IPA password policy is ignored.

3.2.2.2.2. Changing Passwords as the IPA Administrator

If you reset a password using `admin` credentials (that is, as part of the `admins` group), the IPA password policy is ignored, but the expiration date is set to "now". This means that the user is forced to change the password at login time, and the password policy is then enforced. This is also true for users who have had password changing rights delegated to them.

Consequently, the IPA Administrator can easily create users with "default" passwords and reset user's passwords, but will not know the actual, final password entered by the user. Further, any password that is transmitted from the IPA Administrator to the user, even over insecure channels, is a temporary password. Consequently, it is not critical if it is accidentally disclosed, provided that the user promptly resets it.

[1] http://www.redhat.com/docs/manuals/dir-server/ag/8.0/Managing_Indexes-Creating_Indexes.html

3.2.2.2.3. Changing Passwords as a Regular User

If you are logged in as a regular user (that is, you are not part of the `admins` group, or possessed of any elevated privileges), then you can only change your own password, and these changes are always subject to the IPA password policy.

3.2.2.3. Editing the Password Policy

You can use either the web interface or the command-line to edit the IPA password policy. However, you can only edit those attributes supported by IPA.

3.2.2.3.1. Using the Web Interface

Procedure 3.3. To edit the password policy using the web interface:

1. Click **Manage Policy** in the **Tasks** list, and then click **IPA Policy** to display the **Manage IPA Policy** page.
2. Click **Edit Policy** to display the **Edit IPA Policy** page, where you can edit different aspects of the IPA policy.
3. In the **Password Policy** section, edit the password attributes as required, and click **Update Policy**.

3.2.2.3.2. Using the Command Line

Use the **`ipa-pwpolicy`** command to modify IPA password policy details. This command uses the following syntax:

`ipa-pwpolicy` `[--maxlife days][--minlife hours][--history number][--minclasses number][--minlength number]`

For example, to update the minimum password length to 10 characters, and to specify that no history of passwords be kept:

```
# /usr/sbin/ipa-pwpolicy --minlength 10 --history 0
```

Refer to *Section 3.2.2.4, "Password Policy Attributes"* for information on password policy attributes.

Refer to the **`ipa-pwpolicy`** man page for more information on this command.

3.2.2.4. Password Policy Attributes

The password policy is enforced by the `pwd_extop` SLAPI plug-in. freeIPA 1.2.1 supports the following password policy attributes:

- **Minimum Password Lifetime** (krbMinPwdLife): The minimum period of time, in hours, that a user's password must be in effect before the user can change it. The default value is one hour.

  You can use this attribute to prevent users from changing their password to a "temporary" value and then immediately changing it back to the original value.

- **Maximum Password Lifetime** (krbMaxPwdLife): The maximum period of time, in days, that a user's password can be in effect before it must be changed. The default value is 90 days.

- **Minimum Number of Character Classes** (krbPwdMinDiffChars): The minimum number of different classes, or types, of character that must exist in a password before it is considered valid. The default value is zero (0).

  For example, setting krbPwdMinDiffChars = 3 requires that passwords contain at least one character from three of the supported classes.

  The following character classes are supported:

  - Upper-case characters
  - Lower-case characters
  - Digits
  - Special characters (for example, punctuation)

  The following special classes also exist:

  - Number of repeated characters

    This weights in the opposite direction, so that if you have too many repeated characters you will not meet the quorum to satisfy the "level" expressed by krbPwdMinDiffChars.

- **Minimum Length of Password** (krbPwdMinLength): The minimum number of characters that must exist in a password before it is considered valid. The default value is eight characters.

- **Password History Size** (krbPwdHistoryLength): The number of previous passwords that IPA stores, and which a user is prevented from using. For example, if you set this value to 10, IPA prevents a user from reusing any of their previous 10 passwords. The default value is zero (0) (disable password history).

Note

If password history checking is enabled, and a user attempts to use one of the passwords in the history list, the error message returned by the system may be misleading. For example, you may see the following error:

```
A database error occurred: Constraint violation: Password fails
 to meet minimum strength criteria
```

This is because *python-ldap* prevents the retrieval of extended information on password policy failures over LDAP.

Note

Even with krbPwdHistoryLength set to zero, users cannot reuse their existing password.

3.2.2.5. Notifying Users of Password Expiration

Future versions of IPA will support the concept of automatic user notification when passwords are due to expire. This feature is not available in freeIPA 1.2.1. You can, however, manually search for passwords that are due to expire by a specified date.

For example, to retrieve all user entries whose password is due to expire before March 1st, 2008, run the following command:

```
$ ldapsearch -Y GSSAPI -b "cn=users,cn=accounts,dc=example,dc=com"
'(krbPasswordExpiration<=20080301000000Z)'
```

3.2.2.6. Using Password Authentication

If you use password authentication (no GSSAPI authentication, no ticket on the client) with a new user or a user whose password has expired, you need to enable Challenge-Response authentication. Otherwise, the password changing dialog will not display.

This is not enabled by default because some older SSL clients may not support Challenge-Response authentication, and it is needed only if the password has expired.

To enable Challenge-Response authentication:

- Set ChallengeResponseAuthentication to **`yes`** in the **`/etc/ssh/sshd_config`** file.

3.2.2.7. Using Local Logins

The default settings specified by the IPA installation script include timeout settings that still allow local logins to succeed if the client cannot access the IPA server. These settings are specified in the **`/etc/ldap.conf`** file, and can be tuned to suit your particular deployment. A typical deployment would normally include two or more servers for redundancy, and so this would not normally be a problem.

Warning

These timeout settings are only set on operating systems that support the IPA installation script. Currently this only includes Red Hat Enterprise Linux 4 and 5. On other operating systems, you need to specify these values manually. Failure to do so can result in the inability to log in to the machine if no IPA servers are available.

3.2.3. Specifying User Settings

3.2.3.1. Using the Web Interface

You can specify a range of attributes that are automatically applied to each new user account that you create. Any changes that you make to the default settings only apply to newly created accounts; existing accounts are not affected.

Procedure 3.4. Using the web interface

1. Click **Manage Policy** in the **Tasks** list, and then click **IPA Policy** to display the **Manage IPA Policy** page.

2. Click **Edit Policy** to display the **Edit IPA Policy** page, where you can edit different aspects of the IPA policy.

3. In the **User Settings** section, edit the user setting attributes as required, and then click **Update Policy**.

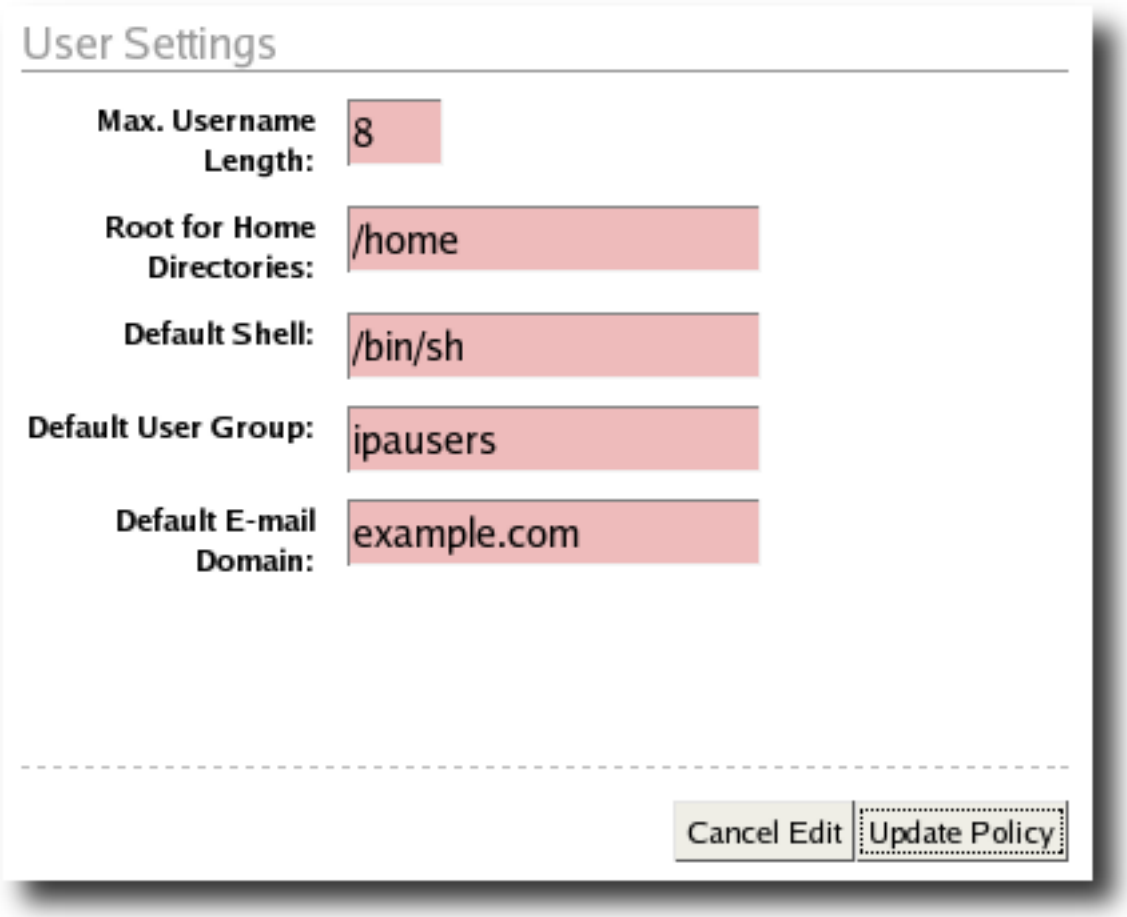

Figure 3.2. Editing the default user settings for the IPA policy.

3.2.3.2. User Setting Attributes

IPA supports the following User Setting attributes:

- **Max. Username Length** (ipaMaxUsernameLength): The maximum length of any username. The default value is 8.
- **Root for Home Directories** (ipaHomesRootDir): The root directory for all home directories. The default value is **`/home`**
- **Default Shell** (ipaDefaultLoginShell): The default shell for all user accounts. The default value is **`/bin/sh`**
- **Default User Group** (ipaDefaultPrimaryGroup): The default group to which all newly created accounts are added. The default value is `ipausers`, which is automatically created during IPA server installation process.

- **Default E-mail Domain** (ipaDefaultEmailDomain): The default domain used to create email addresses for all newly created accounts. The default is the domain to which the IPA server belongs.

Note

The default root directory for all home directories is **/home**, but it is the responsibility of the system administrator to ensure that whatever value is specified for this attribute is actually available.

Red Hat Enterprise Linux and most other Linux distributions include a PAM module called pam_mkhomedir that can automatically create a home directory if one does not exist for the user authenticating against the system. IPA does not force the use of this module because it may try to create home directories even when the shared storage is simply not available. It is the responsibility of the system administrator to activate this module on the clients if needed.

Configuring Applications to use Kerberos with IPA

IPA supports more than just user authentication. You can also configure different applications to use Kerberos to authenticate against your IPA installation. The following sections describe how to configure different applications to authenticate against your Kerberos server.

4.1. Configuring Apache for Kerberos Authentication

Refer to *Kerberos Module for Apache*[1] for information on how to configure Apache for Kerberos authentication against your IPA installation.

After you have configured Apache, you can test the configuration by attempting to retrieve an appropriate keytab, as follows:

```
# kinit admin

# ipa-addservice HTTP/web.example.com

# ipa-getkeytab -s ipa.example.com -p HTTP/web.example.com -k /etc/httpd/
httpd.keytab
```

[1] http://modauthkerb.sourceforge.net/index.html

Customizing Your IPA Deployment

5.1. Extending the Directory Schema

You can add site-specific fields to the IPA schema to provide for the specific requirements of your organization. This requires updates to the ipacustomfields attribute, which is overwritten every time it is modified. That is, if you add two elements to this attribute, and later need to add further elements, you need to include the original two elements in the new definition.

The ipacustomfields attribute is part of *cn=ipaConfig,cn=etc,$SUFFIX* in Directory Server. This attribute delimits each element with a dollar symbol ($), and each element consists of three, comma-delimited parts:

- label: The label displayed to the user
- field: The attribute name
- required: Whether or not the attribute requires a value (true or false)

Each element is displayed as a separate field in the web interface in the Custom Fields section.

The following is an example ipacustomfields attribute containing two new elements:

```
"See Also,seealso,false$Country,c,false"
```

Refer to the *Extending the Directory Schema*[1] section of the *Directory Server Administration Guide* for more information on how to extend and customize the schema.

5.2. Modifying the IPA Directory Information Tree (DIT)

You can add new branches to the IPA DIT at any time. Removing the default branches is not supported.

When you create new users using the web interface, they are always added to the *cn=users,cn=accounts,$SUFFIX* container. If you want to create your own hierarchical organization, you need to use the command-line, or some other tool that supports direct manipulation of the DIT.

[1] http://www.redhat.com/docs/manuals/dir-server/ag/8.0/Extending_the_Directory_Schema.html

Backup and Recovery

6.1. Backing Up Your IPA Deployment

Your IPA deployment represents critical infrastructure and so full backups should always be available. In particular, you should make a backup of the self-signed CA that is created during the initial installation of the IPA server. This CA issues all of the certificates for the server and for any replicas that are created. Without this CA, recovering from a critical failure would be extremely difficult.

It is the responsibility of the site to develop and implement a backup strategy to suit the needs of the deployment. This includes backing up user information and other data.

6.2. Recovering From a Failure

In the event of a hardware or similar failure, you should perform a restoration from your backup library using your own policies and procedures. The extent and success of your system recovery depends on the effectiveness of your backup procedures and their implementation. IPA does not provide tools to facilitate backup and recovery operations.

Troubleshooting

This chapter provides a range of troubleshooting tips and ideas related to different aspects of your IPA deployment. These are based on a simple IPA deployment, using standard operating system configurations as much as possible. If you have made extensive changes to your deployment, or if you have customized your operating systems, more extensive troubleshooting may be required if your IPA installation does not function as expected.

7.1. Kerberos Problems

7.1.1. Basic Kerberos Testing

To retrieve your `Kerberos` ticket and ensure that `Kerberos` is working, run the following command, and enter the Administrator's password when prompted:

`$ kinit admin`

You can use the **`klist`** command to verify that you received an appropriate ticket. The following is an example session of this process:

```
$ kinit admin
Password for admin@EXAMPLE.COM:
$ klist
Ticket cache: FILE:/tmp/krb5cc_0
Default principal: admin@EXAMPLE.COM

Valid starting     Expires            Service principal
06/03/08 22:57:17  06/04/08 22:57:13  krbtgt/EXAMPLE.COM@EXAMPLE.COM

Kerberos 4 ticket cache: /tmp/tkt0
klist: You have no tickets cached
```

Now run the following command:

`$ ldapsearch -Y GSSAPI -b "dc=example,dc=com" uid=admin`

If you receive an error message similar to the following, ensure that you have installed Fedora Directory Server version 1.1 or later.

```
SASL/GSSAPI authentication started
ldap_sasl_interactive_bind_s: Invalid credentials (49)
additional info: SASL(-1): generic failure: GSSAPI Error: Unspecified
GSS failure.  Minor code may provide more information (No such file or
directory)
```

7.1.2. Changing Kerberos Password Problems

Some versions of the *krb5** package set, including those currently shipped in Red Hat Enterprise Linux 5, reset the secret for the kadmin/changepw@REALM principal and store the keytab in the **`/var/kerberos/krb5kdc/kadm5.keytab`** file if you restart the `kadmin` daemon.

This causes the ipa_kpasswd service to lose access to the credentials cache and also causes password changes using either **kinit** or **kpasswd** to fail. This is because ipa_kpasswd reads the **/var/kerberos/krb5kdc/kpasswd.keytab** file when it searches for credentials.

Should this problem occur, you need to copy the new **kadm5.keytab** file over the original **kpasswd.keytab** file. You need to copy this new keytab file to all replicas, because they all share the same keytab.

For example, on the IPA master:

```
# cp /var/kerberos/krb5kdc/kadm5.keytab /var/kerberos/krb5kdc/
kpasswd.keytab
```

7.2. SSH Connection Problems

7.2.1. System Appears to Hang

You may find that the system appears to hang if you attempt to use SSH to connect to a host for which you do not have a Kerberos ticket.

This can occur when the host where the KDC is running has failed. After three attempts to contact the KDC, SSH falls back to other authentication methods. This can take several minutes.

To resolve this issue, restart the host where the KDC is running, and ensure that the KDC restarts as well.

If the KDC has failed, but the actual host is still running, the timeout is relatively quick. In this case, you only need to restart the KDC.

7.2.2. New User Cannot Log in Using SSH

Before a new user can log in to IPA, they need to change their initial password. To do this using SSH, you need to enable Challenge-Response authentication so that the password changing dialog is displayed to the user.

Procedure 7.1. To enable Challenge-Response authentication:

- Set ChallengeResponseAuthentication to **yes** in the **/etc/ssh/sshd_config** file.

7.3. Problems Using the IPA Tools

If Kerberos is working properly, but the **ipa-*** tools fail, ensure that your **/etc/hosts** file is correctly configured. Refer to the *Configuring the /etc/hosts File* section in the IPA Installation and Deployment Guide for more information.

If the **ipa-*** tools still fail, enable debug output in **Apache**, as follows:

1. Edit the **/etc/httpd/conf/httpd.conf** file, and set *LogLevel* to **debug**.
2. Restart the httpd service (# **service httpd restart**)

   This provides more verbose output that might help to identify the problem.

7.4. Firefox Problems

If you have trouble connecting to the IPA server using your browser, ensure that you have followed all of the instructions in the Configuring Your Browser chapter of the *IPA Client Configuration Guide*.

If you continue to have issues, refer to the following resources:

- *Common Error Messages*[1]

The following sections discuss issues specifically related to using Firefox with IPA.

7.4.1. Negotiate Authentication Problems

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 7.2. 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 enables verbose logging, and all information is 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 a `Kerberos` ticket, and need to run **kinit**.

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.

7.4.2. Certificate Authority Problems

Invalid Security Certificate Error

The first time you try to connect to the IPA web interface, **Firefox** will not be able to verify the server certificate, because it uses a self-signed Certificate Authority (CA). Instead, you will receive a Secure Connection Failed message:

```
ipaserver.example.com uses an invalid security certificate.
The certificate is not trusted because it was issued by an invalid CA
 certificate.
(Error code: sec_error_ca_cert_invalid)
```

To overcome this problem, you need to add an exception for the IPA server.

Procedure 7.3. To configure Firefox 3 to accept the IPA server certificate:

1. Click the **Or you can add an exception** link in the Connection Failed dialog in the browser.
2. Click **Add Exception**.
3. In the **Add Security Exception** dialog, click **Get Certificate**.
4. Click **Confirm Security Exception**.

You should now be able to connect to the IPA web interface.

Figure 7.1. Adding a security exception for the IPA server.

Reused Issuer and Serial Number Error

If you are running **Firefox** version 3 and you have uninstalled and reinstalled IPA, you may find that you cannot connect to the IPA web interface, and that **Firefox** returns a sec_error_reused_issuer_and_serial error. This can occur if you load the CA certificate from the previous installation.

To circumvent this problem, remove the CA from the list of known CAs in **Firefox**, and then restart **Firefox**. If you receive a sec_error_untrusted_issuer error when you navigate to the IPA homepage, add an exception for this CA to add it to the list of known CAs.

7.5. Service Principal Problems

If you are trying to install a service principal and keytab on a client, you might see the following error after issuing the **`ipa-getkeytab`** command:

```
SASL/GSSAPI authentication started
SASL Bind failed!
```

A SASL bind failure indicates either that you do not have a `Kerberos` ticket or that you have an invalid ticket (it has expired or is for the wrong user). Ensure that your ticket is valid and retry the operation.

If this does not solve the problem, check the **`/var/log/krb5kdc.log`** file on the IPA server for entries that contain the following (or similar) messages:

```
admin@EXAMPLE.COM for ldap/192.168.1.1@EXAMPLE.COM, Server not found in
 Kerberos database
```

This indicates a DNS reverse look-up failure. Service principals must be constructed from host names, and the existence of an IP address indicates that the system was unable to resolve the IP address to a host name.

To address this problem, ensure that your reverse look-up zone on the DNS is correctly configured, and retry the operation.

7.6. Other Possible Errors

- The IPA command line tools might report the following error:

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

  This can occur if you have multiple entries for the same host created by different KDCs.

- The IPA command line tools might report "Connection refused".

  This can occur when the **Apache** server is not accepting connections on the SSL port.

  - Verify that the **/etc/httpd/conf.d/nss.conf** file exists. If this file does not exist, ensure that you have installed the *mod_nss* rpm package.
  - Verify that all references to 8443 have been changed to 443 in the **/etc/httpd/conf.d/nss.conf** file.
  - Restart `httpd` after you have installed *mod_nss* and edited the **/etc/httpd/conf.d/nss.conf** file.

- The Web browser might return a Service Temporarily Unavailable message, when all indications are that the service is functioning normally. You may also see Connection Refused entries in the **/var/log/httpd/nss_error_log** file.

  This can occur if you have an incorrectly configured **/etc/hosts** file. Refer to the *Configuring the /etc/hosts File* section in the IPA Installation and Deployment Guide for more information.

7.7. Performing a Re-Install

If the installation fails, or if you want to run the installation script again, you first need to remove the existing installation. Run the server installation command using the `--uninstall` switch, as follows:

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

You also need to remove the `Kerberos` keytab before you begin the reinstallation process:

```
# rm -f /var/kerberos/krb5kdc/kpasswd.keytab
```

7.8. DNS and Service Discovery Problems

As discussed in *How IPA and DNS Work Together* in the IPA Administration Reference, to take full advantage of its Service Discovery capabilities, IPA relies on a functional DNS. If the DNS fails or is not fully functional, Service Discovery may also fail to operate correctly. This can lead to problems in a high-availability deployment.

The following sections detail some common mistakes that could prevent the DNS from functioning correctly.

7.8.1. Zone Files

- Ensure that you increment the serial number when editing a zone file.

  If you do not increment the serial number, slave nameservers will never be notified of the changes that have occurred on the master nameserver. Consequently, they will not attempt to refresh their data for that zone.

  After editing a zone file, either reload the file (**`# service named reload`**) or restart the named service (**`# service named restart`**).

- Ensure that you use ellipses and semi-colons correctly in the **`/etc/named.conf`** file. An omitted semi-colon or unclosed ellipsis can prevent named from starting.

- Remember to place periods (.) in zone files after all FQDNs and omit them on hostnames.

  A period at the end of a domain name denotes a fully qualified domain name. If the period is omitted, then named appends the name of the zone or the `$ORIGIN` value, resulting in an invalid value.

7.9. Firewall Problems

- If a firewall is blocking connections from the named program to other nameservers, edit the firewall's configuration file to allow the connection.

  By default, **BIND** version 9 uses random ports above 1024 to query other nameservers. Some firewalls, however, expect all nameservers to communicate using only port 53. To force named to use port 53, add the following line to the options statement of the **`/etc/named.conf`** file:

  ```
  query-source address * port 53;
  ```

  Refer to the *Required Ports* section in the IPA Installation and Deployment Guide for more information.

7.10. IPA Server Boot Problems

If your newly-installed IPA server hangs during system initialization, at the point where "Starting system message bus" is displayed, it indicates a networking problem. The most likely cause is not changing from using `NetworkManager` to static networking.

Procedure 7.4. To remedy this problem:

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

   ```
   # chkconfig NetworkManager off ; chkconfig network on
   ```

2. Ensure that static networking is correctly configured.

3. Restart the system.

Appendix A. Revision History

Revision 1.1 6 Jan, 2009 David O'Brien *davido@redhat.com*

BZ 475295. Document limit to number of groups for Solaris and AIX.
Update troubleshooting section for SSH connectivity.
BZ 473155. Document need to rename keytab file if kadmin is started.
BZ 453782. Clarify section on deleting user accounts.
BZ 469599. Updates from tech review.
BZ 471511. Update Password History information.
BZ 469981. Added backup & recovery information.
BZ 470420. Troubleshooting web browser connectivity problems.
Added Configuring Apache for Kerberos Authentication.
Modified section on Creating User Accounts to account for UNIX username restrictions.
Added definition for supported usernames in IPA.
BZ 469790. Updates to troubleshooting.

Revision 1.0 3 June, 2008 David O'Brien *davido@redhat.com*

Created.